Capacity Utilization in Production Networks

Kian Abbas Nejad*

12th November 2026

[MOST RECENT VERSION]

Abstract

This paper quantifies how capacity constraints shape the propagation of sectoral shocks through the production network and their long-run cost. I build a dynamic production network model in which output is capped by a capacity proportional to capital, solve it as a graphon mean-field game and calibrate it to 63 US industries. Capacity constraints lower GDP per unit of capital by about 10 per cent, mostly through wedges paid along every production chain, and dampen the propagation of sectoral shocks to about three quarters of the unconstrained response, most for upstream and tight sectors.

Keywords: capacity utilization, capacity constraints, production networks, propagation of sectoral shocks, rationing, graphon mean-field games

JEL Codes: E22, E23, E32, C67, L16

*Tilburg University. Email: k.abbasnejad@tilburguniversity.edu

1 Introduction

A production plant's capacity is the maximum quantity its machines could sustainably produce if variable inputs, such as material and labour, were to be freely available. Aggregate data show that on average a significant proportion of plants operate at their capacity constraints[1]. However, a plant's output is bounded not only by its own capacity but also by that of its suppliers: in a normal year, about 6% of US manufacturing plants that operate *below* capacity, and as many as 43% of them in late-2021, name a shortage of materials as a reason for underutilization. This suggests that while imposing a substantial bottleneck on production at the plant level, capacity constraints can block or dampen the propagation of shocks through the input-output linkages. In this paper I model and quantify how capacity constraints shape the propagation of sectoral shocks through the production network.

I build a dynamic production network model with capacity constraints. Sectors buy intermediate inputs from one another through the input–output network, and within each sector many firms produce close substitutes from their own capital and a bundle of the other sectors' goods, each facing an idiosyncratic demand shock. Firms choose their capital in advance, and capital sets capacity: a firm's output is capped by a capacity proportional to its capital stock. Prices are posted before the demand shock is known, so a firm whose demand exceeds its capacity sells its capacity at the posted price, and the excess demand is redirected, at a scarcity premium, to the firms of the same sector that have capacity to spare. Firms post their prices knowing that they may be rationed, and so charge a higher markup the more often they expect to reach their cap. One statistic per sector, the demand draw at which capacity starts to bind, thus determines the sector's utilization rate, its markup, the share of its firms that are rationed and its expenditure shares, and it makes the sector's supply curve state-dependent. Capital is accumulated over time under installation costs and sectoral productivity shocks, so capacity itself, and not only utilization, responds to the state of the network. The standard production network model is nested as the case in which the cap never binds, and every result is read against it: the calibrated economy with its capacity constraint switched off and everything else unchanged, which I call the economy without capacity constraints.

Solving such an economy with many sectors is the technical contribution of the paper. In a production network with sector-specific capital, every sector's investment depends on the whole cross-section of capital and productivity, because each sector's shock moves aggregate prices and the household's discount factor, so the state of the problem has as many dimensions as there are sectors. I take the limit in which each sector consists of a continuum of industries that are small in the economy and connected to the other sectors through a graphon, the continuum analogue of the input–output

[1] In the US about 16% of total capacity is located in firms operating above 95% capacity utilization rate (Census Bureau's Survey of Plant Capacity, 1976–2016). In Sweden about 15% of manufacturing firms report running at full capacity on average, fewer than 7% in the Great Recession and more than 25% in the recovery after the Covid-19 pandemic (Amberg, Friberg, Liu & Syverson, 2026)

matrix. In this limit cross-sectional aggregates are deterministic, each industry solves a two-dimensional problem in its own capital and productivity coupled to the rest of the economy through a few moments, and the equilibrium is a graphon mean-field game in the sense of Caines and Huang (2021), extended to a payoff that runs through general-equilibrium prices. The limit removes aggregate risk; in exchange it delivers an economy with a full cross-section of heterogeneous firms in every sector, whose balanced growth path is solved in minutes and whose transitions are solved in about an hour. This is what allows me to calibrate the model to the US economy at the level of the 63 industries of its input–output accounts, with the capacity margin disciplined by the Census share of plants close to their capacity and by the convexity of industry supply curves in utilization.

The first result is that capacity constraints are costly in the long run. On the balanced growth path of the calibrated economy about one firm in five is rationed at any time and the scarcity rents on rationed goods amount to 0.9 per cent of GDP, yet GDP per unit of capital is about 10 per cent lower than in the economy without capacity constraints, consumption about 13 per cent lower, and the economy grows about 0.6 percentage points a year more slowly, because rationing and the markups it induces lower the return that capital earns for its owners. Most of the gap in GDP, 8 of its 11 log points, is the cost of the wedges that capacity puts on the links of the network at a given allocation of capital: rationing, state-dependent markups and idle capacity. The remaining 2.5 log points are the reallocation of capital that the constraint induces, and this reallocation is a result in its own right. About 7 per cent of the economy's capital sits in different sectors with capacity constraints than without, and it moves not towards the sectors in which it earns the highest return but towards the sectors that need it to produce their sales, service sectors with small capital shares and large sales. A network model without capacity constraints fitted to the same accounts cannot see this: it attributes the rents that capacity earns to the technology, and in doing so assigns the tight, capital-light sectors capital exponents up to almost three times the calibrated economy's.

The second result is that capacity constraints dampen the propagation of sectoral shocks, and that the dampening happens in the network rather than at the shocked sector. A productivity shock to one sector moves GDP by 0.76 of its response in the economy without capacity constraints on average over the 63 sectors, and by 0.71 when sectors are weighted by their size. The shocked sector's own output responds almost as it would without the constraint; what capacity reduces is the pass-through of the shock to the sector's customers and suppliers, whose output responds at less than half of its unconstrained value. The dampening is strongest for shocks to upstream sectors, whose output passes through many production stages before it reaches final use, and to tight sectors; shocks to the downstream sectors that sell to households directly pass through almost undamped. A network model without capacity constraints fitted to the same accounts overstates the response of GDP to a sectoral shock by about a sixth and the response of the other sectors by about a factor of two. Capacity also makes the response of the economy depend on the sign and the size of a shock, which it cannot

in the standard model: the economy responds more to a negative common shock than to a positive one of the same size, and the asymmetry grows with the size of the shock, because a positive shock pushes the tight sectors against their cap while a negative one frees capacity. Over the transition that follows a sectoral shock, utilization absorbs the shock first and capital follows. The dampening recedes only slowly as capital adjusts, from about 0.7 of the unconstrained response at impact to 0.8 along the transition, and the same shock reallocates up to 1.3 times more capital across sectors with capacity constraints than without.

The third result is where capacity is worth most. Raising one sector's capital, and with it its capacity, by one per cent at a time and recording the change in GDP measures the value of the sector's capacity to the economy against its value to the sector's owners. The two differ because a sector's capacity also spares its customers, and through them their customers, the rationing they would otherwise suffer, and this part of its value does not enter the sector's own investment decision. In every sector a unit of capital is worth more to the economy than to its owners, and more so the further upstream the sector. This holds in the economy without capacity constraints as well, because markups compound along the chain, so it is a property of networks with market power rather than of capacity. What capacity changes is the ranking of the sectors. The value of a sector's capital to the economy with capacity constraints over its value without runs from 0.7 to 5.4 across sectors: it is high in tight sectors that use little capital per unit of sales, such as computer systems design, trucking and administrative services, and low in capital-intensive sectors such as utilities, pipelines, oil and gas and real estate. A unit of capital is an input, worth its share of the sector's cost, and it is also capacity, worth the rationing it relieves, and the second part does not scale with the first. Putting a given amount of extra capital into the five sectors where it is worth most raises GDP by about 1.7 per cent, against 1 per cent for a uniform expansion; most of that gain is the markup wedge, which the economy without capacity constraints shares, and what is specific to capacity is which sectors to target.

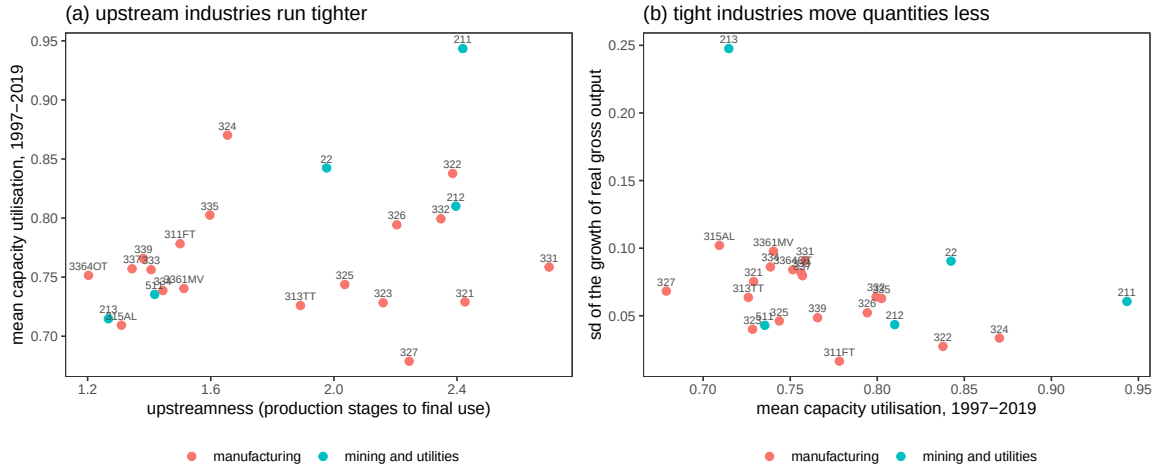
The mechanism behind the dampening is a state-dependent supply curve. Without capacity constraints, a sector that receives cheaper inputs or more demand expands along a supply curve whose elasticity is the inverse of its capital share, which is very flat at the calibrated shares, so almost all of a shock that reaches a sector becomes output. With the constraint, the firms of the sector that are at their cap cannot expand, and the demand they turn away is rationed rather than met at a higher price. The sector's supply elasticity is then that of its firms with spare capacity, weighted by their share of its sales: about a quarter of its unconstrained value at the median sector, and lower the tighter the sector. A shock therefore travels as it would in the standard model until it reaches a tight sector, where part of the quantity response is replaced by rationing while the price response is unchanged, because the rents on rationed units are not recorded as sales. This is also why the dampening is a property of the network and not of the shocked sector. A sector's own productivity raises its capacity together with its demand and does not tighten it; the sectors that receive the shock meet it with the capacity they have, and the benefit of a shock to an upstream sector reaches

households only through the sectors it supplies, whose capacity limits how much of it they can pass on.

The long-run cost is large relative to the rents for the same reason: the wedges sit on the links of a network and are paid along every chain. Gross output is about four times GDP, so every euro of final output is produced through about four euros of sales, three of them between sectors, each crossing a link on which a wedge sits, and a wedge of a few per cent per link is paid several times over on the way to every unit of final output. Rationing also shortens the chains, because the part of each purchase that is scarcity rent buys no inputs, so gross output falls by more than GDP. The allocation of capital, finally, follows from the two roles that capital plays. Without the constraint a sector's capital is determined by its Cobb–Douglas return, so sectors with small capital shares hold little; with it, a sector's capital also bounds what the sector can produce, so a sector with a small capital share and large sales must hold enough capital to serve those sales. Capital moves out of mining and utilities and into the service sectors, and a model without the constraint reads the rent that this capital earns as a larger capital share.

Motivating Figure

Figure 1: Capacity utilization and network position in the data.



Notes: The 24 industries of mining, utilities and manufacturing with a Federal Reserve utilization series, 1997–2019. (a) Mean utilization against upstreamness, the average number of production stages between the industry's output and final use, from the input–output accounts. (b) The standard deviation of the growth of real gross output against mean utilization. Labels are BEA industry codes.

Two patterns in the data motivate the exercise, and the model turns out to produce both (Figure 1). Across the 24 industries of mining, utilities and manufacturing for which the Federal Reserve publishes a utilization rate, the industries that sit further upstream in the input–output network run tighter: the correlation of an industry's mean utilization over 1997–2019 with its upstreamness is 0.3. And the industries that run tighter move their quantities less: the correlation of mean utilization with the

volatility of real gross-output growth is -0.3 (rank correlation -0.4).[2] In the model both patterns have one source. A sector at its cap meets demand with rationing rather than output, so its quantities move less than a slack sector's, and the sectors whose output passes through the most stages before reaching final use are the ones whose cap is reached first when the economy expands. The paper's main result, that capacity dampens the propagation of shocks and does so most for shocks to upstream and tight sectors, is the same statement read along the network.

[Q1 | here | tag: Later] First a note: this motivating figure was made a while ago and its accompanying text in the introduction was attached to a much older version of the introduction so its prose style or even some facts in it might be obsolete.

I have a few questions about the two motivating figures:

1. How do these two motivating figures relate to (the final version) of the "A Regression the Model Implies" results? Is what the regression(s) show already cover this figure (thus making the two figures unneccessary)? if not, i.e. if the two motivating figures suggest a relationship that the regression(s) haven't tested, can we indeed extend the regresssion(s) to cover this too? If yes, now? If no, why?
2. In panel (a) sector 211 seems like an outlier and in panel (b) sector 213 seems like an outlier. How do these motivating figures change if these two are excluded?
3. In case (and only in that case) we decide to import some high-digit (but fewer years, maybe just one year) data for the "A Regression the Model Implies" stuff, can we use that data for this too?

Related literature. The paper builds on four strands of the literature. The first is the theory of production networks and the propagation of sectoral shocks. From Long and Plosser (1983) and Acemoglu, Carvalho, Ozdaglar and Tahbaz-Salehi (2012) to the nonlinear economies of Baqaee and Farhi (2019), the distorted economies of Bigio and La'O (2020) and the dynamic networks with capital of Vom Lehn and Winberry (2022) and Liu and Tsyvinski (2024), sectors supply along smooth cost curves: a sector produces whatever its buyers demand at a cost that rises with output at a constant elasticity, so propagation does not depend on how much slack the sectors have, and capacity, a stock, never responds to the state of the network. The evidence that shocks travel through input linkages, and travel furthest where inputs are hard to replace, is in Acemoglu, Akcigit and Kerr (2016), Boehm, Flaaen and Pandalai-Nayar (2019) and Carvalho, Nirei, Saito and Tahbaz-Salehi (2021); Carvalho and Tahbaz-Salehi (2019) survey the field. This paper adds to these models a supply margin that is state-

[2] The sample is small and the correlations are what they are, a tendency rather than a regularity of the data. The price side, that tight industries move their prices more, holds in the raw data (correlation 0.7) but rests on the two oil industries, extraction and refining, and is 0.25 without them. The utilization series are the Federal Reserve's G.17 rates; output and prices are the BEA's gross-output quantity and price indexes by industry; upstreamness is computed from the input–output accounts as the average number of production stages between an industry's output and final use.

dependent, nonlinear and directional, and whose cap is chosen by firms that know they may reach it, and it shows what the margin does to propagation, to the allocation of capital and to the value of capacity by sector in an economy calibrated at the scale of the accounts.

The second is the macroeconomics of capacity utilization and convex supply. Capacity that is fixed before demand is known, with prices posted in advance and excess demand rationed, goes back to Prescott (1975) and Eden (1990); Fagnart, Licandro and Portier (1999) put firm-level demand uncertainty and the capacity choice it induces into a one-sector business-cycle model, and the paper's derivation of a firm's capacity from its capital follows theirs. The recent evidence is that supply curves are convex in utilization at the industry and the plant level (Boehm & Pandalai-Nayar, 2022; D. A. Comin, Johnson & Jones, 2023; Bohr, 2024) and that firms choose how much slack to hold (Amberg et al., 2026); the paper's firm-level block follows Bohr (2024) in the way demand in excess of capacity is rationed and redirected. These models have one sector, or sectors whose capacity is given, so the questions this paper asks, how the margin shapes propagation through a network and how capacity is allocated across sectors over time, do not arise in them.

The third is the literature on bottlenecks and industrial policy in networks. Liu (2019) shows that a wedge between price and marginal cost in one sector is paid along every chain that passes through it, so that relieving a distortion is worth most in the sectors through which many chains pass; Baqaee and Farhi (2020) measure the aggregate cost of such wedges, and Acemoglu, Autor and Patterson (2023) and di Giovanni, Kalemli-Özcan, Silva and Yildirim (2022) read the US productivity slowdown and the inflation of 2021–22 as the work of sectoral bottlenecks. In these papers the wedge is given or is a policy; here it is the equilibrium outcome of capacity choices, it depends on the state of the economy, and relieving it means building capacity, which the paper prices sector by sector against what the capacity is worth to its owners.

The fourth is technical: mean-field games and their graphon extension. Mean-field games, from Lasry and Lions (2007) and Huang, Malhamé and Caines (2006), describe a continuum of small agents coupled through the distribution of their states; in macroeconomics they are the continuous-time heterogeneous-agent economies of Achdou, Han, Lasry, Lions and Moll (2022), which descend from Aiyagari (1994). Graphon mean-field games, introduced by Caines and Huang (2021) and developed by Caines, Ho, Huang, Jian and Song (2022) and Bayraktar, Chakraborty and Wu (2023), put the agents on a network whose limit is a graphon (Lovász, 2012); static graphon games are the subject of Parise and Ozdaglar (2023). The paper brings the graphon formulation to a production network in which the coupling runs through general-equilibrium prices rather than through an additively separable cost, and it uses the exact law of large numbers of Sun (2006) to make its cross-sectional aggregates deterministic. In macroeconomic terms it solves a heterogeneous-agent economy whose agents are the industries of the input–output table.

The paper proceeds as follows. Section 2 states the model with finitely many sectors and characterises its static and dynamic equilibrium. Section 3 takes the con-

tinuum limit and defines the graphon mean-field game that the paper solves. Section 4 describes the data and the calibration. Sections 5 to 7 report the results: the long-run cost of capacity, the propagation of shocks, and the value of capacity by sector. Section 8 concludes. The appendices collect the proofs for the two model sections, the model without capacity constraints, the numerical method (Appendix D), every source and step of the calibration (Appendix E) and the results by node (Appendix F).

2 A Dynamic Network Model with Capacity Constraints

[Q2 | text | tag: Later0] I have a request whose main goal is to shorten the appendix section **without** removing information. This request is probably rather large in scope *but* I believe it does not require super careful thought (it's mainly procedurial...)

Now, what I want to do is to replace the equations in the appendix sections (mainly the lemmas/theorems in app:fs-details and app:g-details, but maybe also some others...) by replacing equations that are exactly (or even almost exactly) the same as the equations that already appear in the main text by references to them. For example if the appendix (e.g. a lemma in it) has a sentence like "the demand function is

$$y_{in} = v_{in} Y_i \left(\frac{p_{in} + \varphi_{in}}{P_i} \right)^{-\varepsilon}, \tag{1}$$

"; and since (3) already appears in the main text, that part of the appendix should be replaced with a shorter "the demand function is (3)" (no need to even write "the demand function is (3) in the main text" – the reader knows where it is).

Small note: since this is a large, but not very difficult task, maybe you can assign agent(s) to do it but since it's not difficult and I'd rather save on tokens, maybe you can use a lower model (e.g. Opus medium or high effort) for it. But you know best how to do this!

[Q3 | text | tag: Later1] Check the entire paper for equations that are numbered (i.e. not using the "*" in Latex) and not cited anywhere else and make them unnumbered.

Small note: since this is a large, but not very difficult task, maybe you can assign agent(s) to do it but since it's not difficult and I'd rather save on tokens, maybe you can use a lower model (e.g. Opus medium or high effort) for it. But you know best how to do this!

In this section I develop a dynamic general-equilibrium model of a production network with sector-specific capital in which every firm's output is bounded from above by a capacity proportional to its capital stock. Each of J sectors is populated by monopolistically competitive firms that combine sector-specific capital with a Cobb–Douglas bundle of intermediate inputs and sectoral goods are aggregated into a final good used for both consumption and investment. Prices are posted before the firm-level demand shock is known, so demand in excess of capacity is rationed. The modelling of capacity is a version of Fagnart et al. (1999). Bohr (2024) and Amberg et al.

(2026) are recent models in the same tradition. Rationing turns the demand-shock distribution from a single moment into a family of tail functions and makes the effective demand elasticity, the markup, the expenditure shares and the rate of capacity utilization functions of one sector-level object: the demand threshold $\overline{v}_i$ at which capacity starts to bind. I derive the static and dynamic equilibrium and the homogeneity-based state reduction for finitely many sectors, and then carry the same model to a continuum of sectors (Sections 3–3.5).

2.1 Model Setup

Sectoral output The economy consists of J industries (indexed by i and j), each populated by a unit measure of monopolistically competitive firms $n \in [0,1]$ producing differentiated products, together with a unit measure of competitive aggregators that bundle the monopolists' output into a sectoral good through a CES aggregator. Sector-i output is

$$Y_i = \left[\int_0^1 v_{in}^{1/\varepsilon} y_{in}^{(\varepsilon-1)/\varepsilon} \mathrm{d}n\right]^{\varepsilon/(\varepsilon-1)}, \qquad \varepsilon > 1, \tag{2}$$

where v_{in} are firm-specific demand shocks drawn from distributions G_i with mean normalised to one[3]. Each monopolist's output is capped by a capacity $\overline{y}_{in}$, specified below as proportional to the firm's capital, and each monopolist sets price p_{in} *before* observing its demand shock v_{in}. As such, the aggregator takes prices and capacities as given and buys the cheapest bundle that delivers Y_i, minimising $\int_0^1 p_{in} y_{in} \mathrm{d}n$ subject to (2) and $y_{in} \leq \overline{y}_{in}$ for every n, as in Bohr (2024). However, because it cannot buy more than $\overline{y}_{in}$ from a firm whose good is cheap, the constraints *reallocate* demand towards the unconstrained firms who have excess capacity. The demand function is

$$y_{in} = v_{in} Y_i \left(\frac{p_{in} + \varphi_{in}}{P_i}\right)^{-\varepsilon}, \tag{3}$$

where p_{in} is the price of firm n's output and φ_{in} is the Lagrange multiplier associated with the capacity constraint. P_i is the shadow price that describes the aggregator's marginal cost for a unit of Y_i, which depends on the degree to which the quantity constraints bind as indicated by φ_{in}. It can be characterized by combining (2) and (3),

$$P_i = \left[\int_0^1 v_{in} (p_{in} + \varphi_{in})^{1-\varepsilon} \mathrm{d}n\right]^{1/(1-\varepsilon)}. \tag{4}$$

The demand function, the price index, the threshold representation below and the expected-output formula that follows are stated and proved as Lemma 1 in Appendix A.

The multiplier φ_{in} is the amount by which one more unit of firm n's capacity would lower the aggregator's cost of producing Y_i. By complementary slackness, it is zero

[3]Mild assumptions about the demand-shock distributions, in order to ensure that expected output and input use are finite and smooth in the firm's price, and that its pricing problem has a unique solution, are stated in Appendix A.1. These conditions are satisfied by the log-normal form used in the calibration of the model in Section 4.

for a firm with spare capacity and for a rationed firm (firm for whom $y_{in} = \overline{y}_{in}$ holds) it is the premium over the posted price that the aggregator would be willing to pay for an additional unit of the good. More precisely, for a rationed firm $p_{in} + \varphi_{in} = P_i(v_{in}Y_i/\overline{y}_{in})^{1/\varepsilon}$, which is the price at which the firm's CES demand would equal its capacity. This is the market-clearing price the firm would have posted had it known v_{in} in advance. The aggregator values the rationed variety at that price $p_{in} + \varphi_{in}$ yet pays the posted price p_{in}, and the difference is the rationing rent collected by the aggregators.

The index P_i is the CES price index at these prices and, by the envelope theorem, the aggregator's marginal cost of Y_i: an additional unit of the sectoral good has to be sourced from the firms that still have spare capacity, and it costs more than the posted prices alone would indicate. The competitive sectoral aggregators sell the sectoral good at this marginal cost, so P_i is the price paid by the producer of the aggregate final good (the consumption good below) and by the other sectors, and at P_i the sectoral aggregators' supply of the sectoral good equals the demand for it.[4] Because the CES aggregator is homogeneous of degree one, an application of Euler's theorem gives

$$P_i Y_i = \int_0^1 p_{in} y_{in} \mathrm{d}n + \int_0^1 \varphi_{in} y_{in} \mathrm{d}n.$$

So the sector's buyers pay the firms' revenue at posted prices plus the value of the multipliers. We denote the second term by *rationing rent* Π_i^{r}, the scarcity value of the rationed firms' capacity. This term arises because prices are set before v_{in} is known: a rationed firm sells its products at p_{in} although its buyers value it at $p_{in} + \varphi_{in}$. The difference is earned instead by the aggregators, who buy at the posted prices and sell at P_i. They have no other costs and are owned by the household, to which they pay the rent as a dividend (see Households, below). The rent is thus part of the sector's value added but not of its firms' revenue[5], a distinction that matters for the firms' investment incentives (Remark 1). As capacity ceases to bind the rent vanishes and P_i reduces to the CES index at posted prices (Lemma 4).

Using the complementary slackness condition, output can be written with a threshold representation

$$y_{in} = \overline{y}_{in} \min\left\{\frac{v_{in}}{\overline{v}_{in}}, 1\right\} \quad \text{with} \quad \overline{v}_{in} := \frac{\overline{y}_{in}}{Y_i}\left(\frac{p_{in}}{P_i}\right)^{\varepsilon}, \tag{5}$$

[4]Note that selling at average cost, $\int_0^1 p_{in} y_{in} \mathrm{d}n / Y_i < P_i$, would remove the rationing rent mentioned but not the scarcity behind it. At that price the buyers demand more of the sectoral good than aggregators facing the marginal cost P_i are willing to supply, and the excess demand would have to be rationed again downstream. Marginal-cost pricing keeps every market-clearing condition of the model without capacity constraints (derivations in Appendix C) in force, with P_i in place of the price index at posted prices.

[5]Competition among the aggregators does not pass the rent to the firms. Firms post p_{in} before v_{in} is drawn and sell at that price up to capacity, and they do not reprice or accept higher bids afterwards. So the rent goes to whoever buys the rationed units at p_{in}, which is the aggregators. Free entry of aggregators does not compete it away either since aggregation is costless, so entry only splits the same rent among more aggregators.

implying that firm n is demand-constrained when $v_{in} < \overline{v}_{in}$ and capacity-constrained when $v_{in} \geq \overline{v}_{in}$, which occurs with probability $1 - G_i(\overline{v}_{in})$. Using the representation (5) expected output of firm n (expectation taken with respect to the only unknown at time t, i.e. demand shock v_{int}) is given by

$$\mathbb{E}_v[y_{in}] = \int_0^{\overline{v}_{in}} v Y_i \left(\frac{p_{in}}{P_i}\right)^{-\varepsilon} \mathrm{d}G_i(v) + \int_{\overline{v}_{in}}^{\infty} \overline{y}_{in} \mathrm{d}G_i(v) \tag{6}$$

and so the effective elasticity of expected demand with respect to own price is $\varepsilon_{in} := -\partial \ln \mathbb{E}_v[y_{in}] / \partial \ln p_{in} = \varepsilon \eta_{i,1}(\overline{v}_{in}) \in (0, \varepsilon)$ where

$$\eta_{i,\theta}(\overline{v}) := \frac{\int_0^{\overline{v}} (v/\overline{v})^{\theta} \mathrm{d}G_i(v)}{\int_0^{\overline{v}} (v/\overline{v})^{\theta} \mathrm{d}G_i(v) + 1 - G_i(\overline{v})} \quad \text{for} \quad \theta \geq 0 \tag{7}$$

The effective elasticity is thus ε scaled by the share of expected sales that comes from the unconstrained lower tail: in the states in which the firm is capacity-constrained a higher price does not lower its sales, since demand exceeds capacity, so its demand elasticity is zero there. Two consequences of this run through the paper. First, G_i no longer enters through a single moment but through tail functions evaluated at the threshold. Second, every object the firm cares about before v_{in} is realised (expected output, expected input use, the elasticity it faces) depends on $(p_{in}, k_{in}, Y_i, P_i)$ only through the scale $\overline{y}_{in}$ and the threshold $\overline{v}_{in}$. As such the threshold is the one state variable of the pricing problem, and, at the symmetric point, of the sector.

[Q4 | notes | tag: ResultToModel] About "First, G_i no longer enters through a single moment but through tail functions evaluated at the threshold.": this is key! These tail functions are convex and drive some of the comparative static results of the model! Maybe worth mentioning this convexity (due to Gaussian tail) here **or** in the results section. As I am not sure exactly what should be written or where, please answer this in a **short** notes box below before I decide what to do in the final version.

Production and capacity When it is not capacity-constrained, monopolist n in sector i produces according to a Cobb-Douglas production function using own capital and inputs from other sectors

$$y_{in} = Z_i k_{in}^{\alpha_i} M_{in}^{1-\alpha_i} \qquad \alpha_i \in (0,1), \tag{8}$$

where Z_i is sector-specific and stochastic productivity and M_{in} is a Cobb-Douglas composite of the inputs of other sectors' outputs,

$$M_{in} = \prod_{j=1}^{J} m_{ijn}^{a_{ij}}, \qquad a_{ij} \geq 0, \quad \sum_{j=1}^{J} a_{ij} = 1. \tag{9}$$

However, the firm's production is capped at a maximum capacity $\overline{y}_{in}$. We assume that each unit of capital can process at most $1/x_i$ units of the intermediate bundle within the instant, where $x_i > 0$ is a sector-specific technological constant, as in Fagnart et

al. (1999): in sector i processing one unit of the bundle ties up x_i units of the firm's capital. The firm's capacity in bundle units is therefore $\overline{M}_{in} = k_{in}/x_i$, and its capacity in output units is

$$\overline{y}_{in} := Z_i k_{in}^{\alpha_i} \overline{M}_{in}^{1-\alpha_i} = Z_i \, k_{in} \, x_i^{\alpha_i - 1}. \tag{10}$$

Capacity is linear in the firm's capital and in sectoral productivity, and x_i is the capital intensity $k_{in}/\overline{M}_{in}$ at full utilization. Realised output is $y_{in} = \overline{y}_{in} \min\{(M_{in} x_i / k_{in})^{1-\alpha_i}, 1\}$: the firm produces along (8) until the bundle it needs exceeds what its machines can process. Figure 2 draws two isoquants of this technology: to the right of the capacity line $M_{in} = k_{in}/x_i$ they are the Cobb–Douglas isoquants of (8), and on the line they turn vertical, because once the bundle reaches k_{in}/x_i more inputs do not raise output. After v_{in} is realised the firm buys the bundle that just meets its, possibly rationed, demand. Inverting (8) and using (5), the conditional input bundle demand is

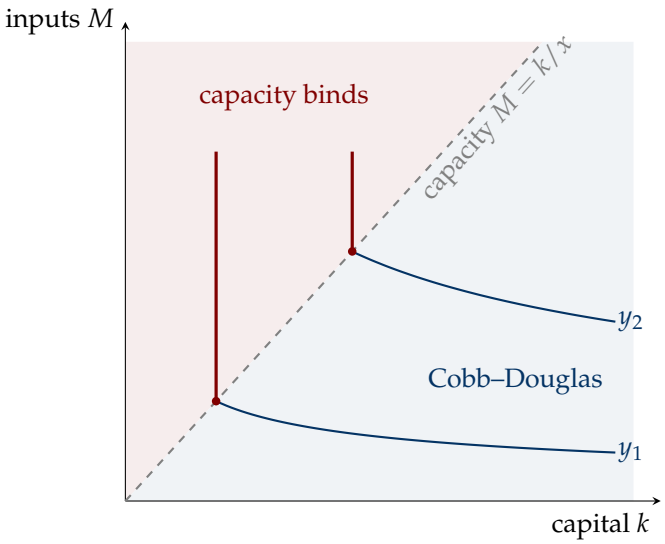
$$M_{in} = \left(\frac{y_{in}}{Z_i k_{in}^{\alpha_i}} \right)^{\frac{1}{1-\alpha_i}} = \frac{k_{in}}{x_i} \min\left\{ \left(\frac{v_{in}}{\overline{v}_{in}} \right)^{\frac{1}{1-\alpha_i}}, 1 \right\} \leq \frac{k_{in}}{x_i}, \tag{11}$$

and the firm's expected bundle use is

$$M_{in}^e := \mathbb{E}_v[M_{in}] = \frac{k_{in}}{x_i} \left[\int_0^{\overline{v}_{in}} \left(\frac{v}{\overline{v}_{in}} \right)^{\frac{1}{1-\alpha_i}} \mathrm{d}G_i(v) + 1 - G_i(\overline{v}_{in}) \right] \tag{12}$$

The upper tail of v is truncated at the threshold, so both expectations depend on the whole lower tail of G_i and on the constrained probability $1 - G_i(\overline{v}_{in})$. When capacity never binds they reduce to the forms in which G_i enters only through a single moment (Appendix C).

Figure 2: Isoquants of the technology with capacity.



Decision sequence. Time is continuous. Each firm owns and invests in its own capital stock k_{in} subject to adjustment costs (below). Within $[t, t + \mathrm{d}t]$:

1. the investment rate has been chosen before t, hence before any time-t shock is known;
2. after observing the sectoral productivity Z_{it} but not the idiosyncratic v_{int}, the firm sets its price p_{int};
3. once v_{int} is realised, production, input use, and the next investment rate are chosen.

Intermediate-input cost minimisation After the realisation of v_{in}, the firm produces the rationed quantity (5) and the intermediate bundle actually used satisfies the capacity-capped requirement (11). The firms' within-period cost-minimisation problem, conditional on total bundle M_{in} is,

$$\min_{m_{ijn}} \sum_{j=1}^{J} P_j m_{ijn} \text{ s.t. (9)}$$

first-order conditions give the unit cost of the input bundle as the log-linear index

$$\mathcal{P}_i(\mathbf{P}) = \prod_{j=1}^{J} \left(\frac{P_j}{a_{ij}} \right)^{a_{ij}}, \quad \mathbf{P} := (P_1, \ldots, P_J) \tag{13}$$

and the standard "expenditure shares" condition for bilateral demands

$$m_{ijn} = a_{ij} \frac{X_{in}}{P_j} \tag{14}$$

where $m_{ijn} = 0$ if $a_{ij} = 0$, and $X_{in} := \sum_{j=1}^{J} P_j m_{ijn} = \mathcal{P}_i(\mathbf{P}) M_{in}$ denotes total intermediate expenditure. The standard equation (14) implies that each firm in sector i spends a_{ij} share of its total input expenditures on the good of sector j.

Firm pricing. Before observing v_{in} the firm chooses p_{in} to maximise expected operating profit

$$\max_{p_{in}} p_{in} \mathbb{E}_v[y_{in}] - \mathcal{P}_i(\mathbf{P}) \mathbb{E}_v[M_{in}] \text{ s.t. (5), (11),}$$

taking $Y_i, P_i, \mathcal{P}_i$ as given. The firm is atomistic, so its own threshold and multiplier move with its price while the sectoral index does not. The optimum is a variable markup over an effective expected marginal cost, both functions of the firm's threshold[6]

$$p_{in} = \frac{\varepsilon_{in}}{\varepsilon_{in} - 1} \widetilde{mc}_{in}, \tag{15}$$

[6]Uniqueness of the maximiser rests on a mild shape condition on G_i, stated as Assumption 2 in Appendix A, which holds for the log-normal specification used in the calibration.

where ε_{in} is, as explained before, the effective demand elasticity of the firm, which is smaller when the firm is often capacity constrained. And $\widetilde{mc}_{in}$ is the effective expected marginal operating cost, or the expected saving in intermediate-input cost associated with a marginal price-induced reduction in expected sales

$$\widetilde{mc}_{in} := \mathcal{P}_i \frac{\partial \mathbb{E}_v[M_{in}]/\partial p_{in}}{\partial \mathbb{E}_v[y_{in}]/\partial p_{in}} = \frac{\mathcal{P}_i}{1-\alpha_i} \frac{\eta_{i,\frac{1}{1-\alpha_i}}(\overline{v}_{in})\, \mathbb{E}_v[M_{in}]}{\eta_{i,1}(\overline{v}_{in}) \mathbb{E}_v[y_{in}]}$$

with η function given in (7). Lemma 2 in Appendix A establishes existence and uniqueness of the maximiser and gives the price, the markup wedge and the optimal threshold in closed form in terms of the tail functions of G_i.

At given capital the only flexible input is the bundle, whose output elasticity is $1-\alpha_i$, so the marginal cost of an extra unit of output is the bundle price divided by the bundle's marginal product,

$$\frac{\mathcal{P}_i}{\partial y_{in}/\partial M_{in}} = \frac{\mathcal{P}_i M_{in}}{(1-\alpha_i) y_{in}},$$

that is, average bundle cost per unit of output, scaled up by $\frac{1}{1-\alpha_i} > 1$ because of diminishing returns to the bundle. $\widetilde{mc}_{in}$ is this object with the realised (M_{in}, y_{in}) replaced by the expectations the firm prices against, since the price is set before v_{in} is realised. Under capacity the relevant expectations are the *marginal* ones, $\eta_{i,\frac{1}{1-\alpha_i}} \mathbb{E}_v[M_{in}]$ and $\eta_{i,1} \mathbb{E}_v[y_{in}]$, because a price cut raises output and input use only in the states in which the firm is not rationed. In other words, the truncated moments appear because the expected saving in input cost from selling one unit less is computed on the unconstrained tail alone.

Inserting the expectations (6)–(12) and the elasticity $\varepsilon_{in} = \varepsilon\, \eta_{i,1}(\overline{v}_{in})$ into (15) gives the pricing rule as,

$$p_{in} = \underbrace{\frac{\varepsilon_{in}}{\varepsilon_{in}-1} \cdot \frac{1}{1-\alpha_i} \cdot \frac{\int_0^{\overline{v}_{in}} v^{\frac{1}{1-\alpha_i}} \mathrm{d}G_i(v)}{\int_0^{\overline{v}_{in}} v \mathrm{d}G_i(v)}}_{=:\, \Gamma_i(\overline{v}_{in})} \frac{\mathcal{P}_i(\mathbf{P})}{Z_i x_i^{\alpha_i}} \overline{v}_{in}^{-\frac{\alpha_i}{1-\alpha_i}}, \tag{16}$$

Written as $\Gamma_i(\overline{v})$, the wedge between price and the marginal-cost index is the product of three factors: the variable markup, the diminishing-returns factor, and a demand-risk factor which is the ratio of the lower-truncated moments $\int_0^{\overline{v}} v^{\frac{1}{1-\alpha_i}} \mathrm{d}G_i$ and $\int_0^{\overline{v}} v \mathrm{d}G_i$. Demand risk therefore raises effective marginal cost, through the convexity of the input requirement on the unconstrained tail, *and* lowers the effective elasticity, through the rationed tail. The second channel is new in this model, and both work through the same threshold (Lemma 2(iii) in Appendix A.3).

Symmetric sectoral equilibrium. Since price and investment are set before v_{in} is realised, there is a symmetric equilibrium $p_{in} = p_i$, $k_{in} = K_i$ for all n, in which all

firms in the sector share the capacity $\overline{y}_i := Z_i K_i x_i^{\alpha_i - 1}$ and the threshold $\overline{v}_i$. Unlike the model without capacity constraints, $P_i \neq p_i$: the shadow index carries the multipliers of the rationed firms.

Using symmetry in the sectoral equilibrium, the price ratio p_i/P_i can be written as a function of $\overline{v}_i$ only,

$$\frac{p_i}{P_i} = \left[\int_0^{\overline{v}_i} v \mathrm{d}G_i(v) + \overline{v}_i^{1-\frac{1}{\varepsilon}} \int_{\overline{v}_i}^{\infty} v^{1/\varepsilon} \mathrm{d}G_i(v) \right]^{\frac{1}{\varepsilon-1}} =: \mathcal{R}_i(\overline{v}_i), \tag{17}$$

which is strictly increasing in the threshold and lies below one: the posted price lies below the shadow index, and the gap closes as capacity of firms in sector i ceases to bind. By the definition of the threshold in (5), output relative to capacity is a function of the threshold as well, $Y_i/\overline{y}_i = \mathcal{R}_i(\overline{v}_i)^{\varepsilon}/\overline{v}_i$, and a strictly decreasing one. As such, output relative to capacity pins the threshold down uniquely, and since the ratio is bounded above, sector output cannot exceed the sector's capacity and this makes capacity utilisation a well-defined object bounded above by 1 (see Definition 1 below). Furthermore, in the symmetric equilibrium, the expected bundle demand in (12) becomes aggregate bundle demand M_i^e. Aggregate and bilateral intermediate spending are given by

$$X_i = \mathcal{P}_i(\mathbf{P})\, M_i^e, \qquad m_{ij} = a_{ij} \frac{X_i}{P_j}. \tag{18}$$

Sectoral goods trade at the shadow index P_i, but due to the rationing rents, which accrue to aggregators and not the monopolists, the quantity of sectoral sales is different from that of firm revenues. Sales at the shadow index, S_i, the revenue of the sector's firms S_i^{f}, and their difference are given by

$$S_i := P_i Y_i, \quad S_i^{\mathrm{f}} := p_i\, \mathbb{E}_v[y_i], \quad \Pi_i^{\mathrm{r}} := S_i - S_i^{\mathrm{f}} = p_i\, \overline{y}_i \int_{\overline{v}_i}^{\infty} \left[\left(\frac{v}{\overline{v}_i} \right)^{1/\varepsilon} - 1 \right] \mathrm{d}G_i(v) \;\geq\; 0, \tag{19}$$

where Π_i^{r} is the rationing rent earned by the competitive aggregators and passed to the household (Households, below), strictly positive whenever some firms are rationed. GDP, $Q = \sum_i (S_i - X_i)$, does not depend on where the rent is booked, but the split of S_i between the firms' capital income, their markup rents and the rationing rent does (Lemma 5).

Within-sector heterogeneity is transient: all firms post the same price and hold the same capital, and capacity adds heterogeneity in *realised* output and in who is rationed, but no persistent firm-level state. The within-sector distribution of firms therefore stays degenerate, and v is a purely transient shock. Lemma 3 in Appendix A.4 collects these results, with the monotonicity properties of $\mathcal{R}_i$, and adds the sector's supply relation, the pricing rule (16) at the symmetric point.

The model without capacity constraints is the limit $x_i \to 0$: capacity $\overline{y}_i = Z_i K_i x_i^{\alpha_i - 1}$ becomes infinite, the probability of being rationed vanishes, and every object of this section converges to its counterpart in that model (Lemma 4 in Appendix A.5).

Without capacity constraints, the pricing first-order condition splits every euro of firm revenue into three constant pieces: a share $\frac{\varepsilon-1}{\varepsilon}(1-\alpha_i)$ paid for intermediates, a share $\frac{\varepsilon-1}{\varepsilon}\alpha_i$ earned by capital, and the markup rent $\frac{1}{\varepsilon}$. Under capacity the markup rent remains the share $\frac{1}{\varepsilon}$ of firm revenue but factor shares move with the threshold. Intermediate spending is a state-dependent share of firm revenue, $X_i = \xi_i(\overline{v}_i)S_i^{\mathrm{f}}$, with ξ_i given in closed form in Lemma 5(i). With $\pi_i(k)$ the maximised expected operating profit of a sector-i firm with own capital k and π_i^k its derivative at $k = K_i$, sectoral sales split into

$$S_i = \underbrace{X_i}_{\text{materials}} + \underbrace{\pi_i^k K_i}_{\text{capital income}} + \underbrace{\frac{S_i^{\mathrm{f}}}{\varepsilon}}_{\text{markup rent}} + \underbrace{\Pi_i^{\mathrm{r}}}_{\text{rationing rent}}, \tag{20}$$

where the markup rent is exactly the share $1/\varepsilon$ of firm revenue at every threshold, as without capacity constraints. Since the firms' share of sales S_i^{f}/S_i is itself state-dependent, even the markup-rent share of GDP is no longer a constant. Here, an extra unit of capital raises a firm's expected profit through two channels: in the states in which the firm is rationed it relaxes the constraint, which pays only if the posted price covers the bundle cost of a unit produced at capacity. In the states in which it is not, output is demand-determined and the extra capital only lowers the bundle needed to produce it, a saving that is always positive. Only the second channel exists in the model without capacity constraints. Lemma 5 in Appendix A.6 states these results with the share functions in closed form and their limits, and decomposes the marginal profit of capital into the two channels; Remark 3 there explains why the markup rent survives.

Utilization compares sector output with the output the sector would produce if every firm were at capacity.

Definition 1 (Capacity and utilization). Sector-i capacity is the sectoral output when every firm produces at capacity, $y_{in} = \overline{y}_i$ for all n in (2):

$$\overline{Y}_i := \overline{y}_i \left[\int_0^\infty v^{1/\varepsilon} \mathrm{d}G_i(v)\right]^{\frac{\varepsilon}{\varepsilon-1}}. \tag{21}$$

Sector utilization is $U_i := Y_i/\overline{Y}_i = U_i(\overline{v}_i)$, with

$$U_i(\overline{v}) := \frac{\mathcal{R}_i(\overline{v})^\varepsilon}{\overline{v}} \left[\int_0^\infty v^{1/\varepsilon} \mathrm{d}G_i(v)\right]^{-\frac{\varepsilon}{\varepsilon-1}}. \tag{22}$$

Aggregate utilization is the value-added-weighted mean of sector utilization,

$$\overline{U} := \frac{\sum_{i=1}^J (S_i - X_i)\, U_i}{\sum_{i=1}^J (S_i - X_i)} = \frac{1}{Q}\sum_{i=1}^J (S_i - X_i)\, U_i. \tag{23}$$

$U_i(\cdot)$ is a strictly decreasing bijection of $(0,\infty)$ onto $(0,1)$: utilization is one when every firm is constrained ($\overline{v}_i \to 0$) and tends to zero as capacity becomes slack ($\overline{v}_i \to \infty$). It depends on the state $(\mathbf{K}, \mathbf{Z})$ and on prices only through the threshold, because

both Y_i and $\overline{Y}_i$ are $\overline{y}_i$ times a function of $\overline{v}_i$. U_i is the object that corresponds to a measured sectoral output-to-capacity ratio.[7]

Final good. A competitive final producer combines sectoral goods according to a CES aggregator with elasticity of substitution $\eta > 0$, $\eta \neq 1$,

$$Q = \left[\sum_{j=1}^{J} \beta_j^{1/\eta} Q_j^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}, \quad \beta_j > 0, \quad \sum_j \beta_j = 1; \tag{24}$$

the Cobb–Douglas aggregator is its limit $\eta \to 1$ and is treated in Appendix A.8. Cost minimisation gives the final demands and the price index,

$$\frac{P_j Q_j}{PQ} = \beta_j \left(\frac{P_j}{P} \right)^{1-\eta} \qquad P = \left[\sum_{j=1}^{J} \beta_j P_j^{1-\eta} \right]^{\frac{1}{1-\eta}}, \tag{25}$$

With the final good as numeraire, $P = 1$ becomes

$$\sum_{j=1}^{J} \beta_j P_j^{1-\eta} = 1, \tag{26}$$

so that $\sum_j \beta_j = 1$ is a normalisation of the units of the final good rather than a restriction on expenditure shares (see Appendix A.8).

Market clearing. Sectoral clearing, in units of good j, is

$$Y_j = Q_j + \sum_{i=1}^{J} m_{ij} \tag{27}$$

with $m_{ij} = a_{ij} X_i / P_j$ and $X_i = \mathcal{P}_i(\mathbf{P}) M_i^e$ of (18). In this model sector i spends the fixed value share a_{ij} of a state-dependent bill X_i on good j, whereas without capacity constraints X_i would itself be a fixed share of sales. The value form of (27), $S_j = P_j Q_j + \sum_i a_{ij} X_i$, is the outer block of the static equilibrium in Section 2.2.

Households. A representative household has isoelastic preferences over the final good,

$$\mathbb{E}_0 \int_0^\infty e^{-\rho t} u(C_t) \mathrm{d}t, \quad u(c) = \begin{cases} \frac{c^{1-\gamma}}{1-\gamma} & , \quad \eta \geq 0, \eta \neq 1 \\ \ln(c) & , \quad \eta = 1 \end{cases}$$

The household owns the firms and the aggregators, so it receives the firms' operating profits net of investment and the aggregators' rationing rents $\sum_i \Pi_i^{\mathrm{r}}$, and consumes the final good net of investment.

[7] The value-added weights in (23) follow the convention of the Federal Reserve Board's aggregate capacity-utilization rate, which averages industry rates with value-added proportions.

Capital accumulation and shocks. We assume that the capital of each firm evolves according to a deterministic law of motion with sector-specific capital adjustment costs,

$$\dot{k}_{in} = [\Phi_i(\iota_{in}) - \delta_i]k_{in} \tag{28}$$

where ι_{in} is investment per unit of capital, so that in the symmetric equilibrium $\iota_{in} = \iota_i$ and sector investment is $I_i = \iota_i K_i$, and Φ_i are the capital adjustment cost functions. We use the following specification

$$\Phi_i(\iota) := (1 + \phi_i \iota)^{1/\phi_i} - 1, \qquad \phi_i > 1.$$

The productivity levels Z_{it} are independent exogenous Ornstein–Uhlenbeck log-productivity processes with mean-reversion speed θ_i and volatility ς^{Z_i}: $\mathrm{d}\ln Z_i = -\theta_i \ln Z_i \,\mathrm{d}t + \varsigma^{Z_i}\,\mathrm{d}B^Z_{i,t}$. Itô's lemma implies[8]

$$\mathrm{d}Z_i = \underbrace{Z_i\left(-\theta_i \ln Z_i + \frac{(\varsigma^{Z_i})^2}{2}\right)}_{=:\mu_{Z_i}(Z_i)}\mathrm{d}t + \underbrace{Z_i\varsigma^{Z_i}}_{=:\sigma_{Z_i}(Z_i)}\,\mathrm{d}B^Z_{i,t}.$$

2.2 Static General Equilibrium

Given sectoral capital $\mathbf{K} = (K_1, \ldots, K_J)$ and productivity $\mathbf{Z} = (Z_1, \ldots, Z_J)$, a static equilibrium is prices $\mathbf{P}$, thresholds $\overline{\mathbf{v}} = (\overline{v}_1, \ldots, \overline{v}_J)$, outputs $\mathbf{Y}$, intermediate expenditures $\mathbf{X}$, bilateral demands $\{m_{ij}\}$, final demands $\{Q_j\}$, and final output Q such that:

1. sectoral prices and thresholds satisfy the pricing rule (16) and the price ratio (17), with $\overline{v}_i \in (\underline{v}_i, \infty)$, and outputs satisfy $Y_i = \overline{y}_i\, \mathcal{R}_i(\overline{v}_i)^{\varepsilon}/\overline{v}_i$;
2. expenditures satisfy $X_i = \mathcal{P}_i(\mathbf{P})M_i^e$ with M_i^e of (18);
3. $m_{ij} = a_{ij}X_i/P_j$;
4. final demands satisfy (25);
5. goods markets clear, (27);
6. the numeraire $P = 1$ holds.

The static general equilibrium is characterised, firstly, by J independent one-dimensional non-linear fixed points given by the pricing rule (16) and the price ratio (17), which combine into one equation per sector in its threshold alone (equation (81) in Appendix A.9). This inner block solves for the thresholds $\overline{v}_i$ given the vector of prices $\mathbf{P}$, and neither the sector's capital nor any other sector's state enters it. Secondly, an outer block solves for the prices $\mathbf{P}$ using $J - 1$ independent goods-market clearing conditions[9] in value form,

$$S_j := P_j Y_j = \beta_j P_j^{1-\eta}\, Q + \sum_{i=1}^{J} a_{ij} X_i, \qquad Q = \sum_{j=1}^{J} S_j - \sum_{j=1}^{J} X_j, \qquad j = 1, \ldots, J, \tag{29}$$

[8] $\mathcal{L}_{Z,i} f := \mu_{Z_i}(Z_i) f_{Z_i} + \frac{1}{2}\sigma_{Z_i}(Z_i)^2 f_{Z_i Z_i}$ denotes the generator of Z_i and $\mathcal{L}_Z := \sum_{i=1}^{J} \mathcal{L}_{Z,i}$ that of the vector process $\mathbf{Z}$, a sum of one-dimensional operators because the Z_i are independent.

[9] Due to Walras' law one clearing equation is redundant.

and the numeraire (26).

Homogeneity Define aggregate capital and capital shares as

$$K := \sum_{i=1}^{J} K_i, \qquad \kappa_i := \frac{K_i}{K}. \tag{30}$$

Since the shares add to unity, we have $\kappa_J = 1 - \sum_{i<J} \kappa_i$. In Proposition 3 in Appendix A.9 I prove three properties of the static equilibrium. First, the threshold equations have a solution for every price vector. Second, scaling every capital stock by $\lambda > 0$ scales every quantity by λ and leaves prices and thresholds unchanged, so that, along any equilibrium selection that respects this scaling, prices, thresholds, utilization and the share functions depend on $(\boldsymbol{\kappa}, \mathbf{Z})$ alone and

$$Q(\mathbf{K}, \mathbf{Z}) = K\, q(\boldsymbol{\kappa}, \mathbf{Z}), \quad S_i = K\, \widehat{s}_i(\boldsymbol{\kappa}, \mathbf{Z}), \quad S_i^{\mathrm{f}} = K\, \widehat{s}_i^{\mathrm{f}}(\boldsymbol{\kappa}, \mathbf{Z}), \quad X_i = K\, \widehat{x}_i(\boldsymbol{\kappa}, \mathbf{Z}), \tag{31}$$

with $q = \sum_i (\widehat{s}_i - \widehat{x}_i)$. That is, final output is homogeneous of degree one in the capital stocks. Therefore the relevant endogenous state is the $(J-1)$-vector of capital shares on the simplex

$$\boldsymbol{\kappa} := (\kappa_1, \ldots, \kappa_{J-1}) \in \Delta^{J-1} := \left\{ \kappa \in \mathbb{R}_+^{J-1} : \sum_{i=1}^{J-1} \kappa_i \leq 1 \right\}. \tag{32}$$

Third, Proposition 3 proves that in every equilibrium the threshold stays above $\underline{v}_i$, so utilization and the rationed share of firms are bounded away from one by constants that depend on $(\varepsilon, \alpha_i, G_i)$ alone.

Two consequences of this must be emphasised. First, at *given prices* the threshold $\overline{v}_i$, and with it utilization, the rationed proportion of firms, the markup and the expenditure shares, do not depend on the sector's capital. Capital moves utilization only through prices, and a sector that receives more capital at unchanged prices scales output and capacity together. Secondly, full utilization is never reached in equilibrium, and neither is universal rationing: at a threshold below $\underline{v}_i$ expected demand is inelastic, so the firm raises its price rather than accept rationing on a larger set of states.

In such models, final expenditure shares respond to relative prices with elasticity $1 - \eta$. Under substitutes ($\eta > 1$) spending moves towards a sector whose price falls, so productive sectors gain influence and shocks are amplified. Under complements ($\eta < 1$) spending moves towards expensive sectors, so adverse shocks raise a sector's influence, which leads to bottlenecks in networks with complementarities. With the existence of capacity constraints the state-contingent reallocation of spending is present even with unit-elastic final demand, because the thresholds, and through them the shares and the wedges, respond to relative prices. what CES final demand adds in my model is the response of the final expenditure shares themselves (Appendix A.10).

[Q5 | tag: Later] what to do with the notes box above?

2.3 Dynamic General Equilibrium

Ownership. As stated before, each monopolist owns and accumulates its own capital stock k and pays out operating profit net of investment as dividends to the representative household. The competitive aggregators pay out their rationing rents Π_i^{r} of Lemma 3(iv). The representative household holds all equity of the monopolists and the aggregators. There are no capital rental markets and there are no financial frictions. The household's marginal utility defines the stochastic discount factor $\Lambda_t = e^{-\rho t} C_t^{-\gamma}$, which at finite J is a *stochastic* process adapted to the aggregate state $(\mathbf{K}, \mathbf{Z})$. Because prices and investment are chosen before the idiosyncratic v_{int} are realised (Section 2.1), all firms in a sector choose the same price and investment and hold the same capital, so sector i is described by its representative stock K_i: the endogenous aggregate state is $\mathbf{K}$, and the exogenous state is $\mathbf{Z}$.

Profit as a function of own capital. The dynamic problem requires firm profit at an arbitrary own stock k, not only at the symmetric point. Under capacity it has no closed form.[10] What the dynamic problem needs are its derivative at the symmetric point and its curvature. The curvature rests on a condition on G_i, stated in Appendix A.11 (Assumption 3), under which profit is strictly concave in own capital.

[Q6 | text | tag: ToApp] The paragraph "The dynamic problem requires".

Let $\pi_i(k; \mathbf{K}, \mathbf{Z})$ denote the expected operating profit of a sector-i firm with own capital k that prices optimally against the symmetric sectoral aggregates, and $\pi_i^k := \partial \pi_i / \partial k \big|_{k=K_i}$ its marginal profit at the symmetric point where all sector-i firms have the same amount of capital. This corresponds to the capital income of the split (20). The flow that prices capital is then given by

$$\begin{gathered} \pi_i^k K_i = \zeta_i(\overline{v}_i)[1 - \xi_i(\overline{v}_i)]\, S_i^{\mathrm{f}}, \\ \text{with } \zeta_i(\overline{v})[1 - \xi_i(\overline{v})] := 1 - \eta_{i,1}(\overline{v}) - \left[\eta_{i,1}(\overline{v}) - \frac{1}{\varepsilon}\right]\left[\frac{1-\alpha_i}{\eta_{i,\frac{1}{1-\alpha_i}}(\overline{v})} - 1\right], \end{gathered} \tag{33}$$

which is the capital share of firm revenue, so that firm revenue splits into the materials share ξ_i, the capital share $\zeta_i[1 - \xi_i]$ and the markup rent $1/\varepsilon$ (Lemma 5(ii) in Appendix A.6). $\zeta_i(\overline{v})$ is the local own-capital elasticity of profit and equals $1 - \frac{1}{\varepsilon}$ when capital is scarce and the firm is rationed often (the threshold near $\underline{v}_i$, where $\xi_i \to 0$). As capacity becomes slack, this tends to $\lim_{\overline{v}\to\infty} \zeta_i(\overline{v}) = \frac{(\varepsilon-1)\alpha_i}{1+(\varepsilon-1)\alpha_i} < 1 - \frac{1}{\varepsilon}$, the constant elasticity of the model without capacity constraints. The return to own capital is therefore highest exactly where the capacity constraint bites most. Concavity of profit in own capital has two sources. Firstly, even without capacity constraints the curvature is inherited from the CES demand curve. Secondly in the model with capacity constraints, the marginal unit of capital is worth most when the firm is rationed often, and as capital grows the rationed states become more rare.

[10] A firm with more capital has a larger capacity, a higher threshold, a smaller effective markup and a lower price, all determined jointly by its pricing problem at its own k.

The firm's problem. Monopolist firms maximise the market value of equity, the Λ-discounted dividend stream. A firm is atomistic *within its sector*: as one of a unit measure of monopolists it takes as given the exogenous $\mathbf{Z}$, the path of $\mathbf{K}$ generated by the symmetric sectoral policies $\iota_j^\star(\mathbf{K}, \mathbf{Z})$, the equilibrium functions $S_i(\mathbf{K}, \mathbf{Z})$ and $C(\mathbf{K}, \mathbf{Z})$, and the SDF, and internalises only its own k:

$$\max_{\{\iota_u\}} \mathbb{E}_t \left[\int_t^\infty \frac{\Lambda_u}{\Lambda_t} [\pi_i(k_u; \mathbf{K}_u, \mathbf{Z}_u) - \iota_u k_u] \mathrm{d}u \right] \text{ s.t. } \dot{k} = [\Phi_i(\iota) - \delta_i]\, k.$$

Working in utility units, let $w^i(k, \mathbf{K}, \mathbf{Z})$ be the equity value multiplied by current marginal utility (the consumption-unit equity value is $w^i / u'(C)$). The firm HJB for w^i is

$$\rho\, w^i = \max_\iota \Big\{ u'(C)[\pi_i(k; \mathbf{K}, \mathbf{Z}) - \iota k] + w_k^i [\Phi_i(\iota) - \delta_i] k \Big\} + \sum_{j=1}^J w_{K_j}^i K_j \Big[\Phi_j(\iota_j^\star) - \delta_j \Big] + \mathcal{L}_Z\, w^i, \tag{34}$$

with first-order condition[11]

$$\Phi_i'(\iota)\, w_k^i = u'(C). \tag{35}$$

Symmetric dynamic equilibrium. A symmetric recursive equilibrium is a collection of value functions and policies $\{w^i, \iota_i^\star\}_{i=1}^J$, a consumption map $C(\mathbf{K}, \mathbf{Z})$, and static maps $\big(\mathbf{P}, \overline{\mathbf{v}}, Q, \mathbf{S}, \mathbf{S}^{\mathrm{f}}\big)(\mathbf{K}, \mathbf{Z})$ such that:

1. *(Static block)* at every $(\mathbf{K}, \mathbf{Z})$, $\big(\mathbf{P}, \overline{\mathbf{v}}, Q, \mathbf{S}, \mathbf{S}^{\mathrm{f}}\big)$ is the static equilibrium of Section 2.2;
2. *(Firm optimality)* given C, the aggregate law generated by $\{\iota_j^\star\}$, and the static maps, w^i solves (34)–(35) for each i, with π_i the profit function of Lemma 6 in Appendix A.11;
3. *(Consistency)* the maximiser of (34) at $k = K_i$ equals $\iota_i^\star(\mathbf{K}, \mathbf{Z})$ for each i;
4. *(Goods market clearing)* $C = Q(\mathbf{K}, \mathbf{Z}) - \sum_i \iota_i^\star K_i$; equivalently the household budget $C = \sum_i \big[\pi_i(K_i) - \iota_i^\star K_i\big] + \Pi^{\mathrm{r}}$, with $\Pi^{\mathrm{r}} := \sum_i \Pi_i^{\mathrm{r}}$ the rationing rents, since $\sum_i [S_i^{\mathrm{f}} - X_i] + \Pi^{\mathrm{r}} = \sum_i [S_i - X_i] = Q$ by Lemma 3(iv);
5. *(SDF consistency)* $\Lambda_t = e^{-\rho t}\, C(\mathbf{K}_t, \mathbf{Z}_t)^{-\gamma}$.

Capital-pricing system. Define the on-diagonal, utility-weighted marginal value $\mathrm{q}_i(\mathbf{K}, \mathbf{Z}) := w_k^i\big|_{k=K_i}$. Weighting by marginal utility absorbs all risk-premium terms into the state-dependence of $u'(C)$, so no explicit covariance corrections appear: $e^{-\rho t} w_t^i$ plus accumulated Λ-weighted dividends is a martingale, so w^i obeys a plain HJB with discount ρ, flow $u'(C)[\pi_i - \iota k]$, and the *physical* generator of the state $(k, \mathbf{K}, \mathbf{Z})$.

[11] Of coursem, only the own-value derivative w_k^i appears in (35), since the firm perceives no effect of its investment on K_i, hence none on (S_i, Q, C, Λ). The flow that prices capital below is the private marginal profit (33), not the derivative Q_{K_i} of the reduced output map.

Proposition 1 (Capital-pricing system). *For each $i = 1, \dots, J$ the on-diagonal marginal value satisfies the envelope equation*

$$\rho\, \mathfrak{q}_i = u'(C)\Big[\pi_i^k - \iota_i\Big] + \mathfrak{q}_i[\Phi_i(\iota_i) - \delta_i] + \sum_{j=1}^{J} \frac{\partial \mathfrak{q}_i}{\partial K_j} K_j\big[\Phi_j(\iota_j) - \delta_j\big] + \mathcal{L}_Z\, \mathfrak{q}_i, \tag{36}$$

and investment is recovered from the on-diagonal first-order condition

$$\Phi_i'(\iota_i)\, \mathfrak{q}_i = u'(C) = C^{-\gamma} \iff \iota_i = \frac{[\mathfrak{q}_i / u'(C)]^{\phi_i/(\phi_i - 1)} - 1}{\phi_i}. \tag{37}$$

Proof in Appendix A.11.

The dynamic equilibrium in levels is the solution $\Big(\{\mathfrak{q}_i, \iota_i\}_{i=1}^J,\, C\Big)$ on $(\mathbf{K}, \mathbf{Z}) \in \mathbb{R}^J_{++} \times \mathbb{R}^J_{++}$ of the J envelope equations (36), the J investment conditions (37) and the resource constraint $C = Q(\mathbf{K}, \mathbf{Z}) - \sum_i \iota_i K_i$, which determines $u'(C)$, one of the two channels (the other being the drift of $\mathbf{K}$) through which every sector's policy enters every other sector's equation. The static block of Section 2.2 supplies Q and the flows π_i^k through the thresholds at each state, and $\mathbf{K}$ and $\mathbf{Z}$ follow (28) and $\mathcal{L}_Z$. Only the static inputs differ from the model without capacity constraints. Q is homogeneous of degree one and $\mathbf{P}$, $\overline{\mathbf{v}}$ and S_i^{f}/K of degree zero in $\mathbf{K}$ (Proposition 3(ii)), which is what permits the state reduction below.

State reduction using homogeneity. The state $(\mathbf{K}, \mathbf{Z})$ has $2J$ dimensions, but one of them is a pure scale. Write $\mu^{K_i}(\iota_i) := \Phi_i(\iota_i) - \delta_i$ for the growth rate of sector i's capital in (28). Aggregate capital and the capital shares of (30) then follow

$$\frac{\dot K}{K} = \mu^K(\boldsymbol{\kappa}, \boldsymbol{\iota}) := \sum_{i=1}^{J} \kappa_i\, \mu^{K_i}(\iota_i), \qquad \dot\kappa_i = \mu_{\kappa_i}(\boldsymbol{\kappa}, \boldsymbol{\iota}) := \kappa_i\Big[\mu^{K_i}(\iota_i) - \mu^K\Big]. \tag{38}$$

The capital share drifts sum to zero, so the state stays on the simplex, and K enters the motion of $\boldsymbol{\kappa}$ nowhere. The detrended sales $\widehat{s}_i$, firm revenues $\widehat{s}_i^{\mathrm{f}}$ and intermediate spending $\widehat{x}_i$ of (31) are functions of $(\boldsymbol{\kappa}, \mathbf{Z})$ alone, as are the thresholds $\overline{v}_i(\boldsymbol{\kappa}, \mathbf{Z})$, and since $Q = Kq(\boldsymbol{\kappa}, \mathbf{Z})$ goods clearing is homogeneous of degree one in $\mathbf{K}$ as well. Define consumption per unit of aggregate capital

$$c(\boldsymbol{\kappa}, \mathbf{Z}, \boldsymbol{\iota}) := q(\boldsymbol{\kappa}, \mathbf{Z}) - \sum_{i=1}^{J} \kappa_i \iota_i, \qquad C = K\, c(\boldsymbol{\kappa}, \mathbf{Z}, \boldsymbol{\iota}). \tag{39}$$

The same scaling runs through the dynamic problem. The profit $\pi_i(k; \mathbf{K}, \mathbf{Z})$ of a firm is homogeneous of degree one in $(k, \mathbf{K})$ jointly, because scaling own and sectoral capital together leaves prices and the firm's threshold unchanged (Proposition 3(ii) and Lemma 6(i)), and C is of degree one. The utility-weighted value w^i is therefore homogeneous of degree $1 - \gamma$, and every object of the capital-pricing system (36)–(37) deflates by $K^{-\gamma}$. What remains is a system on the $(2J - 1)$-dimensional state $(\boldsymbol{\kappa}, \mathbf{Z})$.

Proposition 2 (Homogeneity reduction of the decentralised equilibrium). *Consider a symmetric equilibrium whose static maps respect the scaling of Proposition 3(ii) and whose policies $\iota_j^\star$ are homogeneous of degree zero in* $\mathbf{K}$.

(i) $w^i(\lambda k, \lambda \mathbf{K}, \mathbf{Z}) = \lambda^{1-\gamma}\, w^i(k, \mathbf{K}, \mathbf{Z})$ *for all* $\lambda > 0$, *and on the diagonal*

$$\mathfrak{q}_i = K^{-\gamma}\, \widehat{\mathfrak{q}}_i(\boldsymbol{\kappa}, \mathbf{Z}), \qquad u'(C) = K^{-\gamma} c^{-\gamma};$$

(ii) the capital-pricing system (36)–(37) *is equivalent to the deflated system on* $(\boldsymbol{\kappa}, \mathbf{Z}) \in \Delta^{J-1} \times \mathbb{R}^J_{++}$: *for each* $i = 1, \dots, J$,

$$\begin{aligned} \left(\rho + \gamma\, \mu^K\right) \widehat{\mathfrak{q}}_i = c^{-\gamma} \left[\zeta_i(\overline{v}_i)[1 - \xi_i(\overline{v}_i)]\, \frac{\widehat{s}_i^{\mathrm{f}}(\boldsymbol{\kappa}, \mathbf{Z})}{\kappa_i} - \iota_i \right] \\ + \widehat{\mathfrak{q}}_i\, \mu^{K_i}(\iota_i) + \sum_{j=1}^{J} \frac{\partial \widehat{\mathfrak{q}}_i}{\partial \kappa_j}\, \mu_{\kappa_j} + \mathcal{L}_Z\, \widehat{\mathfrak{q}}_i, \end{aligned} \tag{40}$$

with $\overline{v}_i(\boldsymbol{\kappa}, \mathbf{Z})$ *and* $\widehat{s}_i^{\mathrm{f}}(\boldsymbol{\kappa}, \mathbf{Z})$ *the threshold and detrended firm revenue delivered by the static block, the drifts* μ^K *and* μ_{κ_j} *of* (38), *and first-order condition* $c^{-\gamma} = \widehat{\mathfrak{q}}_i\, \Phi_i'(\iota_i)$, *equivalently* $\Phi_i'(\iota_i)\, \tilde{q}_i = 1$ *for Tobin's* $\tilde{q}_i := \widehat{\mathfrak{q}}_i c^{\gamma}$;

(iii) the J *equations* (40) *are coupled at every state point only through the scalar* c *and the shared drifts,*

$$\begin{aligned} 0 &= c + \sum_{j=1}^{J} \kappa_j\, \iota_j(c) - q(\boldsymbol{\kappa}, \mathbf{Z}), & \iota_j(c) &= \frac{\left[\widehat{\mathfrak{q}}_j c^{\gamma}\right]^{\phi_j/(\phi_j - 1)} - 1}{\phi_j}, \\ \mu^K &= \sum_{l=1}^{J} \kappa_l\, \mu^{K_l}(\iota_l), & \mu_{\kappa_j} &= \kappa_j \left[\mu^{K_j}(\iota_j) - \mu^K\right], \end{aligned} \tag{41}$$

and at every state the first equation of (41) *has exactly one root* $c \in (0,\, q + \sum_i \kappa_i/\phi_i)$;

(iv) the reduction is exact: level objects are recovered by $\mathfrak{q}_i = K^{-\gamma}\widehat{\mathfrak{q}}_i$, $C = Kc$ *and* $\dot{K}/K = \mu^K$.

Proof in Appendix A.12.

The effective discount $\rho + \gamma\mu^K$ is state-dependent[12], and sector i's policy enters sector j's equation only through the two aggregates generated at each state point by (41), the marginal utility $c^{-\gamma}$ and the drift pair $(\mu^K, \mu_{\boldsymbol{\kappa}})$ The network enters through the static block that delivers the thresholds $\overline{v}_i$, the detrended firm revenues $\widehat{s}_i^{\mathrm{f}}$ and q. The matrix $\mathbf{A}$ acts in three distinct places: through the bundle prices $\mathcal{P}_i(\mathbf{P})$ inside each sector's threshold equation (equation (81) in the appendix), hence inside its private flow $\zeta_i(\overline{v}_i)[1 - \xi_i(\overline{v}_i)]\widehat{s}_i^{\mathrm{f}}/\kappa_i$; through the functional form of the common q, via the clearing system in which each sector's spending $\widehat{x}_i$ reaches its suppliers; and indirectly through

[12] It is the finite-J counterpart of the constant balanced-growth rate $\rho + \gamma g$ of Section 3.3, where aggregate growth is deterministic; likewise $\Phi_i'(\iota_i)\tilde{q}_i = 1$ is the exact analogue of the investment condition there.

(41), where the flows shape every sector's policy and the policies feed back into the shared aggregates $(c, \mu^K, \mu_{\boldsymbol{\kappa}})$ entering every other sector's equation. A fourth place is new in this model: the thresholds, and with them utilization and the shares, respond to the whole price vector, so a shock anywhere in the network moves the rationing margin everywhere.

Remark 1 (Private and social returns to capital) The flow that prices capital in (40) is the firm's marginal profit π_i^k, not the marginal product of sector-i capital in final output, Q_{K_i}. Summing the revenue split (20) over sectors, with $Q = \sum_i (S_i - X_i)$ and the Euler identity $Q = \sum_i K_i Q_{K_i}$ (Proposition 3(iii) in Appendix A.9) shows

$$\sum_{i=1}^{J} \pi_i^k K_i = Q - \frac{1}{\varepsilon} \sum_{i=1}^{J} S_i^{\mathrm{f}} - \Pi^{\mathrm{r}} < Q = \sum_{i=1}^{J} K_i Q_{K_i}.$$

Private investment incentives are therefore too weak on average, by the markup rents plus the rationing rents. Without capacity constraints only the markup rents separate the two sides, and the sign and cross-section of the wedge are established there (Appendix C.3). Under capacity constraints the rationing rent is a second aggregate wedge, and both wedges move with the state through the thresholds.

3 The Continuum Model

3.1 Model Setup on $[0, 1]$

From sectors to nodes. In the finite-sector model the state of the economy is $(\boldsymbol{\kappa}, \mathbf{Z}) \in \Delta^{J-1} \times \mathbb{R}_+^J$, a vector of $2J - 1$ dimensions, and every sector's capital-pricing equation is defined on this domain. The continuum model replaces the index set $\{1, \ldots, J\}$ by the vertex set $[0, 1]$ and the input–output matrix $\mathbf{A} = (a_{ij})$ by a kernel $a(i, j)$ on the unit square. The payoff is a collapse of dimension: instead of J equations on a $(2J - 1)$-dimensional state we obtain, for each node $i \in [0, 1]$, a pair of *two-dimensional* equations in the node's own capital and productivity, a Hamilton–Jacobi–Bellman-type equation and a Fokker–Planck–Kolmogorov equation, coupled across nodes only through a small set of kernel-weighted moments. The mechanism at work is a form of the law of large numbers: with a continuum of nodes carrying independent shocks, every cross-sectional average, and with it every aggregate and every equilibrium price functional, is deterministic.

In this continuum model the logic of classical mean field games is transplanted to a network. Consider N households in an incomplete-market economy whose interest rate clears against the empirical capital stock. In this setting, each household's value function carries the entire cross-section as a state: a household of mass $1/N$ moves the capital stock and hence the interest rate, and with $N < \infty$ the idiosyncratic shocks do not average out, so the interest rate is stochastic and forecasting it requires the distribution as a state. The limit $N \to \infty$ removes both, such that the distribution

remains an unknown of the problem, but it is no longer an argument of any agent's value function: in a stationary equilibrium its whole decision-relevant content is one scalar, the interest rate. In my model the same occurs node by node, with network position surviving as a persistent type. The equilibrium object is an *ensemble* of node distributions coupled through kernel-weighted moments, and the single interest rate becomes a few functions on $[0,1]$. Appendix B.9 develops the relation to graphon mean field games.

Input-output kernels and operators.

Definition 2 (Input–output kernel). An *input–output kernel* is a bounded measurable function $a : [0,1]^2 \to [0,\bar{a}]$, $\bar{a} < \infty$, that is *row-stochastic*,

$$\int_0^1 a(i,j)\,\mathrm{d}j = 1 \qquad \text{for a.e. } i \in [0,1]. \tag{42}$$

Here $a(i,j)$ is the Cobb–Douglas weight of node-j goods in node i's intermediate bundle, and constitutes the continuum analogue of the row-stochastic matrix $\mathbf{A}$.

Definition 3 (Graphon operators). For $f \in L^\infty([0,1])$ and $g \in L^1([0,1])$ define the operators

$$(\mathbb{A}f)(i) := \int_0^1 a(i,j)\,f(j)\,\mathrm{d}j, \qquad (\mathbb{A}^*g)(i) := \int_0^1 a(j,i)\,g(j)\,\mathrm{d}j. \tag{43}$$

The operator $\mathbb{A}$ averages *along rows* (what i buys) and $\mathbb{A}^*$ averages *along columns* (what is sold *to* i). They are adjoint to each other: $\int (\mathbb{A}f)g = \int f(\mathbb{A}^*g)$ whenever both sides are defined. The finite-graph counterparts are the matrix $\mathbf{A}$ and its transpose. $\mathbb{A}$ corresponds to $\mathbf{f} \mapsto \mathbf{A}\mathbf{f}$, the row-weighted average $\sum_j a_{ij} f_j$ of node values, and $\mathbb{A}^*$ to $\mathbf{g} \mapsto \mathbf{A}^\top \mathbf{g}$, the column sum $\sum_j a_{ji} g_j$ that routes node masses back to their suppliers. These operators have a Leontief inverse: for any profile of intermediate expenditure shares bounded away from one, the operator $\mathrm{Id} - \mathbb{A}^*\Xi$, with Ξ pointwise multiplication by the shares, is invertible, the continuum analogue of $[\mathbf{I} - \mathbf{A}^\top \boldsymbol{\Xi}]^{-1}$. Lemma 9 in Appendix B.1 proves this and bounds the Domar density the inverse generates; under capacity every share the model generates lies below $1 - \frac{1}{\varepsilon}$, so the lemma applies with that bound. **[Q7 | notes | tag: OperatorBound]** why should a reader who's not deeply interested in theory care about "...bounds the Domar density the inverse generates; under capacity every share the model generates lies below $1 - \frac{1}{\varepsilon}$, so the lemma applies with that bound."? Can you write 1 sentence on the economic meaning of this bound or why it matters for the model? Write this wrapped in notes environment (even if you write it inside the above paragraph) so that I can check what you added.

Step kernels. A finite J-sector economy can be embedded as a *step kernel*: partition $[0,1]$ into intervals (which we refer to as "blocks") $I_1, \ldots, I_J$ of length $1/J$ and set

$$a^J(i,j) := J\,a_{ij}, \quad \beta^J(i) := J\,\beta_i, \qquad (i,j) \in I_i \times I_j,$$

so that (42) holds and block-integrated weights reproduce the matrix entries[13],

$$\int_{I_j} a^J(i,j)\mathrm{d}j = a_{ij}.$$

Using this, an empirical J-sector input–output table can be mapped directly into a step kernel with J blocks, with all primitives $(\alpha, x, \sigma_v, \beta, \phi, \delta, \theta, \varsigma^Z)$ becoming block-constant functions of the node index. On step functions the two operators reduce to the matrices $\mathbf{A}$ and $\mathbf{A}^\mathsf{T}$ acting on block values and block masses, so every kernel-weighted object below is $J \times J$ matrix algebra, and the graphon Leontief inverse collapses exactly to the finite one: the Domar density is the finite Domar vector in density units (this is shown Lemma 10 in Appendix B.2, which also records a constant by which continuum and finite price levels differ).

On a step kernel, then, only the network layer of the continuum model collapses to a finite object. Every object the kernel enters becomes a $J \times J$ matrix product of block values and block masses, so the whole system reduces to J coupled per-block problems. What does not collapse is the economy inside a block: each block still holds a continuum of industries with their own shocks, so its aggregates are deterministic, whereas a sector of the finite model receives one common shock and its aggregates are random (Section 3.5). For the computation this means that the kernel introduces no discretisation error; only the grid on a node's own state does.

One consequence of Definition 2,[14] which shapes the reading of the model, is that the bundle-weight measure $\nu_i(\mathrm{d}j) := a(i,j)\mathrm{d}j$ is *diffuse*. Diffuseness means that no *individual* supplier carries positive weight in a buyer's bundle: only sets of suppliers of positive measure do. This is the supply-side counterpart of the atomism that makes the node problem two-dimensional. As such, a kernel in this model cannot represent a situation where node i sources a fixed share $a_0 > 0$ from one named supplier j_0.[15] What is *not* excluded by Definition 2 is a buyer sourcing heavily from its own industry group. That is, under the step-kernel embedding a node $i \in I_n$ has $\nu_i(I_n) = \int_{I_n} J\, a_{nn}\mathrm{d}j = a_{nn}$, so the finite table's diagonal weight survives in full as weight on the node's own *block*, which has positive measure $1/J$. Only the weight on the single node i itself vanishes.[16]

The model on $[0,1]$. A node in the continuum model is the sector of Section 2.1 with the kernel row in place of the matrix row. Node i holds a unit measure of monopolistically competitive firms facing CES demand of elasticity ε and firm-specific demand

[13] The factor J converts matrix *weights* into *densities*.

[14] This and another quantitative restriction implied by the definition are stated as Lemma 11 in Appendix B.3.

[15] That specification would be the mixed measure $\nu_i(\mathrm{d}j) = a_0\,\delta_{j_0}(\mathrm{d}j) + [1-a_0]\,\tilde{a}(i,j)\,\mathrm{d}j$, a well-defined probability measure on $[0,1]$, but not of the form $a(i,j)\mathrm{d}j$ for any bounded a, and so excluded by Definition 2.

[16] A positive weight on the node's own good, the continuum analogue of $a_{ii} > 0$, is an atom and therefore not a kernel. It can be added as an explicit extension, at the cost of the node's own price entering its own marginal cost and of one additional scalar, monotone fixed-point per node (Appendix B.3).

shocks drawn from G_i, aggregated by competitive bundlers into the node good with price index P_i. Each firm produces with capital and an intermediate bundle that is Cobb–Douglas over the continuum of node goods with weights $a(i,\cdot)$, up to the capacity $\overline{y}_{in} = Z_i k_{in} x(i)^{\alpha(i)-1}$ of (10). The unit cost $\mathcal{P}_i$ of the bundle is the continuum version of the finite index (13): a kernel average of the suppliers' log prices plus a constant that depends on the row alone. All primitives are bounded measurable functions of the node index $i \in [0,1]$: the capital share $\alpha(i)$, the capital intensity $x(i)$, the demand-shock distribution G_i, the final-demand density $\beta(i)$ with $\int_0^1 \beta(i)\mathrm{d}i = 1$ and $\beta \geq \underline{\beta} > 0$, the adjustment-cost curvature $\phi(i)$ and the depreciation rate $\delta(i)$, and the mean-reversion speed $\theta(i)$ and volatility $\varsigma^Z(i)$ of the node's productivity, whose logarithm $z_i := \ln Z_i$ follows[17] the Ornstein–Uhlenbeck process of Section 2.1. Capital accumulates deterministically, as in (28). The bounds, the within-sector elasticity $\varepsilon > 1$, common to all nodes, and the household's preferences are listed in Assumption 4 in Appendix B.3.

Timing, pricing before the shock is realised, and the capacity block (rationing within the node, the threshold, the variable markup, the state-dependent revenue shares and utilization) are node-local and identical to the finite model. The pricing, sales and spending relations of Section 2.1 hold at every node with its own parameters $(\alpha(i), x(i), G_i)$, and every object in them is node-local except $\mathcal{P}_i$. Since investment and prices are chosen before the demand shock is realised, all firms at a node choose identically: the node state is (K_i, z_i), with no persistent firm-level state. The displays are collected in Appendix B.5, where Lemma 12 derives the bundle price index and value clearing on $[0,1]$.

A competitive final producer combines node goods through a CES technology with elasticity $\eta > 0$, $\eta \neq 1$, and a density $\beta(i)$ of final-demand weights, and the final good is the numeraire. Nominal final spending on node i is then $\beta(i)P_i^{1-\eta}Q$, which moves with the node's relative price and hence with its own state[18]. Node goods clear: node i's sales are its final demand plus the intermediate bills of its buyers routed back through the kernel, $\int_0^1 a(j,i)X_j\mathrm{d}j = (\mathbb{A}^*X)(i)$, a fixed value share of each buyer's state-dependent bill; and aggregate final expenditure equals gross output net of intermediate use. Appendix B.4 collects the dictionary between the objects of the finite model and their continuum counterparts.

Aggregation across nodes. The node state (K_i, z_i) is random, but the economy's cross-sectional aggregates are not. With a continuum of nodes carrying independent shocks, every cross-sectional integral of node-level random variables is almost surely equal to the integral of its mean (Lemma 13 in Appendix B.5).[19] This is the ingredient

[17] Throughout the continuum part we write the node's productivity state as z.

[18] Under Cobb–Douglas final demand it would be the constant $\beta(i)Q$: a node's productivity would never reach its payoffs, and the continuum model degenerates, as Appendix B.7 shows.

[19] An independent continuum of Brownian motions is not jointly measurable on the ordinary product of the node and sample spaces, so the model lives on a *Fubini extension* of that product, on which Fubini's theorem holds and such families exist (Sun, 2006). Appendix B.5.1 explains why an ordinary

that removes aggregate risk, makes the stochastic discount factor deterministic (Section 3.3), and reduces the coupling between nodes to a small set of moment integrals. The averaging operates *across* nodes, not within them: the firms at a node all share its shock, so within-node averaging has nothing to work on and the node state remains genuinely random. Determinism arrives one level up, where economy-wide integrals run over the continuum of nodes.

Definition 4 (Node distribution and measure ensemble). Let $\mathfrak{M}(t,i,\cdot)$ denote the distribution of the node state $(K_i(t), z_i(t))$ across realisations. The collection $\mathfrak{M}_G(t) := \{\mathfrak{M}(t,i,\cdot)\}_{i\in[0,1]}$ is the *measure ensemble*. By Lemma 13, for any jointly measurable f with $\int |f(i,K,z)|\,\mathfrak{M}(t,i,\mathrm{d}K,\mathrm{d}z) < \infty$ for a.e. i,

$$\int_0^1 f(i,K_i(t),z_i(t))\mathrm{d}i = \int_0^1 \int f(i,K,z)\,\mathfrak{M}(t,i,\mathrm{d}K,\mathrm{d}z)\,\mathrm{d}i \qquad \mathbb{P}\text{-a.s.}$$

Cross-sectional aggregates are therefore deterministic functionals of the ensemble $\mathfrak{M}_G$. This is the precise sense in which the mean field replaces the $(2J-1)$-dimensional state of the finite model: $\mathfrak{M}_G(t)$ is the only object beyond its own (K_i, z_i) that a node's problem needs, and because the aggregates it generates are deterministic it enters that problem as a given path of functions, not as extra arguments of the node's value function. In the finite model $(\boldsymbol{\kappa}, \mathbf{Z})$ had to be carried inside the value function precisely because it was random and every sector moved it.

3.2 Static Equilibrium Given the Cross-Sectional Distribution

Fix a date t and the measure ensemble $\mathfrak{M}_G(t)$. The static block determines prices, thresholds, sales and profits at every node as functions of the node's own (K,z) and of deterministic aggregate objects, which we define first and then show to be deterministic.

Aggregates. Given candidate node price, spending and sales functions P, X and S of (t,i,K,z), write $\mathbb{E}_i[f](t) := \int f(t,i,K,z)\,\mathfrak{M}(t,i,\mathrm{d}K,\mathrm{d}z)$ for a moment against the node distribution. This gives the three moments $m_P(t,i) := \mathbb{E}_i[\ln P](t)$, $m_X(t,i) := \mathbb{E}_i[X](t)$ and $m_S(t,i) := \mathbb{E}_i[S](t)$, and these generate the aggregates

$$\ln \mathcal{P}(t,i) = (\mathbb{A}m_P)(t,i) - \int_0^1 a(i,j)\ln a(i,j)\,\mathrm{d}j, \tag{44}$$

$$\mathcal{D}(t,i) := (\mathbb{A}^* m_X)(t,i) = \int_0^1 a(j,i)\,m_X(t,j)\,\mathrm{d}j, \tag{45}$$

$$Q(t) = \int_0^1 [m_S(t,j) - m_X(t,j)]\,\mathrm{d}j. \tag{46}$$

Each aggregate is the *moment form* of a realised cross-sectional object: take the realised formula and replace the random node objects inside the kernel integral by their

product space will not do.

node-level means (m_P, m_X, m_S), the replacement that the law of large numbers justifies (Lemma 14 in Appendix B.6).

- The bundle price (44) is the bundle price index (95) of Appendix B.5 with $\ln P_j$ replaced by $m_P(t,j) = \mathbb{E}_j[\ln P](t)$. Realised supplier prices fluctuate with each supplier's own (K_j, z_j); but node i buys from a *continuum* of suppliers whose states are essentially pairwise independent, so the cost-share-weighted average $\int a(i,j) \ln P_j \, \mathrm{d}j$ washes out all idiosyncratic deviations and only the profile of *mean* log prices survives.
- Inbound demand (45) is the inbound-demand term of the clearing condition (102) with X_j replaced by $m_X(t,j) = \mathbb{E}_j[X](t)$: node i's order book $\int a(j,i) X_j \, \mathrm{d}j$ aggregates a continuum of independent buyers, each spending the fixed value share $a(j,i)$ of its own state-dependent bill, so it too is deterministic. Note which slot of the kernel each aggregate uses: $\mathcal{P}$ averages node i's *row* (whom i buys from, via $\mathbb{A}$), $\mathcal{D}$ its *column* (who buys from i, via $\mathbb{A}^*$). Without capacity constraints $X = \xi(j)S$, the spending moment is $\xi(j) m_S(t,j)$ and (45) is $(\mathbb{A}^* \Xi m_S)(t,i)$ (Appendix C.4); under capacity spending is the state-dependent share $\xi_j(\overline{v}) S^{\mathrm{f}}/S$ of sales, so m_X is a genuine third moment of the node distribution, which moves with the buyers' utilization.
- GDP (46) is the GDP identity of (102) in moment form: aggregate value added, an integral of independent node objects, hence deterministic. It is the nominal resource base that final demand draws on.

So we first *define* the aggregates by (44)–(46), which are deterministic by construction, and then prove (Lemma 14) that the *realised* kernel integrals coincide with them almost surely. This ordering avoids circularity: the node block can be solved against deterministic aggregates, and the law of large numbers then guarantees that no realisation of the economy ever deviates from them. A single node's realisation therefore never moves $\mathcal{P}$, $\mathcal{D}$ or Q: node i's marginal cost depends on everyone else's prices only through the average along its own row.

The node problem. Fix the aggregates $(\mathcal{P}(t,i), \mathcal{D}(t,i), Q(t))$ and a node state (K, z). The within-node block involves three node-level unknowns, the price P, the threshold $\overline{v}$ and sales S. The pricing relation is the finite inner-root equation (81) at node i, with the marginal-cost side $\mathcal{P}(t,i)/(e^z x(i)^{\alpha(i)})$: it fixes the price as a function of the threshold. Sales are the node's bundle bill at full capacity, $\mathcal{P}(t,i)K/x(i)$, times a function of the threshold alone, productivity having cancelled. The demand relation is the node's value clearing: sales equal final spending $\beta(i) P^{1-\eta} Q(t)$ plus inbound demand $\mathcal{D}(t,i)$. Substituting the first two relations into the third leaves one scalar equation in the threshold, in which capital enters only the supply side and productivity only the final-demand term; the finite model's sector-by-sector inner block of Section 2.2 reappears node by node, with the aggregates in place of the price vector. It has a solution at every state and for every $\eta > 0$, and the solution is unique under a monotonicity condition on the threshold map of (81), stated as Assumption 5 in Appendix B.6, where

Lemma 15 collects the properties of the solution $(P, \overline{v}, S)(t, i, K, z)$. The threshold rises with capital, so utilization falls: more capital at given aggregates raises capacity faster than sales. It moves with productivity by the sign of $1 - \eta$: higher productivity lowers the price, which under substitutes attracts more than proportional final spending and raises utilization, and under complements loses spending and lowers it. At fixed capital and productivity, a higher price goes with a lower threshold and higher utilization, so sales rise more than proportionally with the price.

The demand relation shows the two customer types of a node. The final producer responds to the node's *relative price* with elasticity η; other nodes' bundlers spend $\mathcal{D}(t, i)$ on node i regardless of its current state, because Cobb–Douglas bundles allocate fixed value shares.[20] What capacity adds is on the supply side and inside $\mathcal{D}$: the bundlers' spending is state-independent for node i but moves with the *buyers'* utilization, and the node's own supply relation is no longer a power of its price but the pair of pricing and sales relations, in which the threshold is the node's operating position. Given the solution, the node's firms earn the operating profit $[1 - \xi_i(\overline{v})]S^{\mathrm{f}}$, its bundlers the rationing rent $S - S^{\mathrm{f}} \geq 0$, and the node's value added $S - X$ is the sum of the two (Appendix B.6).

Definition 5 (Static graphon equilibrium given $\mathfrak{M}_G(t)$). Fix a date t and a measure ensemble $\mathfrak{M}_G(t) = \{\mathfrak{M}(t, i, \cdot)\}_{i \in [0,1]}$. A *static equilibrium at t* consists of

1. jointly measurable node maps $(P, \overline{v}, S)(t, \cdot, \cdot, \cdot) : [0,1] \times (0, \infty) \times \mathbb{R} \to (0, \infty)^3$ with $\overline{v}(t, i, K, z) \in (\underline{v}_i, \infty)$, together with the spending map $X = \xi_i(\overline{v})S^{\mathrm{f}}$, integrable against $\mathfrak{M}_G(t)$ in the sense of Lemma 14; and
2. measurable aggregates $\mathcal{P}(t, \cdot) : [0,1] \to (0, \infty)$ and $\mathcal{D}(t, \cdot) : [0,1] \to [0, \infty)$, and a scalar $Q(t) > 0$;

such that

1. *(Node optimality and market clearing.)* For a.e. node i and every state (K, z), the triple $(P, \overline{v}, S)(t, i, K, z)$ solves the node block of Appendix B.6, (107)–(109), given $(\mathcal{P}(t, i), \mathcal{D}(t, i), Q(t))$.
2. *(Consistency of aggregates.)* The aggregates are generated by the moments (m_P, m_X, m_S) of (P, X, S) through (44)–(46). This is the mean-field fixed-point requirement: the marginal-cost index, inbound intermediate demand and GDP that nodes take as given must be the kernel aggregates of the very price, spending and sales functions the nodes' behaviour generates.
3. *(Numeraire.)* The final-good price equals one in moment form,

$$\int_0^1 \beta(i) \int P(t, i, K, z)^{1-\eta} \, \mathfrak{M}(t, i, \mathrm{d}K, \mathrm{d}z) \, \mathrm{d}i = 1, \tag{47}$$

which by Lemma 14 is the realised normalisation of the final-good price.

[20] The whole difference between CES and Cobb–Douglas final demand, Appendix B.7, therefore lives in the final-demand term.

Structure of the coupling. The cross-node coupling is a fixed point in two moment *functions* on $[0,1]$, (m_P, m_X), and one scalar, Q; the sales moment m_S is regenerated alongside them but enters only through Q. Given a candidate (m_P, m_X, Q), the map that a static equilibrium must leave invariant runs in three steps:

$$
\begin{aligned}
&\text{(aggregate)} && \ln \mathcal{P}(t,i) = (\mathbb{A} m_P)(t,i) - \int_0^1 a(i,j) \ln a(i,j) \mathrm{d}j, \\
& && \mathcal{D}(t,i) = (\mathbb{A}^* m_X)(t,i); \\
&\text{(solve nodes)} && (P, \overline{v}, S)(t,i,K,z) \text{ solves the node block given } (\mathcal{P}(t,i), \mathcal{D}(t,i), Q(t)), \\
& && X = \xi_i(\overline{v})\, S^{\mathrm{f}}; \\
&\text{(regenerate)} && m_P(t,i) = \int \ln P \,\mathrm{d}\mathfrak{M}(t,i,\cdot), \qquad m_X(t,i) = \int X \,\mathrm{d}\mathfrak{M}(t,i,\cdot), \\
& && m_S(t,i) = \int S \,\mathrm{d}\mathfrak{M}(t,i,\cdot), \qquad Q(t) = \int_0^1 [m_S(t,j) - m_X(t,j)] \mathrm{d}j,
\end{aligned}
$$

with (47) pinning the numeraire. Each node therefore needs to know two functions on $[0,1]$ and one number, not the ensemble of node distributions; what capacity changes is which second function, the spending moment in place of the sales moment (Remark 6 in Appendix B.6).

Individual profit and utilization. The object that drives investment in Section 3.3 is the individual monopolist's profit as a function of its own capital k, holding the node's state (K, z) and the aggregates fixed; it is the finite profit function of Section 2.3 with the node's parameters and the node objects of the static block in place of the sector's, and its marginal value at the symmetric point $k = K$ is the flow (33) at the node's threshold (Lemma 16 in Appendix B.6). Node utilization is the finite U_i of Definition 1 at the node's threshold, random across realisations according to the node distribution, and aggregate utilization is its value-added-weighted mean across nodes, deterministic like every cross-sectional aggregate (Appendix B.6).

3.3 Dynamic Equilibrium: The Graphon Mean-Field Game

The generic agent of the dynamic problem is the individual monopolist, atomistic within its node. It sees the rest of the economy only through deterministic paths, and that is what turns the finite model's $(2J-1)$-dimensional problem into one backward and one forward equation per node, coupled through the two moment functions and the scalar of Section 3.2.

Households. A representative household consumes the final good and holds its financial wealth W_t in a fully diversified portfolio: the equity of the continuum of monopolists and of the competitive aggregators at all nodes, and a locally riskless bond in zero net supply paying r_t. Preferences are CRRA with curvature γ and discount rate ρ over final-good consumption, maximised subject to $\dot{W}_t = r_t W_t + D_t - C_t$, where

D_t is the portfolio's dividend flow: operating profits net of investment spending plus the rationing rents, summed over the cross-section. Each individual firm's dividend is risky, loading on its node's state through its sales and its threshold, and so is each node's rent; but the household owns the entire cross-section, and by the law of large numbers the realised flow coincides almost surely with its deterministic mean $D_t = \int_0^1 \mathbb{E}[\pi_i - \iota_i K_i + \Pi_i^{\mathrm{r}}]\mathrm{d}i$. Node-level risk is fully diversified and never priced. The same holds for every aggregate the household cares about (Lemma 14), in particular for the consumption path, since final-good clearing $Q_t = C_t + I_t$ makes C_t deterministic. The household therefore faces an ordinary deterministic consumption–savings problem, its stochastic discount factor is the deterministic $\Lambda_t = e^{-\rho t} C_t^{-\gamma}$, and its Euler equation prices the bond:

$$r_t = \rho + \gamma \frac{\dot{C}_t}{C_t}. \tag{48}$$

With a unique deterministic discount factor Fisher separation is immediate: every shareholder instructs each firm to maximise the present value of its expected dividends at the riskless rate, and there are no risk premia in the limit economy (Lemma 18 in Appendix B.8). The fixed point of the continuum model thus never involves a stochastic pricing kernel. The cost is the loss of aggregate risk, which the numerical work handles by MIT shocks (Appendix D).

The firm's problem. The generic agent is the individual monopolist at node i. Its problem has three kinds of arguments: its own capital k, which it controls; the node state (K, z), common to the node's firms and exogenous to each of them, since a firm is atomistic within its own sector; and the economy-wide state, which enters only through the deterministic paths $\Theta_t := (\mathcal{P}(t,\cdot), \mathcal{D}(t,\cdot), Q(t), r_t)$. This is the structure of Section 3.1: an individual state, a cross-section the agent cannot move, and aggregates it takes as given. The firm maximises the expected flow of operating profit net of investment spending, $\pi(t,i,k;K,z) - \iota k$ with the profit function of Lemma 16, discounted at r_t, subject to $\dot{k} = [\Phi_i(\iota) - \delta(i)]k$ and taking the node's symmetric policy $\bar{\iota}(t,i,K,z)$ as given ((132) in Appendix B.8). Its value $w(t,i,k,K,z)$ solves

$$r_t\, w = \max_{\iota} \{\pi(t,i,k;K,z) - \iota k + w_k\, k[\Phi_i(\iota) - \delta(i)]\} + w_K\, K[\Phi_i(\bar{\iota}) - \delta(i)] + \mathcal{L}_z^{(i)} w + \partial_t w. \tag{49}$$

Set beside the finite HJB (34), the flow is the same, the first-order condition $w_k \Phi_i'(\iota) = 1$ is the same (there the marginal value is weighted by marginal utility, here it is directly in final-good units, because the discount already uses the SDF; (133)), and the firm cares about the rest of the economy in the same way, through the prices and the demand it faces. What changes is the list of arguments: own k, the node's (K, z) and the deterministic paths Θ_t, in place of the whole $(\mathbf{K}, \mathbf{Z})$ of the finite value function. There is no sum over the other sectors' drifts and no generator of the other sectors' shocks, because no other node's state is a state variable for this firm.

All monopolists at a node are identical and face the common z_i, so in equilibrium $k = K_i$ identically and $\bar{\iota}$ is the maximiser of (49) on the diagonal. The equilibrium is

fully characterised by the *on-diagonal marginal value*

$$\mathfrak{q}(t,i,K,z) := w_k(t,i,k,K,z)\Big|_{k=K}, \tag{50}$$

which solves the node-indexed version of the finite capital-pricing equation (36): the required return $r_t\mathfrak{q}$ equals the marginal dividend $\pi_k - \iota^*$, plus the marginal unit's own growth in value, plus the capital gains from the drift of the node state and of z (Proposition 7 in Appendix B.8; the equation is displayed once below, detrended and stationary, in the equilibrium definition). The marginal dividend is the capital share of firm revenue per unit of node capital, the flow (33) at the node's threshold; given the bundle price it is a decreasing function of the threshold alone, so a node whose capacity is slack earns less on the marginal unit than one whose firms are rationed often.

One shortcut is wrong. It is tempting to write a "representative firm of node i" HJB directly in (K,z), with the node's operating profit as its flow, and to read investment off its K-derivative. That derivative, $[w_k + w_K]_{k=K}$, is not the marginal value that governs investment: $w_K \leq 0$ is the demand-curve externality, the fact that more node capital depresses the node price, which an atomistic firm rightly ignores. Investing on it would let the node internalise its own demand curve, a planner-of-the-node distortion; the equilibrium flow is the private marginal profit, as in the finite model (Remark 1), and the node-level equation is needed only afterwards, for equity values (Remark 10 in Appendix B.8).

Node dynamics. In equilibrium the node's capital grows at the rate $\Phi_i(\iota^*) - \delta(i)$ set by its own policy $\iota^*(t,i,K_i,z_i)$, and its log productivity follows the OU process (93). Randomness enters capital only through the policy's dependence on z: a diffusion in z and pure transport in K, a hypoelliptic pair whenever the policy depends on z (where it does not, as under Cobb–Douglas demand, the capital marginal is a moving point mass). The node distribution $\mathfrak{M}(t,i,\cdot)$ solves a Fokker–Planck–Kolmogorov equation, one per node, in the two state variables, displayed once below in the equilibrium definition ((136)–(137) in Appendix B.8). In the terminology of Caines et al. (2022) the model is a graphon mean field game with mean-field-dependent performance functions: the graphon acts on the firm's problem through its payoff, the marginal dividend with its firm revenue and its threshold, via $\mathcal{P} = \mathbb{A}m_P$ and $\mathcal{D} = \mathbb{A}^* m_X$, not through the drift of its state; and the mean field enters through a static general-equilibrium map rather than an additively separable cost.

Balanced growth. The economy is homogeneous of degree one in the capital profile: scaling every node's capital by a common factor leaves all prices, thresholds and utilization rates unchanged and scales every nominal quantity, so that the marginal profit and the marginal value $\mathfrak{q}$ are invariant (Lemma 19 in Appendix B.9). This is the continuum analogue of the finite reduction (Proposition 3(ii)), with one more ingredient under capacity: the node's threshold equation is homogeneous of degree zero in $(K,Q,\mathcal{D})$ at fixed prices, so a node whose capital and whose customers' spending scale

together keeps its operating position, its utilization and its shares, and only the euro amounts move. The natural detrending is by aggregate capital,

$$\overline{K}_t := \int_0^1 \mathbb{E}[K_i(t)]\mathrm{d}i, \qquad g_t := \frac{\dot{\overline{K}}_t}{\overline{K}_t}, \qquad \kappa := \frac{K_i}{\overline{K}_t}, \tag{51}$$

so κ is the node's capital *density*, the mass-one analogue of the finite share $\kappa_i = K_i/K$, and g_t is the growth rate of aggregate capital, deterministic by the law of large numbers, which is why it carries the balanced-growth letter g rather than the drift notation of the finite model. Hatted objects are nominal quantities per unit of aggregate capital, as the finite hats of (31): $\widehat{s} = S/\overline{K}_t$, $\widehat{s}^{\mathrm{f}} = S^{\mathrm{f}}/\overline{K}_t$, $\widehat{x} = X/\overline{K}_t$, $\widehat{\mathcal{D}} = \mathcal{D}/\overline{K}_t$ and $q_t = Q_t/\overline{K}_t$; prices, thresholds and utilization need no detrending. The node's capital density drifts at the node's own growth rate net of the aggregate's, $\Phi_i(\iota^*) - \delta(i) - g_t$, the continuum analogue of the finite share dynamics (38), and the marginal value and the node distribution become functions of (κ, z) (Proposition 8 in Appendix B.9).

The GMFG system. The equilibrium is a Caines–Huang-type system: a backward equation per node for the marginal value, a forward equation per node for the node distribution, coupled through graphon-weighted moments and closed by the general-equilibrium conditions. We state it on the stationary balanced growth path, on which $\partial_t \equiv 0$, $g_t \equiv g$, $r = \rho + \gamma g$, and every function drops its time argument.[21]

Definition 6 (Stationary GMFG equilibrium). A stationary graphon mean-field equilibrium is a collection

$$\{\mathfrak{q}(i,\cdot,\cdot),\ \mathfrak{m}(i,\cdot,\cdot)\}_{i\in[0,1]}, \qquad \left(m_P, \widehat{m_X}, \widehat{m_S}, \mathcal{P}, \widehat{\mathcal{D}}\right) : [0,1] \to \mathbb{R}, \qquad (q, g, r) \in \mathbb{R}^3,$$

of marginal values and node distributions, moment and aggregate functions, and scalars, such that:

1. **Per-node statics.** For a.e. i and all (κ, z), $(P, \overline{v}, \widehat{s})(i, \kappa, z)$ solves the detrended node block of Appendix B.9, (139)–(141), given $\left(\mathcal{P}(i), \widehat{\mathcal{D}}(i), q\right)$ (Lemma 15), with detrended firm revenue $\widehat{s}^{\mathrm{f}}(i, \kappa, z)$ and spending $\widehat{x} = \xi_i(\overline{v})\,\widehat{s}^{\mathrm{f}}$.
2. **HJB(i): marginal $\mathfrak{q}$.** $\mathfrak{q}(i,\cdot,\cdot)$ solves

$$\begin{aligned}(\rho + \gamma g)\,\mathfrak{q} = \zeta_i(\overline{v})[1 - \xi_i(\overline{v})]\,\frac{\widehat{s}^{\mathrm{f}}(i,\kappa,z)}{\kappa} - \iota^*(\mathfrak{q}) + \mathfrak{q}[\Phi_i(\iota^*) - \delta(i)] \\ + \mathfrak{q}_\kappa\,\kappa[\Phi_i(\iota^*) - \delta(i) - g] + \mathcal{L}_z^{(i)}\mathfrak{q},\end{aligned} \tag{52}$$

   with $\iota^*(\mathfrak{q})$ from the first-order condition $\Phi_i'(\iota^*)\,\mathfrak{q} = 1$ and $\overline{v} = \overline{v}(i, \kappa, z)$ from item 1.

[21] Well-posedness of values requires $\rho + (\gamma - 1)g > 0$ at the equilibrium g, so that $r - g > 0$ and detrended firm values and household utility are finite (Assumption 6 in Appendix B.9); it is verified ex post.

3. **FPK(i).** $\mathfrak{m}(i,\cdot,\cdot) \geq 0$ with $\int \mathfrak{m}(i,\kappa,z)\mathrm{d}\kappa\,\mathrm{d}z = 1$ solves

$$0 = -\frac{\partial}{\partial\kappa}\{\kappa[\Phi_i(\iota^*(\mathfrak{q})) - \delta(i) - g]\,\mathfrak{m}\} + \frac{\partial}{\partial z}\{\theta(i)\,z\,\mathfrak{m}\} + \frac{\varsigma^Z(i)^2}{2}\frac{\partial^2}{\partial z^2}\mathfrak{m}, \tag{53}$$

with the policy from item 2.

4. **Graphon consistency (moments).** With $\mathbb{E}_i[\cdot] := \int \cdot\ \mathfrak{m}(i,\kappa,z)\mathrm{d}\kappa\,\mathrm{d}z$,

$$\begin{aligned} m_P(i) &= \mathbb{E}_i[\ln P(i,\kappa,z)], & \widehat{m_X}(i) &= \mathbb{E}_i[\hat{x}(i,\kappa,z)], \\ \widehat{m_S}(i) &= \mathbb{E}_i[\hat{s}(i,\kappa,z)], & \ln \mathcal{P}(i) &= (\mathbb{A}m_P)(i) - \textstyle\int_0^1 a(i,j)\ln a(i,j)\mathrm{d}j, \\ \widehat{\mathcal{D}}(i) &= (\mathbb{A}^*\widehat{m_X})(i), & q &= \textstyle\int_0^1 [\widehat{m_S}(j) - \widehat{m_X}(j)]\mathrm{d}j. \end{aligned} \tag{54}$$

5. **Numeraire.** $\int_0^1 \beta(i)\,\mathbb{E}_i\big[P(i,\kappa,z)^{1-\eta}\big]\mathrm{d}i = 1$.

6. **Aggregation of capital and growth.** $\int_0^1 \mathbb{E}_i[\kappa]\mathrm{d}i = 1$ and

$$g = \int_0^1 \mathbb{E}_i[(\Phi_i(\iota^*) - \delta(i))\,\kappa]\mathrm{d}i, \tag{55}$$

equivalently node-wise stationarity of $\mathbb{E}_i[\kappa]$ under (53) plus the normalisation.

7. **Household consistency.** $r = \rho + \gamma g$, $\widehat{C} = q - \int_0^1 \mathbb{E}_i[\iota^*\kappa]\mathrm{d}i > 0$, and Assumption 6 holds.

On a step kernel with N blocks the system is exactly the finite system(156)–(159) of Appendix B.10: N pairs of equations in (κ, z), N block moments and the three scalars. That is the system the paper computes.

Items 2 and 3 are the HJB–FPK pair of the Caines–Huang system, one two-dimensional pair per node; item 4 is the graphon coupling, through exactly two kernel-weighted moment functions $(\mathbb{A}m_P, \mathbb{A}^*\widehat{m_X})$ and one scalar, the sales moment being regenerated alongside them and entering only through q; items 5–7 are the general-equilibrium closures, with no analogue in the game-theoretic formulation. The whole across-node interaction is summarised by two functions on $[0,1]$ and a number. This is the tractability gain over the $(2J-1)$-dimensional finite problem, and it is delivered jointly by the law of large numbers and by the Cobb–Douglas value shares of the bundle, which make each buyer's spending on a node a fixed share of a bill that depends on the buyer's state but not on the node's (Remark 6). Constancy of the expenditure share within nodes is not needed; what capacity changes is which second function is carried, the spending moment in place of the sales moment. The productivity shock can be switched off without changing the format: at $\varsigma^Z(i) = 0$ the node's z is deterministic, equal to zero on the balanced growth path, $\mathcal{L}_z^{(i)}$ drops out, and the node's problem is a deterministic one in κ whose stationary distribution is a point mass.

3.4 What Is Established, What Is Assumed, What Is Open

[Flag: set aside on 30 October 2026] This subsection is held, unchanged and with its heading and label, inside a maroon box until what to do with it is decided. The heading stays so that the references to it, all of them now inside maroon or dark-teal boxes, print the right number while the boxes are on; with the maroon boxes switched off the subsection and every reference to it vanish together, and the subsection after it moves up one number.

Definition 6 is an equilibrium concept, not a theorem, and it is worth recording in one place what the paper proves about it, what it assumes, and what it leaves open. Appendix B.9 sets the system against the graphon mean field games of Caines and Huang (2021) and explains why their existence and ε-Nash theorems cannot simply be cited: the firm's payoff is unbounded, the coupling runs through general-equilibrium prices rather than through an additively separable cost, and the kernel is directed.

The static layer is characterised up to three shape assumptions on the demand-shock distribution. Cross-sectional aggregates are deterministic functionals of the ensemble (Lemma 14, on the exact law of large numbers of Lemma 13, which is cited from Sun (2006) and Bayraktar et al. (2023) rather than proved here); the per-node static block has a solution at every state, unique under Assumption 5, with the threshold increasing in the node's capital (Lemma 15); node-level sales are state-independent under Cobb–Douglas final demand and state-relevant under CES, so that the Cobb–Douglas degeneracy in productivity survives capacity while the closed forms of the model without capacity constraints do not (Lemma 17, Theorem 1, Propositions 4 and 6); and the static block is homogeneous of degree one in the capital profile (Lemma 19). The dynamic layer is characterised conditionally on equilibrium: the household's discount factor is deterministic (Lemma 18), the individual firm's marginal value solves the marginal-q equation (Proposition 7), and the detrended system is Proposition 8. The step-kernel embedding is exact (Lemma 10): for a finite input–output table the continuum system *is* the finite system of Appendix B.10, with no kernel approximation. What is assumed rather than proved falls into two groups. The three shape restrictions, single-peakedness of the pricing problem (Assumption 2), concavity of profit in own capital (Assumption 3) and monotonicity of the threshold map (Assumption 5), are verified numerically for the log-normal specification; and regularity sufficient for the firm's equation and for the symmetric node fixed point, together with the integrability of the node distributions that Lemma 14 needs, is treated as an equilibrium regularity condition rather than derived from primitives.

Two statements are conjectured, not proved, and two questions are left open. The first conjecture is existence: under the standing assumptions a stationary equilibrium in the sense of Definition 6 exists, which for a step kernel with N blocks is a finite-dimensional fixed point in the block moments and the three scalars; the natural route is the Schauder-type argument of the graphon mean field game literature, and what is missing is the compactness and regularity of the per-node problem under unbounded payoffs (Appendix B.9). The second is ε-optimality: that the step-kernel policies, used by every firm of the finite J-sector economy on the same table, are approximate equilibria of that economy, with an error that vanishes as $J \to \infty$ and that contains,

beside the within-block granularity term of the mean-field literature, an aggregate-fluctuation term of order $J^{-1/2}$, because the finite economy carries aggregate risk that the limit policies ignore (Appendix B.9). The first open question is new to the capacity model: without capacity constraints the Cobb–Douglas static block is solved in closed form, so existence concerns the dynamic layer only, whereas under capacity even the Cobb–Douglas static block is a fixed point over two functions on $[0,1]$ whose existence and uniqueness are not established (Appendix B.7.1), and the same fixed point sits inside the CES system at every η. The second is uniqueness, which should not be expected in general, because the monotonicity condition of Lasry and Lions (2007) has no obvious analogue when the mean field enters through prices; multiplicity is probed numerically and reported rather than assumed away (Appendix D).

3.5 Comparison of the Finite and Continuum Models

The computational appeal of the continuum model is that other sectors' states no longer enter a node's problem. This section makes the price of that reduction explicit. Fix a *matched pair*: the finite J-sector CES economy of Section 2 with row-stochastic matrix $\mathbf{A}$, and the CES GMFG on the step kernel of Section 3.1 built from the *same* matrix, with equal blocks $\Delta_n = 1/J$, block-integrated weights $\int_{I_m} a^J(i,j)\mathrm{d}j = a_{nm}$ and block-constant primitives. Appendix B.10 states the three stationary systems in full, the finite one, the general-kernel one and the step-kernel one, reads their backward equations line by line and tabulates the correspondence object by object. This section takes stock of the *mechanisms* that the passage to the limit removes and of what it buys. Levels are not directly comparable across the pair, because step-kernel quantities are densities and the row entropy carries the $-\ln J$ shift of Section 3.1, so the comparison is run on unit-free objects: ι, $\Phi(\iota) - \delta$, prices, q, Domar weights, elasticities.

The analogy of Section 3 supplies the dictionary: finite *sectors* play the role of finite *agents*, the step-kernel GMFG plays the role of the mean-field limit, and the general-kernel GMFG is the limit's closure. The two reasons why a finite agent's value function carries the whole cross-section reappear at the sector level. In the finite J-sector economy, sector i's realised shock Z_i moves aggregate output, consumption and the stochastic discount factor with weight of order $1/J$, and the J sector shocks do not wash out, so the aggregates are stochastic and every sector's policy conditions on the whole realised state $(\boldsymbol{\kappa}, \mathbf{Z})$. In the GMFG a node is atomistic and the exact law of large numbers across nodes makes every aggregate deterministic, so a node's policy conditions on its own (κ, z) only; what remains of the cross-section is the deterministic ensemble of node distributions, coupled through kernel-weighted moments. The one feature with no classical analogue is that network position survives the limit as a persistent type: the mean field becomes *function-valued*, an ensemble indexed by node, and the single market-clearing condition becomes a continuum of consistency conditions, one per node. There is also one respect in which the agent-level analogy must not be pushed: neither the finite nor the continuum economy has strategic interaction or internalisation, because in both the decision unit is an atomistic monopolist within

its sector, so what changes between the two is *statistical*, not strategic.

The objects: firms, industries, blocks. Fix a step kernel with N blocks $I_1, \dots, I_N$. The population hierarchy has three tiers:

$$\underbrace{\text{firm } n \in [0,1]}_{\text{atomistic in its industry}} \subset \underbrace{\text{industry (node) } i \in [0,1]}_{\text{atomistic in the economy}} \subset \underbrace{\text{block } I_m,\ |I_m| = \Delta_m}_{\text{non-atomistic}}.$$

The classical mean-field object, one HJB equation (in marginal-$\mathfrak{q}$ form) paired with one FPK equation in the node's own (κ, z), is indexed by the *node*, that is, by the individual industry. Within block I_m there is a continuum of ex-ante identical industries, each carrying its own essentially pairwise independent productivity shock z_i, and each hosting its own within-industry population of monopolists. Because primitives are block-constant, the i-indexed HJB–FPK pairs are *identical* for all $i \in I_m$, so only one representative pair per block need ever be written down or solved. What is aggregated to block level is therefore not the games but their *moments*, of log prices, spending, sales and utilization: by the exact law of large numbers across the industries of a block, the block's contribution to any coupling functional is

$$\int_{I_m} \mathbb{E}[f(\kappa_i, z_i)]\mathrm{d}i = \Delta_m \, \mathbb{E}_m[f(\kappa, z)], \qquad \mathbb{E}_m[\cdot] := \int \cdot \, \mathfrak{m}_m(\kappa, z)\, \mathrm{d}\kappa\, \mathrm{d}z,$$

which is deterministic. The block distribution $\mathfrak{m}_m$ carries a dual interpretation: it is simultaneously the distribution of a single industry's state across realisations and, by the exact law of large numbers, the *realised cross-sectional distribution* of industries within the block on almost every path.

A sceptical reader may ask why the block is not simply one mean-field game played at the block level. "One classical mean-field game played directly at the block level" would make the block a single industry, with a common shock z_m shared by everything inside I_m, hence stochastic block aggregates and a stochastic discount factor. That is precisely the finite N-sector model, not its GMFG limit. The distinction between the two descriptions is therefore not pedantic: it *is* the $O(N^{-1/2})$ aggregate-risk gap between the finite economy, with one common shock per sector, and the step-kernel GMFG, with a continuum of independently shocked industries per block and deterministic block aggregates.

The atomism involved sits at two of the three tiers. The firm is atomistic within its industry: it takes (K_i, z_i), P_i and Y_i as given and influences only its own k. The industry is atomistic in the economy: its (K_i, z_i) moves no block moment and no aggregate, which is exactly why only own states appear in the node problem, and why the Cobb–Douglas degeneracy bites in the limit but not at finite J. The block is *not* atomistic, since it enters the coupling with weight Δ_m, but no decision is taken at block level, so its non-atomism has no strategic content: blocks only weight the moment integrals.

Reading the backward equations. Set the finite capital-pricing equation (40) against the node equation (52) on a step kernel. The flow is the *same* private capital-income

flow, and the discount, drift and coupling terms are not. The marginal-profit flow $\zeta_n(\overline{v}_n)[1-\xi_n(\overline{v}_n)]\widehat{s}^{\mathrm{f}}_n/\kappa_n$ is identical on the two sides: both firms are priced at the private rental flow (33), the capital share of firm revenue per unit of capital, so the GMFG is the limit of exactly this firm system, and the comparison is pure granularity, with no change in what is internalised. What differs is the argument of the threshold and of the revenue: the finite $(\overline{v}_n, \widehat{s}^{\mathrm{f}}_n)(\boldsymbol{\kappa}, \mathbf{Z})$ come from the static system of Section 2.2 solved on the full realised state, where every sector's price, and hence every sector's rationing margin, responds to every sector's state, whereas the GMFG $(\overline{v}_n, \widehat{s}^{\mathrm{f}}_n)(\kappa, z)$ come from the scalar node block solved against the deterministic aggregates $(\mathcal{P}_n, \widehat{\mathcal{D}}_n, q)$. Without capacity constraints the flow is the constant coefficient $\frac{\varepsilon-1}{\varepsilon}\alpha_n$ times $\widehat{s}_n/\kappa_n$ on both sides (Appendix C.6). The finite equation weights this flow by marginal utility $c^{-\gamma}$ and discounts at the state-dependent rate $\rho+\gamma\mu^K$; rewritten in goods units it discounts at the state-dependent risk-free rate and carries an explicit risk-premium term, and the GMFG discounts at the constant $r=\rho+\gamma g$ with no premium(Lemma 20 in Appendix B.10). The two systems are also coupled differently. The finite system is coupled *pointwise in the state*: at each $(\boldsymbol{\kappa}, \mathbf{Z})$ the scalar c and the drift pair $(\mu^K, \mu_{\boldsymbol{\kappa}})$ tie the J equations together, no distribution appears anywhere, and the ergodic stationary density $\mathfrak{m}_\infty(\boldsymbol{\kappa}, \mathbf{Z})$ is computed downstream from one FPK equation on the shared state. The GMFG is coupled *in distribution*: the N per-block distributions $\mathfrak{m}_n(\kappa, z)$ feed back into the backward equations through the moments of Definition 6, making the stationary problem genuinely forward–backward; but given those moments the N two-dimensional pairs are completely decoupled.

What the limit removes.

(i) *Aggregate risk and risk premia.* At finite J the sector shocks do not wash out: $(\widehat{s}_n, c, \mu^K)$ and hence the stochastic discount factor are stochastic, firm values covary with consumption, and those covariances are priced, explicitly in the risk-premium term of the goods-units form of (40), implicitly through the utility weighting of $\widehat{\mathfrak{q}}_n$ in (40) itself. In the limit the exact law of large numbers (Lemma 13) makes every aggregate deterministic: no aggregate fluctuations, no consumption-smoothing dynamics, no risk premia, discounting at the risk-free rate. The approximation error this induces against a finite economy is of order $J^{-1/2}$, the aggregate-fluctuation term, on top of any graphon-approximation term(Section 3.4). MIT shocks (Appendix D) restore *perfect-foresight* aggregate transitions, but not priced aggregate risk.

(ii) *Own-shock general-equilibrium feedback.* A finite sector has mass $1/J$: its own Z_n moves aggregate output, consumption, the stochastic discount factor and, through the static block, every other sector's prices, and part of the state-relevance of its flow runs through exactly this channel. In the GMFG a node is atomistic: its own (κ, z) moves no moment and no aggregate. The starkest consequence is the Cobb–Douglas degeneracy(Appendix B.7): at $\eta=1$ the finite model remains a nondegenerate stochastic system *only* because of the $1/J$ own-shock feedback, and in the limit that feedback dies, so the node's *stochastic* state-relevance dies

with it: z becomes payoff-irrelevant and the node distributions degenerate to point masses in κ(Theorem 1). Capacity does not change this: at $\eta = 1$ the node's threshold, and with it its utilization and its revenue shares, is free of z(Lemma 17), so the rationing margin opens no channel through which productivity could re-enter the payoff. At $\eta \neq 1$ the CES relative-price channel $\beta_n P_n^{1-\eta} q$ survives as the *sole* source of own-state relevance, which is precisely why this paper insists on CES final demand.

(iii) *State-contingent cross-sector comovement.* Finite policies are functions $\iota_n(\boldsymbol{\kappa}, \mathbf{Z})$: sector n's investment responds to the *realised* state of every other sector, propagation of a given shock through the network is resolved state by state, and the joint distribution $\mathfrak{m}_\infty(\boldsymbol{\kappa}, \mathbf{Z})$ encodes realised cross-sector correlations, comovement and granular concentration dynamics. GMFG policies are functions $\iota^*(n, \kappa, z)$ of own states alone; other nodes enter only through deterministic moments, and across nodes the equilibrium randomness is essentially pairwise independent by construction. First moments of the cross-section survive exactly; realised cross-node second moments, comovement and its asset-pricing and volatility implications, are zero in the limit by fiat.

(iv) *Not lost: strategic interaction and internalisation.* Neither economy has them. In both, the decision unit is an atomistic monopolist that takes its sector's price and quantity, all aggregates and, at finite J, the stochastic discount factor as given, and internalises the effect of its investment on nothing: capital is priced at the private rental flow in both systems, not at the internalised derivative Q_{K_n} (Remark 10 in Appendix B.8). The GMFG firm problem uses the private flow $\zeta_i(\overline{v})[1 - \xi_i(\overline{v})]\hat{s}^{\mathrm{f}}/\kappa$, the continuum analogue of (33), so the GMFG is the large-J limit of the firm-type system (36), which prices capital at the private flow and in the aggregate falls short of the social return by the markup and rationing rents (Remark 1); that wedge is a property of the static block and passes to the limit node by node, untouched by the comparison. Sectors as units take no decisions in either model, so there is no oligopolistic margin to lose. What changes between the two models is *statistical*, not strategic: which objects in the generic firm's environment are random, and which states its policy is allowed to condition on.

What the limit buys. The corresponding gains: dimension collapses from J coupled PDEs on a $(2J-1)$-dimensional state, exponential in J, to N decoupled two-dimensional pairs plus a moment fixed point, linear in N; the discount factor is deterministic, so no marginal-utility weighting or consumption map enters the state; and the limit adds an object the finite model does not have, a nondegenerate *within-block cross-section*. By the exact law of large numbers $\mathfrak{m}_n(\kappa, z)$ is the realised distribution of the block's continuum of industries, delivering stationary dispersion, firm-size distributions, the dispersion of utilization and of the rationed share of firms across the block's industries, and within/between decompositions, whereas the finite sector sits at a single realised point at each date and only its time-average is available.

Rule of thumb. Moving from the finite model to the GMFG trades the equilibrium's entire *covariance* structure, that is, aggregate fluctuations and their pricing, own-shock general-equilibrium feedback and state-contingent comovement, for a moment-coupled system whose cost is linear in the number of blocks and which resolves the cross-sectional *distribution* exactly. Questions about means, dispersion, network incidence and balanced-growth objects are the GMFG's home turf; questions about aggregate volatility, risk premia, comovement, or the general-equilibrium footprint of an individual sector's shock require the finite model, or MIT-shock transitions as the deterministic middle ground.

Exactness and the two error terms. Two distinct approximations are easily conflated, and the step kernel makes exactly one of them. Representing a finite input–output table by a kernel is exact: by Section 3.1, Lemma 10 and (90) the step kernel reproduces the table's block weights, its Leontief inverse and its entropy up to the bookkeeping shift, and both the sectional and the cut distance between the table and its kernel are zero (Appendix B.2.1), so the graph-to-graphon approximation error that the graphon mean field game literature controls by a sectional-convergence hypothesis is identically zero here. What separates the step-kernel GMFG from the finite J-sector economy on the same table is therefore a single thing, granularity: in the finite economy each sector is one industry carrying one common shock, so aggregates are stochastic and each sector's shock moves them with weight of order $1/J$, whereas in the step-kernel GMFG each block contains a continuum of independently shocked industries, so block moments are deterministic and a block's state moves no aggregate. The gap is the aggregate-fluctuation term of the ε-optimality conjecture, the $O(J^{-1/2})$ sampling noise of J independent sector shocks that the limit policies ignore, together with the within-block granularity term familiar from the mean-field literature; under Cobb–Douglas final demand the gap is not a correction term at all but the entirety of the finite model's dynamics, since the step-kernel Cobb–Douglas GMFG is the deterministic skeleton that the finite Cobb–Douglas economy fluctuates around (Proposition 5). The general-kernel GMFG is one step further removed from any finite economy: it requires no approximation of the model either, but if it is to be read as the limit of a growing sequence of finite networks, the convergence has to be sectional, not merely in the cut metric, and the kernel has to satisfy the conditions of Appendix B.2.2.

The one-sentence comparison. Splitting the finite equilibrium into J firm-type equations does not by itself reduce dimensionality: each equation still lives on the full state because a sector's shock moves aggregates with weight $\sim 1/J$; the GMFG's collapse to N two-dimensional pairs comes entirely from industry-level atomism plus the exact law of large numbers, which turn every cross-industry object a firm cares about from a stochastic process on the shared state into a deterministic moment of the ensemble, at the price of replacing the sector's general-equilibrium own-shock feedback, which kept even the Cobb–Douglas finite model non-degenerate, with the CES own-price channel as the sole source of own-state relevance.

4 Data and Calibration

The model is calibrated to the industries of the United States at a yearly frequency. Appendix E gives every source, every mapping from the accounts to the model's objects and every step of the computation; this section states what is matched to what and defines the economy against which every result is read.

Nodes. A node is a summary industry of the input–output accounts of the Bureau of Economic Analysis. The five government industries and housing are removed, their purchases counted as their suppliers' final demand, and the three finance industries are merged into one node, which leaves $N = 63$ nodes of equal mass. The network a_{nm} is the matrix of domestic direct requirements, each buyer's shares rescaled to sum to one; the final-demand weights β_n match the nodes' shares in final expenditure; depreciation δ_n is depreciation over the net stock of equipment and structures. All accounts are means over 2007–2019. Labour is outside the model and the model's capital is physical capital, equipment and structures. Of a node's value added the model carries the operating surplus of physical capital, net of the imputed pay of the self-employed, and the ratio of that surplus to surplus plus domestically bought intermediate inputs, ω_n, is the accounts' counterpart of the model's capital share of sales.

The capacity block. Two parameters govern how often capacity binds and how much it matters when it does, and two pieces of evidence on US manufacturing identify them. The dispersion of firms' demand shocks, σ_v, is set so that the share of manufacturing capacity in plants operating above 95 per cent of capacity equals its mean in the Census Bureau's Survey of Plant Capacity, 16.2 per cent. Given σ_v, the elasticity of substitution between firms, ε, is set so that the model's industry supply curves are as steep as those that Boehm and Pandalai-Nayar (2022) estimate: the inverse supply elasticity by bin of lagged utilization, which rises from 0.12 in the slackest bin to 0.57 in the tightest. The model's counterpart of each estimate is computed on the solved model by moving each manufacturing node along its own supply curve through the industry-years of the authors' sample, with the estimator's own weights, and ε minimises the χ^2 distance to the four estimates with their covariance. The two moments identify the two parameters separately, and Appendix E.5.1 shows why: the Census share responds to σ_v and hardly to ε, the supply curves to ε given σ_v.

utilization targets. A node's capital intensity x_n is set so that its mean utilization on the balanced growth path equals a target. In the baseline calibration, the calibration of this paper, capacity is on in the 49 nodes for which a utilization rate exists and off in the 14 without one. The 23 core nodes (manufacturing, mining, utilities) are at the long-run means of the Federal Reserve's capacity utilization rates, 1972–2019, and 26 service nodes at rates transferred from the European Commission's business surveys through the European ratio of the service to the manufacturing rate of the same industry. A

second calibration in which every node is as tight as manufacturing is described in Appendix E and is deferred to the robustness analysis.

Capital shares, the interest rate and the rest. The capital exponent α_n is the value at which the model's split of a node's sales at its utilization target reproduces the accounts' ratio ω_n; where no such value exists above 0.02 the exponent is floored and the node is flagged (12 of the 49 nodes). The interest rate is the value at which the model reproduces the capital stock over sales of all nodes, 1.356, and $\rho = r - \gamma g$ follows. The growth rate $g = 0.021$ is that of the nodes' real value added; $\eta = 0.5$ and $\gamma = 1$, log utility, come from the literature and convention; the productivity process within a node is measured on the detailed manufacturing industries of the NBER–CES database with the model's own weights. The installation curvature is set at $\phi = 2$ by convention, the order of the convex adjustment costs that structural estimates find; the regression of capacity growth on lagged utilization that the Federal Reserve's data allow (coefficient 0.167 over 21 manufacturing industries, 1958–2021) is the moment that would discipline it inside the model, and Appendix E says what is known about it. Table 1 lists the parameters.

Table 1: The parameters of the baseline calibration.

	meaning	how it is set	value
N	number of nodes	BEA summary industries	63
a_{nm}	cost share of m in n's bundle	domestic direct requirements	matrix
β_n	weight of n in final demand	share in final expenditure	63 targets
δ_n	depreciation	equipment and structures	0.029 to 0.150
g	growth rate	real value added, 2007–2019	0.021
η	elasticity across nodes in final demand	literature	0.5
γ	relative risk aversion	convention	1 (log utility)
Corr, Var	productivity process within a node	NBER–CES industries	0.97, 0.057
σ_v	dispersion of demand shocks	share of capacity above 95 %, 0.162	0.185
ε	elasticity between firms' goods	supply curves by utilization bin	9.6
α_n	capital exponent	ω_n at the node's threshold; floor 0.02	0.065 (mean in manufacturing)
x_n	capital intensity	mean utilization, 49 targets	0.72 to 0.93
r, ρ	interest rate, discount rate	capital stock over sales, 1.356	0.039, 0.018
ϕ	installation curvature	convention; structural estimates of adjustment costs	2

Notes: Appendix E, Table 10, gives the sources and the targets in full.

[Flag, 3 November 2026: the baseline calibration is now $\phi = 2$, $\gamma = 1$] Until 3 November 2026 the baseline calibration had $\phi = 13.0$ from the simulated regression and $\gamma = 2$ ($r = 0.033$, $\rho = -0.009$). Kian's decisions of 31 October and 2 November 2026: $\phi = 2$ by convention with ε and σ_v re-estimated by the outer loop (9.614 and 0.1851 against 9.604 and 0.1845; the capacity block does not identify ϕ), then log utility

with $\rho = r - g = 0.018$, run first as trials and elevated on 3 November 2026. What moved against the previous baseline: the interest rate 3.3 $\rightarrow$ 3.9 per cent and ρ from -0.009 to $+0.018$; Tobin's q from 1.8–6.7 to 1.05–1.17; the growth gap 0.19 $\rightarrow$ 0.60 points (equal to the interest gap under log utility) and the welfare gap 26.5 $\rightarrow$ 37.6 per cent; the capital reallocation along a transition four times larger and twice as fast, the GDP paths unchanged, the consumption response a third smaller; the targeted expansion 2.3 $\rightarrow$ 1.7 per cent with capacity and 2.5 $\rightarrow$ 1.5 without, so the comparison of Section 7 reverses. Unchanged to the printed digits: the capital shares (ten cells of the node table by one unit in the third decimal), the floors, utilization, the rationed share, the rent, the markups, the impact and sign-size responses. Not re-run on the new baseline: the comparative statics in ε and σ_v (the notes box at the end of Section 5, and the identification figure of Appendix E), the fixed-threshold static, the simulated-panel fit numbers, and the ϕ moment itself (about nineteen hours; whether to compute it at $\phi = 2$ is Kian's decision, to be asked later). The abstract, the conclusion and the talk are not updated.

Fit. The targeted moments are reproduced exactly except the four supply-curve estimates, which the model matches at $\chi^2 = 1.1$ with three degrees of freedom, and the intermediate shares of the floored nodes, which miss by about 4 points at the median. Three untargeted moments say how the capacity block behaves. The dispersion of utilization across the firms of a manufacturing node is 0.13, between the within-industry and the whole-manufacturing dispersions measured on Swedish plant data (Amberg et al., 2026); the share of manufacturing firms that are rationed at any time is 10 per cent, against about 8 per cent of European manufacturers who name equipment or material as the factor limiting their production; and the slope of the share of capacity above 95 per cent on utilization is 1.7 in the model against 0.74 in the data, a miss that one σ_v cannot repair. On the balanced growth path the rationing rent is 0.9 per cent of GDP, the aggregate markup 1.155 against 1.116 without capacity, and value-added-weighted utilization 0.74. The model's capital stock over sales matches the accounts in total and misses by group: with one interest rate and one markup the model cannot reproduce the spread of recorded returns across industries, and it is calibrated to the total.

The comparison economy. Every result is read against the *economy without capacity constraints*: the calibrated economy with its capacity constraint switched off, that is, the limit $x_n \rightarrow 0$ of Section 2.1 at the same a_{nm}, β_n, α_n, δ_n, ε, η, γ, ρ and productivity process. It is the standard production-network model with capital that the calibrated economy nests, so the difference between the two economies is the effect of capacity at a given technology, preferences and network. Because it keeps ρ, its growth rate and its interest rate differ from the calibrated economy's. In tables and figures it is the *no-capacity economy*. Where it is informative a second comparison is added: a no-capacity economy recalibrated to the same accounts, which fits them with different capital shares. It says what a network model without capacity, fitted to the same data,

would predict, and it is called the *recalibrated economy*.

5 The Long-Run Cost of Capacity

[Q8 | mynotes | tag: AAAA | ANSWERED] I want to talk about the return to capital, interest rate and growth. You don't have to answer my questions below in order that I ask them, but rather give me a well-structured (and well thought-out!) text that answers them (and my general concerns about this whole business) in an order that you best see fit. You can refer to the theory (equations of model) or data/code to motivate your answers.

1. What is the economic intuition/mechanism behind lower interest rate and growth rate result of the main calibration? Why is the return to capital lower? Because, very naively, shouldn't capacity constraints make capital more important (in order to build "buffer capacity"), and thus lead to higher investment? or is the lower return to capital purely a network thing (which would be great because right now my supervisors think that I should have more on the Network part of "Capacity utilization in Production Networks". That is, as they see it, right now my paper has a lot on capacity and its costs and effects but doesn't talk enough about the role that the network plays in all this)? A bit more precise questions also: (1) does this lower return to capital and lower investment hold for *all* parameters or could there be, potentially, some parameter constellations where it's the opposite (higher return and higher growth)? (2) how about a one-sector economy? If you cannot answer this question precisely without solving an actual one-sector economy, then *don't* solve a one-sector economy: either ansewr with intution that you strongly think is correct or, if really unsure, just don't answer.

2. Related to this perhaps, a long time ago while doing a very preliminary (now deleted) calibrations I actually found that in that calibration capacity mildly amplified aggregate resonse to upstream shocks and if I rememeber correctly that it also increased growth. I think this "opposite" result – when I thought about it – came down to the balance between two forces whose magnitude would be decided by the parameters. The first force was that capacity by itself increases the return on installed capital and partially reverses the under-investment in capital under monopolistic competition and this effect is linear in ε (and/or σ_v...I don't remember exactly *what* it is linear in). The second force was that rationing destroys surplus: demand turned away is a wedge between posted price and shadow price of node's good and this one is convex in tightness: very small at first but then rising sharply (Gaussian tail). And thus, as it's normal in economics, when you have two economic forces that are opposite and one is linear and another is convex, you usually get a hump-shaped graph where things reverse based on the exact parameter values. Another note that I had written about this but I struggle to fully make sense of right now was "deadweight loss of rationing scales with how imperfectly-rationed demand is redirected to other firms" (seems to be just a part of the explanation). So I am very curious if a more rigorous version of the mechanism above makes sense and is actually correct, and if so if it can be (or already

is) shown with the model equations and/or some graphs. If this "hump-shaped" result is correct in some form I think it could be a very nice result to have in the paper.

Can you please write your answer in a private box at the very beginning of section 5 (here)? I actually first wanted to ask you to give your answer to me in chat but then I thought maybe you'd like to refer to or write some equations or maybe even show some graphs and that's why I think it's best to give the answer in tex.

[The return to capital, the interest rate and growth: the mechanism, the numbers, the sign, the hump. 29 October 2026] [Flag, 3 November 2026: written on the previous baseline, $\phi = 13$ and $\gamma = 2$; its numbers are that calibration's and are not updated. On the baseline since 3 November 2026 ($\phi = 2$, $\gamma = 1$) the interest gap is 0.60 points and equals the growth gap, Tobin's q is 1.05–1.17, and the accounting of item 2 is to be redone.] This answers the box above (question tagged AAAA) in the order of the argument rather than of the questions: (1) the equation that ties the growth rate and the interest rate to the return to capital; (2) why that return is lower in the baseline calibration, with the exact accounting; (3) whether this is a network result and what a one-sector economy would show; (4) whether the sign holds for all parameters, with everything computed so far that bears on it; (5) the "two forces" and the hump, made precise; (6) two corrections to the current text and the computations that would settle what is open. Every number is re-read from the stored output of the long-run block, of the sweep at fixed targets and of the static at fixed thresholds, or from the two-node economy of Appendix D; the few objects computed for this box from those outputs are defined where they appear. Nothing was re-solved. The basis of each claim is stated.

1. The interest rate and the growth rate are one fact, and both follow the private return to capital.

- The two economies share ρ and γ, and on a balanced growth path the household's Euler equation is $r = \rho + \gamma g$ (145); the stored interest rates equal it to the last digit. So the lower interest rate is not a second result: $r_{\text{cap}} - r_{\text{twin}} = \gamma\,(g_{\text{cap}} - g_{\text{twin}}) = 2 \times (-0.19) = -0.37$ points, the table's number. The question is why g is lower.
- Growth is set by the return that firms earn on capital. On a deterministic balanced growth path the capital-pricing equation (40) at its rest point (equivalently the backward equation (52) with no time dependence) reads, node by node,

$$(r-g)\,\tilde{q}_i = R_i - \iota_i, \quad \Phi_i'(\iota_i)\,\tilde{q}_i = 1, \quad \Phi_i(\iota_i) = g + \delta_i, \quad R_i := \frac{\pi_i^k K_i}{K_i} = \zeta_i(\overline{v}_i)[1 - \xi_i(\overline{v}_i)]\,\frac{S_i^{\text{f}}}{K_i},$$

  where R_i is capital income per unit of the node's capital, the private flow of Lemma 5(ii), and $\tilde{q}_i$ is Tobin's q. Without installation costs ($\phi_i \to 1$, so $\tilde{q}_i = 1$ and $\iota_i = g + \delta_i$) it is the textbook $r = R_i - \delta_i$. At $\phi = 13$ Tobin's q is 1.8 to 6.7 across nodes (median 2.8), so the return is diluted by the installation cost, but the logic is unchanged: the required return rises with g (through r, ι_i and $\tilde{q}_i$), the return the technology delivers, R_i, is scale-free in this AK economy, and g is where the

two meet. *Higher return, higher growth*: on the baseline calibration the local slope is about 3 basis points of g per one per cent of the aggregate return (computed for this box from the equation above with the stored δ_i, ϕ and capital shares).

- **Check (numerical).** Aggregating the equation with the capital shares, $\sum_i \kappa_i[\iota_i + (r-g)\tilde{q}_i]$ at the stored (g, r) is 0.1814 per unit of aggregate capital in the capacity economy and 0.1904 in the economy without capacity constraints; the accounting identity of item 2 gives the same two numbers to four decimals, and node by node the ratio of the two economies' required returns (0.93 to 0.96, the depreciation rates differ) equals the ratio of their capital incomes per unit of capital. The stationary solution carries the productivity risk and the node distributions, so the equality is not forced: the deterministic reading is exact in the aggregate on the stored economies.

2. Why the return is lower: the exact accounting.

- Sum the four-way split of sales (20) over nodes and use $Q = \sum_i (S_i - X_i)$ (the display of Remark 1): the aggregate private return per unit of capital is GDP per unit of capital net of the two rents that do not price capital,

$$\bar{R} := \frac{\sum_i \pi_i^k K_i}{K} = \bar{q} - \frac{1}{\varepsilon}\frac{\sum_i S_i^{\mathrm{f}}}{K} - \frac{\Pi^{\mathrm{r}}}{K} = \bar{q}\left[1 - \frac{\mathcal{D}^{\mathrm{f}}}{\varepsilon} - \frac{\Pi^{\mathrm{r}}}{Q}\right], \qquad \mathcal{D}^{\mathrm{f}} := \frac{\sum_i S_i^{\mathrm{f}}}{Q},$$

  with $\mathcal{D}^{\mathrm{f}}$ the Domar sum of firm revenue. The markup rents are the share $1/\varepsilon$ of firm revenue at every threshold (Lemma 5(iii)); they are two fifths of GDP here only because revenue is four times GDP. The rationing rent is the aggregators' dividend. Neither prices capital.

- On the baseline calibration, per unit of aggregate capital (the long-run block; the split by the identity):

per unit of aggregate capital	with capacity	without	difference
GDP, $\bar{q}$	0.3150	0.3488	-0.0338
markup rents, $\sum_i S_i^{\mathrm{f}}/(\varepsilon K)$	0.1307	0.1584	-0.0277
rationing rent, Π^{r}/K	0.0029	0	$+0.0029$
capital income, $\bar{R}$	0.1814	0.1904	-0.0090
capital income as a share of GDP	0.576	0.546	$+0.030$
firm revenue, $\sum_i S_i^{\mathrm{f}}/K$	1.256	1.522	-19 log points
capital income as a share of firm revenue	0.144	0.125	$+14$ log points
growth rate g	2.10 %	2.28 %	-0.19 points

- **Reading.** Capacity removes 3.4 cents of GDP per unit of capital. Of these, 2.8 cents come out of the markup rents, because firm revenue per unit of capital falls by 19 log points (gross output falls by more than GDP, the chains shorten, and a slice of sales is diverted to the rationing rent); 0.3 cents go to the aggregators as that rent; 0.9 cents come out of capital income. The return falls by 4.7 per cent while GDP per unit of capital falls by 10: *the loss falls mostly on the rents, which is*

why the growth effect (0.19 points) is small next to the level effect. At the local slope of item 1 the 4.7 per cent account for 15 of the 19 basis points; the remainder is that the two economies hold their capital in nodes with different depreciation rates and Tobin's q (computed for this box from the same equation).

- **Both halves of the intuition in the question are right, and they are the same fact.** Per euro of firm revenue, capital *is* worth more with capacity: its share of firm revenue rises from 0.125 to 0.144 in the aggregate, and by node from parity in the capital-heavy nodes to 2.7 to 4 times the no-capacity share $\frac{\varepsilon-1}{\varepsilon}\alpha_n$ in the tight, capital-light service nodes (the first channel of (76), relaxing the rationing, which has no counterpart without capacity). Firms respond as the intuition says: they hold more capital per unit of sales, 0.79 against 0.66 in the aggregate, two and a half times the economy without capacity constraints's stock in the floored nodes as a group and up to 3.5 times in the tightest service nodes. That is the buffer. But the buffer runs at 74 per cent and relaxes a constraint in one state out of five, so the revenue a unit of capital generates falls by as much as its share rises: by node, the ratio of the capital shares of revenue (with over without) times the ratio of firm revenues per unit of capital (with over without) is 0.93 to 0.96 in every capacity-on node, which is the ratio of the two required returns of item 1, as the equalisation of returns across nodes demands. The buffer shows up in the capital–output ratio, a level; growth is set by the average return per unit of capital, which the buffer lowers. "Capacity makes capital more important" is true of the marginal unit at a given revenue and false of the average unit once the firm has chosen its buffer, and growth follows the average. (By node: computed for this box from the node table of the long-run block; the aggregates are the table above.)

3. Is it the network? Not for the sign; yes for the size and the incidence.

- *What is within a node.* Every term of the identity moves inside a node: the truncation of output at capacity and the idle capital lower revenue per unit of capital, the lower effective elasticity raises the markup and cuts the bundle bought per euro of revenue, and the rationing rent leaves the firm. A one-sector roundabout economy, one node buying its own good as its bundle with the same intermediate share (so with a Domar sum near 4, as here), has all of them. At the calibrated tightness its GDP per unit of capital would fall, its firm revenue per unit of capital would fall by more, and its return would fall by the difference net of the cushion; the sign would be the baseline calibration's. *Basis: reasoned from the identity and from the firm-level anatomy; not solved, as the question asks.* So the answer to the first question is the uncomfortable one: the lower return, the lower growth rate and the lower interest rate are not a network result.
- *What the network does.* (a) The size of the loss at a given allocation of capital: at the economy without capacity constraints's allocation the wedges cost 38 log points, at the capacity economy's own 8 (Table 3); the reallocation across nodes that recovers 28 is the network's response and has no one-sector counterpart. (b)

The convexity: the cost is convex in a node's shortfall of capital, so at a given aggregate utilization an economy whose nodes run at 0.72 to 0.93 loses more than a one-sector economy at the mean (Jensen; reasoned). (c) The incidence: where the private return rises per euro of revenue and where capital goes, the 6 per cent reallocation and the 2.5-fold stock of the floored service nodes, is a cross-sectional statement with no one-sector counterpart. (d) *Not* network-specific, to be clear about it: the cushion of item 2 comes from the intermediate share, which a one-sector roundabout economy with the same share also has.

- The line this supports, for the point the supervisors made: *the sign of the growth and interest effects is a capacity result; their size, their incidence across sectors and most of the level loss are network results* (the 38 against the 8, the reallocation, the upstream gradient of the dampening). The number that would make the sentence precise is the no-network end of the same calibration, the input-output matrix dialled towards the identity with the utilization and growth targets re-matched at every setting.

4. Does the sign hold for all parameters? No. Everything computed so far that bears on it.

economies compared	what moves	$g_{\text{twin}} - g_{\text{cap}}$, basis points	$\ln(\bar{q}_{\text{twin}}/\bar{q}_{\text{cap}})$	evidence
baseline calibration and its twin	—	$+18.6$	$+0.102$	reproduced
sweep at fixed targets, ε from 7.5 to 21	price of rationing	$+39$ to -0.4	$+0.246$ to -0.007	once; both cross between 17 and 21
sweep at fixed targets, σ_v from 0.105 to 0.22	mass of rationing	$+1.6$ to $+35$	$+0.015$ to $+0.170$	once; no crossing
static at fixed thresholds, ε from 7 to 30	both, same direction	$+9$ to $+56$	$+0.056$ to $+0.182$	once; no crossing
two-node economy, every x_n scaled by s	tightness	-0.5 at $s = 0.71$; positive past $s = 0.78$	-0.0007 at $s = 0.5$; positive past $s = 0.62$	four grids

- Along ε at the data's utilization the capacity economy's g is held at the data's 2.10 (it is re-fitted at every point) and the economy without capacity constraints's falls from 2.49 per cent at $\varepsilon = 7.5$ to 2.09 at 21; the gap in g and the gap in $\bar{q}$ cross zero together between 17 and 21, both above the 95 per cent interval of ε. So yes, there are parameter constellations in which the capacity economy grows faster and produces more per unit of capital, and the stored sweep contains two of them. The effect there is under half a basis point in g and under one log point in $\bar{q}$; the loss region is where the numbers are large. The crossing is flagged in the box after the sweep figure (grid unchecked).
- Along σ_v the mass of rationing alone produces no crossing down to a rationed share of 0.07: the gap shrinks towards zero from above. At fixed thresholds nothing crosses either, because a higher ε there raises the mass faster than it lowers

the price.

- The two-node economy of Appendix D ($\varepsilon = 3$, capital exponents 0.5 and 0.45, utilization 0.80 in both nodes; far from the baseline calibration) is the hump remembered in the question, measured. Scaling every capital intensity by s from the slack limit, *a little capacity is expansionary*: g is 4.8×10^{-5} above its no-capacity value at $s = 0.71$ (aggregate utilization 0.71) and $\ln \bar{q}$ is 7×10^{-4} above at $s = 0.5$; the two cross back at $s = 0.78$ and $s = 0.62$, and past the calibrated point the loss rises steeply (at $s = 2.4$, utilization 0.90, g falls to 1.2 per cent and the aggregate markup reaches 2.3). The hump was four times the solver's sensitivity to the capital window and was reproduced to four significant digits on doubled grids. So the earlier finding was real; it is tiny, and it sits on the slack side of the calibrated tightness. Whether the 63-node economy has one has never been computed (item 6).

5. The two forces, made precise, and why a hump can arise.

- *Force (a), "capacity raises the return on installed capital".* In the model this is the first term of (76): $[1 - G_i(\overline{v}_i)] \cdot (\mathcal{P}_i K_i / x_i) \cdot [p_i Z_i x_i^{\alpha_i} / \mathcal{P}_i - 1]$, the probability of being rationed, times the bundle a unit of capital can process, times the margin on a unit produced at capacity. It is a tail probability times a bounded margin: tail-order in tightness, linear in nothing, and the margin is larger the larger the markup, so it falls with ε rather than rises. The aggregate counterpart, which is where "linear in ε" belongs, is the cushion of item 2: capital's share of GDP, $1 - \mathcal{D}^{\mathrm{f}}/\varepsilon - \Pi^{\mathrm{r}}/Q$, rises because the markup rents' share falls, and that term is linear in $1/\varepsilon$. "Capacity partially reverses the under-investment of monopolistic competition" is right in exactly this form: the wedge between the private and the social return in Remark 1 is the rents, and capacity shrinks the markup rents relative to GDP.
- *Force (b), "rationing destroys surplus, convex in tightness".* The loss in $\bar{q}$ at a given allocation. It is the product of two tail objects, the mass $1 - G_i(\overline{v}_i)$ and the price of rationing, $(p_{in} + \varphi_{in})/p_{in} = (v_{in}/\overline{v}_{in})^{1/\varepsilon}$ (the note in the question, "how imperfectly the turned-away demand is redirected", is this price), and the price compounds through the network because the redirected demand of a node's customers meets rationing again. On the two-node economy the loss vanishes faster than any power of s as capacity slackens, the Gaussian tail.
- *Why the sign can turn.* At the onset both forces are tail-order, so the sign at the onset is a competition of coefficients, not of functional forms. By the identity, $\Delta\bar{R} = \Delta\bar{q} + \frac{1}{\varepsilon}\left|\Delta \sum_i S_i^{\mathrm{f}}/K\right| - \Pi^{\mathrm{r}}/K$: the return rises even while GDP per unit of capital falls if the cushion exceeds the loss, roughly if $\mathcal{D}^{\mathrm{f}} m/\varepsilon > 1$ with m the ratio of the gross-output loss to the GDP loss ($m \approx 1.9$ on the baseline calibration). On the baseline calibration the cushion is 0.82 of the loss, so the return falls; at $\varepsilon = 3$ the markup rents are a third of firm revenue, so the cushion is of the order of the loss and the sign at the onset is open, and in the two-node economy it came out positive. Past the onset the loss grows faster than the cushion, because its two tail factors both rise as the threshold falls and the price factor compounds,

while the cushion is proportional to one of them: the hump. *Basis: the identity is exact; the ordering of the two sides past the onset is reasoned, not derived.* The two-node economy also had $\ln \bar{q}$ itself above its no-capacity value at mild tightness, a static gain this argument does not produce. A candidate: with prices set before the draw and diminishing returns to the bundle at given capital, the economy without capacity constraints's high-draw firms serve demand at a marginal cost that grows like $v^{\alpha_i/(1-\alpha_i)}$ while the price is fixed, so truncating the highest draws moves the ex-post allocation of output across firms towards the planner's when $\alpha_i/(1-\alpha_i)$ is large relative to the love of variety $1/\varepsilon$; at $\alpha = 0.5$ and $\varepsilon = 3$ the economy without capacity constraints's firms sell at a marginal loss in about one state in a hundred, at $\alpha = 0.065$ and $\varepsilon = 9.6$ in none. Reasoned only; it would need its own derivation.

6. Two corrections to the text (the first applied on 30 October 2026), and what would settle the open points.

- The comparative-statics subsection (since 30 October 2026 a notes box at the end of this section) explained the crossing at high ε by "the rents earned by rationed firms raise the return to capital and with it the rate of accumulation" until 30 October 2026, when the sentence was replaced by a statement of the facts (the size of the gain in $\bar{q}$ and in g, the rents' recipient, the direction of causation, and that the static channel is not decomposed). Two things were wrong as stated. The rationing rents are not earned by the firms (they are the aggregators' dividend, (19)); what a rationed firm earns on a marginal unit of capacity is the posted-price margin of (76). And the comparison is per unit of capital, so faster accumulation cannot raise $\bar{q}$: at $\varepsilon = 21$ GDP per unit of capital is itself higher with capacity, a static fact about the allocation and the wedges, and the growth gain follows from it through the identity, not the other way round. Which static channel produces it is not derived; the text now says so.
- Section 5.1 says the rates are lower "as the capacity constraint lowers the return to capital that accrues to the firm". That is right, and item 2 is what it rests on; a clause could add that the loss falls mostly on the rents, which is why the growth effect is small next to the level effect.
- *What would settle the hump on the* 63*-node economy:* the tightness sweep, every x_n scaled by a common s over about twelve points from 0.25 to 2.5 with nothing else moved, about one to two hours of compute. It gives the loss, the growth gap and the dampening against aggregate utilization and against the rationed share, hence the onset and the crossing, if any, at the calibrated dispersion; it is the graph the question asks for. Its economies are not fitted to the accounts, so it would sit in a box or an appendix like the static at fixed thresholds. *What would settle the sign at high* ε*:* the two fixed-allocation solves of Table 3 at the $\varepsilon = 21$ point of the sweep (a re-fit at that point, then seconds), which say whether the gain sits in the wedge term or in the reallocation term.

[Evidence status of this section, 3 November 2026] Every number is the solved

model on the baseline calibration at $\phi = 2$ and $\gamma = 1$, computed once on 2–3 November 2026 by the results chain on the stored calibration; the capacity economy's balanced growth path is the one solved on 31 October 2026 at $\gamma = 2$, reproduced by a second route (warm from that solution instead of cold) to seven digits, since the stationary economy sees γ and ρ only through r. The recalibrated economy is computed once. The comparative statics (the notes box at the end of this section) and the identification figure of Appendix E were computed on the previous baseline ($\phi = 13$, $\gamma = 2$) and are not re-run; every number inside that box is that calibration's.

Sections 5 to 7 report the results. This section compares the balanced growth paths with and without capacity: the cost of capacity and its sources (Section 5.1). Section 6 turns to shocks: their propagation on impact, the role of their sign and size, and two of them over time. Section 7 asks in which sectors capacity is worth most. Throughout, "without capacity" means the economy without capacity constraints of Section 4, the calibrated economy with its constraint switched off, so that every comparison measures the effect of capacity, not of the network.

[Flag: the comparative statics are set aside] Until 30 October 2026 the first sentence of this paragraph continued ", and what each capacity parameter controls" and pointed to a second subsection, the comparative statics in ε and σ_v. That subsection is now a notes box at the end of this section, kept unchanged, to be looked at more closely later.

5.1 Cost and Allocation of Capacity

Table 2 compares the balanced growth path of the calibrated economy with that of the economy without capacity constraints. The two economies have the same technology, preferences and network, differing only in whether the capacity constraint is on. Every row of the table therefore reads as the long-run effect of the capacity constraint. On the balanced growth path of the calibrated economy, value-added-weighted utilization is 74%, about one firm in five is rationed at any time, and the scarcity rents on rationed goods amount to 0.9% of GDP. Relative to the economy without capacity constraints, GDP, consumption and gross output per unit of aggregate capital are lower. The decrease in gross output is more than GDP because rationing cuts the flows of intermediate goods between nodes. Specifically, the share of intermediate inputs in gross output falls from 0.77 to 0.75. In this sense capacity shortens the production chains. The reason is that, of every euro a buyer spends on an intermediate good, less becomes spending on that good's own inputs than without the constraint: part of the euro is the scarcity rent on rationed goods, which the aggregators collect and which buys no inputs, and the firm that receives the rest prices at a wider wedge over its input cost, because a higher price loses it no sales in the states in which it is rationed. Since this cut is taken at every stage a purchase passes through, the chain of input purchases behind a unit of final demand dies out sooner, and gross output becomes a smaller multiple of GDP. Furthermore, the growth rate and the interest rate are lower as the capacity constraint lowers the return to capital that accrues to the firm, so the

households accumulate less. Finally, with capacity, about 7% of aggregate capital sits in different nodes than without.[22]

Table 2: The long run with and without capacity.

	baseline calibration
the capacity economy	
value-added-weighted utilisation $\overline{U}$	74.3 %
share of firms rationed (value-added weighted)	18.6 %
rationing rent / GDP	0.92 %
aggregate markup, with capacity / without	1.155 / 1.116
the capacity economy against the no-capacity economy	
GDP per unit of capital	−10.1 %
consumption per unit of capital	−13.4 %
gross output per unit of capital	−17.7 %
welfare gap, consumption equivalent	−37.6 %
of which from the level of consumption	−13.4 %
of which from the growth rate	−28.0 %
growth rate	−0.599 p.p.
interest rate	−0.599 p.p.
share of aggregate capital in different nodes	6.6 %

Notes: The comparison block is in percentage differences with the no-capacity economy as the base.

The long-run GDP gap has two sources, which are decomposed in Table 3. Firstly, capacity puts a set of wedges on every link of the network at a *given* allocation of capital across nodes: rationing, state-dependent markups and idle capacity. The second is the reallocation of capital across nodes itself, which the economy with capacity constraints undertakes in response to the constraint. The latter leads to some sectors accumulating a higher or lower share of aggregate capital than the counterfactual scenario without capacity constraints. These two can be decomposed in the model with a counterfactual static solve: we take the economy without capacity constraints and move its capital across nodes to the capacity economy's allocation, and re-solve its static general equilibrium at that allocation. The resulting GDP is what the economy without capacity constraints would produce if it held its capital where the capacity economy holds it. At this allocation, GDP per unit of aggregate capital is 2.5 log points below the no-capacity economy's own GDP. The remaining 8 log points of the gap are the cost of the wedges at the capacity economy's allocation of capital.

The two parts add up exactly because the gap is a ratio. Write $q^{\text{cap}}(\boldsymbol{\kappa})$ and $q^{\text{no cap}}(\boldsymbol{\kappa})$ for GDP per unit of capital of the two technologies at capital shares $\boldsymbol{\kappa}$, the function q of (31), and $\boldsymbol{\kappa}^{\text{cap}}$, $\boldsymbol{\kappa}^{\text{no cap}}$ for their balanced-growth-path allocations. Inserting the

[22]This is calculated as $\frac{1}{2}\sum_n \left|\kappa_n^{\text{cap}} - \kappa_n^{\text{no cap}}\right|$, with $\kappa_n = K_n/K$ node n's share of aggregate capital on the balanced growth path (30): the share of the economy's capital that would have to change node to turn the no-capacity allocation into the capacity one. The shares sum to one in both economies, so the nodes that gain capital gain exactly what the other nodes lose, and the $1/2$ counts each unit of capital that moves once rather than twice.

intermediate economy splits the ratio into a product,

$$\frac{q^{\text{no cap}}\left(\boldsymbol{\kappa}^{\text{no cap}}\right)}{q^{\text{cap}}\left(\boldsymbol{\kappa}^{\text{cap}}\right)} = \underbrace{\frac{q^{\text{no cap}}\left(\boldsymbol{\kappa}^{\text{no cap}}\right)}{q^{\text{no cap}}\left(\boldsymbol{\kappa}^{\text{cap}}\right)}}_{\text{reallocation, at the no-capacity technology}} \times \underbrace{\frac{q^{\text{no cap}}\left(\boldsymbol{\kappa}^{\text{cap}}\right)}{q^{\text{cap}}\left(\boldsymbol{\kappa}^{\text{cap}}\right)}}_{\text{wedges, at the capacity economy's allocation}},$$

and the logarithm turns the product into a sum, so the two log differences add to the total. The remaining 8 log points are thus not a residual but the second factor itself, the gap between two economies that hold their capital in the same place and differ only in whether the constraint is on.

Most of the long-run cost of capacity is therefore the cost of the wedges at the allocation the economy actually chooses, and the wedges are costly because they sit on the links of a network: a wedge on every link is paid along every chain. Gross output, the sum of all sales in the economy, intermediate and final, is about four times GDP: every euro of final output is produced through about four euros of sales, three of them between nodes, and each sale crosses a link on which a wedge sits. A wedge of a few per cent per link is thus paid several times over on the way to every euro of GDP, which is how wedges of that size can cost the economy 8 log points of output.

The reverse ordering is also informative, where we take the capacity economy and move its capital to the counterfactual economy's allocation. Its GDP falls by 28 log points, so that at the no-capacity allocation the wedges cost 38 log points, of which the reallocation of capital recovers 28. At the no-capacity allocation some nodes do not have the capital that their sales require, so a large share of their firms are rationed, and the rationing propagates to their customers through the network. The two orderings give different numbers because the cost of capacity is strongly convex in how far a node's capital falls short of what its sales require: small shortfalls cost little, large shortfalls cost a great deal. Read together, the two orderings say that the cost of capacity in the calibrated economy is mostly the cost of the wedges, and that the reallocation of capital is what keeps the wedges from being much larger.

Table 3: Decomposition of the long-run GDP gap.

	wedges measured at the allocation of capital of	
	the capacity economy	the no-capacity economy
wedges: rationing, markups and idle capacity	−8.22	−38.29
reallocation of capital across nodes	−2.48	+27.60
total, $\ln\left(q_{\text{cap}}/q_{\text{no cap}}\right)$	−10.70	−10.70

Notes: Log points. The wedge effect switches the capacity constraint on at a fixed allocation of capital across nodes, the column's; the reallocation effect moves capital between the two economies' allocations at a fixed technology. Each column sums to the total.

Figure 3 takes the reallocation of capital and the output gap to the node level. Network position does not explain the direction of the reallocation: the log change in a

node's capital is almost uncorrelated with the node's upstreamness[23] and only weakly correlated with its utilization. What explains it is the role that capital plays. Without capacity constraints, a node's capital is determined by its Cobb–Douglas return, so nodes with a small capital exponent α_n hold little capital. With capacity, a node's capital also bounds how much the node can produce, so a node with a small capital exponent but a large volume of sales must hold enough capital to serve those sales. Figure 3(a) shows the result. The log change in a node's capital falls steeply with its capital exponent[24]. By group, capacity moves capital out of mining and utilities, where the capital exponent is large, and into the service nodes, where it is small; manufacturing gains a little (the capital shares by node are in Table 17).

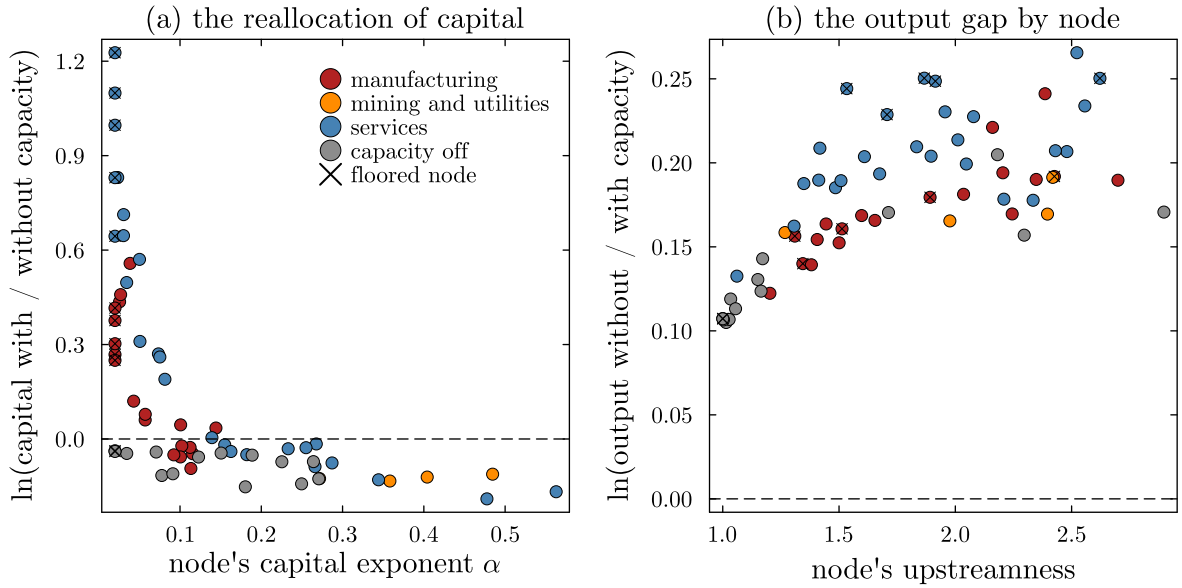
Figure 3(b) and Table 17 show the long-run output gap node by node. Real output per unit of aggregate capital is lower with capacity than without in every node, by 20 log points at the sales-weighted mean. The gap is larger the more upstream the node (the correlation with upstreamness is 0.68) and the tighter the node (the correlation with its utilization is 0.56 among the 49 nodes with capacity on). Upstream nodes lose most because they sell to other nodes, and the demand of other nodes for inputs is what rationing and markups cut.

A no-capacity model fitted to the same accounts. The economy without capacity constraints keeps the capacity economy's capital shares, so it does not reproduce the accounts: in the calibrated economy part of each node's capital income is the rent that capacity earns, and a model without capacity has to attribute that income to the technology instead. The last result of this section is what that attribution looks like. Recalibrating the no-capacity model to the same accounts, by the first and second stage of Appendix E with capacity off in every node, gives each node the capital exponent that explains its capital income without the rent. Figure 4 plots that exponent over the capacity economy's, node by node. In the 14 nodes without capacity the ratio is one. In the 49 nodes with capacity it runs from 1.0 to 2.9: 1.02 on average in mining and utilities, 1.23 in manufacturing and 1.45 in services, and largest in administrative and support services (2.9), miscellaneous professional services and warehousing (2.4 each). What the no-capacity model reads as technology is the rent relative to the node's capital income, so the misreading is largest where capital income is smallest: among the 39 nodes with capacity whose exponent is at the floor in neither economy, the correlation of the log of the ratio with the log of the capital exponent is -0.83 (panel (a)), and at a given exponent the ratio is larger the tighter the node (panel (b); utilization alone

[23]Upstreamness is the measure of Antràs, Chor, Fally and Hillberry (2012) and Fally (2012), written $\mathcal{U}_n$ here because U_n is utilization. It solves $\mathcal{U}_n = 1 + \sum_m d_{nm}\mathcal{U}_m$, where d_{nm} is the share of node n's gross output that node m buys as an intermediate input. A node that sells only to final use has $\mathcal{U}_n = 1$, and every further node that its output passes through before reaching final use adds one, so $\mathcal{U}_n$ is the average number of production stages between the node's output and final use, each path weighted by the share of the node's output that takes it. It is computed from the domestic input–output flows of the calibration (Appendix E; means over 2007–2019), is the same in both economies, and runs from 1.0 to 2.9 across the 63 nodes.

[24]The correlation with $\ln \alpha_n$ is -0.86 among the capacity-on nodes whose exponent is not floored.

Figure 3: The long run by node.



Notes: (a) The log of each node's capital with capacity over its capital without, against the node's capital exponent α_n; floored nodes are crossed. (b) The log of each node's real output without capacity over with, against the node's upstreamness (the average number of production stages between the node's output and final use).

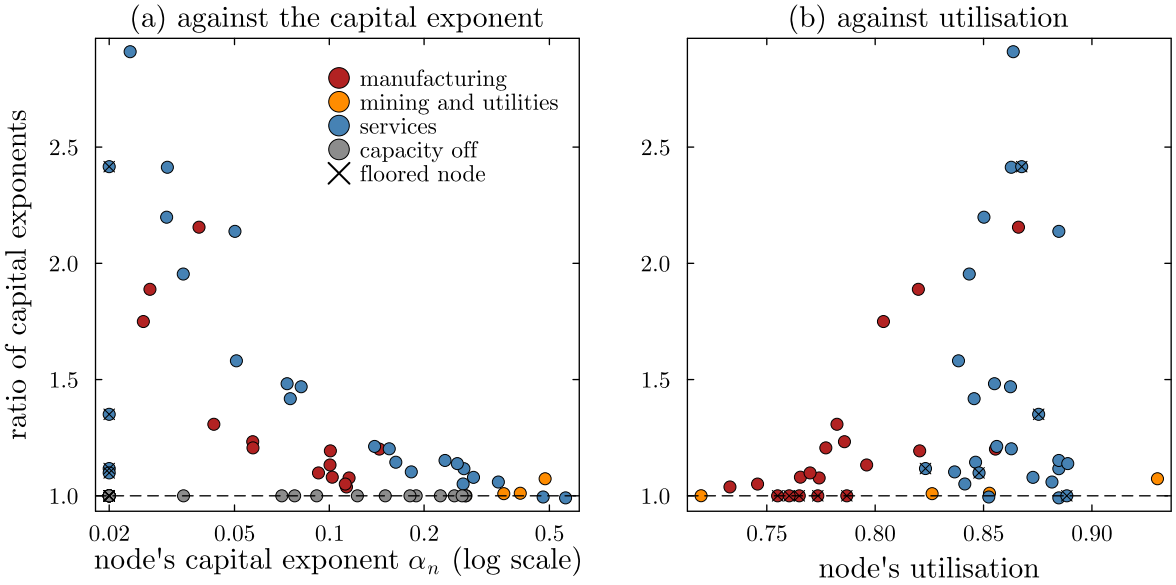
explains little of the ratio, but the exponent and utilization together account for 88 per cent of the variation of its log). In the aggregate the recalibration raises the mean capital exponent from 0.136 to 0.151, floors 8 nodes instead of 12, and puts 2 per cent of aggregate capital in different nodes than the capacity economy does. Its GDP per unit of capital is 2 per cent below the capacity economy's, because it fits the same capital stock with a technology that earns less on it. This recalibrated economy hits the same long-run targets by construction, so it has no long-run column in the tables. It returns in Section 6.1 as the second comparison for the impact responses, where the difference between it and the capacity economy reads as the error of a network model without capacity fitted to the same data.

[Q9 | chat | ANSWERED] the paragraph "A no-capacity model fitted to the same accounts."'s purpose here is unclear to me. Can you please suggest either (1) where to move it, if at all, or (2) perhaps adding some figure/graph to it and make it an additional result in this subsection.

[Reading the table: the 6.6 per cent, the decomposition, the γ rows]

- **6.6 per cent** is $\frac{1}{2}\sum_n \left|s_n^{\text{cap}} - s_n^{\text{no cap}}\right|$, with s_n node n's share of aggregate capital on the balanced growth path: the share of the economy's capital that would have to change node to turn the no-capacity allocation into the capacity one. Since the evening of 28 October 2026 it is the last row of the comparison block of Table 2 (Kian's decision; for an hour that evening it headed a separate table with the three group shares under it); the group shares are no longer printed, the capital shares by node are in Table 17.

Figure 4: What a model without capacity reads as technology.



Notes: The capital exponent of the no-capacity model recalibrated to the same accounts over the capacity economy's, $\alpha_n^{\text{recal}}/\alpha_n$, by node: (a) against the node's capital exponent α_n, on a log scale; (b) against the node's utilization. The 14 nodes without capacity have a ratio of one and no utilization, and appear in (a) only. Nodes whose exponent is at the floor in either economy are crossed; in the 8 nodes floored in both, the ratio is one by construction.

- **The decomposition.** Let q be GDP per unit of aggregate capital. Four economies: C, the capacity economy at its own allocation of capital ($q = 0.193$); B, the economy without capacity constraints at its own (0.214); $B(\kappa_C)$, the economy without capacity constraints's technology with its capital forced to the capacity economy's node shares and the static block re-solved (0.209); $C(\kappa_B)$, the capacity technology at the economy without capacity constraints's shares (0.146). Since the afternoon of 28 October 2026 the decomposition is its own table (Table 3), two columns in the cost orientation, from B to C, negative a cost: first column $B \to B(\kappa_C)$ -2.48 then $B(\kappa_C) \to C$ -8.22; second column $B \to C(\kappa_B)$ -38.29 then $C(\kappa_B) \to C$ $+27.60$; each sums to -10.70. Until then it was a block of Table 2 in the gain orientation, each row $100 \ln$ of a ratio of two of the four economies, positive a gain on the way from C to B, two routes each summing to 10.70:
  - wedges first: $C \to B(\kappa_C)$, the constraint switched off with capital where it is, $+8.22$; then $B(\kappa_C) \to B$, the economy without capacity constraints moves its capital to its own allocation, $+2.48$;
  - reallocation first: $C \to C(\kappa_B)$, the capacity economy forced onto the economy without capacity constraints's allocation, -27.60 (the service nodes then lack the capital their sales need); then $C(\kappa_B) \to B$, the constraint switched off there, $+38.29$.

  The wedge is 8 at one allocation and 38 at the other because the cost of the constraint is convex in how far a node's capital falls short of its sales. The bullet

"Where the cost comes from" reads the first route.

- **The welfare rows** (the comparison block is in percentage differences with the economy without capacity constraints as base since the afternoon of 28 October 2026, Kian's decision). The stored number is the compensating variation of the capacity economy's household, $1+\lambda = (\widehat{C}_{\text{no cap}}/\widehat{C}_{\text{cap}}) \cdot G = 1.155 \times 1.388 = 1.604$, with $G = \exp\left[(g_{\text{no cap}} - g_{\text{cap}})/\rho\right]$ the growth part under log utility (at $\gamma \neq 1$ it is $(\nu_{\text{cap}}/\nu_{\text{no cap}})^{1/(1-\gamma)}$ with $\nu = \rho + (\gamma-1)g$); the table prints the equivalent variation of the economy without capacity constraints's household on the block's base, $1/(1+\lambda) - 1 = -37.6$ per cent, whose level part $\widehat{C}_{\text{cap}}/\widehat{C}_{\text{no cap}} - 1 = -13.4$ per cent is the consumption row and whose growth part $1/G - 1 = -28.0$ per cent compounds with it. The growth part is large because the growth gap of 0.60 points is discounted at $\rho = 0.018$, and it moves with γ (at $\gamma = 2$ the gap is 0.30 points, and on the previous baseline the growth part was 13 per cent), so the welfare row is the one not to lead with; the body's footnote that once said so was removed in v3D, so the row is explained here only.
- **The γ rows** are out of the table and out of the text (the private box after the paragraph on the cost, below): growth effect, points, -0.599 at $\gamma = 1$ ($\rho = 0.018$, the calibration); -0.302 at $\gamma = 2$ ($\rho = -0.003$); -0.202 at $\gamma = 3$ ($\rho = -0.024$).

[Is a 10 per cent loss too large? What the number is, and what would lower it]

- **No property of the model makes it smaller than it reads.** The row is the log gap in GDP per unit of aggregate capital between two balanced growth paths: a TFP-type gap at equal capital. The usual long-run numbers let the capital stock respond as well and are larger for the same wedge; here that response is a separate number, the growth gap. So the level gap and the growth gap are two numbers, and the welfare row (37.6 per cent as the table states it, 60 measured from the capacity side), which compounds them, is the one not to lead with.
- **Two readings do make it less than it seems.**
  - **Reading 1: the number depends on the comparison economy.** The 10 per cent compares the capacity economy with the economy without capacity constraints: the calibrated economy in the limit $x_n \to 0$, every other primitive the same, so that every firm serves all its demand with the capital it holds. The economy without capacity constraints is not fitted to the data; it is the calibrated economy with one feature switched off.
  - The no-capacity network model that *is* fitted to the same accounts (the recalibrated economy) produces 2 per cent *less* per unit of capital than the capacity economy. So the 10 per cent is the cost of the constraint at a given technology; it is not what the data would gain against a frictionless network model fitted to them, which is the comparison that makes the number sound large.
  - **Reading 2: most of the level loss is a markup, compounded by the network.** The constraint raises the average markup from 1.116 (the economy

without capacity constraints) to 1.155: a wedge of 3.5 per cent on every link. A wedge on every link is paid along every chain, and gross output is about four times GDP (the Domar weights sum to 4.0 with capacity and 4.4 without), so a wedge of this size compounds to a loss of this order.

  - The model confirms this. A no-capacity economy with a uniform markup of 1.154 (the economy without capacity constraints at $\varepsilon = 7.5$, every other primitive at its calibrated value) produces 11 log points less than the economy without capacity constraints, close to the capacity economy's 10. But it does not dampen shocks at all: its response to a common shock is 7.00 per unit, the same as the economy without capacity constraints's at every ε. (One point of the fixed-threshold comparative static, computed once; a run at the exact markup, by node and for the growth rate, would repeat the comparison.)
  - The literature finds losses of the same size from markups in networks: Baqaee and Farhi (2020) find that removing markups would raise US TFP by about 15 per cent; Edmond, Midrigan and Xu (2023) put the consumption-equivalent cost of markups at about 9 per cent in an economy whose aggregate markup is 1.15. Both numbers moved across drafts and are to be checked against the published versions.
  - So if the level loss is what a markup produces, the result that is specific to capacity is not the level loss. It is the damping of shocks and its dependence on the state of the economy: a markup lowers the level of output but does not change how a shock propagates.

- **What would lower it.** ε: the loss is 7 log points at $\varepsilon = 10.5$ and 4 at 12 (the sweep figure in the comparative-statics box at the end of this section; the sweep is the previous baseline's, $\phi = 13$), the one-standard-error interval being 8.2–11.8; the honest fix is to report the range (a comparative static re-solved at the ends of the interval), not to move the point. The level of the utilization targets: the secondary calibration's looser service sectors gave a rough 2 per cent (the secondary calibration, deferred), and this lever changes the least elsewhere. Not available: σ_v (0.13 would give 3 points but puts the Census share at 0.087 against 0.162); γ (does not enter the stationary economy). Not ϕ: at $\phi = 2$, the baseline since 3 November 2026, the loss is 10.1 per cent against 9.7 at $\phi = 13$.

[The growth effect at other values of γ, kept out of the text and of the table] The growth gap in the text and in Table 2 is the one at the calibrated $\gamma = 1$ ($\rho = 0.018$): 0.60 points, equal to the interest-rate gap because under log utility $r - g = \rho$ in both economies. At $\gamma = 2$ ($\rho = -0.003$) it is 0.30 points and at $\gamma = 3$ ($\rho = -0.024$) 0.20 points; the capacity economy's own growth rate is the same at all three values, because the stationary economy depends on γ and ρ only through $r = \rho + \gamma g$. The size of the growth effect is therefore a convention: the economy without capacity constraints keeps the calibrated ρ, so its interest rate $r = \rho + \gamma g$ moves with its own growth rate, and the larger is γ, the less its growth rate has to fall to bring the interest rate down to

the return that capital earns. The three values are computed once by the static results on the baseline calibration (a generated table holds them, input nowhere), and they stay in this box until their place is decided.

[Q10 | ANSWERED] - The sentence "The long-run GDP gap has two parts: the wedges that capacity puts on every link, rationing, state-dependent markups and idle capacity, at a given allocation of capital, and the reallocation of capital that the economy undertakes in response" is badly-written grammatically: it's not clear what these two parts are because of so many commas.

- the sentence "the difference from its own GDP is what the reallocation costs a technology that has no use for it, 2 per cent, and the rest of the gap, 8 per cent, is the wedge at that allocation." is not very clear.

- In "Shifting the capacity economy's cross-section to the economy without capacity constraints's allocation, that is, taking away the capital that the service nodes hold for capacity's sake, lowers its GDP by 28 per cent: at the no-capacity allocation the nodes that lose capital become severely rationed and the rationing propagates through the network. The two orderings differ bec" how did service suddenly show up? You need to be clearer! A lot of your sentences read like riddles!

[Evidence and files]

- Every number is the solved model on the baseline calibration ($\phi = 2$, $\gamma = 1$), computed once on 2–3 November 2026 by the static results script, from its stored output tables (the aggregates, the decomposition and the welfare gap; the growth effect at $\gamma = 1, 2, 3$; the per-node results; the recalibrated economy); the capacity economy was solved by two routes (cold at $\gamma = 2$ on 31 October 2026, warm from that solution at $\gamma = 1$ on 2 November 2026) with the same balanced growth path to seven digits; the sweep's and the fixed-threshold script's base points (0.1018–0.1019) are the previous baseline's. The correlations in the text, the ratios of capital exponents in the paragraph on the no-capacity model fitted to the same accounts, their group means and their correlations were recomputed from the new baseline's by-node tables on 3 November 2026 (they move by at most 0.01 against the previous baseline), and Figure 4 was drawn from the same tables.
- The comparison block of the table and the cost paragraph are percentage differences with the economy without capacity constraints as base since the afternoon of 28 October 2026 (Kian's decisions of that day; for a few hours that afternoon the text said log points, and that morning per cent read off the log differences), the decomposition block is in log points and headed so, the welfare rows are the economy without capacity constraints's household's equivalent variation (37.6 per cent; the stored compensating variation is 60), the footnote of the cost paragraph gives the conversions, and the correlation of the output gap with utilization is over the 49 capacity-on nodes since the same date; over all 63 nodes, with the capacity-off nodes at a placeholder utilization of 0.25, it is 0.66.
- The decomposition's two orderings and the convexity reading are computed

facts; "the reallocation keeps the wedges from being much larger" is the reading, reasoned.

[The price and the mass of rationing: moving ε and σ_v] [Flag: set aside on 30 October 2026] This box holds, unchanged, what was the second subsection of this section, the comparative statics in ε and σ_v, set aside for a closer look later. No cross-reference in the main text points into it; the abstract, the road map of the introduction and the conclusion still describe the comparative statics in words, as they stand until their final rewrite. The box keeps its equation, its figure and its three private boxes. Since 3 November 2026 the baseline calibration is $\phi = 2$, $\gamma = 1$; the sweep and the fixed-threshold static were not re-run, so every number, both figures and the three private boxes inside this box are the previous baseline's ($\phi = 13$, $\gamma = 2$).

Two parameters set how much capacity matters, and they act through two different channels. Separating the channels says what the long-run cost of Section 5.1 is made of and where the calibrated economy sits; that the two pieces of evidence of Section 4 identify the two parameters separately is shown in Appendix E (Figure 11).

Two channels. Consider a firm whose demand draw exceeds its capacity. By the threshold representation (5) and the multiplier of (3), the price at which its buyers would have demanded exactly its capacity, had the firm known its demand draw, exceeds its posted price by the factor

$$\frac{p_{in} + \varphi_{in}}{p_{in}} = \left(\frac{v_{in}}{\overline{v}_{in}} \right)^{1/\varepsilon} \quad \text{for} \quad v_{in} \geq \overline{v}_{in}. \tag{56}$$

We call this factor the *price of rationing*. The demand that the firm cannot serve is redirected to the other firms of its sector through the CES demand system, and the factor measures what a buyer loses by being sent to another firm. It falls with ε: the better the substitutes, the smaller the loss. The *mass of rationing* is the share of firms above their capacity, $1 - G_i(\overline{v}_i)$, a tail probability of the distribution of firm-level demand. Both objects depend on the sector's threshold alone, and so does the sector's utilization rate (22). A sector's utilization rate therefore pins down the share of its firms that are rationed, given the shape of G_i. In a calibration that targets utilization, the mass of rationing is set by the dispersion σ_v of the demand draws, and the elasticity ε acts only on its price. This is the sense in which the two parameters are the two channels, and the two moments of the calibration measure them separately: the Census share of capacity in plants close to their cap measures the mass, and the convexity of supply curves in utilization measures the price.

The comparative static at fixed utilization targets. Figure 5 moves one parameter at a time and re-does the calibration at every point: the utilization targets, the final-demand shares, the capital stock and the growth rate are re-matched, so that every point is an economy consistent with the same accounts and the same utilization data. Along ε the share of rationed firms stays at 0.19 over the whole range, as the argument

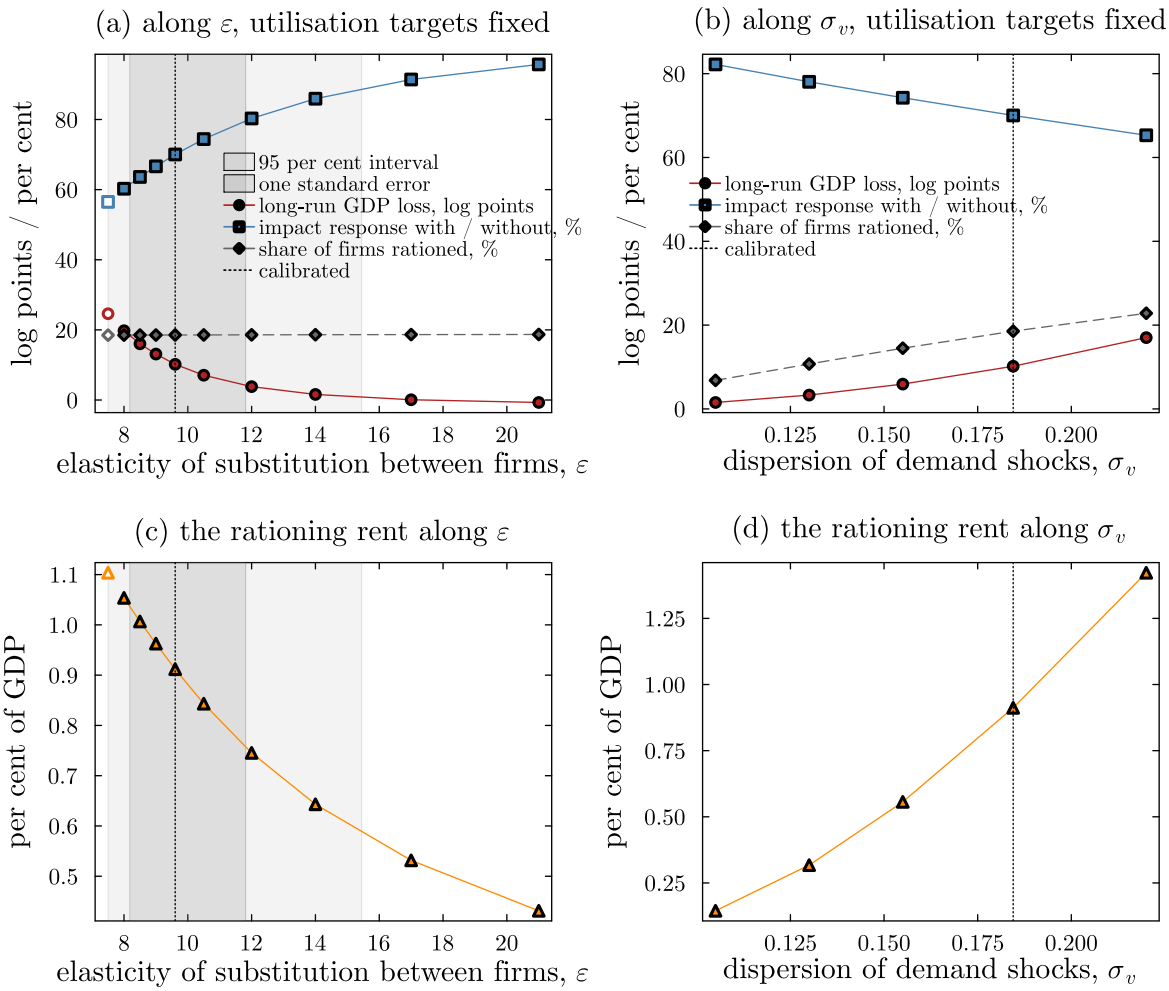
above says it should, and everything else moves with the price of rationing. The long-run GDP loss falls from 10 log points at the calibrated $\varepsilon = 9.6$ to 7 at $\varepsilon = 10.5$, 4 at 12, 2 at 14 and zero at 17. The dampening of a common shock at impact, 0.70 at the calibrated point, weakens to 0.80 at $\varepsilon = 12$ and to 0.96 at 21. Beyond $\varepsilon \approx 17$ the sign of the long-run effect turns: at $\varepsilon = 21$ the capacity economy produces 0.7 log points more than the economy without capacity constraints and grows faster by less than half a basis point. The gain is small, and it does not come from accumulation. The comparison is per unit of capital; the rationing rent is 0.4 per cent of GDP at that point and accrues to the aggregators, not to the firms; and on a balanced growth path the growth rate follows the private return to capital, which is GDP per unit of capital net of the markup rents and the rationing rent (Remark 1), so the slightly higher growth rate is a consequence of the static allocation at $\varepsilon = 21$, not its cause. Which part of that allocation produces the gain is not decomposed here. What the sweep establishes is that the loss from rationing shrinks with the price of rationing and vanishes at $\varepsilon \approx 17$, and that the calibrated economy sits well inside the region where the loss dominates. Below the calibrated value the loss rises steeply: 13 log points at $\varepsilon = 9$, 16 at 8.5, 20 at 8 and about 25 at 7.5, which is close to the lowest value of ε at which the model admits a solution; there the dampening is 0.57, and the re-fitting of the economy stops just short of its tolerance. Along σ_v it is the mass that moves. At $\sigma_v = 0.13$ the share of rationed firms is 0.11 and the long-run loss 3 log points; at 0.155 they are 0.145 and 6; at the calibrated 0.184 they are 0.19 and 10; and at 0.22 they are 0.23 and 17. The two calibration moments confirm this reading: along ε the Census share does not move, and along σ_v the distance to the supply-curve estimates is smallest at the calibrated value (Figure 11 in Appendix E).

What is held fixed decides the sign. The static above holds the utilization data fixed. With the technology held fixed instead, that is, the calibrated thresholds, capital shares and productivity levels, a higher ε lowers every markup, raises every firm's demand and pushes more firms over an unchanged cap, so the loss rises in ε instead of falling (the experiment is in the private box at the end of this section). The results rely on the comparison at fixed targets: it is the one that answers a reader who doubts the estimate of ε, because every point of it is an economy fitted to the same accounts.

[The range of the headline, read off the sweep; not yet in the text]

- The estimate of ε has a one-standard-error interval of 8.2–11.8 and a 95 per cent interval from the admissibility bound (7.3–7.5) to 15.5 (Appendix E), the shaded bands of Figure 5. Read off the sweep, the long-run loss is 4–18 log points within one standard error and 1–25 within the 95 per cent interval, and the impact response of GDP to a common shock, with capacity over without, 0.62–0.80 and 0.57–0.89. Sections 5 and 6 state neither the interval nor these ranges yet.
- **Interpolated, not re-solved.** The loss is 19.8 at $\varepsilon = 8$ and 16.0 at 8.5, so about 18 at 8.2; 7.1 at 10.5 and 3.8 at 12, so about 4 at 11.8; 1.6 at 14 and 0.1 at 17, so about 1 at 15.5; the lower end of the 95 per cent interval is read at the 7.5 point. The response is interpolated the same way (0.603 at 8, 0.636 at 8.5; 0.744 at 10.5,

Figure 5: Comparative statics at fixed utilization targets.



Notes: The calibration is re-done at every point. (a) Along ε at the calibrated σ_v and (b) along σ_v at the calibrated ε: the long-run loss $\ln(q_{\text{no cap}}/q_{\text{cap}})$, the impact response of GDP to a common productivity shock with capacity over without, and the share of rationed firms. (c), (d) The rationing rent as a share of GDP along the same two sweeps. The dotted line marks the calibrated value. The shaded bands in (a) and (c) are the one-standard-error and the 95 per cent intervals of the estimate of ε, the lower end of the latter being the admissibility bound. The open markers at $\varepsilon = 7.5$ are the point whose re-fitting stopped at a miss of 3×10^{-4}.

0.803 at 12; 0.859 at 14, 0.914 at 17). Re-solving the sweep at the four ends would replace these numbers.

[Evidence and flags for the comparative static at fixed targets]

- The sweep is computed once (about four hours on the default capital grid, then nearly four on the wide grid), with its per-node results stored; the base point reproduces the static results' numbers to four decimals. **The grid.** The points $\varepsilon \leq 9$ and $\sigma_v = 0.22$ failed on the default capital grid and were solved on the wide grid (201 points, log-width 1.0); $\varepsilon = 7.5$ stopped there at a worst re-fitting miss of 3×10^{-4} and is the open marker of Figure 5; $\sigma_v \geq 0.26$ was not obtained. Not plotted because flat by construction: aggregate utilization along the sweeps (0.742–0.745 at every point, the targets being re-matched). The Census share along both sweeps is in Figure 11.
- **Flag: the sign at high ε.** The loss is 0.1 at $\varepsilon = 17$ and -0.7 at 21, both above the 95 % interval's upper end (15.5), and whether a gap of that size survives a finer capital grid is unchecked. The text gives no mechanism for the gain: the sentence that attributed it to the rents of rationed firms raising the return to capital was replaced on 30 October 2026 by a statement of the facts (the rents are the aggregators', the comparison is per unit of capital, the growth gap follows the return), and which part of the static allocation produces the gain is not decomposed.
- The statement that the rationed share does not depend on ε at a given utilization is exact for one industry at a given threshold and approximate for a node (the sweep shows 0.185–0.187).

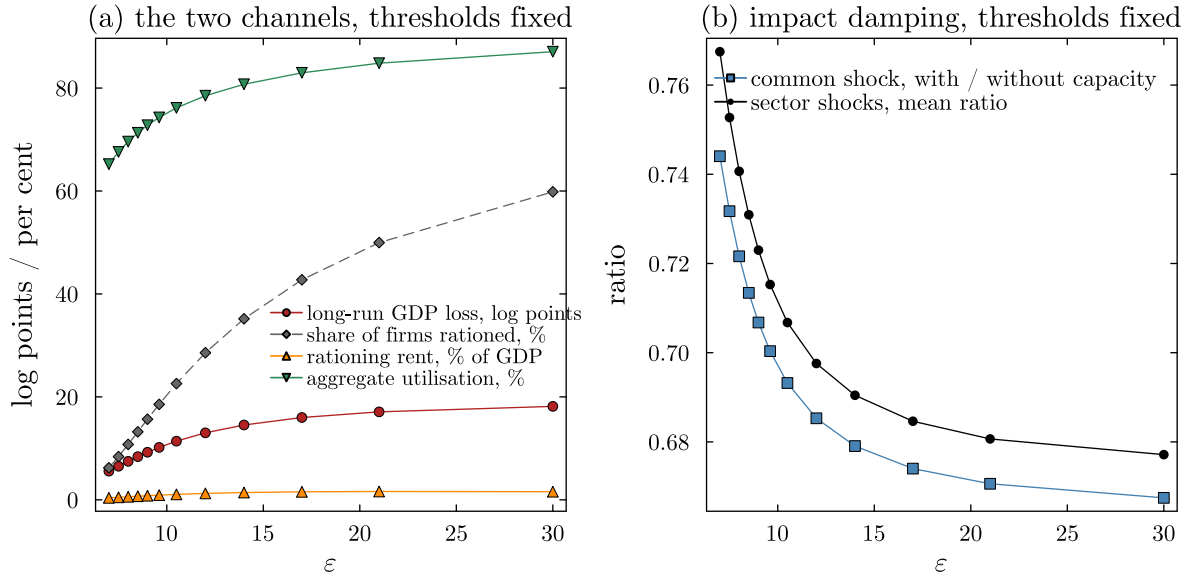
[The comparative static at fixed thresholds: the experiment behind the paragraph "What is held fixed decides the sign"]

- **The experiment.** This static holds the calibrated capacity thresholds, capital shares and productivity levels fixed and moves ε alone (Figure 6). It changes both channels at once, in opposite directions. A higher ε lowers the markup, raises every firm's demand and pushes more firms over their unchanged cap. The share of rationed firms rises from 0.06 at $\varepsilon = 7$ to 0.19 at the calibrated 9.6, 0.35 at 14 and 0.60 at 30, and aggregate utilization rises with it, from 0.65 to 0.87. At the same time the price of rationing falls, as (56) says it should: the rent per rationed firm falls by more than half over the range. The product of the two channels is the total rent that capacity earns, which rises from 0.4 per cent of GDP at $\varepsilon = 7$ to 1.6 per cent at $\varepsilon = 21$ and is flat beyond. The loss does not turn around: the mass channel dominates, so the long-run GDP loss rises throughout, from 6 log points at $\varepsilon = 7$ through 10 at the calibrated value to 18 at $\varepsilon = 30$, flattening as it goes, and the dampening of a common shock at impact strengthens from 0.74 to 0.67.
- **Reading.** The sign of the comparative static in ε therefore depends on what is held fixed. With the technology fixed, a higher ε makes the economy tighter and capacity matters more. With the utilization data fixed, a higher ε makes rationing cheaper for the same share of rationed firms and capacity matters less.

Both statics are correct, and they answer different questions. The static at fixed thresholds says what capacity does in an economy with a given technology as the substitutability of firms' goods changes. The static at fixed utilization targets says what the long-run cost and the dampening would be if the utilization data were what they are and the price of rationing were different. The second is the question that a reader who doubts the estimate of ε asks, and it is the one the rest of the paper relies on.

- **Why it is in a box.** The economies along this static no longer fit the accounts (their utilization, rationed share and capital income all move), so it is not a comparison the paper relies on. It is kept as the record of the experiment and as the evidence for the sign statement in the text. Whether its figure goes to an appendix is open.
- **Evidence.** Computed once (two hours); the points $\varepsilon = 7$ and 7.5 were solved on the wide capital grid. Units: this experiment shocks productivity by ± 1 per cent in levels, where the static results and the sweep shock it in logs; the responses per unit differ by half a per cent and the ratios agree (0.700 at the base point in both).

Figure 6: Comparative static at fixed thresholds.



Notes: Only ε moves; the capacity thresholds, capital shares and productivity levels stay at their calibrated values. (a) The long-run GDP loss, the share of rationed firms, the rationing rent and aggregate utilization. (b) The dampening of a common and of a sector shock at impact.

6 Capacity Dampens the Propagation of Shocks

[Q11 | text | ANSWERED] write a very short opening: 1 sentence intro and 2-3 sentences roadmap.

The long-run comparison holds the shocks fixed; this section asks how capacity changes the economy's response to a shock. Section 6.1 shocks the productivity of one node at a time, with capital given, and compares the impact responses of the capacity economy and the economy without capacity constraints node by node. Section 6.2 asks whether the response depends on the sign and the size of the shock, which it cannot without capacity constraints. Section 6.3 follows two of the shocks over time, with investment and the distribution of industries within every node adjusting.

6.1 Impact: How Capacity Changes the Propagation of a Sectoral Shock

This subsection shocks one node at a time and reads the response of the economy at the date of the shock. The shock raises the productivity of every industry of one node by 1 per cent. Capital is given and the cross-section of industries within every node is held at its balanced-growth distribution, so the response is that of the static block: prices, utilization, rationing and the flows between nodes adjust while capital does not. The same shock is run in the capacity economy and in the economy without capacity constraints, node by node, and the statistic reported is the ratio of the two responses. A ratio below one means that capacity dampens the effect of the shock, while a ratio above one means that it amplifies it. Throughout, 'GDP' refers to GDP per unit of aggregate capital in units of the final good, a node's real output is its sales deflated by its shadow price index; and "the other nodes" are the 62 nodes that are not shocked, weighted by sales. The levels of the elasticities are not informative, because the model has no labour and its GDP is the income of physical capital, so a node's productivity moves GDP by several times the node's Domar weight. The ratios between the two economies are informative.

The first result is that capacity dampens propagation, and that the dampening happens in the network rather than at the shocked node. Table 4 gives the means over the 63 shocked nodes. The response of GDP with capacity is 0.76 of the response without, and the response of real gross output is 0.53 of it. The gross-output ratio can be split between the shocked node's own output and the output of the other nodes. The shocked node's own output responds almost as it would without capacity (ratio 0.95); the other nodes' output responds at only 0.45 of its no-capacity value. In other words, the shocked node sells more and at a lower price in both economies; what capacity changes is how much of that reaches the node's customers and suppliers. Weighting the shocked nodes by their Domar weights, so that large nodes count for more, the ratios are 0.71 for GDP and 0.35 for the other nodes' output. Against the recalibrated economy the ratios are 0.85 and 0.49: a network model without capacity, fitted to the same data, overstates the response of GDP to a sectoral shock by about a sixth and the response of the other sectors by about a factor of two.

The dampening comes from the quantity response of the other nodes, not from their price response. The third block of Table 4 shows this by decomposing the response of the other nodes into a price part and a quantity part. Their sales respond

Table 4: Impact responses to a sectoral productivity shock.

	baseline calibration	
	against the no-capacity economy	against the recalibrated economy
response with capacity over response without, mean over the 63 shocked nodes		
GDP	0.76	0.85
real gross output	0.53	0.58
the shocked node's own output	0.95	0.95
the other nodes' output	0.45	0.49
the same, ratio of the sums over the shocked nodes (Domar-weighted)		
GDP	0.71	0.81
the other nodes' output	0.35	0.41
elasticity of the other nodes' sales to the shock, mean over shocked nodes, with capacity / without		
nominal sales	0.052 / 0.109	0.052 / 0.097
of which price	0.021 / 0.021	0.021 / 0.021
of which quantity	0.031 / 0.088	0.031 / 0.075
the capacity margin with capacity, mean over shocked nodes, change per 1 % shock		
utilisation of the shocked node, points	0.014	0.014
utilisation of the other nodes (sales-weighted), points	0.021	0.021
aggregate utilisation $\overline{U}$, points	0.022	0.022
share of firms rationed (value-added weighted), points	0.042	0.042
aggregate markup, $\times 100$	0.012	0.012

Notes: Impact responses to a $+1$ per cent productivity shock at one node, means over the 63 shocked nodes: the capacity economy's response over the no-capacity economy's and over the recalibrated economy's. GDP and the other nodes' real output are per unit of aggregate capital; the other nodes' response is decomposed into its price and its quantity part.

to the shock with an elasticity about half as large with capacity as without. The price part of that response, the response of their shadow price indices, is the same in the two economies; the whole of the difference is in quantities, which respond at about a third of their no-capacity value. The statement about prices is about the sales-weighted sum of the other nodes' price responses, which nets the shocked node's customers, whose prices fall, against the rest, whose prices rise; at a single receiving node the pass-through of cheaper inputs into its price is damped when the node is tight(Appendix G). Without capacity, a node that receives cheaper inputs or more demand expands along a supply curve with elasticity $1/\alpha_n$, which at the calibrated capital shares is very flat, so almost all of the shock that reaches the node becomes output. With capacity, the firms of the node that are at their cap cannot expand. The demand they turn away is rationed rather than met at a higher price, so the node's recorded sales and output rise by about a third as much as without capacity, while its shadow price index rises by the same amount as without. This is the mechanism in one line: the capacity constraint replaces part of the quantity response of the other nodes with rationing, and the price index does not make up the difference because the rents on rationed units are not recorded as sales.

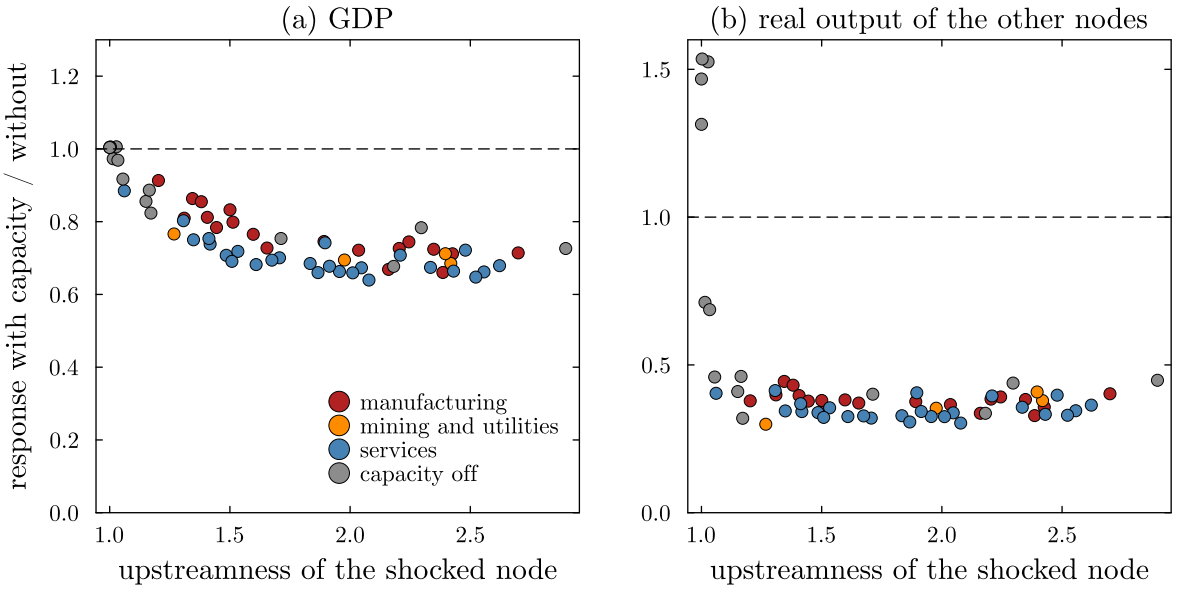
The same mechanism can be summarised in one number per node, the node's supply elasticity: the response of its sales to its price at given capital, which the model gives in closed form as a function of the threshold (Lemma 15); the response of its output is one less. Without capacity it is $1/\alpha_n$, about 12 at the median capacity-on node. With capacity it is the elasticity of the firms with spare capacity, weighted by their share of the node's sales. It is 2.9 at the median, ranges from 1.3 to 8.4 across the capacity-on nodes, and falls steeply with the node's utilization (correlation -0.87). A node at its cap cannot pass a shock on to its customers, and the ratio of the two elasticities, about one to four, is what cuts the quantity response of the other nodes.

The capacity margin shows the same pattern as the output responses: it moves little at the shocked node and more everywhere else. A 1 per cent shock at one node raises that node's own utilization by only about 0.01 percentage points. The reason is that the node's capacity is proportional to its productivity, so the shock raises the node's capacity together with its demand and does not tighten the node that receives it. The shock does tighten the rest of the network: aggregate utilization rises by 0.02 points, the share of rationed firms by 0.04 points and the aggregate markup by 0.01 per cent. This tightening of the other nodes is where the dampening comes from. Section 6.3 shows the same pattern over time.

The means hide a gradient in network position and in tightness. Figure 7 plots the ratios by shocked node. The GDP ratio runs from 0.64 to 1.01 and falls with the position of the shocked node in the network: its correlation with the node's upstreamness is -0.77, and with the node's own utilization -0.59 among the 49 nodes with capacity on. Shocks to upstream and tight nodes are dampened most. Shocks to the capacity-off nodes at the end of the chain (hospitals, nursing, social assistance, retail), which sell to final demand and buy from the rest of the network, pass through almost undamped: their mean ratio is 0.89, against 0.73 for the 49 nodes with capacity on.

For four of these nodes the response is marginally larger with capacity than without, because the shocked node's extra demand for inputs falls on suppliers that are tight. The ratio for the other nodes' output shows the same gradient less sharply (correlation -0.44 with upstreamness), because once the shock has left its origin, what it meets is the tightness of the nodes around it rather than the tightness of the node it came from. The two nodes followed over time in Section 6.3, primary metals and food services, are an upstream manufacturing node and a downstream service node.

Figure 7: Impact by shocked node.



Notes: The response of GDP with capacity over without, against (a) the shocked node's upstreamness and (b) its utilization.

[Evidence and files]

- The impact responses of both economies to the 63 node shocks and to the common shocks, the responses of the recalibrated economy, the four full response matrices and the per-node supply elasticities are stored output tables of the static results; computed once on the previous baseline and again on 2–3 November 2026 on the baseline calibration at $\phi = 2$, $\gamma = 1$, where every ratio of this subsection is unchanged to two decimals (no impact response depends on ϕ, γ or ρ; the capacity block's re-estimation moved ε by 0.01). The correlations and the gradient statistics were recomputed on 3 November 2026; the correlation with utilization is over the 49 capacity-on nodes (over all 63, with the capacity-off nodes at a placeholder utilization of 0.25, it is -0.69).
- The node-level counterpart of "the price response is the same" (the pass-through of cheaper inputs into a receiving node's price about 0.74 of its no-capacity value on average, of an arriving demand shift 0.98) is computed once in Appendix G.

6.2 Sign and Size

Without capacity the response of the economy cannot depend on the sign or the size of the shock, to first order: every node supplies along a constant-elasticity curve, so the response is close to linear in the log of the shock, and the response to a shock of $-x$ is minus the response to $+x$. With capacity it must depend on both. A positive shock pushes the tight nodes into their cap, where extra demand is rationed rather than met, while a negative shock frees capacity, so the fall in demand is met in quantities. The economy's supply curve is therefore concave in the shock, and the same shock does more harm when it is negative than good when it is positive.

The asymmetry shows in the response to a common shock. Table 5 and Figure 8 shock the productivity of every node at once, by ± 1, ± 2 and ± 5 per cent. Without capacity the response of GDP per unit of shock is the same number at every size and sign to three decimals. With capacity it falls as the positive shock grows and rises as the negative shock grows: the response to $-x$ over the response to $+x$ is 1.04 at 1 per cent, 1.08 at 2 per cent and 1.22 at 5 per cent. The dampening of a positive shock also grows with its size, from 0.70 of the no-capacity response at 1 per cent to 0.65 at 5 per cent. Panel (b) of the figure shows the capacity margin behind this. A common shock of $+1$ per cent raises aggregate utilization by 1.3 percentage points and a shock of -1 per cent lowers it by 1.5 points; at ± 5 per cent the two numbers are $+5.4$ and -9.4. utilization falls faster than it rises because, as firms approach their cap, more and more of the extra demand is turned away rather than produced, so the same change in demand buys less utilization on the way up than it loses on the way down. The share of rationed firms moves in the same way.

The same exercise one node at a time finds almost no asymmetry. The response to -1 per cent over the response to $+1$ per cent at one node is one to two decimals at every node, and the response per unit of shock at ± 5 per cent differs from that at ± 1 per cent by less than one per cent, with capacity and without alike. The reason was given in Section 6.1: a shock at one node moves aggregate utilization by 0.02 points per per cent of shock, against 1.3 points for the common shock, so the economy stays where it was on its supply curves. In this model state dependence is a property of how far the economy as a whole moves along its capacity margin, and a shock to one of 63 nodes does not move it far. The same fact explains why the dampening of Section 6.1 is a ratio that holds for any small shock: it is the slope of the economy's supply curve at the calibrated point, not its curvature.

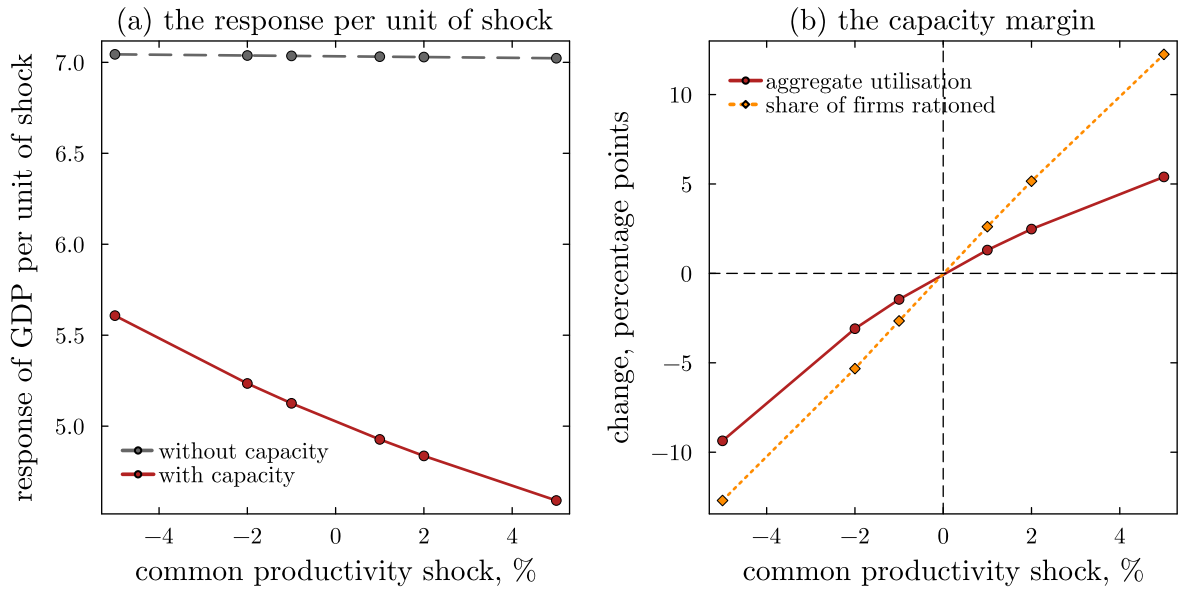
[Evidence and a flag] The common shocks and the node-by-node sign and size responses (-1 per cent at every node, ± 5 at the 23 core nodes) are stored output tables of the static results; computed once. **Flag.** The asymmetry grows about linearly in the size of the shock (1.040, 1.083, 1.222, about 4 per unit) and no derivation of this exists; the common-shock and sector-shock dampening are about equal while the asymmetry is the common shock's alone.

Table 5: Common shocks by sign and size.

common shock to all nodes	with capacity $+x$	with capacity $-x$	without $+x$	without $-x$
$x = 1\,\%$: GDP response per unit shock	4.927	5.126	7.031	7.036
$x = 2\,\%$: GDP response per unit shock	4.836	5.235	7.029	7.038
$x = 5\,\%$: GDP response per unit shock	4.591	5.608	7.023	7.044
asymmetry: response to $-x$ over response to $+x$				
$x = 1\,\%$	1.040		1.001	
$x = 2\,\%$	1.082		1.001	
$x = 5\,\%$	1.221		1.003	
one node shocked at a time, 63 nodes (the size rows: the 23 core nodes)				
sign asymmetry at $\pm 1\,\%$, with capacity (mean / median / max)	1.003 / 1.002 / 1.011			
sign asymmetry at $\pm 1\,\%$, without (mean / median / max)	1.000 / 1.000 / 1.004			
size: per-unit response at $+5$ / at $-5\,\%$ over at $+1\,\%$, with capacity (core nodes)	0.995 / 1.007			
size: per-unit response at $+5$ / at $-5\,\%$ over at $+1\,\%$, without (core nodes)	0.996 / 1.003			

Notes: The response of GDP per unit of aggregate capital to a common productivity shock of ± 1, ± 2 and ± 5 per cent, per unit of shock, with capacity and without; the asymmetry is the response to $-x$ over the response to $+x$.

Figure 8: Common shocks by sign and size.



Notes: (a) The response of GDP per unit of shock, with capacity and without. (b) The change in aggregate utilization and in the share of rationed firms.

6.3 Transition Dynamics

The last exercise follows a shock over time. A 1 per cent productivity shock with a half-life of 3.5 years hits one node, and the economy's transition back to its balanced growth path is computed in full, with investment, the distribution of industries within every node and all prices adjusting, in the capacity economy and in the economy without capacity constraints. Two nodes are shocked: primary metals, an upstream manufacturing node, and food services, a downstream service node. Table 6 and Figure 9 report the paths.

The first date of the transition is the impact response of Section 6.1. The response of GDP at the date of the shock to primary metals is 0.08 per cent with capacity, and the first date of the transition reproduces the static block's response for the same node. This must be so: the response at the shock date is a property of the static block at the state in place before the shock, and investment, adjustment costs and the persistence of the shock cannot move it. Everything that the dynamics add is in what follows.

What the dynamics add is that utilization absorbs the shock first and that the dampening fades slowly as capital adjusts. With capacity, aggregate utilization rises on impact, by 0.03 percentage points after the shock to primary metals and by 0.02 points after the shock to food services, and the share of rationed firms rises by 0.04 points in both cases; both decay with the shock. The response of GDP decays with the shock in both economies, with a half-life of about four years against the shock's three and a half, so capacity changes the size of the response and not its persistence. The ratio of the GDP responses with and without capacity is 0.71 at impact after the upstream shock and 0.75 after the downstream one. It rises along the path as capital has time to move: to 0.77 and 0.79 at year ten and to 0.78 and 0.79 at year twenty. The ratio for the other nodes' output rises faster, from 0.40 and 0.34 at impact to about 0.55 at year ten and 0.6 at year twenty. The dampening is largest when the network has to answer the shock with its existing capacity, and it fades slowly as capital moves. The cumulative response of GDP over ten years, which is the one number to take for the dynamic effect of capacity, is 0.74 of the no-capacity economy's after the upstream shock and 0.77 after the downstream one, within a few hundredths of the impact ratios. Over twenty years the ratios are 0.75 and 0.78, and for the other nodes' output over ten years they are 0.49 and 0.45. What the static block says at impact therefore survives the dynamics almost unchanged.

Capital follows utilization, and it goes to different nodes with capacity than without. Aggregate investment rises with output on impact, by 0.14 per cent with capacity against 0.20 per cent without after the upstream shock, in the same proportion as GDP, and it decays with the shock. A 1 per cent shock with a half-life of 3.5 years moves little capital in either economy. What differs is where the capital goes. The share of aggregate capital that changes node peaks around year four: after the upstream shock at 0.0048 per cent with capacity against 0.0037 per cent without, and after the downstream shock at 0.0048 per cent with capacity against 0.0054 per cent without. The same shock therefore reallocates 1.3 times more capital across nodes with capacity

than without when it hits the upstream node, and about a tenth less when it hits the downstream one. The reason is that the return to capacity depends on the state of the node. In a node that the shock tightens, more firms are rationed and the rent that capacity earns rises, so the node attracts capital; a node that the shock loosens does not. The paths show this. After the shock to primary metals, the node itself and fabricated metals, its main customer, gain capital with capacity, whereas without capacity primary metals sheds capital, as a node that has become more productive and needs less of it. After the shock to food services the node sheds capital in both economies and the business-service nodes gain. This is the mechanism of Section 5.1 in motion: the long-run allocation of Table 17 is the accumulated response of investment to this return.

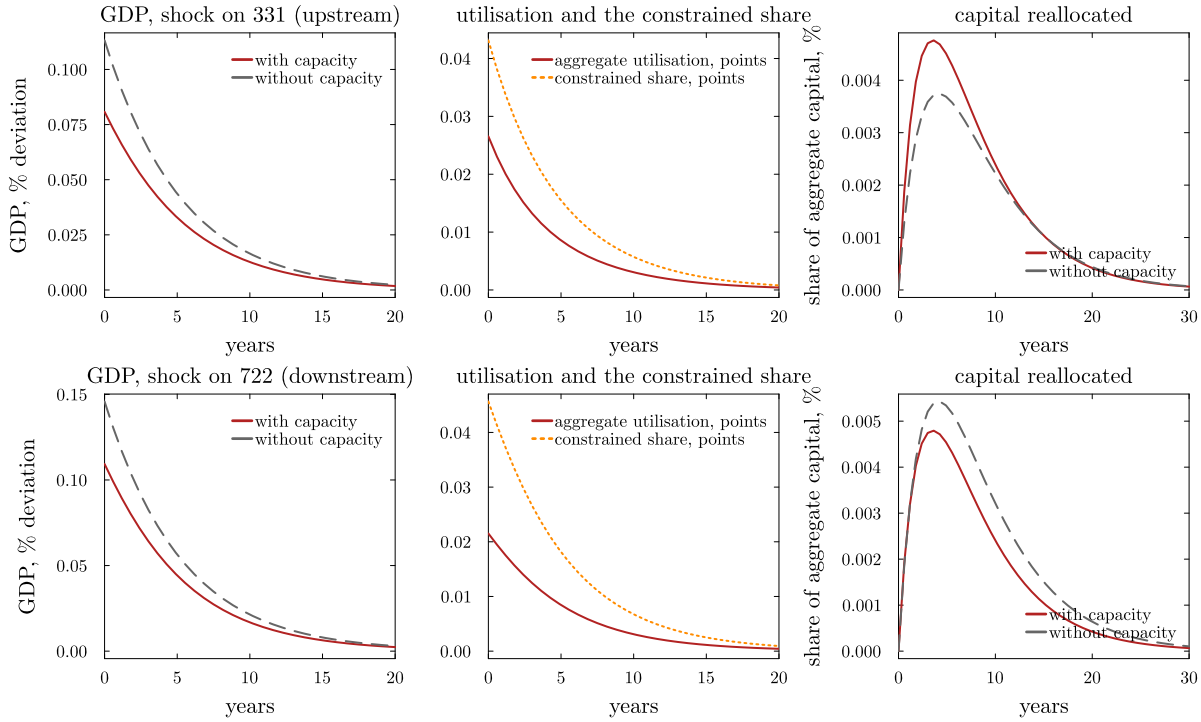
Table 6: Transitions after a sectoral productivity shock.

shock of +1 % on	upstream (331)		downstream (722)	
	with capacity	without	with capacity	without
GDP, impact response, %	0.081	0.113	0.109	0.146
GDP, peak response, %	0.081	0.113	0.109	0.146
year of the peak	0.00	0.00	0.00	0.00
GDP, cumulative response over 10 years, %-years	0.369	0.498	0.497	0.643
GDP, cumulative response over 20 years	0.430	0.578	0.579	0.747
consumption, cumulative response over 10 years	0.122	0.195	0.164	0.255
other nodes' real output, cumulative over 10 years	0.214	0.439	0.208	0.462
aggregate utilisation, impact, points	0.027	–	0.022	–
share of firms rationed, largest change, points	0.043	–	0.046	–
capital reallocated, largest along the path, % of capital	0.0048	0.0037	0.0048	0.0054
capital reallocated at year 10, % of capital	0.0023	0.0022	0.0023	0.0031
path iterations	34	40	45	48
with capacity over without				
GDP, impact	0.714		0.750	
GDP, cumulative over 10 years	0.741		0.774	
GDP, cumulative over 20 years	0.745		0.776	
other nodes' real output, cumulative over 10 years	0.486		0.451	
capital reallocated, largest along the path	1.272		0.883	

Notes: Transitions after a +1 per cent productivity shock with a half-life of 3.5 years at one node, with capacity and without: the impact and peak responses of GDP, the cumulative responses over ten and twenty years, the movement of aggregate utilization, and the share of aggregate capital that changes node.

[Evidence] Four 63-node transitions (relative tolerance 2×10^{-4}; 34 to 48 path iterations; 31 to 51 minutes each), all converged; the paths, the per-node paths and the summary statistics are stored; computed once on 2–3 November 2026 on the baseline calibration at $\phi = 2$, $\gamma = 1$ (on the previous baseline, $\phi = 13$, the reallocation peaked at

Figure 9: Transition paths after a sectoral productivity shock.



Notes: A +1 per cent productivity shock at one node. For each shocked node: the response of GDP, the responses of aggregate utilization and of the share of rationed firms, and the share of aggregate capital that changes node, with capacity and without.

year seven at a quarter of the size and the consumption responses were a third larger; the GDP paths are the same to three digits). The ratios along the path (years 10 and 20) and the investment responses were recomputed from the stored paths on 3 November 2026. No transition after a common shock exists; no transition-based welfare measure exists.

7 Where Capacity Matters: The Value of Capacity by Sector

The results so far compare economies. This section asks a question inside the calibrated economy that a policymaker would ask: in which sectors is capacity worth most? The model gives a sharp answer, because the value of a sector's capacity to the economy and its value to the sector's owners are different objects. To the owners, one more unit of capital in sector n earns the private rental flow of (20), the capital share of firm revenue per unit of capital, which includes the relief of the firm's own constraint in the states in which the firm is rationed. To the economy, one more unit of capacity in sector n also spares the sector's customers the rationing they would otherwise suffer, and through them their customers, and it lowers the shadow price index that the rest of the network pays for the sector's good. The difference is a network externality of

capacity. It is the model's counterpart of the distortion-centrality logic of Liu (2019): a wedge compounds along the chain, so relieving it where many chains pass is worth more than its local cost.

The experiment that measures the two values is the following. Hold the distribution of industries within every node at its balanced-growth form, expand one node's capital, and with it its capacity, by one per cent, re-solve the static block and read the change in GDP. The elasticity of GDP to the node's capital is the social return to its capacity. The private return is the node's capital income as a share of GDP, that is, its state-dependent capital share of revenue times its Domar weight; this is what the elasticity would be in an economy without wedges. The ratio of the social to the private return is the value of the node's capacity to the economy per unit of its value to its owners. The same experiment in the economy without capacity constraints gives the part of the ratio that the markup wedge alone produces, and the difference between the two economies is the part due to capacity. The experiment is static, it holds the distribution of industries fixed, and it does not say how the extra capital is paid for. It ranks sectors; it does not solve a planner's problem.

The results are in Figure 10. Panel (a) plots the ratio of the social to the private return by node against the node's upstreamness, in the capacity economy and in the economy without capacity constraints. The ratio exceeds one in every node in both economies, because the markup wedge compounds along the chain and makes a unit of capital worth more to the economy than to its owners, and it rises with upstreamness in both, from about 1.1 in the downstream nodes to 2 and more in the upstream ones. With capacity the mean ratio across the capacity-on nodes is 1.9 against 1.8 without, and the range is wider: from 1.2 to 3.1 with capacity against 1.25 to 2.1 without. The aggregate is pinned down by constant returns: the elasticities of GDP to the nodes' capital sum to one in both economies, so capacity does not change how much the economy's capital is worth in total; it changes which sectors' capital is worth most. Panel (b) isolates the part due to capacity: the elasticity of GDP to a node's capital with capacity over the same elasticity without, plotted against the node's capital per unit of sales. This ratio runs from 0.7 to 5.4 and is above one in about half of the capacity-on nodes. Two characteristics of a node explain it, with an R^2 of 0.75 across the capacity-on nodes: the ratio is higher where more of the node's firms are rationed, and it is higher where the node uses little capital per unit of sales. The second characteristic is the capacity logic of the model. A unit of capital in a node is an input, worth its capital share, and it is also capacity, worth the rationing it relieves, and the second part does not scale with the capital share. Capacity therefore raises the value of capital most in tight sectors that use little capital per unit of sales: computer systems design, management of companies, professional and administrative services, warehousing, trucking, paper and plastics, where the elasticity with capacity is two to five times the one without. It lowers the relative value of capital in the capital-intensive sectors, utilities, pipelines, oil and gas, rental and leasing, and real estate, where the input part dominates and the adding-up to one tilts the value away from them.

The ranking also says what targeting is worth. A uniform expansion of every

node's capital by one per cent raises GDP by one per cent in both economies. The same amount of aggregate capital put into the five nodes with the largest gain in GDP per unit of capital raises GDP by 1.7 per cent with capacity and by 1.5 per cent without. Targeting pays in both economies because the markup wedge makes capital worth more than its private return everywhere; most of the gain from targeting is therefore not a result of capacity, and the part that is, 0.2 percentage points, comes from the nodes that capacity singles out. The five nodes with capacity are computer systems design, administrative and support services, truck transportation, miscellaneous professional services, and plastics and rubber products; without capacity, data processing and rental and leasing replace the last two. The targeted expansion raises the impact response of GDP to a common shock with capacity, as a ratio of the response without, from 0.701 to 0.714: relieving the tightest sectors makes the economy answer a common shock a little more with output and a little less with rationing. A one per cent expansion of capital is a small change to a margin of 0.7, however, and the paper claims the direction of this effect and not more.

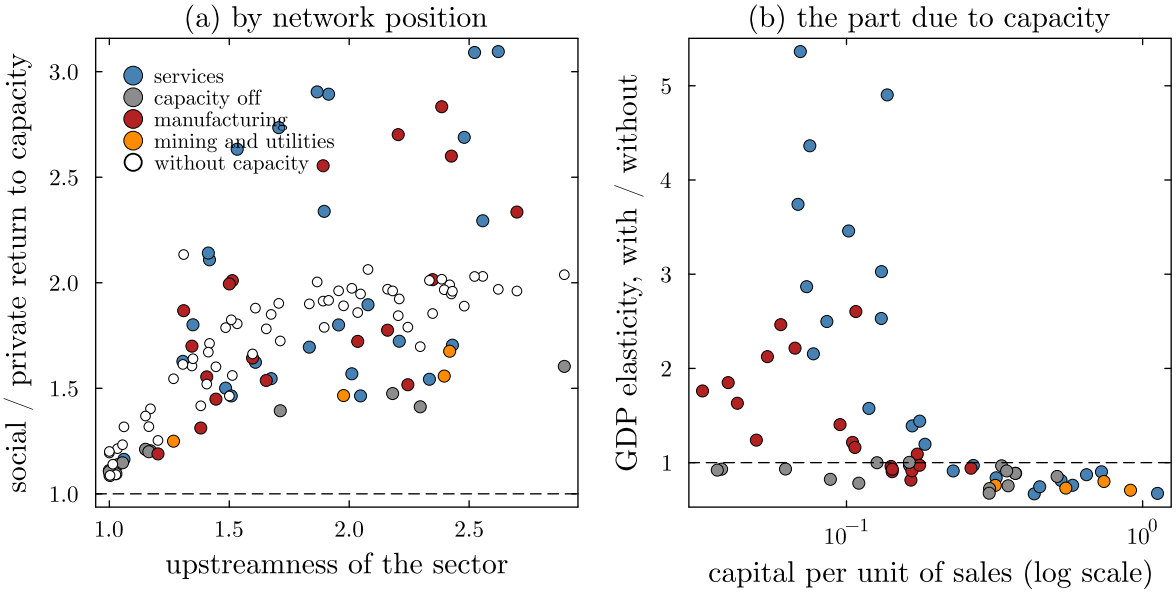
[What this section delivers, and the decision it needs] The capacity-specific result is panel (b): capacity re-ranks the sectors by the value of their capital to the economy, and the re-ranking has a two-factor explanation (tightness and capital per unit of sales) that follows from the model's logic of capital as capacity. The two headline policy numbers are not capacity-specific: the gain from targeting is nearly as large without capacity (1.5 against 1.7 per cent; it is mostly the markup wedge; on the previous baseline, $\phi = 13$, it was larger without, 2.5 against 2.3), and the effect of a one per cent targeted expansion on the dampening of a common shock is small (0.701 to 0.714). Three ways to go: (i) keep the section short, as written, with the re-ranking as its result; (ii) fold panel (b) into the results section as "the value of capacity by sector" and drop the targeting paragraph; (iii) change the instrument to one that is dynamic, a subsidy to capacity in the tight sectors along a transition after a common negative shock, which would measure what capacity buys in a downturn (cost: four transitions, about four hours of compute). The section as written claims no optimality, no welfare number and no dynamic effect.

[Evidence] The per-node and summary results of the capacity expansion are stored; computed once on 3 November 2026 on the baseline calibration at $\phi = 2$, $\gamma = 1$ (24 minutes); the experiment shocks in levels of ± 1 per cent (the units note in the box on the fixed-threshold static, inside the comparative-statics box at the end of Section 5). The means, ranges, the R^2 and the sector lists were recomputed on 3 November 2026 (the R^2 of 0.75 is for the log of the ratio; for the ratio itself it is 0.61).

8 Conclusion

This paper puts an occasionally binding capacity constraint into a dynamic production network and solves the resulting economy as a graphon mean-field game. The constraint changes the economy in one way that shows up three times. Capacity is a

Figure 10: The value of capacity by sector.



Notes: (a) The elasticity of GDP to a node's capital, the distribution of industries held fixed, over the node's capital income as a share of GDP, in the capacity economy (filled) and in the economy without capacity constraints (open), against the node's upstreamness. (b) The part due to capacity: the elasticity with capacity over the same elasticity without, against the node's capital per unit of sales (log scale); grey, nodes with capacity off.

state-dependent supply margin in each sector: a tight sector meets extra demand with rationing and a higher markup rather than with output. In the long run this costs output: GDP per unit of capital is about 10 per cent lower in the calibrated economy than in the same economy without the constraint, and most of the loss is due to the wedges that rationing puts on constrained links. It also reallocates capital, towards the sectors that need it to produce their sales rather than towards the sectors in which it earns the largest Cobb–Douglas return. At the date of a sectoral shock, capacity dampens the propagation of the shock to the rest of the network: it replaces the quantity response of the other sectors with rationing while leaving their price response unchanged, and the dampening is strongest for shocks that originate upstream and in tight sectors. The response is asymmetric in sign and nonlinear in size when the whole economy moves along its capacity margin, as under a common shock, and linear otherwise. Over time the utilization margin absorbs the shock first and capital follows.

Two parameters set the size of all of this, and the model separates what each one does. The elasticity of substitution between a sector's firms sets the price of rationing, the loss from redirecting unmet demand to the firms with spare capacity. The dispersion of firm-level demand sets its mass, the share of firms over their cap, which the utilization data pin down almost regardless of the elasticity. The two pieces of evidence that identify them, the share of plants close to capacity and the convexity of supply curves in utilization, measure exactly these two channels, and the calibrated economy sits where the price of rationing is high. The same separation says where capacity is worth most. A sector's capacity spares its customers the rationing they

would otherwise suffer, so its value to the economy exceeds its value to its owners, and capacity re-ranks the sectors by this value: capital is worth most to the economy in tight sectors that use little of it per unit of sales.

Three things the results rest on should be said plainly. First, the size of every long-run and impact effect depends on how tight the service sectors are, and their utilization rates are transferred from European surveys rather than measured in the United States; the calibration in which every sector is as tight as manufacturing, deferred here to the robustness analysis, gives smaller effects. Second, the model has no labour and no overhead cost, so its GDP is the income of physical capital and a sector's capital share is the whole of its non-input cost; the floored capital shares of twelve sectors are the mark this leaves, and an overhead cost would be the natural extension. Third, the economy has no aggregate risk, by construction of the limit that makes it tractable; the common shocks of Section 6.2 are one-time experiments, not a business cycle, and the systematic comparison of the finite network with its mean-field limit is left to later work.

What the results say for the questions the paper asks is this. Capacity constraints change propagation not by amplifying shocks but by changing what a shock becomes as it travels through the network: rationing and rents take the place of quantities, and most of all for shocks that originate upstream. Capacity is allocated across the network on a different principle from capital in the standard model, so a model without the constraint, fitted to the same accounts, reads the accounts with the wrong technology and the network with the wrong supply curves. And the mean-field formulation is what makes it possible to say so for an economy with as many sectors as the accounts have.

References

Acemoglu, D., Akcigit, U. & Kerr, W. (2016). Networks and the macroeconomy: An empirical exploration. *NBER Macroeconomics Annual*, *30*(1), 273–335.

Acemoglu, D., Autor, D. & Patterson, C. (2023). *Bottlenecks: Sectoral imbalances and the US productivity slowdown* (NBER Working Paper No. 31427). Cambridge, MA: National Bureau of Economic Research.

Acemoglu, D., Carvalho, V. M., Ozdaglar, A. & Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, *80*(5), 1977–2016.

Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L. & Moll, B. (2022). Income and wealth distribution in macroeconomics: A continuous-time approach. *Review of Economic Studies*, *89*(1), 45–86.

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, *109*(3), 659–684.

Amberg, N., Friberg, R., Liu, X. & Syverson, C. (2026). *Have we got news for you: Theory and evidence on firms' optimal expected capacity utilization* (Working Paper No. 466). Stockholm: Sveriges Riksbank.

Antràs, P., Chor, D., Fally, T. & Hillberry, R. (2012). Measuring the upstreamness of production and trade flows. *American Economic Review*, *102*(3), 412–416.

Baqaee, D. R. & Farhi, E. (2019). The macroeconomic impact of microeconomic shocks: Beyond hulten's theorem. *Econometrica*, *87*(4), 1155–1203.

Baqaee, D. R. & Farhi, E. (2020). Productivity and misallocation in general equilibrium. *The Quarterly Journal of Economics*, *135*(1), 105–163.

Barkai, S. (2020). Declining labor and capital shares. *The Journal of Finance*, *75*(5), 2421–2463.

Bayraktar, E., Chakraborty, S. & Wu, R. (2023). Graphon mean field systems. *Annals of Applied Probability*, *33*(5), 3587–3619.

Bigio, S. & La'O, J. (2020). Distortions in production networks. *The Quarterly Journal of Economics*, *135*(4), 2187–2253.

Boehm, C. E., Flaaen, A. & Pandalai-Nayar, N. (2019). Input linkages and the transmission of shocks: Firm-level evidence from the 2011 Tōhoku earthquake. *The Review of Economics and Statistics*, *101*(1), 60–75.

Boehm, C. E. & Pandalai-Nayar, N. (2022). Convex supply curves. *American Economic Review*, *112*(12), 3941–3969.

Bohr, C. E. (2024). Capacity buffers: Explaining the retreat and return of the phillips curve.

Caines, P. E., Ho, D., Huang, M., Jian, J. & Song, Q. (2022). On the graphon mean field game equations: Individual agent affine dynamics and mean field dependent performance functions. *ESAIM: Control, Optimisation and Calculus of Variations*, *28*, 24.

Caines, P. E. & Huang, M. (2021). Graphon mean field games and their equations. *SIAM Journal on Control and Optimization*, *59*(6), 4373–4399.

Carvalho, V. M., Nirei, M., Saito, Y. U. & Tahbaz-Salehi, A. (2021). Supply chain disruptions: Evidence from the Great East Japan Earthquake. *The Quarterly Journal of Economics*, *136*(2), 1255–1321.

Carvalho, V. M. & Tahbaz-Salehi, A. (2019). Production networks: A primer. *Annual Review of Economics*, *11*, 635–663.

Comin, D., Lashkari, D. & Mestieri, M. (2021). Structural change with long-run income and price effects. *Econometrica*, *89*(1), 311–374.

Comin, D. A., Johnson, R. C. & Jones, C. J. (2023). Supply chain constraints and inflation.

Cooper, R. W. & Haltiwanger, J. C. (2006). On the nature of capital adjustment costs. *The Review of Economic Studies*, *73*(3), 611–633.

di Giovanni, J., Kalemli-Özcan, Ş., Silva, A. & Yildirim, M. A. (2022). *Global supply chain pressures, international trade, and inflation* (NBER Working Paper No. 30240). Cambridge, MA: National Bureau of Economic Research.

Eden, B. (1990). Marginal cost pricing when spot markets clear sequentially. *Journal of Political Economy*, *98*(6), 1293–1306.

Edmond, C., Midrigan, V. & Xu, D. Y. (2023). How costly are markups? *Journal of Political Economy*, *131*(7), 1619–1675.

Fagnart, J.-F., Licandro, O. & Portier, F. (1999). Firm heterogeneity, capacity utilization, and the business cycle. *Review of Economic Dynamics*, *2*(2), 433–455.

Fally, T. (2012). *Production staging: Measurement and facts* (Working Paper). Boulder, CO: University of Colorado Boulder.

Gomme, P., Ravikumar, B. & Rupert, P. (2011). The return to capital and the business cycle. *Review of Economic Dynamics*, *14*(2), 262–278.

Hall, R. E. (2004). Measuring factor adjustment costs. *The Quarterly Journal of Economics*, *119*(3), 899–927.

Herrendorf, B., Rogerson, R. & Valentinyi, Á. (2013). Two perspectives on preferences and structural transformation. *American Economic Review*, *103*(7), 2752–2789.

Huang, M., Malhamé, R. P. & Caines, P. E. (2006). Large population stochastic dynamic games: Closed-loop McKean–Vlasov systems and the Nash certainty equivalence principle. *Communications in Information and Systems*, *6*(3), 221–252.

Kreyszig, E. (1978). *Introductory functional analysis with applications*. New York: John Wiley & Sons.

Lasry, J.-M. & Lions, P.-L. (2006a). Jeux à champ moyen. II – Horizon fini et contrôle optimal. *Comptes Rendus Mathématique*, *343*(10), 679–684.

Lasry, J.-M. & Lions, P.-L. (2006b). Jeux à champ moyen. I – Le cas stationnaire. *Comptes Rendus Mathématique*, *343*(9), 619–625.

Lasry, J.-M. & Lions, P.-L. (2007). Mean field games. *Japanese Journal of Mathematics*, *2*(1), 229–260.

Liu, E. (2019). Industrial policies in production networks. *Quarterly Journal of Economics*, *134*(4), 1883–1948.

Liu, E. & Tsyvinski, A. (2024). A dynamic model of input–output networks. *The Review of Economic Studies*, *91*(6), 3608–3644.

Long, J. B. & Plosser, C. I. (1983). Real business cycles. *Journal of Political Economy*, *91*(1), 39–69.

Lovász, L. (2012). *Large networks and graph limits* (Vol. 60). Providence, RI: American Mathematical Society.

Parise, F. & Ozdaglar, A. (2023). Graphon games: A statistical framework for network games and interventions. *Econometrica*, *91*(1), 191–225.

Prescott, E. C. (1975). Efficiency of the natural rate. *Journal of Political Economy*, *83*(6), 1229–1236.

Shea, J. (2002). Complementarities and comovements. *Journal of Money, Credit and Banking*, *34*(2), 412–433.

Sun, Y. (2006). The exact law of large numbers via Fubini extension and characterization of insurable risks. *Journal of Economic Theory*, *126*(1), 31–69.

Sun, Y. & Zhang, Y. (2009). Individual risk and Lebesgue extension without aggregate uncertainty. *Journal of Economic Theory*, *144*(1), 432–443.

Vom Lehn, C. & Winberry, T. (2022). The investment network, sectoral comovement, and the changing us business cycle. *The Quarterly Journal of Economics*, *137*(1), 387–433.

A Model Details and Proofs for Section 2

This appendix collects, in the order of Section 2, the objects that the main text uses in words or in reduced form, the assumptions it only refers to, and the results it states without proof, each result followed by its proof.

A.1 Tail Functions of the Demand-Shock Distribution

Write

$$\mu_{i,\theta} := \mathbb{E}_i[v^\theta] \equiv \int_0^\infty v^\theta \mathrm{d}G_i(v)$$

for the θ-th moment of the sector-i demand shock. The tail functions defined below require the following of G_i.

Assumption 1 (Demand shocks). For each sector i, G_i is a continuous distribution on $(0,\infty)$ with a continuous, strictly positive density g_i, mean $\mu_{i,1} = 1$, and finite moments $\mu_{i,\theta} < \infty$ for every $\theta \in [0, \frac{1}{1-\alpha_i}]$.

The positive density keeps $\eta_{i,\theta}(\overline{v})$ strictly inside $(0,1)$ and makes the constrained probability $1 - G_i(\overline{v})$ positive at every finite threshold; its continuity makes the tail functions defined below continuously differentiable. Continuity of g_i is needed only where a density term survives: it cancels in $\mathcal{R}_i$ and in U_i, which are continuously differentiable for any density, but not in Γ_i or in the threshold map Y_i of Section B.6, whose derivatives carry g_i explicitly and which the implicit function theorem is applied to there. The moment condition is what the expectations (6)–(12) and the limits of Lemma 4 need, and it is the same condition as in the model without capacity constraints, where $\mu_{i,\frac{1}{1-\alpha_i}}$ must be finite. The assumption is otherwise silent on the shape of G_i; Assumption 2 is the only shape restriction.

The following tail functions of G_i describe how much of the demand-shock distribution lies on either side of a cut-off $x > 0$. For $\theta \geq 0$,

$$\mathcal{G}^-_{i,\theta}(x) := \int_0^x \left(\frac{v}{x}\right)^\theta \mathrm{d}G_i(v), \qquad \mathcal{G}^+_{i,\theta}(x) := \int_x^\infty \left(\frac{v}{x}\right)^\theta \mathrm{d}G_i(v), \tag{57}$$

$$\Psi_{i,\theta}(x) := \mathbb{E}_i\left[\min\left\{\left(\frac{v}{x}\right)^\theta, 1\right\}\right] = \mathcal{G}^-_{i,\theta}(x) + 1 - G_i(x), \qquad \eta_{i,\theta}(x) := \frac{\mathcal{G}^-_{i,\theta}(x)}{\Psi_{i,\theta}(x)} \in (0,1), \tag{58}$$

where $\eta_{i,\theta}$ is the lower-tail share defined in (7). The two tails add up to the moment, $\mathcal{G}^-_{i,\theta}(x) + \mathcal{G}^+_{i,\theta}(x) = x^{-\theta}\mu_{i,\theta}$. $\mathcal{G}^-_{i,\theta}(x)$ is the θ-th moment of v/x over the lower tail $v \leq x$, $\Psi_{i,\theta}(x)$ is the same moment with the upper tail capped at one, and $\eta_{i,\theta}(x)$ is the share of $\Psi_{i,\theta}(x)$ contributed by the lower tail. $\Psi_{i,0} \equiv 1$; for $\theta > 0$, $\Psi_{i,\theta}$ is strictly decreasing with $\Psi_{i,\theta}(0^+) = 1$ and $x^\theta \Psi_{i,\theta}(x) \to \mu_{i,\theta}$ as $x \to \infty$ (Lemma 4), and it satisfies the derivative identity

$$\Psi'_{i,\theta}(x) = -\frac{\theta}{x}\,\mathcal{G}^-_{i,\theta}(x) = -\frac{\theta}{x}\,\eta_{i,\theta}(x)\,\Psi_{i,\theta}(x), \tag{59}$$

shown in Step 1 of the proof of Lemma 1 below. The lower-tail share $\eta_{i,\theta}(x)$ is therefore minus the elasticity of $\Psi_{i,\theta}$ with respect to x, divided by θ.

Log-normal specification. The working specification, used in every numerical statement of this paper, is $\ln v \sim \mathcal{N}(-\sigma_{v,i}^2/2, \sigma_{v,i}^2)$, so that $\mu_{i,1} = 1$ and $\mu_{i,\theta} = \exp\left[\frac{\sigma_{v,i}^2}{2}\theta(\theta-1)\right]$ for every θ. Writing $\Phi_{\mathcal{N}}$ for the standard normal distribution function, the tail functions are available in closed form:

$$\begin{aligned} G_i(\overline{v}) &= \Phi_{\mathcal{N}}\left(\frac{\ln \overline{v}}{\sigma_{v,i}} + \frac{\sigma_{v,i}}{2}\right), \\ \mathcal{G}_{i,\theta}^{-}(\overline{v}) &= \exp\left[\frac{\sigma_{v,i}^2}{2}\theta(\theta-1)\right] \overline{v}^{-\theta}\, \Phi_{\mathcal{N}}\left(\frac{\ln \overline{v}}{\sigma_{v,i}} - \left(\theta - \tfrac{1}{2}\right)\sigma_{v,i}\right), \\ \mathcal{G}_{i,\theta}^{+}(\overline{v}) &= \exp\left[\frac{\sigma_{v,i}^2}{2}\theta(\theta-1)\right] \overline{v}^{-\theta}\left[1 - \Phi_{\mathcal{N}}\left(\frac{\ln \overline{v}}{\sigma_{v,i}} - \left(\theta - \tfrac{1}{2}\right)\sigma_{v,i}\right)\right], \end{aligned} \tag{60}$$

and the capacity constant of Definition 1 is $\mu_{i,1/\varepsilon}^{\varepsilon/(\varepsilon-1)} = \exp\left(-\sigma_{v,i}^2/2\varepsilon\right)$. The closed forms follow from $\int_0^{\overline{v}} v^\theta \mathrm{d}G_i = \mathbb{E}_i[v^\theta\, \mathbb{1}\{v \le \overline{v}\}]$ for a log-normal v and are checked against quadrature in the numerical work. Under this specification $\sigma_{v,i}$ is the one distributional parameter of sector i, and it enters the no-capacity model only through $\mu_{i,\frac{1}{1-\alpha_i}} = \exp\left[\frac{\sigma_{v,i}^2}{2}\frac{\alpha_i}{(1-\alpha_i)^2}\right]$.

A.2 Rationing within a Sector

Lemma 1 (Rationing within a sector: demand, shadow price index, threshold). *Fix $Y_i > 0$ and measurable, strictly positive prices $\{p_{in}\}$ and capacities $\{\overline{y}_{in}\}$. The aggregator's cost minimisation has the following solution.*

(i) (Demand and the shadow price index.) *There are multipliers $\varphi_{in} \ge 0$ on the capacity constraints, with $\varphi_{in} = 0$ whenever $y_{in} < \overline{y}_{in}$, such that* (3) *and* (4) *hold. P_i is the aggregator's marginal cost of a unit of Y_i, and $P_i Y_i = \int_0^1 (p_{in} + \varphi_{in}) y_{in} \mathrm{d}n$ exceeds the amount actually paid to the firms, $\int_0^1 p_{in} y_{in} \mathrm{d}n$, by $\int_0^1 \varphi_{in} y_{in} \mathrm{d}n \ge 0$.*

(ii) (Threshold representation.) *With the threshold $\overline{v}_{in}$ of* (5)*, the solution is*

$$y_{in} = \overline{y}_{in} \min\left\{\frac{v_{in}}{\overline{v}_{in}}, 1\right\}, \qquad \varphi_{in} = p_{in}\left[\left(\frac{v_{in}}{\overline{v}_{in}}\right)^{1/\varepsilon} - 1\right]^+ : \tag{61}$$

firm n is demand-constrained when $v_{in} < \overline{v}_{in}$ and capacity-constrained when $v_{in} \ge \overline{v}_{in}$, which happens with probability $1 - G_i(\overline{v}_{in})$.

(iii) (Expected output and the effective demand elasticity.) *Before v_{in} is realised, $\mathbb{E}_v[y_{in}^\theta] = \overline{y}_{in}^\theta \Psi_{i,\theta}(\overline{v}_{in})$ for every $\theta \ge 0$; in particular $\mathbb{E}_v[y_{in}] = \overline{y}_{in} \Psi_{i,1}(\overline{v}_{in})$, and the elasticity of expected demand with respect to the own price, holding (Y_i, P_i) fixed, is*

$$\varepsilon_{in} := -\frac{\partial \ln \mathbb{E}_v[y_{in}]}{\partial \ln p_{in}} = \varepsilon\, \eta_{i,1}(\overline{v}_{in}) \in (0, \varepsilon). \tag{62}$$

Proof. Throughout, i is fixed and $F(y) := [\int_0^1 v_n^{1/\varepsilon} y_n^{(\varepsilon-1)/\varepsilon} \mathrm{d}n]^{\varepsilon/(\varepsilon-1)}$ denotes the aggregator (2) as a function of the profile $y = \{y_n\}$.

Step 1 (the derivative identity (59)*).* Write $\mathcal{G}^-_{i,\theta}(x) = x^{-\theta} \int_0^x v^\theta \mathrm{d}G_i(v)$. Under Assumption 1 G_i has a density g_i, so both factors are differentiable and

$$\frac{\mathrm{d}}{\mathrm{d}x}\mathcal{G}^-_{i,\theta}(x) = -\frac{\theta}{x}\, x^{-\theta} \int_0^x v^\theta \mathrm{d}G_i(v) + x^{-\theta} x^\theta g_i(x) = -\frac{\theta}{x}\mathcal{G}^-_{i,\theta}(x) + g_i(x),$$

and $\mathrm{d}[1 - G_i(x)]/\mathrm{d}x = -g_i(x)$. Adding the two gives $\Psi'_{i,\theta}(x) = -\frac{\theta}{x}\mathcal{G}^-_{i,\theta}(x)$, and $\mathcal{G}^-_{i,\theta} = \eta_{i,\theta}\Psi_{i,\theta}$ by definition. For $\theta > 0$ the derivative is strictly negative because $\mathcal{G}^-_{i,\theta}(x) > 0$ when $g_i > 0$ on $(0, x)$; $\Psi_{i,\theta}(0^+) = 1 - G_i(0^+) = 1$ since $\mathcal{G}^-_{i,\theta}(x) \le G_i(x) \to 0$.

Step 2 (the aggregator's problem and its multipliers). The aggregator solves $\min_y \int_0^1 p_n y_n \mathrm{d}n$ subject to $F(y) \ge Y_i$ and $0 \le y_n \le \overline{y}_n$. The objective is linear, F is concave (it is homogeneous of degree one and quasi-concave, being a CES aggregator with $\varepsilon > 1$) so the constraint set is convex, and the problem has a solution because the feasible set is nonempty when $F(\overline{y}) \ge Y_i$, which we assume, and closed. The Karush–Kuhn–Tucker conditions are therefore necessary and sufficient: with a multiplier $\lambda \ge 0$ on the output constraint and multipliers $\varphi_n \ge 0$ on the capacity constraints,

$$p_n + \varphi_n = \lambda \frac{\partial F}{\partial y_n} = \lambda\, v_n^{1/\varepsilon}\, y_n^{-1/\varepsilon}\, F(y)^{1/\varepsilon}, \qquad \varphi_n(\overline{y}_n - y_n) = 0, \qquad \lambda(F(y) - Y_i) = 0,$$

where the first-order condition uses $\partial F/\partial y_n = v_n^{1/\varepsilon} y_n^{-1/\varepsilon} F^{1/\varepsilon}$. Since prices are strictly positive, $\lambda > 0$ and the output constraint binds, $F(y) = Y_i$; the lower bound $y_n \ge 0$ is never active because $\partial F/\partial y_n \to \infty$ as $y_n \to 0$. Solving the first-order condition for y_n gives $y_n = v_n Y_i [(p_n + \varphi_n)/\lambda]^{-\varepsilon}$.

Step 3 (the multiplier on output is the price index). Substituting the solution into $F(y) = Y_i$,

$$Y_i^{\frac{\varepsilon-1}{\varepsilon}} = \int_0^1 v_n^{1/\varepsilon} (v_n Y_i)^{\frac{\varepsilon-1}{\varepsilon}} \left(\frac{p_n + \varphi_n}{\lambda}\right)^{1-\varepsilon} \mathrm{d}n \iff 1 = \lambda^{\varepsilon-1} \int_0^1 v_n (p_n + \varphi_n)^{1-\varepsilon} \mathrm{d}n,$$

so $\lambda = [\int_0^1 v_n (p_n + \varphi_n)^{1-\varepsilon} \mathrm{d}n]^{1/(1-\varepsilon)} =: P_i$, which is (3). As the multiplier on the output constraint of a cost minimisation, P_i is the marginal cost of Y_i. Multiplying the demand by $p_n + \varphi_n$ and integrating, $\int_0^1 (p_n + \varphi_n) y_n \mathrm{d}n = Y_i P_i^\varepsilon \int_0^1 v_n (p_n + \varphi_n)^{1-\varepsilon} \mathrm{d}n = P_i Y_i$, so $P_i Y_i - \int_0^1 p_n y_n \mathrm{d}n = \int_0^1 \varphi_n y_n \mathrm{d}n \ge 0$. This proves (i).

Step 4 (the threshold). By complementary slackness there are two cases. If $\varphi_n = 0$, then $y_n = v_n Y_i (p_n/P_i)^{-\varepsilon}$, which is feasible if and only if $v_n Y_i (p_n/P_i)^{-\varepsilon} \le \overline{y}_n$, i.e. $v_n \le \overline{v}_n$ with $\overline{v}_n$ as in (5); note that $v_n Y_i (p_n/P_i)^{-\varepsilon} = \overline{y}_n\, v_n/\overline{v}_n$. If $\varphi_n > 0$, then $y_n = \overline{y}_n$ and the first-order condition gives $\overline{y}_n = v_n Y_i [(p_n + \varphi_n)/P_i]^{-\varepsilon}$, i.e.

$$\frac{p_n + \varphi_n}{P_i} = \left(\frac{v_n Y_i}{\overline{y}_n}\right)^{1/\varepsilon} = \left(\frac{v_n}{\overline{v}_n}\right)^{1/\varepsilon} \frac{p_n}{P_i} \iff \varphi_n = p_n \left[\left(\frac{v_n}{\overline{v}_n}\right)^{1/\varepsilon} - 1\right],$$

which is positive if and only if $v_n > \overline{v}_n$. The two cases are exhaustive and consistent, and together they are (61). Under Assumption 1 the event $v_n = \overline{v}_n$ has probability zero, and the constrained probability is $1 - G_i(\overline{v}_n)$. This proves (ii).

Step 5 (expectations and the effective elasticity). From (61), $y_n^\theta = \overline{y}_n^\theta \min\{(v_n/\overline{v}_n)^\theta, 1\}$, so $\mathbb{E}_v[y_n^\theta] = \overline{y}_n^\theta \Psi_{i,\theta}(\overline{v}_n)$ by the definition (58). Holding $(Y_i, P_i, \overline{y}_n)$ fixed, $\partial \ln \overline{v}_n / \partial \ln p_n = \varepsilon$ by (5), and by Step 1

$$\frac{\partial \ln \mathbb{E}_v[y_n]}{\partial \ln p_n} = \frac{\overline{v}_n \Psi'_{i,1}(\overline{v}_n)}{\Psi_{i,1}(\overline{v}_n)} \cdot \varepsilon = -\varepsilon\, \eta_{i,1}(\overline{v}_n),$$

which is (62); $\eta_{i,1}(\overline{v}_n) \in (0,1)$ because both $\mathcal{G}^-_{i,1}(\overline{v}_n)$ and $1 - G_i(\overline{v}_n)$ are strictly positive under the positive-density assumption. This proves (iii). ■

A.3 Pre-Shock Pricing

Uniqueness of the maximiser of the firm's pricing problem rests on a regularity condition on G_i, which is not implied by Assumption 1.

Assumption 2 (Single-peaked pricing). For each sector i, the function $\overline{v} \mapsto (\varepsilon - 1)\mathcal{G}^-_{i,1}(\overline{v}) - 1 + G_i(\overline{v})$ has exactly one zero, $\underline{v}_i \in (0, \infty)$, and

$$\overline{v} \;\longmapsto\; \overline{v}^{1/\varepsilon}\, \frac{(\varepsilon - 1)\mathcal{G}^-_{i,1}(\overline{v}) - 1 + G_i(\overline{v})}{\mathcal{G}^-_{i,\frac{1}{1-\alpha_i}}(\overline{v})}$$

is strictly increasing on $(\underline{v}_i, \infty)$.

The numerator is positive exactly when the effective demand elasticity (62) exceeds one, so the assumption says that the set of thresholds at which the firm faces elastic expected demand is an upper interval, and that on it the displayed ratio is monotone. It is a condition on G_i and (ε, α_i) alone. It is automatically satisfied far in the elastic region: as $\overline{v} \to \infty$ the ratio behaves like $\frac{\varepsilon - 1}{\mu_{i,\frac{1}{1-\alpha_i}}} \overline{v}^{\frac{1}{\varepsilon} + \frac{\alpha_i}{1-\alpha_i}}$ (Lemma 4), a strictly increasing power, which is the form the condition takes in the model without capacity constraints. **[Flag: Assumption 2 is verified numerically, not proved.]** For the log-normal specification of G_i below the assumption was checked on a grid of 294 parameter triples, $\sigma_{v,i} \in \{0.1, \ldots, 2\}$, $\varepsilon \in \{1.2, \ldots, 15\}$, $\alpha_i \in \{0.1, \ldots, 0.85\}$, on a fine logarithmic grid of thresholds $\overline{v} \in [e^{-10}, e^{12}]$: in every triple the numerator has a single zero and the ratio is strictly increasing on $(\underline{v}_i, \infty)$. The ratio is *not* monotone on $(0, \underline{v}_i)$ for some triples (large $\sigma_{v,i}$ with ε close to one), which is why the assumption is stated on the elastic region only; that region is the only one that matters, because profit is increasing in the price wherever expected demand is inelastic (Step 4 of the proof of Lemma 2). Independently, in 720 random pricing problems the profit function was single-peaked in the price and the first-order condition had exactly one root in the elastic region. **[Q12 | text | tag: LogNormal]** Prove Assumption 2 for the log-normal specification, or give a sufficient condition on G_i (a monotone-hazard or log-concavity

condition on the lower tail is the natural candidate) under which it holds; then either replace the assumption by a lemma or keep it and cite the sufficient condition in the text after it.

Lemma 2 (Pre-shock pricing under capacity: a variable markup over effective expected marginal cost). *Let $\varepsilon > 1$ and $\alpha_i \in (0,1)$, fix $(Y_i, P_i, \mathcal{P}_i(\mathbf{P}), Z_i)$ and own capital $k_{in} > 0$, and let G_i satisfy Assumptions 1 and 2. Then:*

(i) The pricing problem has a unique maximiser $p_{in} \in (0,\infty)$, at which expected profit is strictly positive and the effective demand elasticity exceeds one, $\varepsilon_{in} = \varepsilon\,\eta_{i,1}(\overline{v}_{in}) > 1$. The maximiser satisfies

$$p_{in} = \frac{\varepsilon_{in}}{\varepsilon_{in}-1}\,\widetilde{mc}_{in},$$
$$\widetilde{mc}_{in} := \frac{1}{1-\alpha_i}\,\mathcal{P}_i(\mathbf{P})\,\frac{\eta_{i,\frac{1}{1-\alpha_i}}(\overline{v}_{in})\,\mathbb{E}_v[M_{in}]}{\eta_{i,1}(\overline{v}_{in})\,\mathbb{E}_v[y_{in}]} = \frac{1}{1-\alpha_i}\,\frac{\mathcal{P}_i(\mathbf{P})}{Z_i x_i^{\alpha_i}}\,\frac{\mathcal{G}^-_{i,\frac{1}{1-\alpha_i}}(\overline{v}_{in})}{\mathcal{G}^-_{i,1}(\overline{v}_{in})}. \tag{63}$$

(ii) Equivalently, with the state-dependent wedge

$$\Gamma_i(\overline{v}) := \underbrace{\frac{\varepsilon\,\eta_{i,1}(\overline{v})}{\varepsilon\,\eta_{i,1}(\overline{v})-1}}_{\text{markup}} \times \underbrace{\frac{1}{1-\alpha_i}}_{\text{diminishing returns}} \times \underbrace{\frac{\overline{v}^{\frac{1}{1-\alpha_i}}\,\mathcal{G}^-_{i,\frac{1}{1-\alpha_i}}(\overline{v})}{\overline{v}\,\mathcal{G}^-_{i,1}(\overline{v})}}_{\text{demand risk on the unconstrained tail}}$$
$$= \frac{\varepsilon\frac{1}{1-\alpha_i}\mathcal{G}^-_{i,\frac{1}{1-\alpha_i}}(\overline{v})}{(\varepsilon-1)\mathcal{G}^-_{i,1}(\overline{v}) - 1 + G_i(\overline{v})}\,\overline{v}^{\frac{\alpha_i}{1-\alpha_i}}, \tag{64}$$

the price satisfies

$$p_{in} = \Gamma_i(\overline{v}_{in})\,\frac{\mathcal{P}_i(\mathbf{P})}{Z_i x_i^{\alpha_i}}\,\overline{v}_{in}^{\,-\frac{\alpha_i}{1-\alpha_i}} = \Gamma_i(\overline{v}_{in})\,\mathcal{P}_i(\mathbf{P})\,Z_i^{-\frac{1}{1-\alpha_i}}\left(\frac{Y_i(p_{in}/P_i)^{-\varepsilon}}{k_{in}}\right)^{\frac{\alpha_i}{1-\alpha_i}}, \tag{65}$$

which, since $\overline{v}_{in} \propto p_{in}^{\varepsilon}$ by (5), pins down p_{in} implicitly. The threshold at the optimum is the unique root in $(\underline{v}_i, \infty)$ of the scalar equation

$$\overline{v}^{1/\varepsilon}\,\frac{(\varepsilon-1)\mathcal{G}^-_{i,1}(\overline{v}) - 1 + G_i(\overline{v})}{\mathcal{G}^-_{i,\frac{1}{1-\alpha_i}}(\overline{v})} = \frac{\varepsilon}{1-\alpha_i}\,\frac{\mathcal{P}_i(\mathbf{P})}{Z_i x_i^{\alpha_i} P_i}\left(\frac{\overline{y}_{in}}{Y_i}\right)^{1/\varepsilon}, \tag{66}$$

whose right-hand side is the only place where k_{in} (through $\overline{y}_{in}$) and the aggregates enter.

(iii) The markup $\frac{\varepsilon_{in}}{\varepsilon_{in}-1} > \frac{\varepsilon}{\varepsilon-1}$ and the wedge $\Gamma_i(\overline{v}_{in})$ depend on the state $(k_{in}, \mathbf{K}, \mathbf{Z})$ only through $\overline{v}_{in}$, and on the demand-shock distribution only through the tail functions of G_i at $\overline{v}_{in}$.

(iv) As $\overline{v} \to \infty$, $\varepsilon_i(\overline{v}) := \varepsilon\,\eta_{i,1}(\overline{v}) \to \varepsilon$ and $\Gamma_i(\overline{v}) \to \Gamma_i := \frac{\varepsilon}{\varepsilon-1}\frac{1}{1-\alpha_i}\mu_{i,\frac{1}{1-\alpha_i}}$ (Lemma 4), and (65) becomes the constant-markup pricing rule of the no-capacity model (Lemma 21).

Proof. Write $p := p_{in}$, $k := k_{in}$, $\overline{y} := \overline{y}_{in} = Z_i k x_i^{\alpha_i - 1}$, drop the sector index on the tail functions, and let $\Pi(p) := p\,\mathbb{E}_v[y_{in}] - \mathcal{P}_i\,\mathbb{E}_v[M_{in}]$ be expected operating profit: capital is fixed within the instant, so the bundle is the only cost.

Step 1 (the objective as a function of the threshold). By (5), $\overline{v}(p) = \frac{\overline{y}}{Y_i}(p/P_i)^{\varepsilon}$ is a strictly increasing bijection of $(0,\infty)$ onto itself, with inverse $p(\overline{v}) = P_i(\overline{v}Y_i/\overline{y})^{1/\varepsilon}$. By (6)–(12),

$$\Pi = p(\overline{v})\,\overline{y}\,\Psi_1(\overline{v}) - \mathcal{P}_i\frac{k}{x_i}\Psi_{\frac{1}{1-\alpha_i}}(\overline{v}),$$

so maximising over p is the same as maximising over $\overline{v}$.

Step 2 (derivatives). Holding (Y_i, P_i) fixed, $\partial \ln \overline{v}/\partial \ln p = \varepsilon$ and, at fixed p, $\partial \ln \overline{v}/\partial \ln k = \partial \ln \overline{y}/\partial \ln k = 1$ by (10). With the derivative identity (59),

$$\frac{\partial \mathbb{E}_v[y_{in}]}{\partial p} = \overline{y}\,\Psi_1'(\overline{v})\,\frac{\varepsilon\overline{v}}{p} = -\varepsilon\,\eta_1(\overline{v})\,\frac{\mathbb{E}_v[y_{in}]}{p}, \qquad \frac{\partial \mathbb{E}_v[M_{in}]}{\partial p} = -\frac{\varepsilon}{1-\alpha_i}\,\eta_{\frac{1}{1-\alpha_i}}(\overline{v})\,\frac{\mathbb{E}_v[M_{in}]}{p}, \tag{67}$$

$$\frac{\partial \mathbb{E}_v[y_{in}]}{\partial k} = \frac{\overline{y}}{k}\Psi_1(\overline{v}) + \overline{y}\,\Psi_1'(\overline{v})\frac{\overline{v}}{k} = [1-\eta_1(\overline{v})]\frac{\mathbb{E}_v[y_{in}]}{k}, \tag{68}$$

$$\frac{\partial \mathbb{E}_v[M_{in}]}{\partial k} = \frac{1}{x_i}\Psi_{\frac{1}{1-\alpha_i}}(\overline{v}) + \frac{k}{x_i}\,\Psi'_{\frac{1}{1-\alpha_i}}(\overline{v})\frac{\overline{v}}{k} = \left[1 - \frac{1}{1-\alpha_i}\eta_{\frac{1}{1-\alpha_i}}(\overline{v})\right]\frac{\mathbb{E}_v[M_{in}]}{k}.$$

The tag (68) refers to the pair of capital derivatives, the tagged line and the untagged one below it. The price derivatives are the no-capacity model's constant elasticities ε and $\frac{\varepsilon}{1-\alpha_i}$ scaled by the lower-tail shares, because a price change moves output and input use only in the states in which the firm is not rationed. The capital derivatives have two terms each: more capital raises capacity one for one (the first term) and raises the threshold, which lowers the truncated moments (the second).

Step 3 (first-order condition in the price). Differentiating Π and inserting (67),

$$\begin{aligned}\Pi'(p) &= \mathbb{E}_v[y_{in}] + p\,\frac{\partial \mathbb{E}_v[y_{in}]}{\partial p} - \mathcal{P}_i\,\frac{\partial \mathbb{E}_v[M_{in}]}{\partial p}\\ &= [1-\varepsilon\,\eta_1(\overline{v})]\mathbb{E}_v[y_{in}] + \frac{\varepsilon}{1-\alpha_i}\,\eta_{\frac{1}{1-\alpha_i}}(\overline{v})\,\mathcal{P}_i\,\frac{\mathbb{E}_v[M_{in}]}{p}.\end{aligned}$$

The second term is strictly positive, so at any critical point $\varepsilon\eta_1(\overline{v}) > 1$, i.e. $\varepsilon_{in} > 1$; dividing $\Pi'(p) = 0$ by $[\varepsilon\eta_1(\overline{v}) - 1]\mathbb{E}_v[y_{in}] > 0$ gives the first equality in (63), with $\widetilde{mc}_{in}$ as displayed there.

Step 4 (existence and uniqueness of the maximiser). Abbreviate $D(\overline{v}) := (\varepsilon - 1)\mathcal{G}_1^-(\overline{v}) - 1 + G_i(\overline{v})$, so that $\Psi_1 - \varepsilon\mathcal{G}_1^- = -D$ and $D > 0$ if and only if $\varepsilon\eta_1 > 1$. Differentiating the Step-1 form with respect to $\overline{v}$, using $p'(\overline{v}) = p/(\varepsilon\overline{v})$ and (59),

$$\begin{aligned}\frac{\mathrm{d}\Pi}{\mathrm{d}\overline{v}} &= \frac{p\,\overline{y}}{\varepsilon\overline{v}}\left[\Psi_1(\overline{v}) - \varepsilon\mathcal{G}_1^-(\overline{v})\right] + \frac{1}{1-\alpha_i}\,\mathcal{P}_i\frac{k}{x_i}\frac{\mathcal{G}^-_{\frac{1}{1-\alpha_i}}(\overline{v})}{\overline{v}}\\ &= \frac{1}{\overline{v}}\left[\frac{1}{1-\alpha_i}\,\mathcal{P}_i\frac{k}{x_i}\mathcal{G}^-_{\frac{1}{1-\alpha_i}}(\overline{v}) - \frac{p(\overline{v})\,\overline{y}}{\varepsilon}D(\overline{v})\right].\end{aligned}$$

Wherever $D(\overline{v}) \leq 0$ the derivative is strictly positive: profit rises with the price as long as expected demand is inelastic. Where $D(\overline{v}) > 0$, inserting $p(\overline{v})$ and $\overline{y}/k = Z_i x_i^{\alpha_i - 1}$, the derivative is positive if and only if

$$\Lambda(\overline{v}) := \overline{v}^{1/\varepsilon} \frac{D(\overline{v})}{\mathcal{G}^-_{\frac{1}{1-\alpha_i}}(\overline{v})} < c := \frac{\varepsilon}{1-\alpha_i} \frac{\mathcal{P}_i}{Z_i x_i^{\alpha_i} P_i} \left(\frac{\overline{y}}{Y_i} \right)^{1/\varepsilon},$$

a strictly positive constant. The set on which $D > 0$ is exactly $(\underline{v}_i, \infty)$: D is continuous with $D(0^+) = -1$, since $0 \leq \mathcal{G}_1^-(\overline{v}) \leq G_i(\overline{v}) \to 0$, and $D > 0$ far out, since $\overline{v}D(\overline{v}) = (\varepsilon - 1)\int_0^{\overline{v}} v \mathrm{d}G_i - \overline{v}[1 - G_i(\overline{v})] \to \varepsilon - 1 > 0$ by Lemma 4(i); so the single zero granted by Assumption 2 is interior and separates $\{D < 0\}$ from $\{D > 0\}$. On $(\underline{v}_i, \infty)$ that assumption also makes Λ strictly increasing; moreover $\Lambda(\underline{v}_i^+) = 0$ because $D(\underline{v}_i) = 0$ and $\mathcal{G}^-_{\frac{1}{1-\alpha_i}}(\underline{v}_i) > 0$ (the density is strictly positive, so $\mathcal{G}_\theta^-(\overline{v}) > 0$ at every $\overline{v} > 0$), and $\Lambda(\overline{v}) \to \infty$ as $\overline{v} \to \infty$ by Lemma 4(i). Hence there is exactly one $\overline{v}^\star \in (\underline{v}_i, \infty)$ with $\Lambda(\overline{v}^\star) = c$, which is (66); $\mathrm{d}\Pi/\mathrm{d}\overline{v} > 0$ on $(0, \overline{v}^\star)$ and < 0 on $(\overline{v}^\star, \infty)$. So Π has the unique global maximiser $p^\star = p(\overline{v}^\star)$, at which $D > 0$, i.e. $\varepsilon_{in} > 1$. Finally $\Pi(p^\star) > 0$: Π is strictly decreasing on $(\overline{v}^\star, \infty)$ and, as $\overline{v} \to \infty$, $p\,\overline{y}\Psi_1 \propto \overline{v}^{1/\varepsilon}\Psi_1(\overline{v}) \to 0$ and $\Psi_{\frac{1}{1-\alpha_i}} \to 0$ by Lemma 4(i), so $\Pi \to 0$ from above. This proves (i) and the root characterisation in (ii).

Step 5 (the explicit forms). From Step 3, $p = \frac{\varepsilon\eta_1}{\varepsilon\eta_1 - 1} \cdot \frac{1}{1-\alpha_i} \mathcal{P}_i \frac{\eta_{\frac{1}{1-\alpha_i}} \mathbb{E}_v[M_{in}]}{\eta_1 \mathbb{E}_v[y_{in}]}$. By (6)–(12) and $\eta_\theta \Psi_\theta = \mathcal{G}_\theta^-$,

$$\frac{\eta_{\frac{1}{1-\alpha_i}}(\overline{v})\mathbb{E}_v[M_{in}]}{\eta_1(\overline{v})\mathbb{E}_v[y_{in}]} = \frac{\mathcal{G}^-_{\frac{1}{1-\alpha_i}}(\overline{v})}{\mathcal{G}_1^-(\overline{v})} \cdot \frac{k/x_i}{Z_i k x_i^{\alpha_i - 1}} = \frac{1}{Z_i x_i^{\alpha_i}} \frac{\mathcal{G}^-_{\frac{1}{1-\alpha_i}}(\overline{v})}{\mathcal{G}_1^-(\overline{v})},$$

which is the second equality in (63). Multiplying and dividing by $\overline{v}^{\frac{\alpha_i}{1-\alpha_i}}$ and using $\frac{1}{1-\alpha_i} = 1 + \frac{\alpha_i}{1-\alpha_i}$ gives the first form of (65) with $\Gamma_i(\overline{v})$ as in (64); the second expression for Γ_i in (64) follows from $\eta_\theta \Psi_\theta = \mathcal{G}_\theta^-$ and $\varepsilon\mathcal{G}_1^- - \Psi_1 = D$. For the second form of (65), write the threshold as $\overline{v} = Z_i k x_i^{\alpha_i - 1} / [Y_i (p/P_i)^{-\varepsilon}]$; then

$$\frac{\mathcal{P}_i}{Z_i x_i^{\alpha_i}} \overline{v}^{-\frac{\alpha_i}{1-\alpha_i}} = \frac{\mathcal{P}_i}{Z_i x_i^{\alpha_i}} Z_i^{-\frac{\alpha_i}{1-\alpha_i}} k^{-\frac{\alpha_i}{1-\alpha_i}} x_i^{\alpha_i} \left[Y_i (p/P_i)^{-\varepsilon} \right]^{\frac{\alpha_i}{1-\alpha_i}} = \mathcal{P}_i Z_i^{-\frac{1}{1-\alpha_i}} \left(\frac{Y_i (p/P_i)^{-\varepsilon}}{k} \right)^{\frac{\alpha_i}{1-\alpha_i}},$$

because $\left(x_i^{\alpha_i - 1}\right)^{-\frac{\alpha_i}{1-\alpha_i}} = x_i^{\alpha_i}$ cancels the $x_i^{\alpha_i}$ in the denominator and $1 + \frac{\alpha_i}{1-\alpha_i} = \frac{1}{1-\alpha_i}$. The capital intensity therefore enters the pricing rule only through the threshold.

Step 6 (parts (iii) and (iv)). Both the markup $\varepsilon\eta_1/(\varepsilon\eta_1 - 1)$ and Γ_i are functions of $\overline{v}$ alone, built from G_i through $\mathcal{G}_1^-$, $\mathcal{G}^-_{\frac{1}{1-\alpha_i}}$ and G_i evaluated at $\overline{v}$; $\eta_1 < 1$ gives $\varepsilon_{in} < \varepsilon$ and hence a markup above $\frac{\varepsilon}{\varepsilon - 1}$. The limits in (iv) are Lemma 4(i): $\eta_1 \to 1$, $\overline{v}^\theta \mathcal{G}_\theta^- \to \mu_\theta$ with $\mu_1 = 1$, so $\Gamma_i(\overline{v}) \to \frac{\varepsilon}{\varepsilon - 1} \frac{1}{1-\alpha_i} \mu_{i, \frac{1}{1-\alpha_i}}$, and the second form of (65) becomes (163) with (164). ■

A.4 Symmetric Sectoral Equilibrium

Lemma 3 (Symmetric sectoral equilibrium under capacity). *In the symmetric equilibrium of sector i, with $\overline{y}_i = Z_i K_i x_i^{\alpha_i - 1}$ and $\overline{v}_i = \frac{\overline{y}_i}{Y_i}(p_i/P_i)^{\varepsilon}$:*

(i) (Price ratio.) $p_i/P_i = \mathcal{R}_i(\overline{v}_i)$, *where*

$$\mathcal{R}_i(\overline{v}) := \Big[\overline{v}\Big(\mathcal{G}^-_{i,1}(\overline{v}) + \mathcal{G}^+_{i,1/\varepsilon}(\overline{v})\Big)\Big]^{\frac{1}{\varepsilon-1}} = \left[\int_0^{\overline{v}} v \mathrm{d}G_i(v) + \overline{v}^{\,1-\frac{1}{\varepsilon}} \int_{\overline{v}}^{\infty} v^{1/\varepsilon} \mathrm{d}G_i(v)\right]^{\frac{1}{\varepsilon-1}} \tag{69}$$

is strictly increasing, with $\mathcal{R}_i(0^+) = 0$ and $\mathcal{R}_i(\overline{v}) < 1$ for every $\overline{v}$, $\mathcal{R}_i(\overline{v}) \to 1$ as $\overline{v} \to \infty$: the posted price lies below the shadow index.

(ii) (Output.) *Sector output is*

$$Y_i = \overline{y}_i \, \frac{\mathcal{R}_i(\overline{v}_i)^{\varepsilon}}{\overline{v}_i} = \overline{y}_i \Big[\overline{v}_i^{1/\varepsilon}\Big(\mathcal{G}^-_{i,1}(\overline{v}_i) + \mathcal{G}^+_{i,1/\varepsilon}(\overline{v}_i)\Big)\Big]^{\frac{\varepsilon}{\varepsilon-1}}. \tag{70}$$

The map $\overline{v} \mapsto \overline{v}/\mathcal{R}_i(\overline{v})^{\varepsilon}$ is a strictly increasing bijection of $(0,\infty)$ onto $\Big(\mu_{i,1/\varepsilon}^{-\frac{\varepsilon}{\varepsilon-1}}, \infty\Big)$, so for given $(\overline{y}_i, Y_i)$ the threshold exists and is unique if and only if $Y_i < \overline{y}_i \, \mu_{i,1/\varepsilon}^{\frac{\varepsilon}{\varepsilon-1}}$: sector output cannot exceed sector capacity (Definition 1).

(iii) (Bundle and spending.) *Average, hence aggregate, bundle demand and intermediate spending are*

$$M_i^e = \frac{K_i}{x_i} \, \Psi_{i,\frac{1}{1-\alpha_i}}(\overline{v}_i), \qquad X_i = \mathcal{P}_i(\mathbf{P}) \, M_i^e, \qquad m_{ij} = a_{ij} \frac{X_i}{P_j}. \tag{71}$$

(iv) (Two revenue concepts.) *Sales at the shadow index and the revenue of the sector's firms are*

$$S_i := P_i Y_i = p_i \, \overline{y}_i \Big(\mathcal{G}^-_{i,1}(\overline{v}_i) + \mathcal{G}^+_{i,1/\varepsilon}(\overline{v}_i)\Big), \qquad S_i^{\mathrm{f}} := p_i \, \mathbb{E}_v[y_i] = p_i \, \overline{y}_i \, \Psi_{i,1}(\overline{v}_i), \tag{72}$$

and $S_i \geq S_i^{\mathrm{f}}$: the difference

$$\Pi_i^{\mathrm{r}} := S_i - S_i^{\mathrm{f}} = \int_0^1 \varphi_{in} y_{in} \mathrm{d}n = p_i \, \overline{y}_i \int_{\overline{v}_i}^{\infty} \left[\left(\frac{v}{\overline{v}_i}\right)^{1/\varepsilon} - 1\right] \mathrm{d}G_i(v) \;\geq\; 0 \tag{73}$$

is the rationing rent earned by the competitive aggregators, strictly positive whenever $G_i(\overline{v}_i) < 1$.

(v) (Supply relation.) *The pricing rule* (65) *in symmetric form,*

$$P_i \, \mathcal{R}_i(\overline{v}_i) = \Gamma_i(\overline{v}_i) \, \frac{\mathcal{P}_i(\mathbf{P})}{Z_i x_i^{\alpha_i}} \, \overline{v}_i^{\,-\frac{\alpha_i}{1-\alpha_i}}, \qquad \textit{equivalently}$$

$$P_i = \frac{\Gamma_i(\overline{v}_i)}{\mathcal{R}_i(\overline{v}_i)^{1+\frac{\varepsilon\alpha_i}{1-\alpha_i}}} \, \mathcal{P}_i(\mathbf{P}) \, Z_i^{-\frac{1}{1-\alpha_i}} \left(\frac{Y_i}{K_i}\right)^{\frac{\alpha_i}{1-\alpha_i}}, \tag{74}$$

together with (70)*, links the sectoral price, output and threshold given $(K_i, Z_i, \mathcal{P}_i(\mathbf{P}))$; it is the sector's supply relation.*

In the model without capacity constraints the sectoral price is a constant wedge times a marginal-cost index that rises with the bundle price, falls with productivity and rises with output per unit of capital with elasticity $\frac{\alpha_i}{1-\alpha_i}$ (Appendix C.1). Under capacity the same index appears in (74), but the wedge $\Gamma_i(\overline{v}_i)/\mathcal{R}_i(\overline{v}_i)^{1+\frac{\varepsilon\alpha_i}{1-\alpha_i}}$ is state-dependent and the supply relation is implicit in the threshold, which is itself determined by output relative to capacity, (70).

[Which markets clear by price, and why the rent stays with the aggregators.] The meaning of the multiplier and of P_i, where the rent comes from and who pays and receives it, and the average-cost alternative are now in the main text and footnotes after Lemma 1. Two points are not stated there.

- *Two levels of market clearing.* The *firm-level* goods markets do not clear by price: p_{in} and $\overline{y}_{in}$ are fixed before v is known, and a rationed firm sells $\overline{y}_{in}$ and turns demand away. The *sectoral* good market does clear in the Walrasian sense, at P_i. The rationing is absorbed one level up, into P_i through the multipliers. So the model has quantity rationing at the firm level inside an otherwise Walrasian static block, and every clearing condition of the no-capacity model holds with P_i the shadow index.
- *Why the rent is not competed away.* A referee may ask why competing aggregators do not bid the rationed goods up to $p_{in} + \varphi_{in}$, which would hand the rent to the firms. The answer is the posted-price assumption: firms sell at p_{in} up to capacity and do not accept higher bids after v is realised, so the rent belongs to whoever buys at p_{in}. How the rationed units are split across the unit measure of identical aggregators affects each aggregator's share, not the aggregate Π_i^{r}. Free entry of aggregators does not dissipate it either, since the scarce factor is the firms' capacity, not aggregation. If firms could reprice after observing v, a firm facing excess demand would charge the market-clearing price, capture the rent and never ration: that is a different model (ex-post pricing), not a different booking of the same one.

Proof. Drop the sector index on the tail functions and write $\overline{v} := \overline{v}_i$, $\overline{y} := \overline{y}_i$. In the symmetric equilibrium all firms post p_i and hold K_i, so by (10) and (5) they share $\overline{y}$ and $\overline{v}$, and by the law of large numbers across the unit measure of firms every integral $\int_0^1(\cdot)\mathrm{d}n$ of a function of v_{in} equals the expectation $\int_0^\infty(\cdot)\mathrm{d}G_i(v)$.

Step 1 (price ratio). By (61), $p_i + \varphi_{in} = p_i \max\{(v_{in}/\overline{v})^{1/\varepsilon}, 1\}$. Inserting into the price index of (3),

$$P_i^{1-\varepsilon} = p_i^{1-\varepsilon}\left[\int_0^{\overline{v}} v\mathrm{d}G_i(v) + \int_{\overline{v}}^\infty v\left(\frac{v}{\overline{v}}\right)^{\frac{1-\varepsilon}{\varepsilon}}\mathrm{d}G_i(v)\right] = p_i^{1-\varepsilon}\,\overline{v}\Big[\mathcal{G}_1^-(\overline{v}) + \mathcal{G}_{1/\varepsilon}^+(\overline{v})\Big],$$

since $\int_0^{\overline{v}} v\mathrm{d}G_i = \overline{v}\mathcal{G}_1^-(\overline{v})$ and $v(v/\overline{v})^{(1-\varepsilon)/\varepsilon} = \overline{v}(v/\overline{v})^{1/\varepsilon}$. Hence $(p_i/P_i)^{\varepsilon-1} = \overline{v}[\mathcal{G}_1^-(\overline{v}) + \mathcal{G}_{1/\varepsilon}^+(\overline{v})]$, which is (69). For the properties of $\mathcal{R}_i$: since $v \geq \overline{v}$ on the upper tail, $\overline{v}^{1-1/\varepsilon}v^{1/\varepsilon} \leq v$ there, so $\mathcal{R}_i(\overline{v})^{\varepsilon-1} \leq \int_0^\infty v\mathrm{d}G_i = \mu_{i,1} = 1$, with strict inequality

whenever $G_i(\overline{v}) < 1$, which is always under Assumption 1; differentiating the second form in (69), the two density terms cancel and

$$\frac{\mathrm{d}}{\mathrm{d}\overline{v}}\mathcal{R}_i(\overline{v})^{\varepsilon-1} = \left(1-\frac{1}{\varepsilon}\right)\overline{v}^{-1/\varepsilon}\int_{\overline{v}}^{\infty} v^{1/\varepsilon}\mathrm{d}G_i(v) > 0;$$

as $\overline{v} \to 0$ both terms of $\mathcal{R}_i^{\varepsilon-1}$ vanish (the second is bounded by $\overline{v}^{1-1/\varepsilon}\mu_{i,1/\varepsilon}$), and as $\overline{v} \to \infty$ the first tends to $\mu_{i,1} = 1$ and the second to zero by Lemma 4(i). This proves (i).

Step 2 (output). The definition of the threshold, $\overline{v} = \frac{\overline{y}}{Y_i}(p_i/P_i)^{\varepsilon}$, with (i) gives $Y_i = \overline{y}\mathcal{R}_i(\overline{v})^{\varepsilon}/\overline{v}$; substituting (69) and $\frac{\varepsilon}{\varepsilon-1} - 1 = \frac{1}{\varepsilon-1}$ gives the second form of (70). The same expression follows directly from the aggregator (2) with $y_{in} = \overline{y}\min\{v_{in}/\overline{v}, 1\}$, the unconstrained and the constrained firms contributing the two terms of

$$Y_i^{\frac{\varepsilon-1}{\varepsilon}} = \left(\frac{\overline{y}}{\overline{v}}\right)^{\frac{\varepsilon-1}{\varepsilon}}\int_0^{\overline{v}} v\mathrm{d}G_i(v) + \overline{y}^{\frac{\varepsilon-1}{\varepsilon}}\int_{\overline{v}}^{\infty} v^{1/\varepsilon}\mathrm{d}G_i(v) = \overline{y}^{\frac{\varepsilon-1}{\varepsilon}}\,\overline{v}^{1/\varepsilon}\Big[\mathcal{G}_1^-(\overline{v}) + \mathcal{G}_{1/\varepsilon}^+(\overline{v})\Big],$$

which is a check on (i)–(ii). For the bijection, write

$$\frac{\overline{v}}{\mathcal{R}_i(\overline{v})^{\varepsilon}} = W(\overline{v})^{-\frac{\varepsilon}{\varepsilon-1}},$$
$$W(\overline{v}) := \overline{v}^{1/\varepsilon}\Big[\mathcal{G}_1^-(\overline{v}) + \mathcal{G}_{1/\varepsilon}^+(\overline{v})\Big] = \overline{v}^{\frac{1}{\varepsilon}-1}\int_0^{\overline{v}} v\mathrm{d}G_i(v) + \int_{\overline{v}}^{\infty} v^{1/\varepsilon}\mathrm{d}G_i(v).$$

Differentiating, the density terms cancel again and $W'(\overline{v}) = -(1-\frac{1}{\varepsilon})\overline{v}^{\frac{1}{\varepsilon}-2}\int_0^{\overline{v}} v\mathrm{d}G_i(v) < 0$; $W(0^+) = \mu_{i,1/\varepsilon}$, because the first term is bounded by $\overline{v}^{1/\varepsilon}G_i(\overline{v}) \to 0$; and $W(\infty) = 0$, because the first term is bounded by $\overline{v}^{\frac{1}{\varepsilon}-1}\mu_{i,1}$ and the second is the tail of a convergent integral. So $\overline{v} \mapsto \overline{v}/\mathcal{R}_i(\overline{v})^{\varepsilon}$ is continuous and strictly increasing from $\mu_{i,1/\varepsilon}^{-\varepsilon/(\varepsilon-1)}$ to ∞, and the threshold equation $\overline{v}/\mathcal{R}_i(\overline{v})^{\varepsilon} = \overline{y}/Y_i$ has exactly one solution when $\overline{y}/Y_i > \mu_{i,1/\varepsilon}^{-\varepsilon/(\varepsilon-1)}$ and none otherwise. This proves (ii).

Step 3 (bundle and spending). By the law of large numbers and (12) at $k_{in} = K_i$, $M_i^e = \int_0^1 M_{in}\mathrm{d}n = \mathbb{E}_v[M_{in}] = \frac{K_i}{x_i}\Psi_{\frac{1}{1-\alpha_i}}(\overline{v})$; $X_i = \mathcal{P}_i M_i^e$ and $m_{ij} = a_{ij}X_i/P_j$ are (14) integrated over n. This proves (iii).

Step 4 (revenues and the rationing rent). $S_i = P_iY_i = P_i\,\overline{y}\,\mathcal{R}_i^{\varepsilon}/\overline{v} = p_i\,\overline{y}\,\mathcal{R}_i^{\varepsilon-1}/\overline{v} = p_i\overline{y}[\mathcal{G}_1^-(\overline{v}) + \mathcal{G}_{1/\varepsilon}^+(\overline{v})]$ by (i)–(ii), and $S_i^{\mathrm{f}} = p_i\mathbb{E}_v[y_i] = p_i\overline{y}\Psi_1(\overline{v}) = p_i\overline{y}[\mathcal{G}_1^-(\overline{v}) + 1 - G_i(\overline{v})]$ by (6). Their difference is $p_i\overline{y}[\mathcal{G}_{1/\varepsilon}^+(\overline{v}) - (1-G_i(\overline{v}))] = p_i\overline{y}\int_{\overline{v}}^{\infty}[(v/\overline{v})^{1/\varepsilon} - 1]\mathrm{d}G_i(v)$, which is nonnegative because the integrand is, and strictly positive when $G_i(\overline{v}) < 1$. By (61), on the constrained tail $\varphi_{in}y_{in} = p_i[(v_{in}/\overline{v})^{1/\varepsilon} - 1]\overline{y}$ and it is zero otherwise, so $\int_0^1 \varphi_{in}y_{in}\mathrm{d}n$ equals the same integral; by Lemma 1(i) it also equals $P_iY_i - \int_0^1 p_iy_{in}\mathrm{d}n = S_i - S_i^{\mathrm{f}}$. This proves (iv).

Step 5 (supply relation). The first form of (74) is (65) at $p_{in} = p_i = P_i\mathcal{R}_i(\overline{v})$ and $k_{in} = K_i$. For the second, substitute $\overline{v} = \frac{Z_iK_ix_i^{\alpha_i-1}}{Y_i}\mathcal{R}_i(\overline{v})^{\varepsilon}$ into $\overline{v}^{-\frac{\alpha_i}{1-\alpha_i}}$:

$$\frac{\mathcal{P}_i}{Z_ix_i^{\alpha_i}}\,\overline{v}^{-\frac{\alpha_i}{1-\alpha_i}} = \mathcal{P}_i\,Z_i^{-1-\frac{\alpha_i}{1-\alpha_i}}\left(\frac{Y_i}{K_i}\right)^{\frac{\alpha_i}{1-\alpha_i}}\mathcal{R}_i(\overline{v})^{-\frac{\varepsilon\alpha_i}{1-\alpha_i}},$$

the powers of x_i cancelling exactly as in Step 5 of the proof of Lemma 2; dividing by $\mathcal{R}_i(\overline{v})$ and using $1+\frac{\alpha_i}{1-\alpha_i}=\frac{1}{1-\alpha_i}$ gives the second form. This proves (v). ■

Remark 2 (The markup wedge along the threshold) Fix (P_i, K_i, Z_i) and let the sector's threshold $\overline{v}$ vary along its supply relation (74), that is, let $\mathcal{P}_i(\overline{v}) := P_i\,\mathcal{R}_i(\overline{v})\,Z_i x_i^{\alpha_i}\,\overline{v}^{\alpha_i/(1-\alpha_i)}/\Gamma_i(\overline{v})$ be the bundle price at which $\overline{v}$ is the sector's equilibrium threshold. Output $Y_i(\overline{v}) = \overline{y}_i\mathcal{R}_i(\overline{v})^{\varepsilon}/\overline{v}$ (70) and bundle use $M_i^e(\overline{v}) = \frac{K_i}{x_i}\Psi_{i,\frac{1}{1-\alpha_i}}(\overline{v})$ (Lemma 3(iii)) both fall in $\overline{v}$, and

$$\frac{P_i\ \mathrm{d}Y_i/\mathrm{d}\overline{v}}{\mathcal{P}_i(\overline{v})\ \mathrm{d}M_i^e/\mathrm{d}\overline{v}} = \frac{\varepsilon_i(\overline{v})}{\varepsilon_i(\overline{v})-1} \qquad \text{for every } \overline{v} > \underline{v}_i,$$

the range on which $\mathcal{P}_i(\overline{v}) > 0$: the marginal value product of expected bundle use at the fixed index P_i exceeds the bundle's price by the factor of the capacity-adjusted gross markup of Lemma 2, at every threshold. The identity is the pricing rule in threshold form. With $\mathcal{P}_i(\overline{v})$ defined by that rule, the firm's own first-order condition reads $p_i\ \mathrm{d}\mathbb{E}_v[y_{in}]/\mathrm{d}\overline{v} = \frac{\varepsilon_i}{\varepsilon_i-1}\,\mathcal{P}_i\ \mathrm{d}M_i^e/\mathrm{d}\overline{v}$; what the sectoral form adds is that sales at the fixed index and firm revenue at the posted price have the same threshold derivative,

$$P_i\,\frac{\mathrm{d}Y_i}{\mathrm{d}\overline{v}} = p_i\,\frac{\mathrm{d}\mathbb{E}_v[y_{in}]}{\mathrm{d}\overline{v}} = -\,p_i\,\overline{y}_i\,\frac{\mathcal{G}^-_{i,1}(\overline{v})}{\overline{v}},$$

the price being evaluated at $\overline{v}$ and not differentiated. A sector that maximised its value added $P_iY_i - \mathcal{P}_iM_i^e$ at given prices would set the ratio to one, that is, price at effective marginal cost; the firms set it at their markup. Capacity changes the size of this one-sector wedge, through $\eta_{i,1}(\overline{v})$, but not its form.

Proof. Output. From (69), $\mathcal{R}_i^{\varepsilon-1} = \overline{v}\,[\mathcal{G}^-_{i,1} + \mathcal{G}^+_{i,1/\varepsilon}]$. Differentiating the two tail integrals, the boundary terms at $v=\overline{v}$ cancel (both integrands equal one there), and

$$\frac{\mathrm{d}}{\mathrm{d}\overline{v}}\left[\mathcal{G}^-_{i,1} + \mathcal{G}^+_{i,1/\varepsilon}\right] = -\frac{1}{\overline{v}}\left[\mathcal{G}^-_{i,1} + \frac{1}{\varepsilon}\mathcal{G}^+_{i,1/\varepsilon}\right],$$

$$\frac{\mathrm{d}\ln Y_i}{\mathrm{d}\ln\overline{v}} = \frac{\varepsilon}{\varepsilon-1}\cdot\frac{(1-\frac{1}{\varepsilon})\mathcal{G}^+_{i,1/\varepsilon}}{\mathcal{G}^-_{i,1}+\mathcal{G}^+_{i,1/\varepsilon}} - 1 = -\frac{\mathcal{G}^-_{i,1}(\overline{v})}{\mathcal{G}^-_{i,1}(\overline{v}) + \mathcal{G}^+_{i,1/\varepsilon}(\overline{v})},$$

i.e. $\overline{v}\ \mathrm{d}Y_i/\mathrm{d}\overline{v} = -\overline{y}_i\,\mathcal{R}_i(\overline{v})\,\mathcal{G}^-_{i,1}(\overline{v})$, using $\mathcal{R}_i^{\varepsilon}/(\overline{v}[\mathcal{G}^-_{i,1}+\mathcal{G}^+_{i,1/\varepsilon}]) = \mathcal{R}_i$. Since $p_i = P_i\mathcal{R}_i$ and $\mathbb{E}_v[y_{in}] = \overline{y}_i\Psi_{i,1}(\overline{v})$ with $\Psi'_{i,1} = -\mathcal{G}^-_{i,1}/\overline{v}$ by (59), this is the second display of the remark.

Bundle. By (59), $\overline{v}\ \mathrm{d}M_i^e/\mathrm{d}\overline{v} = -\frac{K_i}{x_i}\,\frac{1}{1-\alpha_i}\,\mathcal{G}^-_{i,\frac{1}{1-\alpha_i}}(\overline{v})$.

Ratio. With $\overline{y}_i = Z_iK_ix_i^{\alpha_i-1}$ and $\mathcal{P}_i(\overline{v})$ as defined,

$$\frac{P_i\ \mathrm{d}Y_i/\mathrm{d}\overline{v}}{\mathcal{P}_i(\overline{v})\ \mathrm{d}M_i^e/\mathrm{d}\overline{v}} = \frac{\mathcal{G}^-_{i,1}(\overline{v})\ \Gamma_i(\overline{v})\,\overline{v}^{-\frac{\alpha_i}{1-\alpha_i}}}{\frac{1}{1-\alpha_i}\mathcal{G}^-_{i,\frac{1}{1-\alpha_i}}(\overline{v})} = \frac{\varepsilon\,\mathcal{G}^-_{i,1}(\overline{v})}{(\varepsilon-1)\mathcal{G}^-_{i,1}(\overline{v}) - 1 + G_i(\overline{v})} = \frac{\varepsilon_i(\overline{v})}{\varepsilon_i(\overline{v})-1},$$

by the explicit form (64) of Γ_i and $\varepsilon_i = \varepsilon\mathcal{G}^-_{i,1}/\Psi_{i,1}$. The denominator of the middle expression vanishes at $\underline{v}_i$ and is positive above it, which is the stated range. ■

A.5 The Capacity Limit

Capacity is removed from the model by letting the capital intensity x_i tend to zero: capacity $\overline{y}_i = Z_i K_i x_i^{\alpha_i - 1}$ becomes infinite, the probability of being rationed vanishes, and every capacity object converges to its counterpart in the model without capacity constraints. At every *finite* x_i the constrained probability $1 - G_i(\overline{v}_i)$ is strictly positive (Assumption 1); it is only in the limit that no firm is rationed. Since $x_i = 0$ is not an admissible parameter value, the no-capacity model is a limit, stated once here and cited from Appendix C.

Lemma 4 (The capacity limit). *Let $\varepsilon > 1$, $\alpha_i \in (0,1)$, and let G_i satisfy Assumption 1.*

(i) (Tail functions.) *As $\overline{v} \to \infty$: $G_i(\overline{v}) \to 1$, $\overline{v}\,[1 - G_i(\overline{v})] \to 0$, and for every $\theta \in [0, \frac{1}{1-\alpha_i}]$*

$$\overline{v}^{\theta}\mathcal{G}^{-}_{i,\theta}(\overline{v}) \to \mu_{i,\theta}, \qquad \overline{v}^{\theta}\mathcal{G}^{+}_{i,\theta}(\overline{v}) \to 0, \qquad \overline{v}^{\theta}\Psi_{i,\theta}(\overline{v}) \to \mu_{i,\theta}, \qquad \eta_{i,\theta}(\overline{v}) \to 1.$$

At $\theta = 1/\varepsilon$ the second limit holds in the stronger form $\overline{v}\,\mathcal{G}^{+}_{i,1/\varepsilon}(\overline{v}) \to 0$, since

$$\overline{v}\,\mathcal{G}^{+}_{i,1/\varepsilon}(\overline{v}) = \overline{v}^{1-\frac{1}{\varepsilon}} \int_{\overline{v}}^{\infty} v^{1/\varepsilon} \mathrm{d}G_i(v) \le \int_{\overline{v}}^{\infty} v \mathrm{d}G_i(v).$$

Consequently $\varepsilon_i(\overline{v}) \to \varepsilon$, $\Gamma_i(\overline{v}) \to \Gamma_i = \frac{\varepsilon}{\varepsilon-1}\frac{1}{1-\alpha_i}\mu_{i,\frac{1}{1-\alpha_i}}$, $\mathcal{R}_i(\overline{v}) \to 1$, and $\Psi_{i,1}(\overline{v}) / [\mathcal{G}^{-}_{i,1}(\overline{v}) + \mathcal{G}^{+}_{i,1/\varepsilon}(\overline{v})] \to 1$; the ratio in Assumption 2 satisfies

$$\overline{v}^{1/\varepsilon}\,\frac{(\varepsilon-1)\mathcal{G}^{-}_{i,1}(\overline{v}) - 1 + G_i(\overline{v})}{\mathcal{G}^{-}_{i,\frac{1}{1-\alpha_i}}(\overline{v})} \;\sim\; \frac{\varepsilon-1}{\mu_{i,\frac{1}{1-\alpha_i}}}\,\overline{v}^{\frac{1}{\varepsilon}+\frac{\alpha_i}{1-\alpha_i}} \;\to\; \infty.$$

(ii) (Firm problem.) *Fix $(Y_i, P_i, \mathcal{P}_i(\mathbf{P}), Z_i, k_{in})$ and let $x_i \downarrow 0$. Under Assumption 2 the optimal threshold of Lemma 2 diverges, $\overline{v}_{in}(x_i) \to \infty$, and the optimal price converges to the unique solution of the pricing rule of the no-capacity model,*

$$p_{in} = \frac{\varepsilon}{\varepsilon-1}\,\frac{1}{1-\alpha_i}\,\mathcal{P}_i(\mathbf{P})\,\mu_{i,\frac{1}{1-\alpha_i}}\,Z_i^{-\frac{1}{1-\alpha_i}}\left(\frac{Y_i (p_{in}/P_i)^{-\varepsilon}}{k_{in}}\right)^{\frac{\alpha_i}{1-\alpha_i}} \qquad \textit{(Lemma 21).}$$

(iii) (Sector.) *Let Assumption 2 hold, fix $(Y_i, K_i, Z_i, \mathcal{P}_i(\mathbf{P}))$ and let $x_i \downarrow 0$. For all x_i small enough the symmetric threshold of Lemma 3(ii) exists and lies in $(\underline{v}_i, \infty)$, and then $\overline{y}_i \to \infty$, $\overline{v}_i(x_i) \to \infty$ and*

$$P_i \to \Gamma_i\,\mathcal{P}_i(\mathbf{P})\,Z_i^{-\frac{1}{1-\alpha_i}}\left(\frac{Y_i}{K_i}\right)^{\frac{\alpha_i}{1-\alpha_i}}, \qquad M_i^{e} \to \mu_{i,\frac{1}{1-\alpha_i}}\left(\frac{Y_i}{Z_i K_i^{\alpha_i}}\right)^{\frac{1}{1-\alpha_i}},$$

$$\frac{S_i^{\mathrm{f}}}{S_i} \to 1, \qquad \frac{\Pi_i^{\mathrm{r}}}{S_i} \to 0, \qquad \frac{Y_i}{\overline{y}_i} \to 0, \qquad U_i(\overline{v}_i) \to 0,$$

the first two being the no-capacity model's supply curve (166) *and bundle demand* (165). *The two smallness requirements are $Y_i < \overline{y}_i\,\mu_{i,1/\varepsilon}^{\frac{\varepsilon}{\varepsilon-1}}$, which makes the threshold exist*

(Lemma 3(ii)), and $\overline{v}_i > \underline{v}_i$, which makes the wedge $\Gamma_i(\overline{v}_i)$ of (64) *positive; both hold for all x_i small enough because $\overline{y}_i \to \infty$ and $\overline{v}_i \to \infty$. Only the P_i limit uses Assumption 2, through the supply relation* (74)*, which is the pricing rule of Lemma 2(ii) at the symmetric point; the remaining limits are aggregator identities and follow from (i) and Lemma 3(ii)–(iv) alone.*

In the limit, utilization vanishes rather than reaching one: the no-capacity model is the model with infinite capacity, not the model at full capacity. All results of Appendix C are to be read as statements about this limit.

Proof. Drop the sector index on the tail functions.

Step 1 (tail functions). For $\theta \in [0, \frac{1}{1-\alpha_i}]$, $\overline{v}^\theta \mathcal{G}_\theta^-(\overline{v}) = \int_0^{\overline{v}} v^\theta \mathrm{d}G_i(v) \uparrow \mu_{i,\theta} < \infty$ as $\overline{v} \to \infty$ by monotone convergence, and $\overline{v}^\theta \mathcal{G}_\theta^+(\overline{v}) = \int_{\overline{v}}^\infty v^\theta \mathrm{d}G_i(v) = \mu_{i,\theta} - \int_0^{\overline{v}} v^\theta \mathrm{d}G_i \to 0$. Since $\overline{v}^\theta \le v^\theta$ on $\{v \ge \overline{v}\}$ for $\theta \ge 0$, $\overline{v}^\theta[1 - G_i(\overline{v})] \le \int_{\overline{v}}^\infty v^\theta \mathrm{d}G_i(v) \to 0$ for every such $\theta > 0$, in particular $\overline{v}[1 - G_i(\overline{v})] \to 0$ (admissible because $\frac{1}{1-\alpha_i} > 1$) and $G_i(\overline{v}) \to 1$. (Markov's inequality alone would give only the bound $\overline{v}^\theta[1 - G_i(\overline{v})] \le \mu_{i,\theta}$; what is used is the truncated moment, whose tail vanishes.) Hence $\overline{v}^\theta \Psi_\theta(\overline{v}) = \overline{v}^\theta \mathcal{G}_\theta^-(\overline{v}) + \overline{v}^\theta[1 - G_i(\overline{v})] \to \mu_{i,\theta}$ and $\eta_\theta(\overline{v}) = \overline{v}^\theta \mathcal{G}_\theta^- / \overline{v}^\theta \Psi_\theta \to \mu_{i,\theta}/\mu_{i,\theta} = 1$ (for $\theta = 0$ trivially $\eta_0 = G_i \to 1$). The consequences: $\varepsilon_i(\overline{v}) = \varepsilon\eta_1 \to \varepsilon$; in the second form of (64), $\Gamma_i(\overline{v}) = \frac{\varepsilon \frac{1}{1-\alpha_i} \overline{v}^{\frac{1}{1-\alpha_i}} \mathcal{G}^-_{\frac{1}{1-\alpha_i}}(\overline{v})}{(\varepsilon-1)\overline{v}\mathcal{G}_1^-(\overline{v}) - \overline{v}[1-G_i(\overline{v})]} \to \frac{\varepsilon \frac{1}{1-\alpha_i} \mu_{i,\frac{1}{1-\alpha_i}}}{(\varepsilon-1)\cdot 1 - 0} = \Gamma_i$; $\mathcal{R}_i(\overline{v})^{\varepsilon-1} = \overline{v}\mathcal{G}_1^-(\overline{v}) + \overline{v}\mathcal{G}_{1/\varepsilon}^+(\overline{v}) \to 1 + 0$ because $\overline{v}\mathcal{G}_{1/\varepsilon}^+(\overline{v}) = \overline{v}^{1-1/\varepsilon} \int_{\overline{v}}^\infty v^{1/\varepsilon} \mathrm{d}G_i \le \int_{\overline{v}}^\infty v \mathrm{d}G_i \to 0$; and $\Psi_1 / [\mathcal{G}_1^- + \mathcal{G}_{1/\varepsilon}^+] = \overline{v}\Psi_1 / \mathcal{R}_i^{\varepsilon-1} \to 1/1$. For the ratio of Assumption 2, write it as $\overline{v}^{\frac{1}{\varepsilon}-1+\frac{1}{1-\alpha_i}} \cdot \frac{\overline{v} D(\overline{v})}{\overline{v}^{\frac{1}{1-\alpha_i}} \mathcal{G}^-_{\frac{1}{1-\alpha_i}}(\overline{v})}$ with D as in the proof of Lemma 2; the second factor tends to $(\varepsilon - 1)/\mu_{i,\frac{1}{1-\alpha_i}}$ and the exponent of the first is $\frac{1}{\varepsilon} + \frac{\alpha_i}{1-\alpha_i} > 0$. This proves (i).

Step 2 (firm problem). Fix $(Y_i, P_i, \mathcal{P}_i, Z_i, k_{in})$. By Step 4 of the proof of Lemma 2, the optimal threshold is $\overline{v}_{in} = \Lambda^{-1}(c)$, where Λ is a strictly increasing bijection of $(\underline{v}_i, \infty)$ onto $(0, \infty)$ that does not depend on x_i, and $c = \frac{\varepsilon}{1-\alpha_i} \frac{\mathcal{P}_i}{Z_i x_i^{\alpha_i} P_i} (\overline{y}_{in}/Y_i)^{1/\varepsilon}$ with $\overline{y}_{in} = Z_i k_{in} x_i^{\alpha_i - 1}$, so $c \propto x_i^{-\alpha_i - \frac{1-\alpha_i}{\varepsilon}} \to \infty$ as $x_i \downarrow 0$. Hence $\overline{v}_{in}(x_i) \to \infty$. For the price, the second form of (65) says that $p_{in}(x_i)$ solves $p = \Gamma_i(\overline{v}_{in}(x_i))\, \mathcal{P}_i Z_i^{-\frac{1}{1-\alpha_i}} [Y_i (p/P_i)^{-\varepsilon} / k_{in}]^{\frac{\alpha_i}{1-\alpha_i}}$, an equation in which x_i appears only through the wedge. For any constant $\Gamma > 0$ the equation $p = \Gamma\, \mathcal{P}_i Z_i^{-\frac{1}{1-\alpha_i}} [Y_i (p/P_i)^{-\varepsilon} / k_{in}]^{\frac{\alpha_i}{1-\alpha_i}}$ has a unique root $p(\Gamma)$, because the left side is increasing and the right side decreasing in p, and $p(\Gamma)$ is continuous in Γ (it is the explicit power $p(\Gamma) \propto \Gamma^{1/\Theta_i}$ of Lemma 21(ii)). Since $\Gamma_i(\overline{v}_{in}(x_i)) \to \Gamma_i$ by (i), $p_{in}(x_i) = p(\Gamma_i(\overline{v}_{in}(x_i))) \to p(\Gamma_i)$, which is the root of (163)–(164). This proves (ii).

Step 3 (sector). Fix $(Y_i, K_i, Z_i, \mathcal{P}_i)$. Then $\overline{y}_i = Z_i K_i x_i^{\alpha_i - 1} \to \infty$, and by Step 2 of the proof of Lemma 3 the threshold is $\overline{v}_i = \omega^{-1}(\overline{y}_i / Y_i)$ with $\omega(\overline{v}) := \overline{v}/\mathcal{R}_i(\overline{v})^\varepsilon$ a strictly increasing bijection onto $(\mu_{i,1/\varepsilon}^{-\varepsilon/(\varepsilon-1)}, \infty)$ that does not depend on x_i; it is defined once

$\overline{y}_i/Y_i > \mu_{i,1/\varepsilon}^{-\varepsilon/(\varepsilon-1)}$, hence for all x_i small enough, and then $\overline{v}_i(x_i) \to \infty$ and $Y_i/\overline{y}_i \to 0$. Since $\overline{v}_i \to \infty$, $\overline{v}_i > \underline{v}_i$ for all x_i small enough as well, so the wedge $\Gamma_i(\overline{v}_i)$ is positive and the supply relation (74) is the one of Lemma 2(ii) at the symmetric point; at moderate x_i neither is guaranteed, which is why the statement is a smallness statement in x_i. utilization is $U_i(\overline{v}_i) = \mu_{i,1/\varepsilon}^{-\frac{\varepsilon}{\varepsilon-1}} Y_i/\overline{y}_i$ by (21)–(22), a fixed multiple of $Y_i/\overline{y}_i$, so $U_i(\overline{v}_i) \to 0$. By the second form of (74), $P_i = \Gamma_i(\overline{v}_i)\mathcal{R}_i(\overline{v}_i)^{-1-\frac{\varepsilon\alpha_i}{1-\alpha_i}} \mathcal{P}_i Z_i^{-\frac{1}{1-\alpha_i}} (Y_i/K_i)^{\frac{\alpha_i}{1-\alpha_i}} \to \Gamma_i \cdot 1 \cdot \mathcal{P}_i Z_i^{-\frac{1}{1-\alpha_i}} (Y_i/K_i)^{\frac{\alpha_i}{1-\alpha_i}}$ by (i), which is (166). For the bundle, (18) and the threshold definition give

$$M_i^e = \frac{K_i}{x_i}\Psi_{\frac{1}{1-\alpha_i}}(\overline{v}_i) = \frac{K_i}{x_i}\,\overline{v}_i^{\frac{1}{1-\alpha_i}}\Psi_{\frac{1}{1-\alpha_i}}(\overline{v}_i) \cdot \left(\frac{Y_i}{Z_i K_i x_i^{\alpha_i-1}\mathcal{R}_i(\overline{v}_i)\varepsilon}\right)^{\frac{1}{1-\alpha_i}}$$
$$= \overline{v}_i^{\frac{1}{1-\alpha_i}}\Psi_{\frac{1}{1-\alpha_i}}(\overline{v}_i)\,\mathcal{R}_i(\overline{v}_i)^{-\frac{\varepsilon}{1-\alpha_i}}\left(\frac{Y_i}{Z_i K_i^{\alpha_i}}\right)^{\frac{1}{1-\alpha_i}},$$

the powers of x_i cancelling because $x_i^{-1} \cdot (x_i^{1-\alpha_i})^{\frac{1}{1-\alpha_i}} = 1$ and those of K_i because $K_i \cdot K_i^{-\frac{1}{1-\alpha_i}} = K_i^{-\frac{\alpha_i}{1-\alpha_i}}$; by (i) the prefactor tends to $\mu_{i,\frac{1}{1-\alpha_i}} \cdot 1$, giving (165). Finally $S_i^{\mathrm{f}}/S_i = \Psi_1(\overline{v}_i)/[\mathcal{G}_1^-(\overline{v}_i) + \mathcal{G}_{1/\varepsilon}^+(\overline{v}_i)] \to 1$ and $\Pi_i^{\mathrm{r}}/S_i = 1 - S_i^{\mathrm{f}}/S_i \to 0$ by (i). This proves (iii). ∎

A.6 Revenue Shares

In the model without capacity constraints, the pricing first-order condition makes intermediate spending a constant share of sales and splits every euro of sales into three constant pieces. Under capacity the shares move with the threshold, but one piece of the split survives exactly.

Lemma 5 (Revenue shares under capacity). *Let G_i satisfy Assumptions 1 and 2, and consider the symmetric equilibrium of Lemma 3, with the threshold $\overline{v}_i > \underline{v}_i$ of Lemma 2. Let $\pi_i(k; \mathbf{K}, \mathbf{Z})$ denote the maximised expected operating profit of a sector-i firm with own capital k that prices optimally against the symmetric sectoral aggregates, and $\pi_i^k := \partial\pi_i/\partial k\big|_{k=K_i}$ its marginal profit at the symmetric point.*

(i) (Materials.) *Intermediate spending is the share $\xi_i(\overline{v}_i)$ of firm revenue,*

$$X_i = \xi_i(\overline{v}_i)\, S_i^{\mathrm{f}}, \qquad \xi_i(\overline{v}) := \frac{\eta_{i,1}(\overline{v}) - \frac{1}{\varepsilon}}{\frac{1}{1-\alpha_i}\,\eta_{i,\frac{1}{1-\alpha_i}}(\overline{v})} > 0. \tag{75}$$

(ii) (Capital.) *The marginal profit of capital satisfies* (33)*, which defines the own-capital elasticity of profit at the symmetric point, $\zeta_i(\overline{v})$, wherever $\xi_i(\overline{v}) \neq 1$; under Assumption 3 of Section 2.3 one has $\xi_i(\overline{v}) < 1 - \frac{1}{\varepsilon}$ on the whole of $(\underline{v}_i, \infty)$ (Lemma 6(iii)), so ζ_i*

is then well defined throughout; the product $\zeta_i(\overline{v})[1-\xi_i(\overline{v})]$, which is what the rest of the paper uses, needs no such proviso. Equivalently,

$$\pi_i^k K_i = \underbrace{[1-G_i(\overline{v}_i)]\,\frac{\mathcal{P}_i K_i}{x_i}\left[\frac{p_i Z_i x_i^{\alpha_i}}{\mathcal{P}_i}-1\right]}_{\text{relaxing the rationing}} + \underbrace{\frac{\alpha_i}{1-\alpha_i}\,\frac{\mathcal{P}_i K_i}{x_i}\,\mathcal{G}^-_{i,\frac{1}{1-\alpha_i}}(\overline{v}_i)}_{\text{input saving when unconstrained}}. \tag{76}$$

(iii) (Markup rent.) *The two share functions of* (75) *and* (33) *satisfy*

$$\xi_i(\overline{v}) + \zeta_i(\overline{v})[1-\xi_i(\overline{v})] = 1-\frac{1}{\varepsilon} \qquad \text{for every } \overline{v} > \underline{v}_i, \tag{77}$$

an identity in $\overline{v}$ that uses only the definitions of the two functions. Evaluated at the equilibrium threshold it says that the residual $S_i^{\mathrm{f}} - X_i - \pi_i^k K_i$ equals $S_i^{\mathrm{f}}/\varepsilon$ exactly; that reading is about the optimally posted price, since it is (i)–(ii) that identify the two shares with X_i/S_i^{f} and $\pi_i^k K_i/S_i^{\mathrm{f}}$, and both rest on the pricing first-order condition. Away from the price optimum the residual is not $S_i^{\mathrm{f}}/\varepsilon$.

(iv) (Four-way split of sales.) *Every euro of sectoral sales S_i splits into*

$$\underbrace{\xi_i(\overline{v}_i)\frac{S_i^{\mathrm{f}}}{S_i}}_{\text{materials}} + \underbrace{\zeta_i(\overline{v}_i)[1-\xi_i(\overline{v}_i)]\frac{S_i^{\mathrm{f}}}{S_i}}_{\text{capital income}} + \underbrace{\frac{1}{\varepsilon}\frac{S_i^{\mathrm{f}}}{S_i}}_{\text{markup rent}} + \underbrace{1-\frac{S_i^{\mathrm{f}}}{S_i}}_{\text{rationing rent}} = 1, \tag{78}$$

$$\frac{S_i^{\mathrm{f}}}{S_i} = \frac{\Psi_{i,1}(\overline{v}_i)}{\mathcal{G}^-_{i,1}(\overline{v}_i) + \mathcal{G}^+_{i,1/\varepsilon}(\overline{v}_i)},$$

the second line being (72).

(v) (Limits.) *As $\overline{v} \to \infty$: $\xi_i(\overline{v}) \to \frac{\varepsilon-1}{\varepsilon}(1-\alpha_i)$, $\zeta_i(\overline{v})[1-\xi_i(\overline{v})] \to \frac{\varepsilon-1}{\varepsilon}\alpha_i$, $\zeta_i(\overline{v}) \to \frac{(\varepsilon-1)\alpha_i}{1+(\varepsilon-1)\alpha_i}$ and $S_i^{\mathrm{f}}/S_i \to 1$, so that* (78) *becomes the three-way split of the no-capacity model* (168).

The capital piece is the capital income of the split and the private flow that prices capital in Section 2.3; the markup-rent piece is the dividend $S_i^{\mathrm{f}}/\varepsilon$; the rationing rent is the aggregators' dividend Π_i^{r}. Decomposition (76) shows the two channels through which an extra unit of capital raises a firm's expected profit: in the states in which the firm is rationed, it adds $Z_i x_i^{\alpha_i-1}$ units of output and $1/x_i$ units of bundle, worth $p_i Z_i x_i^{\alpha_i-1} - \mathcal{P}_i/x_i$, a term whose sign depends on whether the posted price covers the bundle cost of a unit produced at capacity; in the states in which it is not, output is demand-determined and the extra capital only lowers the bundle needed to produce it, a saving that is always positive. Without capacity constraints only the second channel exists and (33) reduces to the constant $\frac{\varepsilon-1}{\varepsilon}\alpha_i$; the no-capacity model's split into $\frac{\varepsilon-1}{\varepsilon}(1-\alpha_i)$ and $\frac{\varepsilon-1}{\varepsilon}\alpha_i$ is the $\overline{v}\to\infty$ case of (77) together with the limit of ξ_i in (v), since the identity pins only the sum of the two shares. Two remarks on $\zeta_i(\overline{v})$: without capacity constraints ζ_i is a global elasticity, because profit is exactly a power of own capital; under capacity $\zeta_i(\overline{v})$ is the elasticity at the symmetric point only, and the curvature of π_i in k is taken up in Section 2.3.

Remark 3 (What capacity breaks in the revenue split, and what it leaves) The constant-share result of the no-capacity model (Remark 12) rests on exactly two ingredients: a constant demand elasticity and a Cobb–Douglas requirement of the flexibly chosen bundle. Capacity keeps the second and removes the first, which is precisely the failure that remark anticipates: the effective elasticity $\varepsilon_i(\overline{v}_i)$ varies with the state, so the markup, the materials share $\xi_i(\overline{v}_i)$ and the capital share $\zeta_i(\overline{v}_i)[1-\xi_i(\overline{v}_i)]$ all move with the threshold, hence with $(\mathbf{K},\mathbf{Z})$ and with prices, under either final-demand system. What survives is the markup rent: the share $1/\varepsilon$ of firm revenue is untouched by capacity, at every threshold and for every G_i. The reason is a homogeneity of the pricing problem that rationing does not disturb: along the ray $(k,p)\mapsto(\lambda k,\lambda^{-1/\varepsilon}p)$ the threshold is constant, expected revenue is homogeneous of degree $1-\frac{1}{\varepsilon}$ and expected input cost of degree one, and Euler's identity at the price optimum is (77). The four-way split (78) also shows that the share of sales accruing to the firms, S_i^{f}/S_i, is itself state-dependent, so that even the aggregate markup-rent share of GDP is no longer a constant.

Positivity of the capital share holds under Assumption 3 of Section 2.3: Lemma 6(iii) shows that $\zeta_i(\overline{v})[1-\xi_i(\overline{v})]>0$, hence $\xi_i(\overline{v})<1-\frac{1}{\varepsilon}$, on the whole of $(\underline{v}_i,\infty)$, because the marginal profit of capital is then strictly decreasing towards zero. An unconditional proof is not available; the open question after that assumption covers it.

Proof. Drop the sector index on the tail functions, write $\overline{v}:=\overline{v}_i$, and abbreviate $\eta_1:=\eta_1(\overline{v})$, $\eta_\gamma:=\eta_{\frac{1}{1-\alpha_i}}(\overline{v})$ and $\gamma:=\frac{1}{1-\alpha_i}$ inside this proof only.

Step 1 (materials share). The first-order condition of Step 3 of the proof of Lemma 2, at the symmetric point, reads $[\varepsilon\eta_1-1]\,p_i\mathbb{E}_v[y_i]=\varepsilon\gamma\eta_\gamma\,\mathcal{P}_i\mathbb{E}_v[M_i]$. Since $X_i=\mathcal{P}_iM_i^e=\mathcal{P}_i\mathbb{E}_v[M_i]$ by Lemma 3(iii) and $S_i^{\mathrm{f}}=p_i\mathbb{E}_v[y_i]$,

$$\frac{X_i}{S_i^{\mathrm{f}}}=\frac{\varepsilon\eta_1-1}{\varepsilon\gamma\eta_\gamma}=\frac{\eta_1-\frac{1}{\varepsilon}}{\gamma\eta_\gamma}=\xi_i(\overline{v}),$$

which is positive because $\varepsilon\eta_1>1$ at the optimum (Lemma 2(i)). This proves (i).

Step 2 (marginal profit of capital by the envelope theorem). Let $\Pi(p,k)$ be the expected operating profit of a firm with own capital k posting p, with competitors at the symmetric point, so that $\pi_i(k)=\max_p\Pi(p,k)$. Under Assumption 1 the tail functions are continuously differentiable, so Π is continuously differentiable in (p,k), and by Lemma 2 the maximiser $p^\star(k)$ is unique and interior for every $k>0$; it is continuous in k, because by Step 4 of that proof the optimal threshold is $\Lambda^{-1}(c(k))$ with Λ a continuous strictly increasing bijection and c continuous in k. Then $\pi_i'(k)=\partial_k\Pi(p^\star(k),k)$ by the envelope theorem in its Danskin form: for $k'\neq k$ the two-sided bound $\Pi(p^\star(k),k')-\Pi(p^\star(k),k)\le\pi_i(k')-\pi_i(k)\le\Pi(p^\star(k'),k')-\Pi(p^\star(k'),k)$ holds by optimality, and dividing by $k'-k$ and letting $k'\to k$ squeezes both sides to $\partial_k\Pi(p^\star(k),k)$ by the joint continuity of $\partial_k\Pi$ and the continuity of $p^\star$. No differentiability of $p^\star$, and hence no second-order condition, is needed: the derivative of the optimally chosen price drops out because the bound never differentiates it. Differentiating $\Pi=p\mathbb{E}_v[y]-$

$\mathcal{P}_i \mathbb{E}_v[M]$ in k at fixed p and inserting (68),

$$k\,\pi_i'(k) = [1-\eta_1]\,p\mathbb{E}_v[y] - [1-\gamma\eta_\gamma]\,\mathcal{P}_i\mathbb{E}_v[M],$$

the tail functions being evaluated at the firm's own threshold. At $k = K_i$ this is $\pi_i^k K_i = [1-\eta_1]S_i^{\mathrm{f}} - [1-\gamma\eta_\gamma]X_i$; substituting $X_i = \xi_i(\overline{v})S_i^{\mathrm{f}}$ from Step 1,

$$\frac{\pi_i^k K_i}{S_i^{\mathrm{f}}} = 1-\eta_1-(1-\gamma\eta_\gamma)\frac{\eta_1-\frac{1}{\varepsilon}}{\gamma\eta_\gamma} = 1-\eta_1-\left(\eta_1-\frac{1}{\varepsilon}\right)\left(\frac{1}{\gamma\eta_\gamma}-1\right),$$

which is (33). (The derivation holds at any own k, not only at K_i: $k\pi_i'(k) = \zeta_i(\overline{v}(k))[1-\xi_i(\overline{v}(k))]\,p^\star(k)\mathbb{E}_v[y](k)$ with the firm's own threshold $\overline{v}(k)$, a fact used in Section 2.3.) For the decomposition (76), use $(1-\eta_1)\Psi_1 = 1-G_i(\overline{v})$ and $(1-\gamma\eta_\gamma)\Psi_\gamma = [1-G_i(\overline{v})] - (\gamma-1)\mathcal{G}_\gamma^-$ with $\gamma-1 = \frac{\alpha_i}{1-\alpha_i}$, together with (6)–(12):

$$\pi_i^k K_i = [1-G_i(\overline{v})]\left(p_i\overline{y}_i - \mathcal{P}_i\frac{K_i}{x_i}\right) + \frac{\alpha_i}{1-\alpha_i}\,\mathcal{P}_i\frac{K_i}{x_i}\mathcal{G}_\gamma^-(\overline{v}),$$

and $p_i\overline{y}_i - \mathcal{P}_iK_i/x_i = \frac{\mathcal{P}_iK_i}{x_i}[p_iZ_ix_i^{\alpha_i}/\mathcal{P}_i - 1]$ by (10). This proves (ii).

Step 3 (the rent identity, by Euler's theorem). Adding Steps 1 and 2 gives $\xi_i + \zeta_i[1-\xi_i] = \frac{\eta_1-1/\varepsilon}{\gamma\eta_\gamma} + 1 - \eta_1 - \frac{\eta_1-1/\varepsilon}{\gamma\eta_\gamma} + \eta_1 - \frac{1}{\varepsilon} = 1-\frac{1}{\varepsilon}$, which is (77). The identity has a structural reason that does not use the tail functions. Along the ray $(k,p)\mapsto(\lambda k, \lambda^{-1/\varepsilon}p)$, $\lambda>0$, the threshold $\overline{v} = \frac{Z_ikx_i^{\alpha_i-1}}{Y_i}(p/P_i)^\varepsilon$ is constant, so expected revenue $R(k,p) := p\,\overline{y}(k)\Psi_1(\overline{v})$ is homogeneous of degree $1-\frac{1}{\varepsilon}$ and expected input cost $C(k,p) := \mathcal{P}_i\frac{k}{x_i}\Psi_\gamma(\overline{v})$ of degree one. Euler's theorem along the ray gives $kR_k - \frac{1}{\varepsilon}pR_p = (1-\frac{1}{\varepsilon})R$ and $kC_k - \frac{1}{\varepsilon}pC_p = C$, hence $k\Pi_k - \frac{1}{\varepsilon}p\Pi_p = (1-\frac{1}{\varepsilon})R - C$. At the price optimum $\Pi_p = 0$, so $k\pi_i'(k) = (1-\frac{1}{\varepsilon})R - C$, i.e. $R - C - k\pi_i'(k) = R/\varepsilon$: the residual after materials and the marginal product of capital is exactly the fraction $1/\varepsilon$ of revenue, for every G_i and every threshold. This proves (iii).

Step 4 (four-way split). Divide $S_i = S_i^{\mathrm{f}} + \Pi_i^{\mathrm{r}}$ (Lemma 3(iv)) by S_i and split S_i^{f} by (i)–(iii): $S_i^{\mathrm{f}} = X_i + \pi_i^kK_i + S_i^{\mathrm{f}}/\varepsilon$. The ratio S_i^{f}/S_i is (72). This proves (iv).

Step 5 (limits). By Lemma 4(i), $\eta_1, \eta_\gamma \to 1$, so $\xi_i(\overline{v}) \to (1-\frac{1}{\varepsilon})/\gamma = \frac{\varepsilon-1}{\varepsilon}(1-\alpha_i)$ and $\zeta_i[1-\xi_i] \to 0 - (1-\frac{1}{\varepsilon})(\frac{1}{\gamma}-1) = \frac{\varepsilon-1}{\varepsilon}\alpha_i$, since $\frac{1}{\gamma}-1 = -\alpha_i$; dividing, $\zeta_i \to \frac{\frac{\varepsilon-1}{\varepsilon}\alpha_i}{1-\frac{\varepsilon-1}{\varepsilon}(1-\alpha_i)} = \frac{(\varepsilon-1)\alpha_i}{1+(\varepsilon-1)\alpha_i}$; and $S_i^{\mathrm{f}}/S_i \to 1$ by Lemma 4(i), so the rationing piece of (78) vanishes and the other three tend to the constants of (168). This proves (v). ∎

A.7 Capacity utilization

In the notation of Appendix A.1, the integral in (21) is the moment $\mu_{i,1/\varepsilon}$, so that sector capacity is $\overline{Y}_i = \overline{y}_i\,\mu_{i,1/\varepsilon}^{\frac{\varepsilon}{\varepsilon-1}} = Z_iK_ix_i^{\alpha_i-1}\,\mu_{i,1/\varepsilon}^{\frac{\varepsilon}{\varepsilon-1}}$, and by (69) the utilization function in (22) reads

$$U_i(\overline{v}) = \mu_{i,1/\varepsilon}^{-\frac{\varepsilon}{\varepsilon-1}}\left[\overline{v}^{1/\varepsilon}\Big(\mathcal{G}_{i,1}^-(\overline{v}) + \mathcal{G}_{i,1/\varepsilon}^+(\overline{v})\Big)\right]^{\frac{\varepsilon}{\varepsilon-1}} = \frac{\mathcal{R}_i(\overline{v})^\varepsilon}{\overline{v}\,\mu_{i,1/\varepsilon}^{\frac{\varepsilon}{\varepsilon-1}}}.$$

That $U_i(\cdot)$ is a strictly decreasing bijection of $(0,\infty)$ onto $(0,1)$ is Lemma 3(ii). Since $v \mapsto v^{1/\varepsilon}$ is concave and $\mu_{i,1} = 1$, $\mu_{i,1/\varepsilon}^{\varepsilon/(\varepsilon-1)} \leq 1$: dispersion in the demand shocks lowers sector capacity below the sum of firm capacities, because the CES aggregator values the firms' goods with weights v_{in} that do not match the uniform allocation of capacity across firms. Since both Y_i and $\overline{Y}_i$ are $\overline{y}_i$ times a function of $\overline{v}_i$, a change in K_i or Z_i that leaves $\overline{v}_i$ unchanged scales output and capacity together. Two other decreasing functions of the same threshold describe the sector's operating position: the constrained share $1 - G_i(\overline{v}_i)$, the fraction of firms turning demand away, and the mean physical utilization $\mathbb{E}_v[y_i]/\overline{y}_i = \Psi_{i,1}(\overline{v}_i) = S_i^{\mathrm{f}}/(p_i\overline{y}_i)$, the unweighted average of firm-level output over firm-level capacity. U_i is the CES-aggregate version of the latter. **[Alternative aggregate utilization measures.]** Two natural alternatives to (23), recorded for comparison with data and for the continuum part, where each becomes a moment of the node distribution:

1. *Sales-weighted*: $\overline{U}^{S} := \sum_i S_i U_i / \sum_i S_i$, the Domar-weighted mean; it overweights upstream sectors relative to (23) in proportion to their gross-output-to-value-added ratio.
2. *Capacity-value-weighted*: $\overline{U}^{\overline{Y}} := \sum_i P_i\overline{Y}_i U_i / \sum_i P_i\overline{Y}_i = \sum_i S_i / \sum_i P_i\overline{Y}_i$, since $P_i\overline{Y}_i U_i = P_i Y_i = S_i$: the value of output over the value of capacity at current prices, the closest analogue of a single output-to-capacity ratio for the whole economy.

All three coincide when utilization is common across sectors, and all three are functions of the sectoral thresholds and of the static block's prices and quantities only.

A.8 Cobb–Douglas Final Demand

The Cobb–Douglas final aggregator, the limit $\eta \to 1$ of (24), is

$$Q = \prod_{j=1}^{J} \left(\frac{Q_j}{\beta_j}\right)^{\beta_j}, \quad \beta_j > 0, \quad \sum_j \beta_j = 1 \tag{79}$$

Cost minimisation gives the constant expenditure shares β_j and the price index,

$$\frac{P_j Q_j}{PQ} = \beta_j, \quad P = \prod_{j=1}^{J} P_j^{\beta_j}, \tag{80}$$

with the final good as numeraire, i.e. $P = 1$. Under Cobb–Douglas the β_j are expenditure shares, so $\sum_j \beta_j = 1$ is an adding-up restriction: a $\boldsymbol{\beta}$ with $\sum_j \beta_j \neq 1$ would violate the budget identity, and rescaling $\boldsymbol{\beta}$ would change relative demands. Under CES the shares $\beta_j(P_j/P)^{1-\eta}$ add to one automatically for any positive $\boldsymbol{\beta}$, because P is constructed to make them do so: $\sum_j \beta_j (P_j/P)^{1-\eta} = P^{\eta-1}\sum_j \beta_j P_j^{1-\eta} = 1$. The budget identity therefore imposes nothing on $\sum_j \beta_j$. Under CES the numeraire (26) replaces the Cobb–Douglas one in (80), and $\sum_j \beta_j = 1$ is a normalisation rather than a share restriction.

As $\eta \to 1$, (24) converges to (79), and its demands and price index to (80); the $\beta_j^{-\beta_j}$ normalisation in (79) is what makes the limit exact rather than exact up to a units constant[25].

A.9 Static Equilibrium: Structure

In the definition of the static equilibrium in Section 2.2, item 1 replaces the supply curve of the model without capacity constraints (Appendix C.1) by the pair of relations that link a sector's price, output and threshold, and adds the thresholds to the unknowns; through Lemma 3(ii) it contains the feasibility condition $Y_i < \overline{Y}_i$. Item 2 replaces the constant expenditure share by the state-dependent bundle demand. Everything else is as in the model without capacity constraints. Sectoral sales $S_i = P_i Y_i$, firm revenue S_i^{f}, the rationing rents Π_i^{r}, the shares $\xi_i(\overline{v}_i)$, $\zeta_i(\overline{v}_i)$ and utilization $U_i(\overline{v}_i)$ are then read off Lemma 3, Lemma 5 and Definition 1.

The equilibrium conditions separate into an inner block, solved sector by sector at given prices, and an outer block in the prices alone. Fix $\mathbf{P} \in \mathbb{R}^J_{++}$, hence the bundle prices $\mathcal{P}_i(\mathbf{P})$ of (13). The supply relation, in the first form of (74), is then a scalar equation in the sector's threshold alone,

$$\frac{\mathcal{R}_i(\overline{v})\,\overline{v}^{\frac{\alpha_i}{1-\alpha_i}}}{\Gamma_i(\overline{v})} = \frac{\mathcal{P}_i(\mathbf{P})}{P_i\,Z_i\,x_i^{\alpha_i}}, \qquad \overline{v} \in (\underline{v}_i, \infty), \tag{81}$$

whose left-hand side is a known function of the threshold, built from the tail functions of G_i at $\overline{v}$: by the second form of (64) it equals

$$\frac{\mathcal{R}_i(\overline{v})\Big[(\varepsilon-1)\mathcal{G}^-_{i,1}(\overline{v}) - 1 + G_i(\overline{v})\Big]}{\varepsilon^{\frac{1}{1-\alpha_i}}\mathcal{G}^-_{i,\frac{1}{1-\alpha_i}}(\overline{v})},$$

and whose right-hand side collects the sector's own price, its bundle price and its productivity. Neither the sector's capital nor any other sector's state appears: given prices, the threshold is a function of $(\mathbf{P}, Z_i)$ only. The restriction to $(\underline{v}_i, \infty)$ is the interior-markup condition $\varepsilon_i(\overline{v}) > 1$, which every optimal price satisfies (Lemma 2(i)). Given a solution $\overline{v}_i$ of (81), the sector's quantities follow from Lemma 3 and Definition 1:

$$\begin{aligned}
&\overline{y}_i = Z_i K_i x_i^{\alpha_i - 1}, \quad Y_i = \overline{y}_i\,\frac{\mathcal{R}_i(\overline{v}_i)^{\varepsilon}}{\overline{v}_i}, \quad S_i = P_i Y_i, \quad U_i = U_i(\overline{v}_i),\\
&S_i^{\mathrm{f}} = S_i\,\frac{\Psi_{i,1}(\overline{v}_i)}{\mathcal{G}^-_{i,1}(\overline{v}_i) + \mathcal{G}^+_{i,1/\varepsilon}(\overline{v}_i)}, \quad \Pi_i^{\mathrm{r}} = S_i - S_i^{\mathrm{f}}, \quad Q_j = \beta_j P_j^{-\eta}\,Q,\\
&M_i^{e} = \frac{K_i}{x_i}\Psi_{i,\frac{1}{1-\alpha_i}}(\overline{v}_i), \quad X_i = \mathcal{P}_i(\mathbf{P})M_i^{e} = \xi_i(\overline{v}_i)S_i^{\mathrm{f}}, \quad m_{ij} = a_{ij}\frac{X_i}{P_j},
\end{aligned} \tag{82}$$

by the final demands (25) with $P = 1$.

[25] Write $\beta_j^{1/\eta} Q_j^{(\eta-1)/\eta} = \beta_j (Q_j/\beta_j)^{(\eta-1)/\eta}$ and expand to first order in $(\eta-1)/\eta$; both aggregators turn the balanced bundle $Q_j = \beta_j Q$ into exactly Q units of the final good.

The value form of clearing, (29), follows by multiplying the sectoral clearing condition (27) by P_j and using $m_{ij} = a_{ij}X_i/P_j$ and the final demands (25) with $P = 1$. Summing its first equation over j and using $\sum_j a_{ij} = 1$ gives $\sum_j S_j = Q\sum_j \beta_j P_j^{1-\eta} + \sum_i X_i$, which is the second equation by the numeraire (26). One clearing equation is therefore redundant (Walras's law), and the outer block consists of $J-1$ independent clearing equations plus the numeraire in the J unknown prices, each price vector being mapped to thresholds and quantities through the inner block (81)–(82); the numeraire is a surface in price space on which the clearing equations are solved.

Proposition 3 (Static equilibrium under capacity: structure and homogeneity). *Let Assumptions 1 and 2 hold.*

(i) (Inner block.) *For every $\mathbf{P} \in \mathbb{R}^J_{++}$ and every sector i, equation* (81) *has at least one solution in $(\underline{v}_i, \infty)$: its left-hand side is continuous, tends to 0 as $\overline{v} \downarrow \underline{v}_i$ and to ∞ as $\overline{v} \to \infty$. Its solution set depends on $(\mathbf{P}, Z_i)$ and not on $\mathbf{K}$. Given a solution, $Y_i, S_i, S_i^{\mathrm{f}}, \Pi_i^{\mathrm{r}}, M_i^e, X_i$ and m_{ij} in* (82) *are linear in K_i, while U_i, S_i^{f}/S_i, ξ_i, ζ_i, ε_i and the markup are functions of $\overline{v}_i$ alone and hence, at given prices, do not depend on K_i at all.*

(ii) (Homogeneity.) *If $\big(\mathbf{P}, \overline{\mathbf{v}}, \mathbf{Y}, \mathbf{X}, \{m_{ij}\}, \{Q_j\}, Q\big)$ is a static equilibrium at $(\mathbf{K}, \mathbf{Z})$ and $\lambda > 0$, then $\big(\mathbf{P}, \overline{\mathbf{v}}, \lambda\mathbf{Y}, \lambda\mathbf{X}, \{\lambda m_{ij}\}, \{\lambda Q_j\}, \lambda Q\big)$ is a static equilibrium at $(\lambda\mathbf{K}, \mathbf{Z})$, and every static equilibrium at $(\lambda\mathbf{K}, \mathbf{Z})$ arises in this way. Consequently, along any equilibrium selection that respects this scaling (in particular whenever the static equilibrium is unique), prices, thresholds, utilization, the share functions and the revenue ratios S_i^{f}/S_i are functions of $(\boldsymbol{\kappa}, \mathbf{Z})$ alone, and* (31) *holds with $q = \sum_i(\widehat{s}_i - \widehat{x}_i)$: final output is homogeneous of degree one in the capital stocks.*

(iii) (Euler identities.) *Wherever such a selection $Q(\mathbf{K}, \mathbf{Z})$ is differentiable, $\sum_i K_i Q_{K_i} = Q$, the marginal products Q_{K_i} are homogeneous of degree zero in $\mathbf{K}$, and, writing q as a function of the independent coordinates $(\kappa_1, \ldots, \kappa_{J-1})$ of* (32),

$$q(\boldsymbol{\kappa}, \mathbf{Z}) = \sum_{i=1}^{J} \kappa_i\, Q_{K_i}, \qquad \frac{\partial q}{\partial \kappa_a} = Q_{K_a} - Q_{K_J}, \quad a = 1, \ldots, J-1. \tag{83}$$

(iv) (Operating range.) *In every static equilibrium and every sector, $\overline{v}_i \in (\underline{v}_i, \infty)$; hence $\varepsilon_i(\overline{v}_i) > 1$, $\xi_i(\overline{v}_i) > 0$, and*

$$U_i < U_i(\underline{v}_i) < 1, \qquad 1 - G_i(\overline{v}_i) < 1 - G_i(\underline{v}_i) < 1 :$$

utilization and the rationed share of firms are bounded away from one by constants that depend on $(\varepsilon, \alpha_i, G_i)$ alone.

Proof. Write $\gamma_i := \frac{1}{1-\alpha_i}$ and $D_i(\overline{v}) := (\varepsilon - 1)\mathcal{G}^-_{i,1}(\overline{v}) - 1 + G_i(\overline{v})$, as in the proof of Lemma 2.

Step 1 (the inner block). By the second form of (64), $\Gamma_i(\overline{v})\overline{v}^{-\frac{\alpha_i}{1-\alpha_i}} = \varepsilon\gamma_i\mathcal{G}^-_{i,\gamma_i}(\overline{v})/D_i(\overline{v})$, so the first form of (74) reads $P_i\mathcal{R}_i(\overline{v})D_i(\overline{v})/[\varepsilon\gamma_i\mathcal{G}^-_{i,\gamma_i}(\overline{v})] = \mathcal{P}_i/(Z_i x_i^{\alpha_i})$, which is (81); K_i appears nowhere in it, since $\mathcal{R}_i$, Γ_i and D_i are functions of the threshold alone and

the aggregates on the right are fixed. Let $F_i(\overline{v}) := \mathcal{R}_i(\overline{v})D_i(\overline{v})/\mathcal{G}^-_{i,\gamma_i}(\overline{v})$ on $(\underline{v}_i, \infty)$. F_i is continuous (Assumption 1); $F_i(\underline{v}_i^+) = 0$ because $D_i(\underline{v}_i) = 0$ while $\mathcal{R}_i(\underline{v}_i) > 0$ and $\mathcal{G}^-_{i,\gamma_i}(\underline{v}_i) > 0$; and as $\overline{v} \to \infty$, $\mathcal{R}_i \to 1$, $\overline{v}D_i(\overline{v}) \to \varepsilon - 1$ and $\overline{v}^{\gamma_i}\mathcal{G}^-_{i,\gamma_i}(\overline{v}) \to \mu_{i,\gamma_i}$ by Lemma 4(i), so $F_i(\overline{v}) \sim (\varepsilon-1)\overline{v}^{\gamma_i - 1}/\mu_{i,\gamma_i} \to \infty$ because $\gamma_i - 1 = \frac{\alpha_i}{1-\alpha_i} > 0$. By the intermediate value theorem F_i takes the value $\varepsilon\gamma_i\mathcal{P}_i/(P_iZ_ix_i^{\alpha_i})$ at least once, and the solution set depends only on $(\mathbf{P}, Z_i)$ through that value. Given a root, $Y_i = \overline{y}_i\mathcal{R}_i(\overline{v}_i)^\varepsilon/\overline{v}_i$ by Lemma 3(ii), M_i^e by (iii), S_i, S_i^{f} and Π_i^{r} by (iv), and the remaining quantities by definition; each of $Y_i, S_i, S_i^{\mathrm{f}}, \Pi_i^{\mathrm{r}}, M_i^e, X_i, m_{ij}$ carries the factor K_i through $\overline{y}_i = Z_iK_ix_i^{\alpha_i - 1}$ or through K_i/x_i, and nothing else in it depends on K_i; $U_i(\overline{v}_i)$, $S_i^{\mathrm{f}}/S_i = \Psi_{i,1}/(\mathcal{G}^-_{i,1} + \mathcal{G}^+_{i,1/\varepsilon})$, $\xi_i(\overline{v}_i)$, $\zeta_i(\overline{v}_i)$, $\varepsilon_i(\overline{v}_i)$ and the markup are functions of $\overline{v}_i$ alone. This proves (i).

Step 2 (homogeneity). Let $(\mathbf{P}, \overline{\mathbf{v}}, \mathbf{Y}, \mathbf{X}, \{m_{ij}\}, \{Q_j\}, Q)$ be a static equilibrium at $(\mathbf{K}, \mathbf{Z})$ and $\lambda > 0$; consider $(\lambda\mathbf{K}, \mathbf{Z})$ with the same $\mathbf{P}$ and $\overline{\mathbf{v}}$. Item 1 of the definition: the supply relation is (81), which does not involve $\mathbf{K}$ (Step 1), so it holds with the same thresholds; the output relation (70) holds with λY_i because $\overline{y}_i$ scales by λ. Item 2: M_i^e scales by λ at the same threshold and bundle price, so $X_i \mapsto \lambda X_i$. Item 3: $m_{ij} \mapsto \lambda m_{ij}$. Item 4: with $Q \mapsto \lambda Q$ and $\mathbf{P}$ unchanged, $Q_j = \beta_j P_j^{-\eta} Q \mapsto \lambda Q_j$, and the price index involves prices only. Item 5: clearing, $Y_j = Q_j + \sum_i m_{ij}$, is linear in the quantities and so holds after scaling, and $Q = \sum_j (S_j - X_j) \mapsto \lambda Q$ consistently. Item 6 involves prices only. Hence the scaled tuple is a static equilibrium at $(\lambda\mathbf{K}, \mathbf{Z})$; applying the same argument with λ^{-1} shows that every equilibrium at $(\lambda\mathbf{K}, \mathbf{Z})$ is the scaled image of one at $(\mathbf{K}, \mathbf{Z})$. Along a selection that respects the scaling, $\mathbf{P}$, $\overline{\mathbf{v}}$ and every function of them are homogeneous of degree zero and $Q, S_i, S_i^{\mathrm{f}}, X_i$ of degree one in $\mathbf{K}$; taking $\lambda = 1/K$ gives (31) with $q(\boldsymbol{\kappa}, \mathbf{Z}) := Q(\boldsymbol{\kappa}, \mathbf{Z})$, the final output at aggregate capital one and composition $\boldsymbol{\kappa}$, and likewise for $\widehat{s}_i, \widehat{s}_i^{\mathrm{f}}, \widehat{x}_i$; $q = \sum_i(\widehat{s}_i - \widehat{x}_i)$ is $Q = \sum_i(S_i - X_i)$ divided by K. This proves (ii).

Step 3 (Euler identities). Differentiating $Q(\lambda\mathbf{K}, \mathbf{Z}) = \lambda Q(\mathbf{K}, \mathbf{Z})$ in λ at $\lambda = 1$ gives $\sum_i K_iQ_{K_i}(\mathbf{K}, \mathbf{Z}) = Q(\mathbf{K}, \mathbf{Z})$ (Euler's theorem), and differentiating it in K_i gives $Q_{K_i}(\lambda\mathbf{K}, \mathbf{Z}) = Q_{K_i}(\mathbf{K}, \mathbf{Z})$: the marginal products are homogeneous of degree zero, hence functions of $(\boldsymbol{\kappa}, \mathbf{Z})$. Dividing Euler's identity by K gives $q = \sum_i \kappa_i Q_{K_i}$. For the second identity, write q as a function of the independent coordinates at aggregate capital one,

$$q(\kappa_1, \ldots, \kappa_{J-1}, \mathbf{Z}) = Q\big(\kappa_1, \ldots, \kappa_{J-1},\, 1 - \textstyle\sum_{i<J}\kappa_i,\, \mathbf{Z}\big);$$

by the chain rule $\partial q/\partial\kappa_a = Q_{K_a} - Q_{K_J}$. This proves (iii).

Step 4 (operating range). In a static equilibrium each sector is at the symmetric equilibrium of Lemma 3, whose threshold is the optimal threshold of Lemma 2 at $k_{in} = K_i$, and by Lemma 2(i) that threshold lies in $(\underline{v}_i, \infty)$ with $\varepsilon_i(\overline{v}_i) > 1$; positivity of $\xi_i(\overline{v}_i)$ is Lemma 5(i). Since $U_i(\cdot)$ is strictly decreasing (Definition 1), $U_i(\overline{v}_i) < U_i(\underline{v}_i)$; and $U_i(\overline{v}) = \mu_{i,1/\varepsilon}^{-\frac{\varepsilon}{\varepsilon-1}}\mathcal{R}_i(\overline{v})^\varepsilon/\overline{v} < 1$ for every $\overline{v} > 0$, because $\overline{v}/\mathcal{R}_i(\overline{v})^\varepsilon > \mu_{i,1/\varepsilon}^{-\frac{\varepsilon}{\varepsilon-1}}$ by Lemma 3(ii). Likewise $1 - G_i(\overline{v}_i) < 1 - G_i(\underline{v}_i) < 1$, as G_i is strictly increasing and $\underline{v}_i > 0$ (indeed $D_i(\overline{v}) \to -1$ as $\overline{v} \downarrow 0$, so the zero of D_i is interior). This proves (iv). ■

The proposition is a statement about structure, not about existence: neither the existence nor the uniqueness of the static equilibrium under capacity is claimed, and the

finite model is not solved in this paper. Two pieces are missing. Given $\mathbf{P}$, uniqueness of the inner root requires the left-hand side of (81) to be strictly increasing on $(\underline{v}_i, \infty)$, a property of $(\varepsilon, \alpha_i, G_i)$ alone that we have not established (it holds numerically for the log-normal specification; see the flag below); the continuum model needs the same property at every node and names it as Assumption 5 in Section B.6, where it also delivers the uniqueness of the per-node block. And the outer clearing system in $\mathbf{P}$ has no known monotone structure: without capacity constraints under Cobb–Douglas final demand it is log-linear (Appendix C.2), whereas here the inner block is nonlinear in prices through the thresholds under either demand system. The continuum model of Part II replaces the outer block by a fixed point in cross-sectional moments and gives the per-node block a scalar structure, whose root is studied there (Section 3.2). The maximal utilization $U_i(\underline{v}_i)$ and the maximal rationed share $1 - G_i(\underline{v}_i)$ are primitives of the sector; for the log-normal specification, $U_i(\underline{v}_i)$ ranges from about 0.7 (dispersed shocks, low elasticity) to close to one (concentrated shocks, high elasticity). **[Flag: uniqueness of the inner root is verified numerically, not proved.]** On the grid of 648 log-normal parameter triples used for Assumption 3 below ($\sigma_{v,i} \in \{0.1, \ldots, 2\}$, $\varepsilon \in \{1.2, \ldots, 15\}$, $\alpha_i \in \{0.1, \ldots, 0.85\}$; 4001 logarithmically spaced thresholds on $(\underline{v}_i, e^{14})$) the left-hand side of (81) was strictly increasing in every triple, so that given $\mathbf{P}$ the inner root is unique. Analytically, its logarithmic derivative is $\frac{(1-1/\varepsilon)\mathcal{G}^+_{i,1/\varepsilon}}{\mathcal{G}^-_{i,1}+\mathcal{G}^+_{i,1/\varepsilon}} + \frac{1}{1-\alpha_i} - \frac{\overline{v}g_i(\overline{v})}{\mathcal{G}^-_{i,\frac{1}{1-\alpha_i}}} + \frac{\overline{v}D_i'(\overline{v})}{D_i(\overline{v})}$ with $D_i := (\varepsilon-1)\mathcal{G}^-_{i,1} - 1 + G_i$ and g_i the density; it tends to $+\infty$ at $\underline{v}_i$ and to $\frac{\alpha_i}{1-\alpha_i} > 0$ at infinity, but the two density terms in between have indeterminate signs. In a three-sector Cobb–Douglas example solved numerically (inner root by bisection, outer block by Newton's method) the clearing residuals vanished to machine precision, the scaling $\mathbf{K} \mapsto \lambda\mathbf{K}$ left $(\mathbf{P}, \overline{\mathbf{v}})$ unchanged and scaled Q exactly, both Euler identities of (83) held to 10^{-9} against finite differences, and every threshold lay in $(\underline{v}_i, \infty)$ with $U_i < U_i(\underline{v}_i)$.

Remark 4 (What capacity removes from the static block) Without capacity constraints, under Cobb–Douglas final demand the static block is log-linear and closed-form. Intermediate spending is the constant share ξ_i of sales, so the sales system is $\mathbf{S} = \mathbf{h}Q$ with the constant Domar weights $\mathbf{h} = [\mathbf{I} - \mathbf{A}'\Xi]^{-1}\boldsymbol{\beta}$; combined with the constant-wedge supply curve this gives a linear system in $(\ln \mathbf{P}, \ln Q)$ whose solution is the AK map $Q = \mathcal{B}\prod_i K_i^{\omega_i}\prod_i Z_i^{\omega_i/\alpha_i}$, with capital elasticities $\boldsymbol{\omega}$ built from $(\boldsymbol{\beta}, \mathbf{A}, \boldsymbol{\alpha})$ alone and a level constant $\mathcal{B}$ that carries all distortions (Proposition 9, Appendix C.2). Under capacity nothing of this survives except homogeneity, and it fails under Cobb–Douglas as much as under CES. Intermediate spending is $X_i = \xi_i(\overline{v}_i)S_i^{\mathrm{f}} = \xi_i(\overline{v}_i)[S_i^{\mathrm{f}}/S_i]\,S_i$, a state-dependent share of sales, so the sales system keeps its Leontief form, $\mathbf{S} = [\mathbf{I} - \mathbf{A}'\Xi(\overline{\mathbf{v}})]^{-1}\mathbf{b}(\mathbf{P})\,Q$ with $\Xi(\overline{\mathbf{v}}) := \mathrm{diag}(\xi_i(\overline{v}_i)S_i^{\mathrm{f}}/S_i)$, whose inverse exists whenever the effective spending shares are below one (under Assumption 3 of Section 2.3, $\xi_i(\overline{v}) < 1 - \frac{1}{\varepsilon}$ and $S_i^{\mathrm{f}} \le S_i$ give the uniform bound $\xi_i(\overline{v}_i)S_i^{\mathrm{f}}/S_i < 1 - \frac{1}{\varepsilon}$), but the Domar weights $\mathbf{h}(\mathbf{P}, \overline{\mathbf{v}})$ now depend on prices and thresholds; the supply relation is implicit in the threshold and not log-linear; the price system is therefore nonlinear, there are no constant capital elasticities $\boldsymbol{\omega}$, no level constant $\mathcal{B}$, and no separability $q =$

$\mathcal{Z}(\mathbf{Z})\mathcal{K}(\boldsymbol{\kappa})$. Both demand systems require a root-find, and the comparative statics of Appendix A.10, which without capacity constraints contrast a closed form with a root-find, are qualitative here under either system.

A.10 CES versus Cobb–Douglas Final Demand

Since η enters only the final-good block, all within-sector objects are unchanged: the sectoral CES block, $\mathcal{P}_i(\mathbf{P})$, the pricing rule and the threshold equation (81), and hence $Y_i, M_i^e, X_i, S_i^{\mathrm{f}}$ and U_i as functions of $(\mathbf{P}, \overline{\mathbf{v}})$ in (82). Only the final-demand loading $b_j(\mathbf{P})$ of (29) and the numeraire change. Without capacity constraints the CES case differs from the Cobb–Douglas case precisely in the price-dependence of that loading, which destroys the log-linearity of the sales system and with it the closed form (Appendix C.2); under capacity the sales system is already nonlinear in prices through the thresholds, so the distinction between the two demand systems is no longer that between a closed form and a root-find but only that between price-invariant and price-dependent final expenditure shares. Degree-one homogeneity holds under both (Proposition 3(ii)), so the AK reduction is intact, with q defined implicitly by the static solve in either case.

Without capacity constraints under Cobb–Douglas final demand the local weights $\omega_i(\boldsymbol{\kappa}, \mathbf{Z})$ of Section 2.2 are the constants $\boldsymbol{\omega}$ and the target composition $\boldsymbol{\kappa}^*(\mathbf{Z})$ equals $\boldsymbol{\omega}$ (Appendix C.2).

What is unchanged from Cobb–Douglas. Everything within sectors: the pricing rule, the threshold equation, the share functions and the rent split, all as functions of the threshold; the propagation logic of the sales system, in which final demand is passed upstream and discounted at every link by the effective spending share $\xi_i(\overline{v}_i)S_i^{\mathrm{f}}/S_i$ (Remark 4); and degree-one homogeneity in $\mathbf{K}$, so the AK reduction is intact.

The new channel. Expenditure shares now respond to relative prices:

$$\frac{P_jQ_j}{PQ} = \beta_j\left(\frac{P_j}{P}\right)^{1-\eta} \implies \frac{\partial \ln\left(P_jQ_j/PQ\right)}{\partial \ln\left(P_j/P\right)} = 1-\eta \neq 0,$$

whereas under Cobb–Douglas ($\eta = 1$) a price fall is exactly offset by a quantity rise and no shock can move any final expenditure share. Consequences of $1-\eta \neq 0$:

- *The direction is signed by η.* A positive shock to Z_i lowers P_i. Under substitutes ($\eta > 1$) quantities respond more than one-for-one, so spending moves *towards* the now-cheap sector: productive sectors gain influence, and shocks are amplified. Under complements ($\eta < 1$) quantities respond less than one-for-one, so spending moves towards *expensive* (scarce or unproductive) sectors: adverse shocks raise a sector's influence. This is the bottleneck logic of networks with complementarities.
- *The state-contingent reallocation channel.* Because final shares, and with them the local weights $\omega_i(\boldsymbol{\kappa}, \mathbf{Z})$, move with the state, relative marginal products Q_{K_i}/Q_{K_j}

depend on $\mathbf{Z}$, and $\boldsymbol{\kappa}^*(\mathbf{Z})$ is state-dependent: with substitutes capital should *chase* high-productivity sectors, with complements it should *hedge* towards potential bottlenecks. Under capacity this channel is present under Cobb–Douglas final demand as well, because the thresholds, and through them the shares and the wedges, respond to relative prices; what CES adds is the response of the final expenditure shares themselves.

- *Level distortions are not composition-neutral.* Without capacity constraints under Cobb–Douglas, ε and $\mu_{i,\frac{1}{1-\alpha_i}}$ are pure level effects on $\mathcal{B}$ and leave $\boldsymbol{\omega}$ untouched, while under CES they shift relative prices and price-dependent shares turn the shift into a movement of the local weights. Under capacity the same is true under either demand system: the wedge $\Gamma_i(\overline{v}_i)$ and the shares are state-dependent, so markups and demand risk generically move the composition as well as the level.

A.11 Dynamic Problem

Assumption 3 (Concave profit in own capital). For each sector i, the function

$$\Phi_i(\overline{v}) := \frac{\zeta_i(\overline{v})[1-\xi_i(\overline{v})]\,\mathcal{G}^-_{i,\frac{1}{1-\alpha_i}}(\overline{v})}{\varepsilon_i(\overline{v})-1}, \qquad \varepsilon_i(\overline{v}) = \varepsilon\,\eta_{i,1}(\overline{v}), \tag{84}$$

is strictly decreasing on $(\underline{v}_i, \infty)$.

Up to the constant $\varepsilon\frac{1}{1-\alpha_i}\mathcal{P}_i/x_i$, Φ_i is the marginal profit of capital written as a function of the firm's threshold (Lemma 6(ii) below), and the threshold is increasing in own capital, so the assumption says exactly that a firm's marginal profit falls with its capital. Like Assumption 2, it is a condition on G_i and (ε, α_i) alone, and its two ends are fixed: $\Phi_i(\overline{v}) \to +\infty$ as $\overline{v} \downarrow \underline{v}_i$, where the effective elasticity tends to one, and $\Phi_i(\overline{v}) \sim \frac{\alpha_i}{\varepsilon}\mu_{i,\frac{1}{1-\alpha_i}}\overline{v}^{-\frac{1}{1-\alpha_i}}$ as $\overline{v} \to \infty$ (Lemma 4), a strictly decreasing power, which is the form the assumption takes in the model without capacity constraints, where profit is exactly a power of own capital with exponent below one. The assumption therefore restricts the shape of G_i only through the way Φ_i descends between two fixed ends. **[Flag: Assumption 3 is verified numerically, not proved.]** For the log-normal specification the assumption was checked on a grid of 648 parameter triples, $\sigma_{v,i} \in \{0.1, \ldots, 2\}$, $\varepsilon \in \{1.2, \ldots, 15\}$, $\alpha_i \in \{0.1, \ldots, 0.85\}$, on 4001 logarithmically spaced thresholds in $(\underline{v}_i, e^{14})$: Φ_i was strictly decreasing in every triple, and the capital share $\zeta_i(\overline{v})[1-\xi_i(\overline{v})]$ was strictly positive throughout, with minimum 0.017 (attained in the no-capacity limit at $\varepsilon = 1.2$, $\alpha_i = 0.1$). Independently, in 300 random pricing problems the maximised profit was strictly concave in own capital at the point tested (second difference), its derivative matched the envelope formula (33) to 5×10^{-9} and the explicit form (86) to 2×10^{-13}, and at large own capital the profit function had the no-capacity model's logarithmic slope $\zeta_i(\infty)$ to four digits. The natural route to a proof is a log-concavity or monotone-hazard condition on the lower tail of G_i, as for Assumption 2; the two assumptions are related but neither is known to imply the

other. **[Q13 | text | tag: LogNormal2]** Prove Assumption 3 for the log-normal specification, or give a sufficient condition on G_i under which it holds, and relate it to Assumption 2 (does one imply the other?); then either replace the assumption by a lemma or keep it and cite the sufficient condition after it. This subsumes the positivity of the capital share, $\xi_i(\overline{v}) < 1 - \frac{1}{\varepsilon}$, which Lemma 6(iii) derives from the assumption.

Lemma 6 (Profit and marginal profit of a firm with own capital k under capacity). *Consider a sector-i monopolist with own capital $k > 0$ that prices optimally against the aggregator's demand (Lemma 1), taking $(Y_i, P_i, \mathcal{P}_i(\mathbf{P}), Z_i)$ as given, while its competitors are at the symmetric point of Lemma 3. Let Assumptions 1 and 2 hold.*

(i) (Profit at own k.) *Its capacity is $\overline{y}_i(k) = Z_i k x_i^{\alpha_i - 1}$; its threshold $\overline{v}_i(k)$ is the unique root in $(\underline{v}_i, \infty)$ of* (66)*, whose right-hand side is proportional to $k^{1/\varepsilon}$, and is a continuous strictly increasing bijection of $(0, \infty)$ onto $(\underline{v}_i, \infty)$; its price is $p_i(k) = P_i[\overline{v}_i(k) Y_i / \overline{y}_i(k)]^{1/\varepsilon}$. Its expected operating profit is*

$$\pi_i(k; \mathbf{K}, \mathbf{Z}) = p_i(k)\, \overline{y}_i(k)\, \Psi_{i,1}(\overline{v}_i(k)) - \mathcal{P}_i(\mathbf{P})\, \frac{k}{x_i}\, \Psi_{i,\frac{1}{1-\alpha_i}}(\overline{v}_i(k)), \tag{85}$$

a continuous function of k with no closed form, equal to $S_i^{\mathrm{f}} - X_i = [1 - \xi_i(\overline{v}_i)] S_i^{\mathrm{f}}$ at the symmetric point $k = K_i$.

(ii) (Marginal profit at any k.) *π_i is continuously differentiable and, with the tail functions evaluated at the firm's own threshold $\overline{v} = \overline{v}_i(k)$,*

$$k\, \pi_i'(k) = \zeta_i(\overline{v})[1 - \xi_i(\overline{v})]\, p_i(k) \mathbb{E}_v[y_{in}] \qquad \text{and} \qquad \pi_i'(k) = \frac{\varepsilon}{1 - \alpha_i} \frac{\mathcal{P}_i(\mathbf{P})}{x_i}\, \Phi_i(\overline{v}), \tag{86}$$

with Φ_i of (84)*; at $k = K_i$ the first form is* (33)*.*

(iii) (Concavity.) *π_i is strictly concave on $(0, \infty)$ if and only if Assumption 3 holds for sector i. In that case π_i' decreases strictly from $+\infty$ as $k \downarrow 0$ to 0 as $k \to \infty$; in particular $\pi_i'(k) > 0$ for every k, so that $\zeta_i(\overline{v})[1 - \xi_i(\overline{v})] > 0$ and $\xi_i(\overline{v}) < 1 - \frac{1}{\varepsilon}$ for every $\overline{v} \in (\underline{v}_i, \infty)$.*

(iv) (Large-capital asymptote.) *As $k \to \infty$ at fixed aggregates, $\overline{v}_i(k) \to \infty$ and $k^{1 - \zeta_i(\infty)} \pi_i'(k)$ converges to a positive constant, where $\zeta_i(\infty) = \frac{(\varepsilon - 1)\alpha_i}{1 + (\varepsilon - 1)\alpha_i}$ is the own-capital elasticity of profit of the model without capacity constraints (Lemma 5(v)): the power form of the profit function without capacity constraints (Appendix C.3) is the large-capital asymptote of* (85)*.*

Proof. Notation as in the proofs of Lemmas 2 and 5: $\gamma := \frac{1}{1 - \alpha_i}$, $D(\overline{v}) := (\varepsilon - 1)\mathcal{G}_1^-(\overline{v}) - 1 + G_i(\overline{v})$, $\Lambda(\overline{v}) := \overline{v}^{1/\varepsilon} D(\overline{v}) / \mathcal{G}_\gamma^-(\overline{v})$, and $\Pi(p, k) := p \mathbb{E}_v[y_{in}] - \mathcal{P}_i \mathbb{E}_v[M_{in}]$ the expected operating profit of a firm posting p with own capital k; the sector index is dropped on the tail functions.

Step 1 (threshold, price and profit as functions of own capital). By Step 4 of the proof of Lemma 2, the optimal threshold at own capital k is the unique solution in $(\underline{v}_i, \infty)$

of $\Lambda(\overline{v}) = c(k)$, with $c(k) := \frac{\varepsilon}{1-\alpha_i}\frac{\mathcal{P}_i}{Z_i x_i^{\alpha_i} P_i}(\overline{y}_i(k)/Y_i)^{1/\varepsilon}$, a continuous strictly increasing function of k proportional to $k^{1/\varepsilon}$ with $c(0^+) = 0$ and $c(\infty) = \infty$. Under Assumption 2, Λ is a continuous strictly increasing bijection of $(\underline{v}_i, \infty)$ onto $(0, \infty)$ (it vanishes at $\underline{v}_i^+$ and diverges at infinity, as shown there), so $\overline{v}_i(k) = \Lambda^{-1}(c(k))$ is a continuous strictly increasing bijection of $(0, \infty)$ onto $(\underline{v}_i, \infty)$. The price follows from the threshold definition (5), and (85) is $\Pi(p_i(k), k)$ written with (6)–(12); it is continuous in k. At $k = K_i$, $p_i(K_i) = p_i$ and $\overline{y}_i(K_i) = \overline{y}_i$, so $\Pi = S_i^{\mathrm{f}} - X_i$ by (72) and (18), and $X_i = \xi_i(\overline{v}_i)S_i^{\mathrm{f}}$ by Lemma 5(i). This proves (i).

Step 2 (marginal profit). Π is continuously differentiable in (p, k) (Assumption 1) and, for each k, has the unique interior maximiser $p_i(k)$, which is continuous in k by Step 1. By the envelope theorem in the Danskin form of Step 2 of the proof of Lemma 5, π_i is differentiable with $\pi_i'(k) = \partial_k \Pi(p, k)\big|_{p=p_i(k)}$, which is continuous in k. By (68), $k\,\pi_i'(k) = [1 - \eta_1]\, p\mathbb{E}_v[y] - [1 - \gamma\eta_\gamma]\, \mathcal{P}_i \mathbb{E}_v[M]$ with the tail functions at the firm's own threshold, and Step 2 of the proof of Lemma 5 turns this into $\zeta_i(\overline{v})[1 - \xi_i(\overline{v})]\, p\mathbb{E}_v[y]$: the substitution used there, $\mathcal{P}_i\mathbb{E}_v[M] = \xi_i(\overline{v})\, p\mathbb{E}_v[y]$, is the price first-order condition, which holds at every k. For the explicit form, write $p\mathbb{E}_v[y] = p\,\overline{y}_i(k)\Psi_1(\overline{v})$ and insert the pricing rule (63), $p = \frac{\varepsilon\eta_1}{\varepsilon\eta_1 - 1} \cdot \gamma \frac{\mathcal{P}_i}{Z_i x_i^{\alpha_i}} \cdot \frac{\mathcal{G}_\gamma^-}{\mathcal{G}_1^-}$: since $\overline{y}_i(k)/(Z_i x_i^{\alpha_i}) = k/x_i$ and $\eta_1\Psi_1 = \mathcal{G}_1^-$,

$$p\,\overline{y}_i(k)\Psi_1(\overline{v}) = \varepsilon\gamma\,\frac{\mathcal{P}_i}{x_i}\,k \cdot \frac{\eta_1\Psi_1}{\varepsilon\eta_1 - 1} \cdot \frac{\mathcal{G}_\gamma^-}{\mathcal{G}_1^-} = \varepsilon\gamma\,\frac{\mathcal{P}_i}{x_i}\,k\,\frac{\mathcal{G}_\gamma^-(\overline{v})}{\varepsilon\eta_1(\overline{v}) - 1},$$

so $k\pi_i'(k) = \varepsilon\gamma\frac{\mathcal{P}_i}{x_i}k\,\zeta_i(\overline{v})[1 - \xi_i(\overline{v})]\mathcal{G}_\gamma^-(\overline{v})/(\varepsilon_i(\overline{v}) - 1)$, which is (86) with Φ_i of (84). At $k = K_i$, $p\mathbb{E}_v[y] = S_i^{\mathrm{f}}$ and the first form is (33). This proves (ii).

Step 3 (concavity). By Step 1, $k \mapsto \overline{v}_i(k)$ is a continuous strictly increasing bijection of $(0, \infty)$ onto $(\underline{v}_i, \infty)$, so by (86) π_i' is strictly decreasing on $(0, \infty)$ if and only if Φ_i is strictly decreasing on $(\underline{v}_i, \infty)$, i.e. if and only if Assumption 3 holds; a differentiable function is strictly concave if and only if its derivative is strictly decreasing. For the two ends of Φ_i: as $\overline{v} \downarrow \underline{v}_i$, $\eta_1 \to 1/\varepsilon$, so $\xi_i(\overline{v}) \to 0$ and $\zeta_i[1 - \xi_i] \to 1 - \frac{1}{\varepsilon}$ by (33), while $\mathcal{G}_\gamma^-(\overline{v}) \to \mathcal{G}_\gamma^-(\underline{v}_i) > 0$ and $\varepsilon_i(\overline{v}) - 1 \to 0^+$, so $\Phi_i \to +\infty$; as $\overline{v} \to \infty$, Lemma 5(v) gives $\zeta_i[1 - \xi_i] \to \frac{\varepsilon - 1}{\varepsilon}\alpha_i$ and Lemma 4(i) gives $\overline{v}^\gamma\mathcal{G}_\gamma^- \to \mu_{i,\gamma}$ and $\varepsilon_i \to \varepsilon$, so $\Phi_i(\overline{v}) \sim \frac{\alpha_i}{\varepsilon}\mu_{i,\gamma}\overline{v}^{-\gamma} \to 0$. Under Assumption 3 Φ_i is strictly decreasing with these ends, hence strictly positive on $(\underline{v}_i, \infty)$; since $\mathcal{G}_\gamma^- > 0$ and $\varepsilon_i - 1 > 0$ there, $\zeta_i[1 - \xi_i] > 0$, and by the rent identity (77) $\xi_i = 1 - \frac{1}{\varepsilon} - \zeta_i[1 - \xi_i] < 1 - \frac{1}{\varepsilon}$. The limits of π_i' follow from those of Φ_i through the bijection of Step 1. This proves (iii).

Step 4 (large-capital asymptote). Nothing in Steps 1–3 involves the final-demand elasticity η (not to be confused with the lower-tail shares $\eta_{i,\theta}$): the aggregates $(Y_i, P_i, \mathcal{P}_i)$ are taken as given. Step 1 gives $\overline{v}_i(k) \to \infty$ as $k \to \infty$, and by Lemma 4(i) $\Lambda(\overline{v}) = \frac{\varepsilon - 1}{\mu_{i,\gamma}}\overline{v}^{\frac{1}{\varepsilon} + \gamma - 1}[1 + o(1)]$. Writing $c(k) = c_1 k^{1/\varepsilon}$ and $\Theta_i := 1 + \frac{\varepsilon\alpha_i}{1 - \alpha_i} = \varepsilon\Big(\frac{1}{\varepsilon} + \gamma - 1\Big)$, the identity $\Lambda(\overline{v}_i(k)) = c(k)$ gives $\overline{v}_i(k) = \left[\frac{\mu_{i,\gamma}c_1}{\varepsilon - 1}\right]^{\varepsilon/\Theta_i} k^{1/\Theta_i}[1 + o(1)]$, and by the asymptotics of Φ_i in Step 3, $\pi_i'(k) = \varepsilon\gamma\frac{\mathcal{P}_i}{x_i}\Phi_i(\overline{v}_i(k)) \sim \gamma\alpha_i\mu_{i,\gamma}\frac{\mathcal{P}_i}{x_i}\overline{v}_i(k)^{-\gamma} \propto k^{-\gamma/\Theta_i}$. Finally $\gamma/\Theta_i = \frac{1}{1 + (\varepsilon - 1)\alpha_i} = 1 - \zeta_i(\infty)$, so $k^{1 - \zeta_i(\infty)}\pi_i'(k)$ converges to a positive constant. This

proves (iv). ■

Reading the profit function. The pieces of (85)–(86) are

- $\pi_i(K_i) = [1 - \xi_i(\overline{v}_i)]S_i^{\mathrm{f}}$ is *sector-level operating profit*: the firms' revenue minus their materials bill. It comprises capital income $\zeta_i(\overline{v}_i)[1 - \xi_i(\overline{v}_i)]S_i^{\mathrm{f}}$ plus pure markup rents $S_i^{\mathrm{f}}/\varepsilon$; since the sector has a unit measure of firms, it is also *average* firm profit at the symmetric point. Unlike the model without capacity constraints, it is not the value-added share of sectoral sales, because the rationing rent Π_i^{r} is part of $S_i - X_i$ but not of the firms' profit.
- An extra unit of own capital raises expected profit through the two channels of (76): in the states in which the firm is rationed it adds output and bundle use at the posted price, and in the states in which it is not it lowers the bundle needed for the demand-determined output. The second channel is the only one without capacity constraints, in which a better-capitalised firm has lower expected marginal cost, passes it into a lower price and gains demand along the sectoral CES curve; the first is new, and it is where the posted price meets the bundle cost of a unit produced at capacity.
- $\zeta_i(\overline{v})$ is the local own-capital elasticity of profit. Without capacity constraints it is the global constant $\zeta_i(\infty)$, increasing in ε (flatter demand converts a cost advantage into more revenue with less price-cutting) and in α_i (a more capital-intensive technology turns own capital into a larger cost advantage). Under capacity it moves with the firm's threshold, from $1 - \frac{1}{\varepsilon}$ when capital is scarce and the firm is rationed often (the threshold near $\underline{v}_i$, where $\xi_i \to 0$) to $\zeta_i(\infty) < 1 - \frac{1}{\varepsilon}$ when capacity is slack: the return to own capital is highest exactly where the capacity constraint bites most. The passage between the two ends need not be monotone — Assumption 3 restricts Φ_i, not ζ_i, and numerically ζ_i can dip below $\zeta_i(\infty)$ in the interior before returning to it — so it is the two ends, not a monotone descent, that the assumption and the limits deliver.
- Concavity has two sources. Without capacity constraints the curvature is purely demand-side, inherited from the CES demand curve through $\zeta_i(\infty) < 1$, since the technology is constant-returns in (k, M). Under capacity a second source appears: the marginal unit of capital is worth most when the firm is rationed often, and as capital grows the rationed states become rare. Assumption 3 combines the two into one monotonicity condition; it is the form the well-posedness of the firm problem below takes under capacity, and the same condition is used node by node in the continuum model (Lemma 16).

Proof of Proposition 1. Let $\mathbb{G}$ be the generator of $(\mathbf{K}, \mathbf{Z})$ under (28) at the symmetric policies, writing $\mu^{K_j}(\iota_j) := \Phi_j(\iota_j) - \delta_j$ for the net growth rate of sector-j capital,

$$\mathbb{G}f := \sum_{j=1}^{J} \mu^{K_j}(\iota_j) K_j \, \partial_{K_j} f + \mathcal{L}_Z f. \tag{87}$$

Proof. Differentiate (34) in k; the $\partial\iota$ terms vanish by (35), leaving

$$\rho\, w^i_k = u'(C)[\partial_k \pi_i - \iota] + \partial_k\Big(\mu^{K_i}(\iota)\, k\, w^i_k + \mathbb{G} w^i\Big),$$

with $\mathcal{L}_Z$ acting on $\mathbf{Z}$ alone and therefore commuting with every capital derivative. Restrict to the symmetric point $k = K_i$, where $\iota = \iota_i$ and $\partial_k \pi_i = \pi^k_i$, and write $\mathfrak{q}_i := w^i_k\big|_{k=K_i}$. The on-diagonal restriction obeys the chain rule $\partial_{K_j}\mathfrak{q}_i = \Big[w^i_{kK_j} + \mathbb{1}\{j=i\}\, w^i_{kk}\Big]_{k=K_i}$, so that

$$\Big[\partial_k\Big(\mu^{K_i} k\, w^i_k + \mathbb{G} w^i\Big)\Big]_{k=K_i} = \mathfrak{q}_i\, \mu^{K_i}(\iota_i) + \Big[\mu^{K_i} K_i w^i_{kk} + \sum_{j=1}^{J} \mu^{K_j} K_j w^i_{kK_j}\Big]_{k=K_i} + \mathcal{L}_Z \mathfrak{q}_i = \mathfrak{q}_i\, \mu^{K_i}(\iota_i) + \mathbb{G}\mathfrak{q}_i,$$

the bracket being $\sum_j \mu^{K_j} K_j \partial_{K_j}\mathfrak{q}_i$ by the chain rule. Inserting this gives (36). The first-order condition (37) is (35) evaluated at $k = K_i$, inverted using $\Phi_i'(\iota) = (1+\phi_i\iota)^{-(\phi_i-1)/\phi_i}$. ∎

A.12 State Reduction

Notation for drifts and diffusions. For an Itô process x we write $\mu_x(\cdot)$ and $\sigma_x(\cdot)$ for its *arithmetic* drift and diffusion and $\mu^x(\cdot) := \mu_x(\cdot)/x$, $\sigma^x(\cdot) := \sigma_x(\cdot)/x$ for the *geometric* ones,

$$\mathrm{d}x = \mu_x(\cdot)\,\mathrm{d}t + \sigma_x(\cdot)\,\mathrm{d}B, \qquad \frac{\mathrm{d}x}{x} = \mu^x(\cdot)\,\mathrm{d}t + \sigma^x(\cdot)\,\mathrm{d}B,$$

where $(\cdot)$ signals that drifts and diffusions are in general functions of the state and of policies (e.g. of $\boldsymbol{\kappa}, \mathbf{Z}, \boldsymbol{\iota}$), not constants; when x loads on several Brownian motions, $\boldsymbol{\sigma}_x \equiv (\sigma_{x,i})_i$ and $\boldsymbol{\sigma}^x \equiv (\sigma^x_i)_i$ denote the row vectors of loadings. Exogenous shock volatilities that are primitives keep their own letters (ς^{Z_i}; $\varsigma^Z(i)$ in the graphon part). The generic process x of this paragraph is unrelated to the capital intensities x_i; the two never meet in a formula. Capital has no diffusion, so its only drift objects are the geometric $\mu^{K_i}(\iota_i) = \Phi_i(\iota_i) - \delta_i$ and μ^K and the arithmetic μ_{κ_i} of (38).

Laws of motion in $(K, \boldsymbol{\kappa}, \mathbf{Z})$. The laws of motion (38) are a statement about the capital process alone, with no optimisation in them. Summing (28) at the symmetric point over i,

$$\dot{K} = \sum_{i=1}^{J} \dot{K}_i = \sum_{i=1}^{J} \mu^{K_i}(\iota_i)\, K_i,$$

and dividing by K and writing $K_i/K = \kappa_i$ gives $\dot{K}/K = \sum_i \kappa_i \mu^{K_i}(\iota_i) = \mu^K$. For the shares, $\kappa_i = K_i/K$ gives

$$\dot{\kappa}_i = \frac{\dot{K}_i}{K} - \kappa_i \frac{\dot{K}}{K} = \kappa_i\Big[\mu^{K_i}(\iota_i) - \mu^K\Big] = \mu_{\kappa_i}.$$

The multiplicative form of (28) is essential: without it these share dynamics, and the homogeneity below, fail. The drifts are tangent to the simplex: summing over i and

using $\sum_i \kappa_i = 1$, $\sum_i \mu_{\kappa_i} = \sum_i \kappa_i \mu^{K_i}(\iota_i) - \mu^K = 0$ by the definition of μ^K, so the adding-up constraint $\sum_i \kappa_i = 1$ is preserved.

Combined with the state generator (87), these laws deliver the change of variables that the deflation of the capital-pricing system needs.

Lemma 7 (Deflation of a homogeneous function). *Fix $\theta \in \mathbb{R}$ and let $f(\mathbf{K}, \mathbf{Z}) = K^\theta g(\boldsymbol{\kappa}, \mathbf{Z})$ be homogeneous of degree θ in $\mathbf{K}$. Then, with $\mathbb{G}$ of* (87) *and the drifts of* (38)*,*

$$K^{-\theta}\mathbb{G}f = \theta\mu^K g + \sum_{l=1}^{J} \mu_{\kappa_l}\, \partial_{\kappa_l} g + \mathcal{L}_Z g. \tag{88}$$

The coefficients are tangent to the simplex, $\sum_{l=1}^{J} \mu_{\kappa_l} = 0$. Consequently the right-hand side of (88) *differentiates g only along the simplex: the first-order term is the derivative of g in the direction $\boldsymbol{\mu}_{\boldsymbol{\kappa}}$, whose components sum to zero. Two functions that agree on Δ^{J-1} therefore give the same right-hand side at every point of Δ^{J-1}, however they are extended off it.*

Proof. Step 1 (the generator in $(K, \boldsymbol{\kappa})$). Write $f(\mathbf{K}, \mathbf{Z}) = \tilde{f}(K, \boldsymbol{\kappa}, \mathbf{Z})$ with $K = \sum_j K_j$ and $\kappa_l = K_l/K$. With $\partial K/\partial K_j = 1$ and $\partial \kappa_l / \partial K_j = (\mathbb{1}\{l = j\} - \kappa_l)/K$, the chain rule gives $\partial_{K_j} f = \partial_K \tilde{f} + \sum_l [\mathbb{1}\{l = j\} - \kappa_l]\partial_{\kappa_l}\tilde{f}/K$, so that in (87)

$$\sum_{j=1}^{J} \mu^{K_j} K_j\, \partial_{K_j} f = \sum_{j=1}^{J} \mu^{K_j}\kappa_j\, K\partial_K \tilde{f} + \sum_{l=1}^{J} \left(\sum_{j=1}^{J} \mu^{K_j}\kappa_j [\mathbb{1}\{l = j\} - \kappa_l] \right) \partial_{\kappa_l}\tilde{f} = K\mu^K\, \partial_K \tilde{f} + \sum_{l=1}^{J} \mu_{\kappa_l}\, \partial_{\kappa_l}\tilde{f},$$

since $\sum_j \mu^{K_j}\kappa_j[\mathbb{1}\{l = j\} - \kappa_l] = \kappa_l \mu^{K_l} - \kappa_l \mu^K = \mu_{\kappa_l}$. Hence $\mathbb{G}f = K\mu^K \partial_K \tilde{f} + \sum_l \mu_{\kappa_l}\partial_{\kappa_l}\tilde{f} + \mathcal{L}_Z \tilde{f}$.

Step 2 (substitute the ansatz). For $\tilde{f} = K^\theta g$, $K\partial_K \tilde{f} = \theta K^\theta g$ and $\partial_{\kappa_l}\tilde{f} = K^\theta \partial_{\kappa_l} g$, while $\mathcal{L}_Z \tilde{f} = K^\theta \mathcal{L}_Z g$. Substituting into Step 1 and dividing by K^θ gives (88).

Step 3 (tangency). $\sum_l \mu_{\kappa_l} = 0$ is the tangency of the laws of motion shown above. ■

Proof of Proposition 2.

Proof. (i) Throughout, w^i is the firm's value function, that is, the solution of (34) that satisfies the transversality condition $\lim_{T\to\infty} \mathbb{E}_t\left[e^{-\rho(T-t)} w^i(k_T, \mathbf{K}_T, \mathbf{Z}_T)\right] = 0$ along every admissible plan; the policies $\iota_j^\star$ are locally Lipschitz in $\mathbf{K}$; and w^i is continuously differentiable in $(k, \mathbf{K})$. Fix $\lambda > 0$ and $\mathbf{Z}$, and consider the scaling $(k, \mathbf{K}) \mapsto (\lambda k, \lambda \mathbf{K})$.

Step 1 (profit). The sector-i firm with own capital k prices against $(Y_i, P_i, \mathcal{P}_i(\mathbf{P}), Z_i)$ (Lemma 6). Along the selection of the statement, $\mathbf{K} \mapsto \lambda\mathbf{K}$ leaves P_i and $\mathcal{P}_i(\mathbf{P})$ unchanged and scales $Y_i = S_i/P_i$ by λ (Proposition 3(ii)). The firm's threshold $\overline{v}_i(k)$ is the unique root in $(\underline{v}_i, \infty)$ of (66) (Lemma 6(i)), whose left-hand side depends on $\overline{v}$ alone and whose right-hand side is proportional to $(\overline{y}_i(k)/Y_i)^{1/\varepsilon}$ with $\overline{y}_i(k) = Z_i k x_i^{\alpha_i - 1}$ (Step 1 of the proof of Lemma 6). Under $(k, Y_i) \mapsto (\lambda k, \lambda Y_i)$ this ratio is unchanged, so the

threshold and the price $p_i(k) = P_i[\overline{v}_i(k)Y_i/\overline{y}_i(k)]^{1/\varepsilon}$ are unchanged, and both terms of (85) scale by λ, through $\overline{y}_i(k)$ and k/x_i. Hence

$$\pi_i(\lambda k; \lambda \mathbf{K}, \mathbf{Z}) = \lambda\, \pi_i(k; \mathbf{K}, \mathbf{Z}).$$

Step 2 (consumption and the capital paths). $C = Q - \sum_j \iota_j^\star K_j$ with Q of degree one and $\iota_j^\star$ of degree zero, so $C(\lambda\mathbf{K}, \mathbf{Z}) = \lambda\, C(\mathbf{K}, \mathbf{Z})$ and $u'(C) = C^{-\gamma}$ is of degree $-\gamma$. Given a path of $\mathbf{Z}$, the aggregate stocks solve $\dot{K}_j = \mu^{K_j}\Big(\iota_j^\star(\mathbf{K}, \mathbf{Z})\Big)K_j$, whose right-hand side is homogeneous of degree one and locally Lipschitz in $\mathbf{K}$, so the solution is unique and the path started from $\lambda\mathbf{K}_t$ is λ times the path started from $\mathbf{K}_t$, the path of $\mathbf{Z}$ being the same. Likewise, under any investment plan $\{\iota_u\}$, the own stock started from λk_t is λk_u by (28).

Step 3 (the value). In utility units the firm's problem reads

$$w^i(k, \mathbf{K}, \mathbf{Z}) = \max_{\{\iota_u\}}\ \mathbb{E}_t \int_t^\infty e^{-\rho(u-t)}\, u'(C_u)[\pi_i(k_u; \mathbf{K}_u, \mathbf{Z}_u) - \iota_u k_u]\mathrm{d}u,$$

since $\Lambda_u/\Lambda_t = e^{-\rho(u-t)}u'(C_u)/u'(C_t)$ and w^i is the equity value times $u'(C_t)$; the transversality condition is what identifies the solution of (34) with this value. Admissibility ($1 + \phi_i\iota_u > 0$, adaptedness) is scale-free. Started from $(\lambda k, \lambda\mathbf{K})$ under the same plan, by Steps 1–2 the integrand at every u is $\lambda^{-\gamma}\cdot\lambda$ times the integrand started from $(k, \mathbf{K})$, and the set of feasible plans is the same. Hence $w^i(\lambda k, \lambda\mathbf{K}, \mathbf{Z}) = \lambda^{1-\gamma}w^i(k, \mathbf{K}, \mathbf{Z})$, and the set of maximising plans is the same; it is a singleton, since π_i is strictly concave in k (Lemma 6(iii)) and Φ_i is strictly concave, so the investment policy is of degree zero in $(k, \mathbf{K})$, consistent with the selection of the statement (the selection is closed under the firm's best response; existence is the content of (ii)–(iii)). Differentiating the identity in k, w^i_k is homogeneous of degree $-\gamma$ in $(k, \mathbf{K})$; on the diagonal $k = K_i$, taking $\lambda = 1/K$,

$$\mathfrak{q}_i(\mathbf{K}, \mathbf{Z}) = K^{-\gamma}\,\mathfrak{q}_i(\mathbf{K}/K, \mathbf{Z}) =: K^{-\gamma}\,\widehat{\mathfrak{q}}_i(\boldsymbol{\kappa}, \mathbf{Z}),$$

where $\widehat{\mathfrak{q}}_i(\boldsymbol{\kappa}, \mathbf{Z})$ is $\mathfrak{q}_i$ evaluated at the share vector $(\kappa_1, \ldots, \kappa_{J-1}, 1 - \sum_{a<J}\kappa_a)$; the J-coordinate derivatives of $\widehat{\mathfrak{q}}_i$ in (ii) are those of this function on the simplex, and by Lemma 7 the first-order term in (ii) does not depend on how $\widehat{\mathfrak{q}}_i$ is extended off it. Finally $u'(C) = C^{-\gamma} = K^{-\gamma}c^{-\gamma}$ by (39) at the equilibrium policies.

(ii) Substitute $\mathfrak{q}_i = K^{-\gamma}\widehat{\mathfrak{q}}_i(\boldsymbol{\kappa}, \mathbf{Z})$ and $u'(C) = K^{-\gamma}c^{-\gamma}$ into (36).

1. *The flow, discount and own-drift terms.* $\pi_i^k = \zeta_i(\overline{v}_i)[1 - \xi_i(\overline{v}_i)]\, S_i^{\mathrm{f}}/K_i$ by (33), and $S_i^{\mathrm{f}}/K_i = (S_i^{\mathrm{f}}/K)/\kappa_i = \widehat{s}_i^{\mathrm{f}}/\kappa_i$ with $\overline{v}_i$ and $\widehat{s}_i^{\mathrm{f}}$ functions of $(\boldsymbol{\kappa}, \mathbf{Z})$ alone, so

$$u'(C)\Big[\pi_i^k - \iota_i\Big] = K^{-\gamma}c^{-\gamma}\left[\zeta_i(\overline{v}_i)[1 - \xi_i(\overline{v}_i)]\,\frac{\widehat{s}_i^{\mathrm{f}}}{\kappa_i} - \iota_i\right],$$

   while $\rho\,\mathfrak{q}_i = K^{-\gamma}\rho\,\widehat{\mathfrak{q}}_i$ and $\mathfrak{q}_i\mu^{K_i} = K^{-\gamma}\widehat{\mathfrak{q}}_i\mu^{K_i}$.

2. *The generator term.* The last two terms of (36) are $\mathbb{G}\mathfrak{q}_i$ by (87), and $\mathfrak{q}_i = K^{-\gamma}\widehat{\mathfrak{q}}_i$ is homogeneous of degree $-\gamma$ in $\mathbf{K}$, so Lemma 7 at $\theta = -\gamma$ gives

$$K^\gamma\,\mathbb{G}\mathfrak{q}_i = -\gamma\mu^K\widehat{\mathfrak{q}}_i + \sum_{j=1}^J \mu_{\kappa_j}\,\partial_{\kappa_j}\widehat{\mathfrak{q}}_i + \mathcal{L}_Z\widehat{\mathfrak{q}}_i.$$

3. *Collect.* Every term carries the common factor $K^{-\gamma}$; cancelling it and moving the $-\gamma\mu^K\widehat{\mathrm{q}}_i$ term of Step 2 to the left-hand side produces the effective discount $\rho+\gamma\mu^K$ of (40), and the rest is (40) term by term. Each step is an identity, so conversely a solution $\widehat{\mathrm{q}}_i$ of (40) yields the homogeneous solution $K^{-\gamma}\widehat{\mathrm{q}}_i$ of (36). The FOC (37) deflates in the same way: $\Phi_i'(\iota_i)\,K^{-\gamma}\widehat{\mathrm{q}}_i = K^{-\gamma}c^{-\gamma}$ gives $c^{-\gamma} = \widehat{\mathrm{q}}_i\Phi_i'(\iota_i)$.

(iii) The first equation of (41) is goods clearing (39); the second inverts the first-order condition of (ii) using $\Phi_i'(\iota) = (1+\phi_i\iota)^{-(\phi_i-1)/\phi_i}$; the third and fourth are the drifts (38). Nothing else links the J equations, since $\overline{v}_i$, $\widehat{s}_i^{\mathrm{f}}$ and q depend on the state only. For the root: since $\iota_i \geq -1/\phi_i$, admissible c lie in $(0,\, q+\sum_i \kappa_i/\phi_i)$. On that interval each $\iota_j(c)$ is strictly increasing in c (as $\widehat{\mathrm{q}}_j > 0$ by the first-order condition of (ii), Φ_j' being positive, and $\gamma > 0$), so $F(c) := c + \sum_j \kappa_j\iota_j(c) - q$ is strictly increasing; $F(c) \to -\sum_j \kappa_j/\phi_j - q < 0$ as $c \downarrow 0$, and at $\bar{c} = q + \sum_i \kappa_i/\phi_i$, $F(\bar{c}) = \sum_j \kappa_j\big[\iota_j(\bar{c}) + 1/\phi_j\big] = \sum_j \kappa_j\big[\widehat{\mathrm{q}}_j\bar{c}^{\gamma}\big]^{\phi_j/(\phi_j-1)}/\phi_j > 0$. Hence exactly one root.

(iv) is immediate from (i)–(ii): nothing is discarded in passing to $(\boldsymbol{\kappa}, \mathbf{Z})$, the aggregate stock being a pure scaling factor. ■

B Model Details and Proofs for Section 3

This appendix collects, in the order of Section 3, the objects that the main text uses in words or in reduced form, the assumptions it only refers to, and the results it states without proof, each result followed by its proof. It also holds what the finite part has no counterpart for: Appendix B.2 states what a general kernel needs in order to be the limit of a growing sequence of networks, and what fails without it.

B.1 Graphon Operators and the Leontief Inverse

An input–output network is directed, $a(i,j) \neq a(j,i)$ in general, so a is a *digraphon* rather than a graphon in the sense of graph-limit theory (Lovász, 2012). Nothing in the paper requires symmetry: the kernel enters only through the two operators of Definition 3, which are defined for any bounded measurable kernel, and the compactness theory of graph limits is not used.The graphon mean field games of Caines and Huang (2021) assume symmetric graphons; Appendix B.9 extends their formulation to directed kernels.

The two operators of Definition 3 are bounded on different spaces, and the asymmetry is the one the economics asks for.

Lemma 8 (Operator bounds). $\mathbb{A} : L^\infty \to L^\infty$ *with* $\|\mathbb{A}\|_{\infty\to\infty} \leq 1$ *by* (42)*, and* $\mathbb{A}^* : L^1 \to L^1$ *with* $\|\mathbb{A}^*\|_{1\to 1} \leq 1$.

Proof. The bound for $\mathbb{A}$. Fix $f \in L^\infty$ and (a.e.) i:

$$|(\mathbb{A}f)(i)| = \left| \int_0^1 a(i,j)\, f(j) \mathrm{d}j \right| \overset{(1)}{\leq} \int_0^1 a(i,j)|f(j)|\mathrm{d}j \overset{(2)}{\leq} \|f\|_\infty \int_0^1 a(i,j)\mathrm{d}j \overset{(3)}{=} \|f\|_\infty.$$

(1) The triangle inequality for integrals, $\left|\int h\right| \leq \int |h|$, together with $a \geq 0$, so that $|a(i,j)f(j)| = a(i,j)|f(j)|$.

(2) Insert the pointwise bound $|f(j)| \leq \|f\|_\infty$ (valid off a null set, which does not affect the integral) and use monotonicity of the integral: a pointwise-larger non-negative integrand has a larger integral.

(3) Row-stochasticity (42).

Taking the essential supremum over i gives

$$\|\mathbb{A}f\|_\infty \equiv \operatorname{ess\,sup}_{i\in[0,1]} |(\mathbb{A}f)(i)| \leq \|f\|_\infty \implies \|\mathbb{A}\|_{\infty\to\infty} \leq 1$$

The bound is attained: $f \equiv 1$ gives $\mathbb{A}f \equiv 1$ by (42), so in fact $\|\mathbb{A}\|_{\infty\to\infty} = 1$. So, $\mathbb{A}$ is a genuine *averaging* operator; it can never inflate a uniform bound, and it preserves constants.

The bound for $\mathbb{A}^*$. Fix $g \in L^1$:

$$\begin{aligned}\|\mathbb{A}^* g\|_1 \equiv \int_0^1 |(\mathbb{A}^* g)(i)|\mathrm{d}i &= \int_0^1 \left| \int_0^1 a(j,i)\, g(j)\mathrm{d}j \right| \mathrm{d}i \\ &\overset{(1)}{\leq} \int_0^1 \int_0^1 a(j,i)|g(j)|\mathrm{d}j\,\mathrm{d}i \overset{(2)}{=} \int_0^1 |g(j)| \left(\int_0^1 a(j,i)\mathrm{d}i \right) \mathrm{d}j \overset{(3)}{=} \|g\|_1.\end{aligned}$$

(1) The same triangle inequality, applied inside the outer integral.

(2) *Tonelli's theorem*: for a jointly measurable *nonnegative* integrand (here $(i,j) \mapsto a(j,i)|g(j)| \geq 0$), the order of integration may always be exchanged, with no integrability hypothesis (both iterated integrals equal the double integral, finite or not). This is why one passes to absolute values *first*: Fubini for signed integrands would require exactly the integrability we are in the middle of establishing.

(3) Here the row-sum enters: at fixed j, the inner integral $\int_0^1 a(j,i)\mathrm{d}i$ runs over the kernel's *second* argument, so it is the row-sum of row j, equal to one by (42). (In the $\mathbb{A}$ computation the output index i sat in the first slot and the integration variable in the second; here the output index sits in the second slot, so after Tonelli it is the output variable that the row-sum integrates out.)

Since the chain gives $\|\mathbb{A}^* g\|_1 \leq \|g\|_1 < \infty$, it incidentally shows $(\mathbb{A}^* g)(i)$ is finite for a.e. i, i.e. $\mathbb{A}^*$ is well defined on all of L^1. Finally, for $g \geq 0$ every inequality in the chain holds with equality (mass is conserved, merely redistributed across nodes), so $\|\mathbb{A}^*\|_{1\to 1} = 1$ exactly. ■

A Bit of Analysis Background.

- *Function norms.*
$$\|f\|_\infty := \operatorname{ess\,sup}_{i\in[0,1]} |f(i)|$$
"Essential" because an L^∞ function is only defined up to null sets, so we ignore violations on sets of measure zero. And the L^∞ norm is
$$\|g\|_1 := \int_0^1 |g(i)|\mathrm{d}i$$

- *How the two norms compare (embedding).* On a space of total measure one such as $[0,1]$, the sup norm dominates the integral norm:
$$\|f\|_1 = \int_0^1 |f(i)|\mathrm{d}i \le \int_0^1 \|f\|_\infty \mathrm{d}i = \|f\|_\infty,$$
so $L^\infty([0,1]) \subset L^1([0,1])$: every essentially bounded function is integrable.

  The reverse does not hold. Counterexample: $f(i) = i^{-1/2}$ has $\|f\|_1 = 2$ but $\|f\|_\infty = \infty$, and even among bounded functions the gap can be arbitrarily large (a spike of height c on a set of measure $1/c^2$ has $\|f\|_\infty = c$ but $\|f\|_1 = 1/c$).

  Two easy-to-misremember points:

  - on a probability space it is the *small* exponents that give the *big* spaces, $L^\infty \subset L^p \subset L^1$ for $1 \le p \le \infty$
  - and the inclusion relies on the measure being finite: on $\mathbb{R}$ neither space contains the other (constants are in $L^\infty \setminus L^1$, integrable spikes accumulate in $L^1 \setminus L^\infty$).

- *Operator norm.* For a linear operator $T : X \to X$ the operator norm is
$$\|T\|_{X\to X} := \sup\{\|Tf\|_X : \|f\|_X \le 1\}$$
so proving $\|Tf\|_X \le \|f\|_X$ for *every* f is exactly proving $\|T\|_{X\to X} \le 1$.

- *Intuition: the maximal stretch factor.* Feed T every input of X-length at most one and record the longest output: that is $\|T\|_{X\to X}$. Equivalently, it is the smallest constant c with $\|Tf\|_X \le c\,\|f\|_X$ for all f (by linearity, scaling f scales Tf, so testing unit-ball inputs, $\|f\|_X = 1$, loses nothing).

  An operator with $\|T\| \le 1$ (a *contraction*) never lengthens any vector; that is what makes *iterating* it safe, since submultiplicativity gives $\|T^n\| \le \|T\|^n$, the engine of the Neumann series below.

- *The finite-dimensional analogy.* For a matrix $\mathbf{T} \in \mathbb{R}^{J\times J}$ the same definition, taken with respect to a vector norm, gives the *induced matrix norms*, and the two cases used here have closed forms worth memorising:
$$\|\mathbf{T}\|_{\infty\to\infty} = \max_n \sum_m |T_{nm}| \quad \text{(max } \textit{row} \text{ sum)},$$
$$\|\mathbf{T}\|_{1\to 1} = \max_m \sum_n |T_{nm}| \quad \text{(max } \textit{column} \text{ sum)}.$$

The Euclidean choice would give the spectral norm, not used here.

A row-stochastic matrix $\mathbf{A}$ has max row sum one, so $\|\mathbf{A}\|_{\infty\to\infty} = 1$: each output entry is a weighted *average* of input entries, so a componentwise bound can never be inflated. That is: if every entry of the input satisfies $|x_m| \le M$, then every output entry $(\mathbf{A}x)_n = \sum_m A_{nm}x_m$, being an average with nonnegative weights summing to one, also satisfies $|(\mathbf{A}x)_n| \le \sum_m A_{nm}M = M$; an average can never leave the range of the numbers being averaged.

Its transpose has $\|\mathbf{A}^\mathsf{T}\|_{1\to1} = \max_m \sum_n A_{mn} = 1$ (columns of $\mathbf{A}^\mathsf{T}$ are rows of $\mathbf{A}$): acting on a nonnegative vector, $\mathbf{A}^\mathsf{T}$ *redistributes* total mass without creating or destroying it.

Lemma 8 is exactly this pair of facts with integrals in place of sums:

- $\mathbb{A} \leftrightarrow \mathbf{A}$ on L^∞ (averaging)
- $\mathbb{A}^* \leftrightarrow \mathbf{A}^\mathsf{T}$ on L^1 (mass-conserving)

and the step-kernel closes the circle: on step functions $\mathbb{A}^*\Xi$ literally *is* $\mathbf{A}^\mathsf{T}\Xi$, whose ℓ^1 induced norm $\max_n \xi_n \le \bar{\xi} < 1$ is the matrix analogue of $\|\mathbb{A}^*\Xi\|_{1\to1} \le \bar{\xi}$.

This asymmetry of natural domains is economically meaningful, and its source is that a is row-stochastic but *not* column-stochastic: rows integrate to one by (42), while the column masses $a^*(i) := \int_0^1 a(j,i)\mathrm{d}j$, which is the total cost-share weight all buyers j place on supplier i, can exceed one. As such, a "hub" supplier bought intensively by many nodes has $a^*(i) > 1$. Each operator is therefore a contraction precisely on the space whose norm the row-sums control: $\mathbb{A}^*$ can inflate a sup-bound by the factor $a^*(i)$, and $\mathbb{A}$ can inflate total mass by the same column weights. The economics matches the functional analysis. Log-prices are *intensive*: the meaningful statement about them is a uniform bound (no node's price explodes), and $\mathbb{A}$, a cost-share-weighted average of suppliers' log-prices, stays within any such bound. Sales are *extensive*: they add across buyers, so the meaningful statement is a finite *total*, and $\mathbb{A}^*$, routing every buyer's spending back to its suppliers, conserves exactly that total ($\int \mathbb{A}^* g = \int g$ for $g \ge 0$); no useful uniform bound on sales exists, since a hub legitimately carries a sales density far above average. Hence log-prices live in L^∞ where $\mathbb{A}$ acts, revenues in L^1 where $\mathbb{A}^*$ acts, and Lemma 8 is exactly what makes the graphon Leontief inverse below, which iterates $\mathbb{A}^*\Xi$ on revenues in L^1, converge.

Lemma 9 (Graphon Leontief inverse and the Domar density). *Let $\xi : [0,1] \to [0,\bar{\xi}]$, $\bar{\xi} < 1$, be measurable (it will be the intermediate expenditure share) and write Ξ for pointwise multiplication by ξ. Then, on $L^1([0,1])$:*

(i) (Contraction.) $\|\Xi\|_{1\to1} \le \bar{\xi}$ *and hence* $\|\mathbb{A}^*\Xi\|_{1\to1} \le \bar{\xi} < 1$.

(ii) (Neumann inverse.) $\mathrm{Id} - \mathbb{A}^*\Xi$ *is boundedly invertible, with*

$$(\mathrm{Id} - \mathbb{A}^*\Xi)^{-1} = \sum_{n=0}^{\infty} (\mathbb{A}^*\Xi)^n \tag{89}$$

convergent in operator norm, and $\left\|(\mathrm{Id} - \mathbb{A}^*\Xi)^{-1}\right\|_{1\to1} \le (1-\bar{\xi})^{-1}$.

(iii) (The Domar density.) *For a final-demand density* $\beta \in L^1_+ := \{f \in L^1([0,1]) : f \geq 0 \text{ a.e.}\}$ *(the cone of nonnegative integrable functions) with* $\|\beta\|_1 = 1$*, the* Domar density $h := (\mathrm{Id} - \mathbb{A}^*\Xi)^{-1}\beta$*, the continuum analogue of* $\mathbf{h} = [\mathbf{I} - \mathbf{A}^\mathsf{T}\boldsymbol{\Xi}]^{-1}\boldsymbol{\beta}$*, satisfies*

$$h \geq \beta \geq 0 \quad \text{a.e.}, \qquad 1 \leq \|h\|_1 \leq \frac{1}{1-\bar{\xi}}.$$

(iv) (Lipschitz dependence.) $\|h - h'\|_1 \leq (1-\bar{\xi})^{-1}\|\beta - \beta'\|_1$ *for any two final-demand densities.*

Proof. Step 1 (the multiplication operator). $(\Xi f)(i) = \xi(i)f(i)$, so with $0 \leq \xi \leq \bar{\xi}$ pointwise,

$$\|\Xi f\|_1 = \int_0^1 \xi(i)|f(i)|\mathrm{d}i \leq \bar{\xi}\int_0^1 |f(i)|\mathrm{d}i = \bar{\xi}\|f\|_1,$$

which is $\|\Xi\|_{1\to 1} \leq \bar{\xi}$. Combining with $\|\mathbb{A}^*\|_{1\to 1} \leq 1$ (Lemma 8) and submultiplicativity of operator norms, $\|\mathbb{A}^*\Xi\|_{1\to 1} \leq \|\mathbb{A}^*\|_{1\to 1}\|\Xi\|_{1\to 1} \leq \bar{\xi} < 1$. This is (i), and it is the *only* place the hypothesis $\bar{\xi} < 1$ is used.

Step 2 (invertibility). $L^1([0,1])$ is a Banach space and $T := \mathbb{A}^*\Xi$ is a bounded linear operator on it with $\|T\| < 1$ by Step 1. The Neumann series theorem (the operator form of the geometric series; Kreyszig (1978), Theorem 7.3-1) then gives that $\mathrm{Id} - T$ is boundedly invertible with $(\mathrm{Id} - T)^{-1} = \sum_{n\geq 0} T^n$, convergent in operator norm. This is (89).

Step 3 (the norm bound on the inverse). The bound is a consequence of the series, not a separate ingredient: taking norms termwise and using submultiplicativity again, $\|T^n\| \leq \|T\|^n \leq \bar{\xi}^n$, so

$$\left\|(\mathrm{Id} - T)^{-1}\right\|_{1\to 1} \leq \sum_{n\geq 0}\|T^n\|_{1\to 1} \leq \sum_{n\geq 0}\bar{\xi}^n = \frac{1}{1-\bar{\xi}},$$

completing (ii). Applying it to β gives the upper bound $\|h\|_1 \leq \|\beta\|_1/(1-\bar{\xi}) = (1-\bar{\xi})^{-1}$ of (iii).

Step 4 (positivity and the lower bound). T maps L^1_+ into itself, because $a \geq 0$ and $\xi \geq 0$ make $(Tg)(i) = \int_0^1 a(j,i)\xi(j)g(j)\mathrm{d}j \geq 0$ whenever $g \geq 0$. Hence every term of $h = \sum_{n\geq 0} T^n\beta$ is nonnegative, and keeping only the $n = 0$ term gives $h \geq \beta \geq 0$ a.e.; integrating, $\|h\|_1 \geq \|\beta\|_1 = 1$. This completes (iii). Economically the two bounds bracket the continuum analogue of $\sum_i h_i > 1$: gross output exceeds GDP because intermediates are counted at every link, but the double-counting is finite precisely because each link retains only the fraction $\xi \leq \bar{\xi} < 1$.

Step 5 (Lipschitz dependence). $\beta \mapsto h$ is the linear map $(\mathrm{Id} - T)^{-1}$, so $h - h' = (\mathrm{Id} - T)^{-1}(\beta - \beta')$ and (iv) is Step 3 applied to $\beta - \beta'$. ■

[Reading Lemma 9: the operators, and why a quantitative inverse]

(i) *The two operators in the statement.* Ξ is pointwise multiplication, $(\Xi f)(i) := \xi(i) f(i)$: it keeps node i's intermediate share of whatever reaches i. Id is the identity operator, $(\mathrm{Id} f)(i) = f(i)$, the operator analogue of $\mathbf{I}$ in the finite formula.

(ii) *The composite operator spelled out.* One round of re-spending is

$$(\mathbb{A}^* \Xi\, g)(i) = \int_0^1 a(j,i)\,(\Xi g)(j)\mathrm{d}j = \int_0^1 a(j,i)\,\xi(j)\,g(j)\mathrm{d}j :$$

every buyer j spends the share $\xi(j)$ of what reaches it on intermediates and routes it to its suppliers with weights $a(j,\cdot)$. The Neumann series (89) is then literally the sum over supply chains of every length, exactly as $\sum_n (\mathbf{A}'\Xi)^n$ is in the finite model.

Under capacity the intermediate expenditure share is not a primitive but a function of the node's state: node i spends $X_i = \xi_i(\overline{v}_i) S_i^{\mathrm{f}}$, the share $\xi_i(\overline{v}_i) S_i^{\mathrm{f}}/S_i$ of its sales, with the threshold $\overline{v}_i$ determined in equilibrium (Lemma 5). The operator $\mathbb{A}^*\Xi$ with a fixed profile ξ therefore characterises the model without capacity constraints (Appendix C.4), and the capacity model's inbound demand is instead the kernel integral of the nodes' spending itself (Section 3.2). The lemma's quantitative content survives nonetheless, because the effective share is uniformly bounded: under Assumption 3, $\xi_i(\overline{v}) < 1 - \frac{1}{\varepsilon}$ and $S_i^{\mathrm{f}} \le S_i$ give $X_i/S_i < 1 - \frac{1}{\varepsilon}$ at every node and in every state (Remark 4), so $\bar{\xi} := 1 - \frac{1}{\varepsilon}$ is an admissible bound for every state-dependent share profile the capacity model can generate, and (i)–(iv) apply to the Leontief inverse built from any such profile. Remark 7 uses this bound to control gross output by GDP.

B.2 Step Kernels and Graphon Limits

Beyond the reduction of the operators to the matrices $\mathbf{A}$ and $\mathbf{A}^{\mathsf{T}}$ (Section 3.1), the step-kernel embedding has two further properties.

1. The graphon Leontief inverse collapses *exactly* to the finite one: the Domar density is the step function $h(i) = J\, h_m^{\mathrm{fin}}$ on I_m with $\mathbf{h}^{\mathrm{fin}} = [\mathbf{I} - \mathbf{A}^{\mathsf{T}}\Xi]^{-1}\boldsymbol{\beta}$ (Lemma 10 below).
2. The row entropy $H^J(i)$ of a^J, defined in Lemma 11(i) below, is the finite model's sourcing entropy $\mathcal{H}_n := -\sum_m a_{nm} \ln a_{nm}$, the constant term of the log of the bundle price index (13), shifted by $-\ln J$,

$$\begin{aligned} H^J(i) &:= -\int_0^1 a^J(i,j) \ln a^J(i,j)\mathrm{d}j \\ &= -\sum_{m=1}^{J} a_{nm} \ln (J\, a_{nm}) = \mathcal{H}_n - \ln J \in [-\ln J, 0], \qquad i \in I_n, \end{aligned} \tag{90}$$

the level shift that appears whenever continuum and finite-model levels are compared (Section 3.5).

(The level shift, and why it makes levels incomparable across different J, is discussed where step kernels are discretised: Appendix D, Step 0.)

When the kernel is instead meant as the *limit* of a growing sequence of finite networks, the natural notion of convergence $a^J \to a$ is *section-wise*,

$$\sup_i \left\| a^J(i,\cdot) - a(i,\cdot) \right\|_{L^1} \to 0 \qquad \text{and} \qquad \sup_j \left\| a^J(\cdot,j) - a(\cdot,j) \right\|_{L^1} \to 0,$$

which block-constant primitives satisfy trivially.

Lemma 10 (Step kernels: exact reduction to the finite-sector Leontief inverse). *Take the step-kernel embedding of Section 3.1, $a^J(i,j) = J\, a_{nm}$ and $\beta^J(i) = J\, \beta_n$ for $(i,j) \in I_n \times I_m$, together with block-constant shares $\xi(j) = \xi_m$ for $j \in I_m$, and write $\Xi := \operatorname{diag}(\xi_1, \dots, \xi_J)$. Then the graphon Domar density is exactly the finite-sector Domar vector in density units: the step function*

$$h(i) = J\, h_m^{\text{fin}} \quad \text{on } I_m, \qquad \mathbf{h}^{\text{fin}} = \left[\mathbf{I} - \mathbf{A}^{\mathsf{T}}\Xi\right]^{-1}\boldsymbol{\beta},$$

with the same J-scaling as the embedding of a and β themselves.

Proof. Write $G_n[g] := \int_{I_n} g(j)\mathrm{d}j$ for the block masses of $g \in L^1$.

(1) *One application of $\mathbb{A}^*\Xi$ lands in step functions.* For $i \in I_m$,

$$(\mathbb{A}^*\Xi g)(i) = \int_0^1 a^J(j,i)\,\xi(j)\,g(j)\mathrm{d}j = \sum_{n=1}^{J} \int_{I_n} J\, a_{nm}\,\xi_n\, g(j)\mathrm{d}j = J \sum_{n=1}^{J} a_{nm}\,\xi_n\, G_n[g],$$

constant on each I_m: whatever g is, only its block masses survive one application.

(2) *On block masses the operator is the matrix $\mathbf{A}^{\mathsf{T}}\Xi$.* The image's block mass over I_m is $1/J$ times its constant value there:

$$G_m[\mathbb{A}^*\Xi g] = \sum_{n=1}^{J} a_{nm}\,\xi_n\, G_n[g] = \left(\mathbf{A}^{\mathsf{T}}\Xi\, G[g]\right)_m.$$

(3) *Induction.* β^J has block masses β_n, so $(\mathbb{A}^*\Xi)^k\beta^J$ is the step density with block-mass vector $(\mathbf{A}^{\mathsf{T}}\Xi)^k\boldsymbol{\beta}$, i.e. pointwise value $J\left[(\mathbf{A}^{\mathsf{T}}\Xi)^k\boldsymbol{\beta}\right]_m$ on I_m.

(4) *Sum the series.* Summing termwise (justified by absolute convergence, $\left\|(\mathbb{A}^*\Xi)^k\right\|_{1\to 1} \le \bar{\xi}^k$) gives, for $i \in I_m$,

$$h(i) = \sum_{k=0}^{\infty} J\left[\left(\mathbf{A}^{\mathsf{T}}\Xi\right)^k\boldsymbol{\beta}\right]_m = J\left[\left(\mathbf{I} - \mathbf{A}^{\mathsf{T}}\Xi\right)^{-1}\boldsymbol{\beta}\right]_m = J\, h_m^{\text{fin}},$$

the finite-sector geometric series converging because the ℓ^1 matrix norm is $\left\|\mathbf{A}^{\mathsf{T}}\Xi\right\|_1 = \max_n \xi_n \sum_m a_{nm} = \max_n \xi_n \le \bar{\xi} < 1$, the exact matrix analogue of $\|\mathbb{A}^*\Xi\|_{1\to 1} \le \bar{\xi}$.

■

Graphon limits of production networks: what is needed, what breaks. The rest of this subsection answers a question the main text defers (Section 3.1): which finite input–output networks does the continuum formulation handle, and what goes wrong for those it does not. Two readers need different answers. A reader who uses the step-kernel GMFG to solve an observed J-sector economy (every numerical result of the paper) needs almost nothing: the embedding of Section 3.1 is exact and every kernel condition of Section 3 is automatic. A reader who wants to view a general kernel as the *limit* of a growing sequence of finite networks needs a genuine assumption stack, and each rung of that stack has a named failure mode. Appendix B.2.1 collects the convergence notions needed to state the stack, Appendix B.2.2 states the stack, and Appendix B.2.3 tabulates which layer of the model breaks under which network pathology; the probabilistic stage on which the exact law of large numbers of Lemma 13 is stated is set out beside that lemma, in Appendix B.5.1.

B.2.1 Convergence of Networks to Kernels

The cut norm. Graph-limit theory measures the distance between two kernels by the cut norm

$$\|W\|_{\square} := \sup_{M,T\subseteq[0,1]} \left| \int_{M\times T} W(i,j)\mathrm{d}i\mathrm{d}j \right|, \tag{91}$$

the supremum running over pairs of measurable sets. Applied to $W = a - a'$ it compares aggregate linkages. Since $\int_{M\times T} a$ is the total cost-share weight that the buyer bloc M places on the supplier bloc T, $\|a - a'\|_{\square} \leq \epsilon$ states that however industries are grouped into a buyer bloc and a supplier bloc, the two networks' bloc-to-bloc linkages differ by at most ϵ; the bloc pair is chosen *after* inspecting both networks, so cut-closeness is a worst-case guarantee over every two-bloc aggregation. Equivalently, the cut norm controls all bilinear aggregates $\int\int f(i)W(i,j)g(j)\mathrm{d}i\mathrm{d}j$ with bounded f, g. It is the canonical norm of graph-limit theory (Lovász, 2012): the space of kernels is compact in the associated cut metric, so every sequence of dense finite networks has a subsequence converging to a kernel. Two cut-close input–output tables are therefore indistinguishable in any coarse aggregate. The guarantee is additive, however, and so informative only where the objects compared are themselves of order one: at a coarse aggregation a table entry is $O(1)$ and ϵ constrains it, whereas at fine disaggregation the entries fall below ϵ and the bound is vacuous.

Why cut convergence is too weak for a mean-field limit. Node i's equilibrium objects depend on the network through its own *sections*: its marginal cost through its row $a(i,\cdot)$, via $\mathbb{A}\ln P$ and the row entropy in (95), and its inbound demand through its column $a(\cdot,i)$, via $\mathbb{A}^*\Xi m_S$. The cut norm constrains only bloc aggregates and places no useful bound on any individual row or column: two networks can be cut-close while *every* row is far in L^1 (Example 1), in which case node-level objects such as $\mathcal{P}_i$ do not converge although all bloc aggregates do. A mean-field limit is a statement about the environment a typical individual node faces, and that environment is precisely the

node's own row and column. This is why Caines and Huang (2021) impose a section-wise convergence hypothesis, their (H11), rather than cut convergence(Appendix B.9).

Sectional convergence. The natural notion is convergence of every node's own interaction profile, uniformly over nodes,

$$\sup_i \left\| a^J(i,\cdot) - a(i,\cdot) \right\|_{L^1} \to 0 \qquad \text{and} \qquad \sup_j \left\| a^J(\cdot,j) - a(\cdot,j) \right\|_{L^1} \to 0, \tag{92}$$

rows for costs and columns for demand. Sectional convergence implies cut convergence: for $W := a^J - a$ and any measurable M, T, Fubini (both kernels are bounded) and the triangle inequality give

$$\left| \int_{M\times T} W \right| \le \int_M \int_T |W(i,j)| \mathrm{d}j\mathrm{d}i \le \int_M \left\| a^J(i,\cdot) - a(i,\cdot) \right\|_{L^1} \mathrm{d}i \le \sup_i \left\| a^J(i,\cdot) - a(i,\cdot) \right\|_{L^1},$$

using $\lambda(M) \le 1$ in the last step, and the supremum over (M,T) bounds the cut norm by the sectional distance. The converse fails.

Example 1 (A quasi-random network is cut-close to the uniform kernel but sectionally far) Let $a \equiv 1$ be the uniform kernel, in which every buyer spreads its intermediate spending evenly over all suppliers. Build a quasi-random network: for each buyer sector $n \in \{1,\dots,J\}$ and supplier sector m draw an i.i.d. link indicator $\chi_{nm} \in \{0,1\}$ with $\mathbb{P}[\chi_{nm} = 1] = \frac{1}{2}$, let $d_n := \sum_m \chi_{nm} \approx J/2$ be row n's link count, row-normalise to $a'_{nm} = \chi_{nm}/d_n$, and embed as in Section 3.1: $a'(i,j) = (J/d_n)\chi_{nm}$ on $I_n \times I_m$, a checkerboard of blocks at height 0 or $J/d_n \approx 2$. Then:

(i) *Every row stays L^1-far from uniform.* Row n differs from 1 by $J/d_n - 1$ on the linked suppliers (measure d_n/J) and by 1 on the unlinked ones (measure $1 - d_n/J$), so $\|a'(i,\cdot) - 1\|_{L^1} = 2(1 - d_n/J) \to 1$ as $J \to \infty$, by the law of large numbers for the link indicators.

(ii) *Hence node-level objects do not converge.* Row n is a density of height J/d_n on a set of measure d_n/J, so its entropy term in (95) is $-\ln(J/d_n) \to -\ln 2$, against 0 for the uniform kernel: at fixed supplier prices, every node's marginal cost $\mathcal{P}_i$ remains a factor $\frac{1}{2}$ away from its uniform-network value, for every J.

(iii) *Yet the cut norm vanishes,* almost surely, at rate $J^{-1/2}\sqrt{\ln J}$.

The mechanism is that a bloc aggregate sums $\sim J^2$ independent centred signs, so it fluctuates like J against a normalisation of J^2 and the noise cancels, whereas a single row sums nothing and its L^1 distance is an average of $O(1)$ deviations that never cancel.

[Proof of Example 1(iii): the cut norm of the quasi-random deviation.] Write $W := a' - 1$, constant at $W_{nm} = J\chi_{nm}/d_n - 1$ on each block.

- *The worst blocs are unions of whole sectors.* For any measurable (M,T), $\int_{M\times T} W = \sum_{n,m} u_n v_m W_{nm}$ with $u_n := \lambda(M \cap I_n)$ and $v_m := \lambda(T \cap I_m)$ in $[0, 1/J]$. The sum is

linear in each u_n and each v_m separately, so the supremum is attained at a vertex of the cube and

$$\|W\|_{\square} = \max_{S,T\subseteq\{1,\dots,J\}} \left| J^{-2} \sum_{n\in S} \sum_{m\in T} W_{nm} \right|.$$

- *The degree-fluctuation piece.* Split $W_{nm} = \varepsilon_{nm} + \chi_{nm}(J/d_n - 2)$ with $\varepsilon_{nm} := 2\chi_{nm} - 1$ an i.i.d. array of fair ± 1 signs. Hoeffding's inequality for each row degree and a union bound over the J rows give $\max_n |J/d_n - 2| = O(\sqrt{\ln J/J})$ almost surely, so this piece contributes $O(\sqrt{\ln J/J})$ to the cut norm.
- *The sign piece.* For fixed (S,T) the sign sum has at most J^2 independent terms, so $\mathbb{P}[|\sum_{n\in S,m\in T} \varepsilon_{nm}| > tJ^{3/2}] \le 2\exp(-t^2 J/2)$. A union bound over the 4^J bloc pairs and Borel–Cantelli give, almost surely for all large J, $\|W\|_{\square} \le tJ^{-1/2} + O(\sqrt{\ln J/J}) \to 0$ for any $t > \sqrt{2\ln 4}$.

For block-constant data converging to a step kernel the two notions coincide trivially, and for the exact embedding of Section 3.1 both distances are zero: representing a finite network by its step kernel introduces no approximation at all.

B.2.2 What the Primitive Network Must Satisfy

Two distinct questions have to be kept apart.

For the step-kernel GMFG: nothing beyond finiteness. Let $\mathbf{A}$ be any finite, row-stochastic input–output table with $\varepsilon > 1$, $\alpha_n \in (0,1)$ and $\beta_n > 0$. Its step-kernel embedding satisfies every condition of Definition 2 and Assumption 4 automatically: boundedness, with $\bar{a} = J \max_{n,m} a_{nm}$; finite bundle entropy, with $H^J(i) = \mathcal{H}_n - \ln J$ by (90); atomlessness, built into the density specification (a diagonal entry $a_{nn} > 0$ is *not* an atom problem: it embeds as a block of positive measure, and only the continuum's within-node own-good consumption would need the self-loop extension of Appendix B.3); and bounded column masses, $a^*(i) \le J \max_m \sum_n a_{nm}$. The embedding is exact (Lemma 10; both distances of Appendix B.2.1 are zero), so no limit of network sequences is ever taken. The substantive requirements are economic, not graph-theoretic: $\varepsilon > 1$ and $\alpha_n \in (0,1)$, so that $\bar{\xi} < 1$ and the Leontief inverse exists, and $\beta_n > 0$(Remark 9). The step-kernel user therefore needs no graphon limit theory at all, which is what licenses the use of the continuum machinery on an observed input–output table without proving limit theorems.

For a general kernel as the limit of a growing sequence $\mathbf{A}^{(J)}$, $J \to \infty$: a ladder of assumptions. Each rung names a condition on the sequence, the layer of the model it protects, and what breaks without it.

1. *Density (bounded weights along the sequence).* If the number of suppliers per buyer stays bounded as J grows, the embedded densities $a^J = Ja_{nm}$ have essential supremum at least J/d for a buyer with d suppliers, and row entropies $\mathcal{H}_n - \ln J \to$

$-\infty$: no bounded kernel limit exists, rungs 3–4 below fail along the sequence, and in graph-limit terms the unweighted adjacency pattern converges to the zero graphon in the cut metric (Lovász, 2012), so no mean-field structure survives. Input–output networks at fine disaggregation are sparse and heavy-tailed, so this rung, and not any of the others, is the binding scope limit of the continuum construction. *Failure layer: the limit object itself.*

2. *Sectional convergence*, (92): the analogue of Caines and Huang's (2021) (H11). Cut convergence alone is not enough: Example 1 is cut-close to the uniform kernel while every row stays L^1-far, so node-level objects ($\mathcal{P}_i$, hence prices, hence everything) do not converge even though all bloc aggregates do. *Failure layer: the limit kernel exists but describes no individual node's environment; every transfer of limit policies back to the finite economy (ε-Nash-type results) fails.*

3. *Uniformly bounded weights*, $a^{(J)} \leq \bar{a}$: no buyer's sourcing density concentrates without limit, equivalently no supplier's column mass $a^*(i) := \int_0^1 a(j,i)\mathrm{d}j$ diverges. The probabilistic layer survives without it (Lemmas 13 and 14 need only row-stochasticity and Tonelli, since $\int_0^1 a^*(i)\mathrm{d}i = 1$ makes a^* finite a.e.), and so do the operator bounds and the Neumann series; what breaks is every *uniform-over-nodes* bound on column-aggregated objects. From the Neumann series, $h \geq \beta + \mathbb{A}^*\Xi\beta$ pointwise and $(\mathbb{A}^*\Xi\beta)(i) \geq \underline{\xi}\,\underline{\beta}\,a^*(i)$ with $\underline{\xi} = \frac{\varepsilon-1}{\varepsilon}(1-\bar{\alpha})$, so $a^*(i) \uparrow \infty$ forces $h(i) \uparrow \infty$: the Domar density stays in L^1 ($\|h\|_1 \leq (1-\bar{\xi})^{-1}$) but $\ln h \notin L^\infty$. An explicit finite-entropy kernel with this property is $a(i,j) = (1-s) + \frac{s}{2}j^{-1/2}$, $s \in (0,1)$: rows are identical, row-stochastic and of finite entropy, but $a^*(i) = (1-s) + \frac{s}{2}i^{-1/2} \to \infty$ as $i \downarrow 0$. Consequently the Cobb–Douglas Fredholm layerof Appendix B.7 breaks as written ($\ln h$ sits on the right-hand side of the Fredholm equation(191), and the L^∞-Neumann solve needs a bounded right-hand side), the deterministic inbound demand $\mathcal{D} = \mathbb{A}^*\Xi m_S$ is finite a.e. but no longer uniformly bounded, and the uniform moment bounds behind the existence routeof Section 3.4 are lost. A supplier whose Domar weight does not vanish along the sequence is a *granular* agent, which the limit is by construction designed to remove; such a node is properly treated in the finite model or by a major-agent extension. *Failure layer: this model's uniform-bound layer, with the boundedness hypotheses of the Caines–Huang theorems failing in parallel.*

4. *Uniformly finite bundle entropy*, $\int_0^1 a(i,j)|\ln a(i,j)|\mathrm{d}j \leq \bar{e}$ for a.e. i: no near-atomic sourcing rows along the sequence. It is implied by rung 3 (Assumption 4(4)) and becomes an independent condition only if rung 3 is relaxed. Its role is the bundle price index: whenever log prices are bounded, $|\ln \mathcal{P}_i| \leq |(\mathbb{A}\ln P)(i)| + \left|\int_0^1 a(i,j)\ln a(i,j)\mathrm{d}j\right| \leq \|\ln P\|_\infty + \bar{e}$ by Lemma 8 and the bound; without it $\mathcal{P}_i$ can be 0 or ∞ on a positive-measure set of nodes, and the per-node static block (Lemma 15 takes $\mathcal{P}(t,i) \in (0,\infty)$ as input) is ill posed there. Atomlessness does not imply it: the atomless row density $a(i,j) \propto 1/(j\ln^2 j)$ near $j=0$ is integrable, yet $a|\ln a| \sim 1/(j|\ln j|)$ is not, so $H(i) = -\infty$. A "near-atom" is thus compatible with diffuseness but not with finite entropy, which is the operative quantitative

anti-concentration condition. *Failure layer: the static block.*

5. *No atoms in the bundle measure.* Automatic for any kernel specification, since $a(i,j)\mathrm{d}j$ is absolutely continuous; only the deliberate self-loop extension of Appendix B.3 introduces an atom, at the cost of one extra scalar and monotone fixed point per node (the node's own price entering its own marginal cost). *Failure layer: none; a modelling choice.*

[Scope of the ladder, and repair strategies if a referee pushes.]

- **Rung 1 is the honest scope limit, and it should not be buried.** Real input–output tables at fine disaggregation are sparse-ish and heavy-tailed, so the dense-graphon construction is a genuine idealisation, not a technical convenience. This is the same disclosure as the firm-level caveat: sparse-graph mean-field theory is open, and nothing here pretends otherwise. Because the paper only ever *computes* the step kernel, where no limit is taken, the disclosure costs nothing operationally.
- **Rung 3, if one wanted to drop it.** The repair would be weighted-L^∞ spaces keyed to $a^*(i)$, carrying the weight through the Neumann series, the Fredholm solve and the moment bounds. That is a different paper, and probably not worth doing unless a referee demands hubs.
- **On the granularity restriction:** it is mild and, I think, exactly right — a dominant hub is a granular agent, and removing granularity is the whole point of the limit. If pushed, the answer is the major-agent GMFG literature, not a weakening of rung 3.

B.2.3 A Failure Taxonomy

Table 7 summarises the ladder by network pathology. Each cell states whether the layer survives (yes), fails (no), or is not established either way (open); the bases are the computations of Appendix B.2.2 and the cited literature.

B.3 Assumptions on Primitives

Assumption 4 (Primitives). All primitives are measurable functions of the node index $i \in [0,1]$:

1. Capital shares $\alpha(i) \in [\underline{\alpha}, \bar{\alpha}] \subset (0,1)$. As in the finite model, the two combinations $\frac{1}{1-\alpha(i)}$ (the returns-to-scale factor of the input requirement) and $\frac{1-\alpha(i)}{\alpha(i)}$ (the supply elasticity of the model without capacity constraints) recur throughout.
2. Within-sector CES elasticity $\varepsilon > 1$ (common across nodes for notational ease; an i-dependent $\varepsilon(i)$ changes nothing), and demand-shock distributions G_i satisfying Assumptions 1 and 2 node by node, with moments $\mu_{i,\frac{1}{1-\alpha(i)}} \in [1, \bar{\mu}]$, so that the tail functions of Section 2.1 are finite at every node; under the log-normal specification (60) this is a dispersion profile $\sigma_v(i) \in [\underline{\sigma}_v, \bar{\sigma}_v] \subset (0,\infty)$. Without capacity

constraints, the markup constants $\Gamma(i)$ and expenditure shares $\xi(i) \le \bar{\xi} < 1$ are built from these (Appendix C.4).

3. Capacity technology: capital intensities $x(i) \in [\underline{x}, \bar{x}] \subset (0, \infty)$, one unit of capital at node i processing at most $1/x(i)$ units of the bundle, so that a firm with capital k has capacity $\overline{y} = Z_i k x(i)^{\alpha(i)-1}$ as in (10).
4. Input–output kernel a as in Definition 2; no further restriction on a is imposed. Two consequences of that definition are used repeatedly below, in exactly the form in which Lemma 11 states them: the bundle entropy is finite, and the bundle-weight measure is diffuse.
5. Final-demand density $\beta : [0,1] \to [\underline{\beta}, \bar{\beta}]$ with $\underline{\beta} > 0$ and $\int_0^1 \beta(i)\mathrm{d}i = 1$. Strict positivity is *substantive*, not technical: see Remark 9.
6. Adjustment-cost curvatures $\phi(i) > 1$, depreciation $\delta(i) \in [\underline{\delta}, \bar{\delta}]$, with $\Phi_i(\iota) = (1 + \phi(i)\iota)^{1/\phi(i)} - 1$ as before.
7. Shocks. Log-productivity $z_i(t) := \ln Z_i(t)$ follows a node-wise OU process
$$\mathrm{d}z_i = -\theta(i)\, z_i\, \mathrm{d}t + \varsigma^Z(i)\, \mathrm{d}B_i^Z(t), \qquad \theta(i) \ge \underline{\theta} > 0, \;\; \varsigma^Z(i) \in [0, \bar{\varsigma}^Z], \tag{93}$$
and node capital carries no shock: it accumulates deterministically, as in (28). The family $\{B_i^Z\}_{i\in[0,1]}$ is essentially pairwise independent across nodes, in the sense made precise in Lemma 13. The generator of the productivity coordinate at node i is $\mathcal{L}_z^{(i)} f = -\theta(i) z f_z + \frac{1}{2}\varsigma^Z(i)^2 f_{zz}$. The shock may be switched off: $\varsigma^Z(i) = 0$ makes $z_i(t) = z_i(0)e^{-\theta(i)t}$ deterministic, decaying to zero.
8. Household preferences: CRRA with curvature $\gamma > 0$ and discount rate $\rho > 0$.

Lemma 11 (Two consequences of the kernel definition). *Let a be an input–output kernel in the sense of Definition 2, i.e. bounded measurable with $0 \le a \le \bar{a} < \infty$ and row-stochastic* (42)*, and adopt the convention $0 \ln 0 = 0$. Then, for a.e. $i \in [0,1]$:*

(i) (Finite bundle entropy.) *with the constant $\bar{e} := e^{-1} + \bar{a} \ln \bar{a} < \infty$,*[26]
$$\int_0^1 a(i,j)|\ln a(i,j)|\mathrm{d}j \le \bar{e}, \tag{94}$$
so the row entropy $H(i) := -\int_0^1 a(i,j) \ln a(i,j)\mathrm{d}j$ is finite, and in fact $H(i) \in [-\bar{e}, 0]$, with $H(i) = 0$ if and only if $a(i,\cdot) = 1$ a.e.

(ii) (Diffuse bundle weights.) *The bundle-weight measure $\nu_i(\mathrm{d}j) := a(i,j)\mathrm{d}j$ is a Borel probability measure on $[0,1]$ that is absolutely continuous with respect to Lebesgue measure; hence it is* diffuse *(atomless), $\nu_i(\{j\}) = 0$ for every $j \in [0,1]$, in particular on the diagonal, $\nu_i(\{i\}) = 0$.*

Proof. (i), Step 1 (split the integral at $a = 1$). Write $E_{\le} := \{j : a(i,j) \le 1\}$ and $E_{>} := \{j : a(i,j) > 1\}$. On $E_{\le}$, $|\ln a| = -\ln a$, so $a|\ln a| = -a \ln a$; on $E_{>}$, $|\ln a| = \ln a$. Hence
$$\int_0^1 a|\ln a|\mathrm{d}j = \int_{E_{\le}} (-a \ln a)\mathrm{d}j + \int_{E_{>}} a \ln a\, \mathrm{d}j.$$

[26] e is Euler's number.

Step 2 (bound each piece by a pointwise maximum). The map $x \mapsto -x \ln x$ on $(0,1]$ has derivative $-(\ln x + 1)$, vanishing at $x = e^{-1}$, so its maximum there is e^{-1}; with the convention $0 \ln 0 = 0$ the bound extends to $x = 0$. The map $x \mapsto x \ln x$ is increasing on $[1,\infty)$, so on $E_>$ it is at most $\bar{a} \ln \bar{a}$, which is nonnegative because $\bar{a} \geq 1$: a nonnegative function integrating to one over a set of measure one cannot have essential supremum below one. Since $E_\leq$ and $E_>$ both have measure at most one,

$$\int_0^1 a|\ln a| \mathrm{d}j \leq e^{-1} + \bar{a} \ln \bar{a} = \bar{e},$$

which is (94), and $|H(i)| \leq \int_0^1 a|\ln a| \mathrm{d}j \leq \bar{e}$.

Step 3 (the sign of H). $\chi(x) := x \ln x$ is strictly convex on $(0,\infty)$ with $\chi(0) = 0$. Lebesgue measure on $[0,1]$ is a probability measure, so Jensen's inequality and row-stochasticity give

$$\int_0^1 \chi(a(i,j)) \mathrm{d}j \;\geq\; \chi\left(\int_0^1 a(i,j) \mathrm{d}j \right) = \chi(1) = 0,$$

i.e. $H(i) \leq 0$; strict convexity makes the inequality an equality only when $a(i,\cdot)$ is a.e. constant, and row-stochasticity then forces that constant to be one. Combined with Step 2, $H(i) \in [-\bar{e}, 0]$.

(ii). ν_i is a measure because $a(i,\cdot) \geq 0$ is measurable, and $\nu_i([0,1]) = \int_0^1 a(i,j) \mathrm{d}j = 1$ by (42), so it is a probability measure. It is defined as an integral against Lebesgue measure, so any Lebesgue-null B has $\nu_i(B) = 0$: it is absolutely continuous. Singletons are Lebesgue-null, so $\nu_i(\{j\}) = 0$ for every j; taking $j = i$ gives the diagonal statement. ■

[Background: absolute continuity with respect to Lebesgue measure, and why it forces diffuseness.]

- *Definition.* A measure ν on $[0,1]$ is absolutely continuous with respect to Lebesgue measure λ, written $\nu \ll \lambda$, if $\lambda(B) = 0$ implies $\nu(B) = 0$ for every Borel set B: ν cannot charge a set that Lebesgue measure regards as negligible. By the Radon–Nikodym theorem this is the same as $\nu(\mathrm{d}j) = f(j)\,\mathrm{d}j$ for some density $f \geq 0$; here $f = a(i,\cdot)$, and the lemma's proof is the easy direction, "defined by a density $\Rightarrow$ absolutely continuous".
- *Diffuse (atomless).* ν is diffuse if $\nu(\{j\}) = 0$ for every single point j. An *atom* is a point (more generally, a set of positive measure that cannot be split into two parts of positive measure) carrying positive mass; the Dirac measure δ_{j_0} is the pure example.
- *Why $\nu \ll \lambda$ implies diffuse.* Every singleton is Lebesgue-null, $\lambda(\{j\}) = 0$, so absolute continuity gives $\nu(\{j\}) = 0$. That is the whole argument, and it is why a kernel density can never express "a fixed share of purchases from one named supplier": that needs an atom (paragraph B.3).

- *The converse fails.* Diffuse does not imply absolutely continuous: the Cantor measure (uniform on the Cantor set) has no atoms yet lives on a Lebesgue-null set, so it is singular with respect to λ. Lebesgue's decomposition splits any finite measure into an absolutely continuous part, a singular-continuous part (Cantor-like) and a pure-point part (atoms); "diffuse" rules out only the last, "$\ll \lambda$" rules out the last two.
- *Reading in the model.* The bundle weights ν_i describe "a set of suppliers, not a list of them": positive-measure sets of suppliers matter, individual suppliers do not. That is the supply-side twin of atomism on the demand side, and it is what makes the coupling run through kernel integrals of moments rather than through named counterparties.
- (Standard measure theory; not re-derived.)

Reading (i): entropy and how concentrated a bundle may be. The integral $H(i)$ is the (differential) entropy of node i's bundle-weight density, and $-H(i)$ is the Kullback–Leibler divergence of $a(i,\cdot)$ from the uniform density. That is the reading behind the two bounds of Lemma 11(i). Entropy measures how far node i's sourcing is from spreading evenly over all suppliers. It is nonpositive $H(i) \leq 0$ and $H(i) = 0$ holds if and only if the row is uniform. Also $H(i) \to -\infty$, as the bundle concentrates on a supplier set of vanishing measure. The constant $\bar{e}$ is therefore an anti-concentration condition: it rules out near-atomic bundles, and it is automatic under Definition 2. What it buys is the bundle price index (95) of the model on $[0,1]$, in which $H(i)$ appears as an additive term, so a node whose sourcing collapsed onto a vanishing set would have a divergent marginal cost. Which of these conditions a general kernel genuinely needs, and what fails without them, is discussed in Appendix B.2.

Self-loops as an extension. In the finite model $a_{ii} > 0$ is allowed and means something sharper than block self-reliance: a sector uses *its own* good. Its continuum analogue is an atom, so by Lemma 11(ii) it is not a kernel and has to be added explicitly, as a bundle

$$M = m_{ii}^{a_0(i)} \exp\left\{[1 - a_0(i)] \int_0^1 \tilde{a}(i,j) \ln m_{ij} \mathrm{d}j\right\}$$

with atomic weight $a_0(i) \in [0,1)$ on the node's own good and a kernel $\tilde{a}$ carrying the rest. Everything below extends to $a_0 > 0$ at the cost of the node's own price entering its own marginal cost, so that the per-node static block of Section 3.2 acquires one additional (still scalar and monotone) fixed point. We set $a_0 \equiv 0$ in the baseline: a self-loop atom is a deliberate extension of the specification, not a violation of an assumption.

B.4 Notation Map: From the Finite Model to the Continuum

Every continuum object of Section 3.1 has a finite counterpart, and it is worth collecting them. The dictionary below reads the finite model of Section 2.1 into the continuum. Its rows on capital shares and aggregate capital look ahead: $\overline{K}_t$ and the node's capital share κ are constructed in Section 3.3, where the model is detrended; the rows on the capacity block anticipate Sections 3.1–3.2, where the threshold, utilization, spending and the rationing rent become functions of the node's own state.

Object	Finite ($i, j = 1, \ldots, J$)	Graphon ($i, j \in [0,1]$)
IO weights	a_{ij}, row-stochastic $\mathbf{A}$	kernel $a(i,j)$, (42)
Final shares	β_i, $\sum_i \beta_i = 1$	density $\beta(i)$, $\int \beta = 1$
Leontief inverse	$[\mathbf{I} - \mathbf{A}^\mathsf{T}\Xi]^{-1}$	$(\mathrm{Id} - \mathbb{A}^*\Xi)^{-1}$, (89)
Sector states	(K_i, Z_i)	(K_i, Z_i), distribution $\mathfrak{M}(t,i,\cdot)$
Sector shocks	$B^Z_{i,t}$, common within sector i	$B^Z_i(t)$, common within node i, e.p.i. across nodes
Shock volatilities	ς^{Z_i}	$\varsigma^Z(i)$
Capital shares	$\kappa_i = K_i/K$	$\kappa := K_i/\overline{K}_t$ (a density)
Aggregate capital	$K = \sum_i K_i$	$\overline{K}_t = \int_0^1 \mathbb{E}[K_i(t)]\mathrm{d}i$
Aggregate shocks	$\mathcal{Z}(\mathbf{Z})$ stochastic	none: aggregates deterministic
Capacity technology	x_i; $\overline{y}_i = Z_i K_i x_i^{\alpha_i - 1}$	$x(i)$; $\overline{y}_i = Z_i K_i x(i)^{\alpha(i)-1}$
Demand shocks	G_i, $\sigma_{v,i}$; tail functions $\mathcal{G}^-_{i,\theta}, \mathcal{G}^+_{i,\theta}$	G_i, $\sigma_v(i)$; the same tail functions, node by node
Threshold	$\overline{v}_i(\mathbf{K}, \mathbf{Z})$, shared by the sector	$\overline{v}(t,i,K,z)$, a function of the node's own state
utilization	$U_i(\overline{v}_i)$; $\overline{U}$	$U(t,i,K,z)$; node mean $\mathbb{E}_i[U]$; $\overline{U}(t)$, (116)
Intermediate spending	$X_i = \xi_i(\overline{v}_i) S^\mathrm{f}_i$	$X(t,i,K,z)$, moment $m_X(t,i)$ (Section 3.2)
Rationing rent	$\Pi^\mathrm{r}_i = S_i - S^\mathrm{f}_i$	$\Pi^\mathrm{r}(t,i,K,z) = S - S^\mathrm{f}$

B.5 The Model on $[0,1]$

Within-node block (unchanged, capacity included). Each node i is a sector: a unit measure of monopolistically competitive firms $n \in [0,1]$ with CES demand of elasticity ε and firm-specific demand shocks v drawn from G_i, with moments $\mu_{i,\theta} := \mathbb{E}_v[v^\theta]$ and

tail functions $\mathcal{G}^-_{i,\theta}$, $\mathcal{G}^+_{i,\theta}$, $\Psi_{i,\theta}$, $\eta_{i,\theta}$ of (57)–(58) at node i,[27] aggregated by competitive bundlers into the node good with shadow price index P_i and quantity Y_i (Lemma 1). Each monopolist produces, when it is not capacity-constrained,

$$y_{in} = Z_i\, k_{in}^{\alpha(i)}\, M_{in}^{1-\alpha(i)}, \qquad \ln M_{in} = \int_0^1 a(i,j)\, \ln m_{in}(j)\, \mathrm{d}j,$$

i.e. the intermediate bundle is the continuum Cobb–Douglas with weights $a(i,\cdot)$, and its output is capped by the capacity $\overline{y}_{in} = Z_i k_{in} x(i)^{\alpha(i)-1}$ of (10), one unit of capital processing at most $1/x(i)$ units of the bundle. Cost minimisation over the bundle gives the bundle price index

$$\ln \mathcal{P}_i = \int_0^1 a(i,j) \ln P_j \,\mathrm{d}j - \int_0^1 a(i,j) \ln a(i,j)\,\mathrm{d}j, \qquad m_{in}(j) = a(i,j)\,\frac{\mathcal{P}_i M_{in}}{P_j}, \tag{95}$$

the exact analogue of the finite log-bundle formula, with the row-sum replaced by $\mathbb{A}$ acting on $\ln P$.

Timing, pricing before v is realised, and the capacity block (rationing within the node, the threshold, the variable markup, the state-dependent revenue shares and utilization) are all node-local and identical to the finite model. In particular, at each node's symmetric equilibrium ($p_{in} = p_i$, $k_{in} = K_i$), with capacity $\overline{y}_i := Z_i K_i x(i)^{\alpha(i)-1}$ and threshold $\overline{v}_i = \frac{\overline{y}_i}{Y_i}(p_i/P_i)^{\varepsilon}$:

$$\text{(pricing/threshold)} \qquad P_i\, \mathcal{R}_i(\overline{v}_i) = \Gamma_i(\overline{v}_i)\, \frac{\mathcal{P}_i}{Z_i x(i)^{\alpha(i)}}\, \overline{v}_i^{\,-\frac{\alpha(i)}{1-\alpha(i)}}, \tag{96}$$

$$\text{(sales and revenue)} \qquad S_i := P_i Y_i = P_i\, \overline{y}_i\, \frac{\mathcal{R}_i(\overline{v}_i)^{\varepsilon}}{\overline{v}_i}, \tag{97}$$

$$S_i^{\mathrm{f}} := p_i \mathbb{E}_v[y_{in}] = S_i\, \frac{\Psi_{i,1}(\overline{v}_i)}{\mathcal{G}^-_{i,1}(\overline{v}_i) + \mathcal{G}^+_{i,1/\varepsilon}(\overline{v}_i)},$$

$$\text{(spending)} \qquad X_i = \mathcal{P}_i M_i^e = \mathcal{P}_i\, \frac{K_i}{x(i)}\, \Psi_{i,\frac{1}{1-\alpha(i)}}(\overline{v}_i) = \xi_i(\overline{v}_i)\, S_i^{\mathrm{f}}, \tag{98}$$

with $\mathcal{R}_i$, Γ_i and ξ_i the functions of Lemmas 3, 2 and 5, built from the tail functions of G_i with the node's parameters $(\alpha(i), x(i), G_i)$, and $\overline{v}_i \in (\underline{v}_i, \infty)$ (Lemma 2(i)). Since these steps repeat the finite-sector algebra verbatim (with $\sum_j a_{ij} \mapsto \int a(i,j)\mathrm{d}j$), we do not re-derive them; only note that *every object in* (96)–(98) *is node-local except* $\mathcal{P}_i$, which is a kernel average of other nodes' prices. Without capacity constraints the pricing relation is the constant-wedge supply curve and spending is the constant share $\zeta(i)$ of sales (Appendix C.4); under capacity both are functions of the node's threshold, and the threshold is a third node-level unknown alongside the price and sales.

Final CES demand. A competitive final producer combines node goods through a CES technology with elasticity $\eta > 0$, $\eta \neq 1$:

$$Q = \left[\int_0^1 \beta(i)^{1/\eta}\, Q_i^{(\eta-1)/\eta}\,\mathrm{d}i\right]^{\eta/(\eta-1)}. \tag{99}$$

[27] $\mathbb{E}_v$ is the expectation over the firm-specific shock, not the moment $\mathbb{E}_i$ against the node distribution of Section 3.2.

Cost minimisation gives final demands and the numeraire condition (final good price $\equiv 1$):

$$Q_i = \beta(i)\, P_i^{-\eta}\, Q, \qquad 1 = \left[\int_0^1 \beta(i)\, P_i^{1-\eta}\, \mathrm{d}i\right]^{1/(1-\eta)}. \tag{100}$$

The Cobb–Douglas final aggregator of the finite model is the $\eta \to 1$ limit, in which nominal final spending on node i is the constant $\beta(i)Q$. For $\eta \neq 1$ nominal final spending $P_iQ_i = \beta(i)P_i^{1-\eta}Q$ moves with the node's *relative price*, and hence, through (96)–(97), with the node's own (K_i, Z_i). This is the whole point of the CES; Appendix B.7 makes it precise.

Remark 5 (More about CES) The CES assumption leaves the entire within-node block (96)–(98) and the Cobb–Douglas intermediate bundle (95) untouched: the expenditure share $\xi_i(\overline{v}_i)$, the markup and utilization are functions of the node's threshold whatever the final-demand system, so the static block below stays a scalar equation per node, in the threshold. Only the demand *level* facing a node acquires a price elasticity. The alternative (firm-level persistent productivity within nodes) would instead create a nondegenerate within-node distribution and a genuinely three-layer model; it is orthogonal to what follows and omitted.

Market clearing. Node goods clear,

$$Y_i = Q_i + \int_0^1 m_j(i)\, \mathrm{d}j \tag{101}$$

where $m_j(i)$ is node j's demand for node-i goods. In value terms, using (95) and (98), node j spends $a(j,i)X_j$ on node-i goods, a fixed value share of a state-dependent bill, so

$$S_i = \underbrace{\beta(i)P_i^{1-\eta}Q}_{\text{final}} + \underbrace{\int_0^1 a(j,i)\, X_j\, \mathrm{d}j}_{=\,(\mathbb{A}^* X)(i)}, \qquad Q = \int_0^1 \left[S_j - X_j\right] \mathrm{d}j, \tag{102}$$

the second identity being GDP = gross output net of intermediate use. Without capacity constraints $X_j = \xi(j)S_j$ and the inbound term is $(\mathbb{A}^*\Xi S)(i)$ (Appendix C.4); under capacity the buyers' bills, not their sales, are what the kernel routes to node i.

Lemma 12 (Within-node block and value clearing on $[0,1]$). *(i)* (Bundle price index.) *The unit cost of node i's intermediate bundle is* (95)*, with expenditure-share densities* $P_j m_{in}(j)/(\mathcal{P}_i M_{in}) = a(i,j)$.

In operator form, with $H(i) := -\int_0^1 a(i,j)\ln a(i,j)\mathrm{d}j$ *the row entropy,* $\ln \mathcal{P}_i = (\mathbb{A}\ln P)(i) + H(i)$.

(ii) (Pricing, sales and spending.) *At each node's symmetric equilibrium the pricing/threshold relation* (96)*, the sales and revenue relations* (97) *and the spending relation* (98) *hold, with the threshold in the elastic region* $(\underline{v}_i, \infty)$.

(iii) (Value clearing.) *Node sales satisfy* (102)*, and aggregate final expenditure equals gross output net of intermediate use,* $Q = \int_0^1 [S_j - X_j]\mathrm{d}j$.

Proof. (i) The bundle problem: choose input densities $m(j) \geq 0$ to attain a bundle level M at least cost,

$$\mathcal{P}_i\, M := \min_{m(\cdot)} \int_0^1 P_j\, m(j)\mathrm{d}j \;\text{ s.t. }\; \int_0^1 a(i,j) \ln m(j)\mathrm{d}j = \ln M.$$

(1) *First-order condition.* With multiplier Λ on the constraint, the Lagrangian

$$\int_0^1 \left[P_j\, m(j) - \Lambda\, a(i,j) \ln m(j)\right]\mathrm{d}j + \Lambda \ln M$$

is separable across j, so minimise pointwise:

$$P_j = \Lambda \frac{a(i,j)}{m(j)} \iff m(j) = \Lambda \frac{a(i,j)}{P_j} \qquad \text{for a.e. } j.$$

(The pointwise objective is convex in $m(j)$, and $\ln m$ forces $m > 0$ wherever $a(i,j) > 0$; where $a(i,j) = 0$, optimally $m(j) = 0$.)

(2) *Total cost equals the multiplier.* Substituting the FOC into the objective,

$$\int_0^1 P_j\, m(j)\mathrm{d}j = \Lambda \int_0^1 a(i,j)\mathrm{d}j = \Lambda,$$

by row-stochasticity (42). So total spending is Λ, and the expenditure *share densities* are $P_j m(j)/\Lambda = a(i,j)$: the kernel is literally the density of Cobb–Douglas expenditure shares.

(3) *Pin down Λ from the constraint.*

$$\ln M = \int_0^1 a(i,j) \ln \frac{\Lambda\, a(i,j)}{P_j}\mathrm{d}j = \ln \Lambda + \int_0^1 a(i,j) \ln a(i,j)\mathrm{d}j - \int_0^1 a(i,j) \ln P_j \mathrm{d}j,$$

again using (42) to pull $\ln \Lambda$ out of the integral.

(4) *Unit cost.* The technology is homogeneous of degree one in $m(\cdot)$ (scaling all m by c adds $\ln c$ to $\ln M$), so cost is linear in M and the unit cost is the price index: $\mathcal{P}_i = \Lambda / M$. Taking logs of step (3),

$$\ln \mathcal{P}_i = \ln \Lambda - \ln M = \int_0^1 a(i,j) \ln P_j \mathrm{d}j - \int_0^1 a(i,j) \ln a(i,j)\mathrm{d}j,$$

which is (95); and substituting $\Lambda = \mathcal{P}_i M$ back into the FOC gives the demand density $m(j) = a(i,j)\, \mathcal{P}_i M / P_j$, the second part of (95).

Operator form. In the graphon operators of Definition 3, writing $H(i) := -\int_0^1 a(i,j) \ln a(i,j)\mathrm{d}j$ for the row entropy,

$$\ln \mathcal{P}_i = (\mathbb{A} \ln P)(i) + H(i),$$

and in moment form this is exactly (44): $\ln \mathcal{P}(t,i) = (\mathbb{A} m_P)(t,i) + H(i)$: the kernel average of supplier log prices plus the (nonrandom) entropy of node i's own bundle weights.

(ii) The pricing/threshold relation is Lemma 3(v) in its first form, the sales and revenue relations are Lemma 3(ii) and (iv), and the spending relation is Lemma 3(iii) together with Lemma 5(i), all at sector i with parameters $(\alpha(i), x(i), G_i)$; the threshold lies in $(\underline{v}_i, \infty)$ by Lemma 2(i). Their derivations in Section 2.1 (Appendices A.2, A.3, A.4 and A.6) are node-local, i.e. they involve the node's own firms, its bundlers and the node objects $(Y_i, P_i, \mathcal{P}_i)$ only, and go through verbatim with $\sum_j a_{ij}(\cdot)$ replaced by $\int_0^1 a(i,j)(\cdot)\mathrm{d}j$; the only non-local object is $\mathcal{P}_i$, which (i) provides.

(iii)

(1) Multiply (101) by P_i:

$$S_i = P_i Y_i = P_i Q_i + \int_0^1 P_i\, m_j(i)\mathrm{d}j.$$

For the final-demand term, CES demand (100) gives $P_i Q_i = P_i \cdot \beta(i) P_i^{-\eta} Q = \beta(i) P_i^{1-\eta} Q$. For the intermediate term, the bundle demand in (95), read with j as the buyer and i as the supplier and integrated over the buyer's firms, is $m_j(i) = a(j,i)\,\mathcal{P}_j M_j^e / P_i$, so

$$P_i\, m_j(i) = a(j,i)\,\mathcal{P}_j M_j^e = a(j,i)\, X_j,$$

using the definition of intermediate spending $X_j = \mathcal{P}_j M_j^e$ of (98). Integrating over buyers j gives $\int_0^1 a(j,i) X_j \mathrm{d}j = (\mathbb{A}^* X)(i)$, hence the first identity: every euro of node i's sales is either final spending (price-sensitive, elasticity η) or a fixed value share $a(j,i)$ of some buyer's intermediate bill X_j, which is itself the state-dependent share $\xi_j(\overline{v}_j) S_j^{\mathrm{f}} / S_j$ of that buyer's sales.

(2) Integrate the first identity over i:

$$\int_0^1 S_i \mathrm{d}i = Q \underbrace{\int_0^1 \beta(i) P_i^{1-\eta} \mathrm{d}i}_{=1 \text{ by (100)}} + \underbrace{\int_0^1 (\mathbb{A}^* X)(i) \mathrm{d}i}_{=\int_0^1 X_j \mathrm{d}j},$$

where the second brace is Tonelli plus row-stochasticity exactly as in the proof of Lemma 8 ($\int_0^1 a(j,i)\mathrm{d}i = 1$ at fixed j): the adjoint operator conserves total mass. Rearranging,

$$Q = \int_0^1 \left[S_j - X_j\right] \mathrm{d}j,$$

that is, nominal final expenditure (at the numeraire price 1) equals gross output net of intermediate use, i.e. aggregate value added. Note that the numeraire condition is an input to this step, which is why Remark 7 treats the identity as a Walras check rather than an independent equation.

■

Capital. Each monopolist owns its capital, which accumulates deterministically,

$$\mathrm{d}k_{in} = [\Phi_i(\iota_{in}) - \delta(i)]k_{in}\,\mathrm{d}t, \tag{103}$$

with investment purchased in units of the final good; z_i follows (93). Since investment and prices are chosen before v is realised, all monopolists at a node choose identically, and the node state is (K_i, z_i) with within-node distribution degenerate at K_i, exactly as in the finite model: capacity adds heterogeneity in realised output and in who is rationed, but no persistent firm-level state (Section 2.1).

Aggregation across nodes. The node state (K_i, z_i) is random, but every cross-sectional integral of node-level random variables is almost surely equal to the integral of its mean. The statement is Sun's exact law of large numbers, and it holds not on the ordinary product of the node and sample spaces (on which an independent continuum of Brownian motions is not even jointly measurable) but on a *Fubini extension* $([0,1] \times \Omega, \mathcal{W}, \lambda \boxtimes \mathbb{P})$ of that product, on which Fubini's theorem holds for $\mathcal{W}$-measurable processes and essentially pairwise independent (e.p.i.) families with prescribed marginals exist. We fix such an extension once and for all and construct $\{z_i\}$ on it; Appendix B.5.1 below explains why an ordinary product space will not do and gives the references.

Lemma 13 (Exact law of large numbers; Sun, 2006). *Let $\psi : [0,1] \times \Omega \to \mathbb{R}$ be $\mathcal{W}$-measurable, integrable, and e.p.i. across i. Then, $\mathbb{P}$-a.s.,*

$$\int_0^1 \psi(i,\omega)\,\mathrm{d}\lambda(i) = \int_0^1 \mathbb{E}[\psi(i,\cdot)]\,\mathrm{d}\lambda(i). \tag{104}$$

The same conclusion holds for kernel-weighted integrals $\int a(j,i)\psi(j,\omega)\mathrm{d}j$ and $\int a(i,j)\psi(j,\omega)\mathrm{d}j$ for a.e. i, since $j \mapsto a(\cdot,j)\psi(j,\cdot)$ inherits measurability and e.p.i. from ψ and boundedness of a.

The averaging operates across nodes, not within them (Section 3.1): the unit measure of monopolists at a node all share the common node shock B_i^Z, so within-node averaging has nothing to work on. This is the first of the six differences from the cluster-based construction of Caines and Huang (2021), set out in Appendix B.9.3.

B.5.1 The Probabilistic Stage: Fubini Extensions

The exact law of large numbers of Lemma 13, on which every deterministic aggregate of the model rests, is stated on a space that is not the ordinary product of the node and sample spaces. This subsection says why, and what the enlarged space is.

If $\{B_i^Z\}_{i\in[0,1]}$ were i.i.d. Brownian motions on an ordinary product space, the sample map $i \mapsto B_i^Z(t,\omega)$ would not be Lebesgue-measurable, so the realised cross-sectional integrals the model is built on, $\int_0^1 f(i, K_i(t), z_i(t))\mathrm{d}i$, would not even be defined. The obstruction is not an artefact of a poor construction: a classical observation going back to Doob shows that a process measurable with respect to the product σ-algebra and

essentially pairwise independent across i is almost surely constant for a.e. i, so a genuinely random independent continuum cannot be jointly measurable in the classical sense.

Sun's (2006) resolution keeps the independence and weakens the measurability requirement. A *Fubini extension* $([0,1] \times \Omega, \mathcal{W}, \lambda \boxtimes \mathbb{P})$ is a probability space whose σ-algebra $\mathcal{W}$ strictly contains the product σ-algebra $\mathcal{B} \otimes \mathcal{F}$, whose measure agrees with $\lambda \otimes \mathbb{P}$ on it, and on which Fubini's theorem holds for every $\mathcal{W}$-measurable integrable process. *Rich* Fubini extensions exist (Sun, 2006; Sun & Zhang, 2009) and carry processes that are $\mathcal{W}$-measurable and *essentially pairwise independent* — e.p.i.: for λ-a.e. i, $\psi(i,\cdot)$ is independent of $\psi(j,\cdot)$ for λ-a.e. j — with any prescribed marginals, in particular the node-wise Ornstein–Uhlenbeck family $\{B_i^Z\}$ of Assumption 4. We fix such an extension throughout and construct $\{z_i\}$ on it. On it the realised cross-section is λ-integrable for $\mathbb{P}$-a.e. ω, and Lemma 13 upgrades "well defined" to "deterministic". Bayraktar et al. (2023) establish the law of large numbers for graphon-interacting particle systems in exactly this setting.

[The Fubini extension $([0,1] \times \Omega,\ \mathcal{W},\ \lambda \boxtimes \mathbb{P})$, element by element.] The two ordinary spaces being combined:

- $([0,1],\ \mathcal{B},\ \lambda)$ is the *node space*: $[0,1]$ is the vertex set, $\mathcal{B}$ the Lebesgue sets, and Lebesgue measure λ plays the role of the uniform "count" of nodes: $\lambda(M)$ is the mass of a node set M, and $\mathrm{d}\lambda(i)$ is what every cross-sectional integral $\int \cdots \mathrm{d}i$ in the text integrates against.
- $(\Omega,\ \mathcal{F},\ \mathbb{P})$ is the *sample space*: a single $\omega \in \Omega$ is one complete realisation of all randomness in the economy at once (every node's Brownian path $\{B_i^Z(\cdot,\omega)\}_{i\in[0,1]}$), $\mathcal{F}$ its σ-algebra, $\mathbb{P}$ the probability measure.
- Points of the product set $[0,1] \times \Omega$ are pairs (i,ω): "node i, in state of the world ω". A function $\psi : [0,1] \times \Omega \to \mathbb{R}$ is a *process indexed by nodes*: freezing i gives node i's random variable $\psi(i,\cdot)$; freezing ω gives the cross-sectional sample function $\psi(\cdot,\omega)$, the object the aggregates integrate over nodes.
- $\mathcal{W}$ is a σ-algebra on $[0,1] \times \Omega$ that *strictly contains* the classical product σ-algebra $\mathcal{B} \otimes \mathcal{F}$: every classically product-measurable set is in $\mathcal{W}$, but $\mathcal{W}$ has more sets. "ψ is $\mathcal{W}$-measurable" is therefore a *weaker* requirement than classical joint measurability, and that slack is the entire point (below).
- $\lambda \boxtimes \mathbb{P}$ is a probability measure on $\mathcal{W}$ with two defining properties: (a) it agrees with the ordinary product measure $\lambda \otimes \mathbb{P}$ on $\mathcal{B} \otimes \mathcal{F}$, so nothing that was already well defined changes; and (b) the *Fubini property* holds on all of $\mathcal{W}$: for every $\mathcal{W}$-measurable, $\lambda \boxtimes \mathbb{P}$-integrable ψ, the sections $\psi(i,\cdot)$ and $\psi(\cdot,\omega)$ are (a.e.) measurable and integrable, and

$$\int_{[0,1]\times\Omega} \psi \,\mathrm{d}(\lambda \boxtimes \mathbb{P}) = \int_0^1 \left(\int_\Omega \psi(i,\omega)\mathrm{d}\mathbb{P}(\omega) \right) \mathrm{d}\lambda(i) = \int_\Omega \left(\int_0^1 \psi(i,\omega)\mathrm{d}\lambda(i) \right) \mathrm{d}\mathbb{P}(\omega).$$

  The symbol $\boxtimes$ (rather than $\otimes$) flags that this is *not* the product measure but a proper extension of it carrying the Fubini property on the enlarged σ-algebra,

hence "Fubini extension".

Why an extension is unavoidable rather than a technicality:

- A classical negative result (Doob's observation; see Sun 2006) says: if ψ is measurable with respect to the *classical* product σ-algebra (even completed) and the family $\{\psi(i,\cdot)\}_i$ is essentially pairwise independent, then $\psi(i,\cdot)$ is a.s. *constant* for a.e. i. A genuinely random independent continuum therefore cannot be classically jointly measurable; and without joint measurability the sample function $i \mapsto \psi(i,\omega)$ need not be λ-measurable, so the realised aggregate $\int_0^1 \psi(i,\omega)\mathrm{d}i$ is not even defined. This is precisely the measurability obstruction described in the text.
- Sun's resolution keeps the independence and *weakens the measurability demand* to $\mathcal{W}$: *rich* Fubini extensions exist (Sun 2006; Sun–Zhang 2009) carrying processes that are $\mathcal{W}$-measurable *and* essentially pairwise independent (e.p.i.) with any prescribed marginals, in particular the OU family $\{z_i\}$ driven by e.p.i. Brownian motions. "E.p.i." means: for λ-a.e. i, the random variable $\psi(i,\cdot)$ is independent of $\psi(j,\cdot)$ for λ-a.e. j. Only *pairwise* independence (off null sets) is demanded, and that is all the exact LLN needs.
- On such a space everything clicks into place: by the Fubini property the realised cross-section $\psi(\cdot,\omega)$ *is* λ-integrable for $\mathbb{P}$-a.e. ω, and the exact LLN (Lemma 13) upgrades "well defined" to "deterministic": the realised aggregate equals its expected value with probability one.

On $\mathbb{P}$: The Physical Measure.

- "Physical" (also "statistical", "real-world", "$\mathbb{P}$-measure") is finance's name for the probability measure under which the data are actually generated, as opposed to a *risk-neutral* (pricing) measure $\mathbb{Q}$, which reweights states by marginal utility so that assets price as discounted $\mathbb{Q}$-expectations at the risk-free rate. Here $\mathbb{P}$ is simply the sample-space measure of the model's randomness (all the Brownian paths), so yes: it is the physical measure.
- No measure change is ever used in this document: all pricing is done under $\mathbb{P}$ with the SDF written explicitly. The two languages are equivalent ($\mathrm{d}\mathbb{Q}/\mathrm{d}\mathbb{P}$ is the normalised SDF), and this model makes the equivalence trivial in an instructive way: by Lemma 18 the SDF is *deterministic* in the graphon limit, so the density $\mathrm{d}\mathbb{Q}/\mathrm{d}\mathbb{P}$ is identically one and the risk-neutral and physical measures *coincide*. That is the measure-theoretic face of "no aggregate risk, no risk premia, discounting at $r = \rho + \gamma g$": node-level risk exists under $\mathbb{P}$ but is idiosyncratic, fully diversified in the household's portfolio, and therefore carries a zero risk premium, i.e. no reweighting.
- In the *finite-J* model the distinction is real: the SDF is stochastic, $\mathbb{Q} \neq \mathbb{P}$, and the utility-weighted $\mathfrak{q}_i$ formulation is exactly the device that avoids ever writing $\mathbb{Q}$ while still pricing the covariances.

B.6 The Static Block: Aggregates, the Node System and Profits

This appendix collects the results behind Section 3.2: the lemma that makes the aggregates deterministic, the node system and its solution, the structure of the coupling and the GDP identity, the individual firm's profit function, and node utilization and its moments. Throughout, t is fixed and $\mathbb{E}_i[f](t) = \int f(t,i,K,z)\,\mathfrak{M}(t,i,\mathrm{d}K,\mathrm{d}z)$ as in Section 3.2.

Aggregates are deterministic.

Lemma 14 (Aggregates are deterministic). *Suppose the equilibrium maps P, X, S are jointly measurable and that $\ln P$, S, X and $P^{1-\eta}$ have finite first moments against $\mathfrak{M}(t,j,\cdot)$, uniformly in j:*

$$\sup_{j\in[0,1]} \int |\ln P(t,j,K,z)|\ \mathfrak{M}(t,j,\mathrm{d}K,\mathrm{d}z) < \infty,$$
$$\sup_{j\in[0,1]} \int S(t,j,K,z)\ \mathfrak{M}(t,j,\mathrm{d}K,\mathrm{d}z) < \infty, \qquad \sup_{j\in[0,1]} \int X(t,j,K,z)\ \mathfrak{M}(t,j,\mathrm{d}K,\mathrm{d}z) < \infty,$$
$$\sup_{j\in[0,1]} \int P(t,j,K,z)^{1-\eta}\ \mathfrak{M}(t,j,\mathrm{d}K,\mathrm{d}z) < \infty \tag{105}$$

(for $\eta > 1$ the last is a lower-tail condition on P, not implied by the others; the condition on X is implied by the one on S whenever $X \le (1-\frac{1}{\varepsilon})S$, which holds under Assumption 3 by Remark 4). Then the realised bundle price $\mathcal{P}_i$, the realised inbound intermediate demand $\int a(j,i)X_j\mathrm{d}j$, the realised numeraire integral $\int \beta P_i^{1-\eta}\mathrm{d}i$, and realised GDP all coincide $\mathbb{P}$-a.s. with (44)–(46).

Proof. Fix t. The four integrability requirements in the hypothesis are exactly the four transformations of (P,X,S) that the aggregates use.

Step 1: the node states form an e.p.i. family. In equilibrium each node's state $s_j(t) = (K_j(t), z_j(t))$ is a measurable functional of node j's own Brownian path B_j^Z and its (nonrandom) initial condition: z_j solves (93), and K_j follows the closed-loop capital dynamics (136), which are driven by z_j alone. The family $\{B_j^Z\}_j$ is $\mathcal{W}$-measurable and e.p.i. by construction (Appendix B.5.1), and node-wise measurable transformations preserve both properties; hence $\{s_j(t)\}_{j\in[0,1]}$ is $\mathcal{W}$-measurable and e.p.i., with $s_j(t) \sim \mathfrak{M}(t,j,\cdot)$.

Step 2: the realised bundle price. Fix i and set $\psi(j,\omega) := a(i,j)\ln P\big(t,j,s_j(t,\omega)\big)$. The three hypotheses of Lemma 13 hold: ψ is $\mathcal{W}$-measurable (a jointly measurable function composed with the state family of Step 1); it is e.p.i. across j, because $\psi(j,\cdot)$ depends on the randomness only through s_j; and it is integrable, since $|\psi(j,\omega)| \le \bar{a}\,\big|\ln P(t,j,s_j)\big|$ with finite expectation by hypothesis. The kernel-weighted form of Lemma 13 then gives, $\mathbb{P}$-a.s. for a.e. i,

$$\int_0^1 a(i,j)\ \ln P\big(t,j,s_j(t)\big)\mathrm{d}j = \int_0^1 a(i,j)\,\mathbb{E}[\ln P(t,j,\cdot)]\mathrm{d}j = (\mathbb{A}m_P)(t,i).$$

Adding the nonrandom entropy term $-\int_0^1 a(i,j)\ln a(i,j)\mathrm{d}j$ to both sides shows that the realised $\ln \mathcal{P}_i$ of (95) equals (44).

Step 3: inbound demand, numeraire integral, GDP. The same three checks (measurability, e.p.i. through s_j only, integrability with a nonrandom bounded factor) apply to

$$\psi(j,\omega) = a(j,i)\,X\big(t,j,s_j(t,\omega)\big) \quad \Rightarrow \quad \int_0^1 a(j,i)\,X_j\mathrm{d}j = (\mathbb{A}^* m_X)(t,i) = \mathcal{D}(t,i),$$

$$\psi(j,\omega) = \beta(j)\,P\big(t,j,s_j(t,\omega)\big)^{1-\eta} \quad \Rightarrow \quad \int_0^1 \beta(j)\,P_j^{1-\eta}\mathrm{d}j = \int_0^1 \beta(j)\,\mathbb{E}\Big[P_j^{1-\eta}\Big]\mathrm{d}j,$$

$$\psi(j,\omega) = [S-X]\big(t,j,s_j(t,\omega)\big) \quad \Rightarrow \quad \int_0^1 \big[S_j - X_j\big]\mathrm{d}j = \int_0^1 [m_S - m_X]\,(t,j)\mathrm{d}j = Q(t),$$

with bounded nonrandom factors $\bar{a}$, $\bar{\beta}$, and 1 respectively, and with the integrability of X either assumed or, under Assumption 3, inherited from that of S through $0 \le X \le (1-\frac{1}{\varepsilon})S$. The first line shows the realised inbound demand equals (45); the second shows the realised numeraire integral equals its moment form (the left-hand side of (47)); the third shows realised GDP equals (46). In each case the realised object is $\mathbb{P}$-a.s. equal to a deterministic functional of the ensemble, as claimed. ■

The node system. Fix aggregates $(\mathcal{P}(t,i), \mathcal{D}(t,i), Q(t))$ and a node state (K,z), $Z = e^z$. The within-node block (96)–(97) involves three node-level unknowns: the price P, the threshold $\overline{v}$ and sales S. It is convenient to write the pricing/threshold relation through the *threshold map*

$$\mathrm{Y}_i(\overline{v}) := \frac{\mathcal{R}_i(\overline{v})\,\overline{v}^{\frac{\alpha(i)}{1-\alpha(i)}}}{\Gamma_i(\overline{v})} = \frac{\mathcal{R}_i(\overline{v})\Big[(\varepsilon-1)\mathcal{G}^-_{i,1}(\overline{v}) - 1 + G_i(\overline{v})\Big]}{\varepsilon^{\frac{1}{1-\alpha(i)}}\mathcal{G}^-_{i,\frac{1}{1-\alpha(i)}}(\overline{v})}, \qquad \overline{v} \in (\underline{v}_i, \infty), \quad (106)$$

the left-hand side of the finite inner-root equation (81). It is continuous, with $\mathrm{Y}_i(\underline{v}_i^+) = 0$ and $\mathrm{Y}_i(\overline{v}) \to \infty$ as $\overline{v} \to \infty$ (Proposition 3(i)), and far in the elastic region it behaves like $\overline{v}^{\frac{\alpha(i)}{1-\alpha(i)}}/\Gamma(i)$, with $\Gamma(i)$ the constant wedge of the model without capacity constraints (Lemma 4(i)). Then (96) and (97), written with the node's capacity $\overline{Y} := e^z K x(i)^{\alpha(i)-1}\mu_{i,1/\varepsilon}^{\frac{\varepsilon}{\varepsilon-1}}$ and utilization $U_i(\overline{v})$ of Definition 1, read

$$\text{(pricing/threshold)} \qquad P\,\mathrm{Y}_i(\overline{v}) = \frac{\mathcal{P}(t,i)}{e^z\,x(i)^{\alpha(i)}}, \quad (107)$$

$$\text{(sales)} \qquad S = P\,\overline{Y}\,U_i(\overline{v}) = \frac{\mathcal{P}(t,i)\,K}{x(i)}\,\mu_{i,1/\varepsilon}^{\frac{\varepsilon}{\varepsilon-1}}\,\frac{U_i(\overline{v})}{\mathrm{Y}_i(\overline{v})}, \quad (108)$$

where the second form of (108) substitutes the price out: productivity cancels, and sales are the node's bundle bill at full capacity, $\mathcal{P}(t,i)K/x(i)$, times a function of the threshold alone. Demand-side clearing at the node, from (102), is

$$S = \beta(i)\,P^{1-\eta}\,Q(t) + \mathcal{D}(t,i). \quad (109)$$

Equations (107)–(109) are three equations in the three node-level unknowns $(P, \overline{v}, S)$ given (K, z) and aggregates. Substituting (107) and the second form of (108) into (109) leaves one scalar equation in the threshold,

$$\frac{\mathcal{P}(t,i)\,K}{x(i)}\,\mu_{i,1/\varepsilon}^{\frac{\varepsilon}{\varepsilon-1}}\,\frac{U_i(\overline{v})}{\Upsilon_i(\overline{v})} = \beta(i)\,Q(t)\left(\frac{\mathcal{P}(t,i)}{e^z x(i)^{\alpha(i)}\,\Upsilon_i(\overline{v})}\right)^{1-\eta} + \mathcal{D}(t,i), \tag{110}$$

in which K enters only the supply side and z only the final-demand term. Without capacity constraints the corresponding equation is a power equation in the price, $cP^{1/\alpha(i)} = \beta(i)QP^{1-\eta} + \mathcal{D}$, whose unique root follows from the ordering of the exponents (Appendix C.4); here the supply side is the ratio of two tail-function combinations of the threshold, and uniqueness of the root rests on the monotonicity of the threshold map, the property that Proposition 3 left open for the finite inner block. Here it is named, because the continuum block uses it at every node and in every state.

Assumption 5 (Monotone threshold map). For each node i, the threshold map Υ_i of (106) is strictly increasing on $(\underline{v}_i, \infty)$.

Like Assumptions 2 and 3, this is a shape restriction on G_i through $(\varepsilon, \alpha(i))$ alone, and its two ends are fixed: Υ_i rises from 0 at $\underline{v}_i$ and grows like the increasing power $\overline{v}^{\frac{\alpha(i)}{1-\alpha(i)}}$ at infinity, the form it takes without capacity constraints. Its logarithmic derivative is written out in the flag after Proposition 3; it contains two density terms of indeterminate sign, and no proof is available. The assumption implies the monotonicity clause of Assumption 2: the ratio there equals $\varepsilon\frac{1}{1-\alpha(i)}\Upsilon_i(\overline{v})\,[\overline{v}/\mathcal{R}_i(\overline{v})^{\varepsilon}]^{1/\varepsilon}$, the product of Υ_i and a function that is strictly increasing by Lemma 3(ii), so under Assumption 5 only the single zero of $(\varepsilon-1)\mathcal{G}_{i,1}^{-} - 1 + G_i$ remains to be assumed. Economically, the assumption says that at given aggregates a node with a higher threshold, hence slacker capacity, posts a lower price: the price $\mathcal{P}(t,i)/(e^z x(i)^{\alpha(i)}\Upsilon_i(\overline{v}))$ is then a strictly decreasing function of the threshold, from $+\infty$ at $\underline{v}_i$ to 0. **[Flag: Assumption 5 is verified numerically, not proved.]** On the grid of 648 log-normal parameter triples of the flag after Proposition 3 ($\sigma_v \in \{0.1, \ldots, 2\}$, $\varepsilon \in \{1.2, \ldots, 15\}$, $\alpha \in \{0.1, \ldots, 0.85\}$; 4001 logarithmically spaced thresholds on $(\underline{v}_i, e^{14})$) Υ_i was strictly increasing in every triple; it is the same object as the finite inner-root map, so the numerical evidence recorded there is the evidence for this assumption. Independently, the scalar equation (110) was scanned on the same threshold grid in 1200 random node problems (400 each at $\eta \in \{0.5, 2\}$ and at the Cobb–Douglas value $\eta = 1$; random triples, aggregates $(\mathcal{P}, Q, \mathcal{D})$, share β, state (K, z) and intensity x, with $\mathcal{D} = 0$ in 15% of the cases): the left-hand side minus the right-hand side changed sign exactly once in every case, from positive to negative, its derivative at the root was negative in every case, the two forms of (108) agreed to 10^{-15} at the root, and the root moved with (K, z) as Lemma 15(iv) states in every case. The elasticity (111) approaches the no-capacity value $1/\alpha(i)$, but not uniformly fast over the grid: the density term that has to vanish, $\overline{v}^{1+\frac{1}{1-\alpha(i)}}g_i(\overline{v})$, only begins to decay past $\ln\overline{v} \approx \sigma_{v,i}^2(\frac{1}{1-\alpha(i)} + \frac{1}{2})$, so at $\overline{v} = 10^2$ the match is to four digits at moderate dispersion but not at all at the top of the grid — at $\sigma_{v,i} = 2$, $\varepsilon = 1.2$, $\alpha = 0.8$

the elasticity is still 1.86 against $1/\alpha = 1.25$, and 1.45 at $\overline{v} = 10^5$. The same caveat applies to every "attained at $\overline{v} = 10^2$" statement about the capacity limit elsewhere in this paper. **[Q14 | text | tag: LogNormal3]** Prove Assumption 5 for the log-normal specification, or give a sufficient condition on G_i under which it holds, and relate it to Assumptions 2 and 3 (it implies the monotonicity clause of the former, see the text above; is it implied by the latter, or the converse?). The weaker property that the proof of Lemma 15(iv) actually uses at a root is that $\beta(i)QP^{1-\eta}[\ell_U + \eta\,\ell_Y] + \mathcal{D}[\ell_U + \ell_Y] > 0$ there, in the notation $\ell_Y = \mathrm{d}\ln Y_i/\mathrm{d}\ln\overline{v}$, $\ell_U = -\,\mathrm{d}\ln U_i/\mathrm{d}\ln\overline{v}$ of that proof: with $\mathcal{D} = 0$ this is $\ell_U + \eta\ell_Y > 0$, and a condition uniform in the aggregates is $\ell_U + \min\{1,\eta\}\ell_Y > 0$, which for small η tolerates a locally decreasing threshold map. Note that as a statement for *every* $\eta > 0$, as the lemma is phrased, this collapses back to $\ell_Y \geq 0$, so the assumption is close to necessary at that generality and is stronger than needed only at a fixed η. If proved, replace the assumption by a lemma and remove the corresponding sentence in the reading after Proposition 3.

Lemma 15 (The per-node static block under capacity). *Fix i and aggregates with $\mathcal{P}(t,i) > 0$, $Q(t) > 0$, $\mathcal{D}(t,i) \geq 0$, and let Assumption 4 hold. Then:*

(i) (Pricing and sales.) *At the node's symmetric equilibrium the price, the threshold and sales satisfy* (107)–(108)*, with $\overline{v} \in (\underline{v}_i, \infty)$.*

(ii) (Demand.) *Node sales satisfy the demand relation* (109)*.*

(iii) (Existence.) *For every $(K,z) \in (0,\infty)\times\mathbb{R}$ and every $\eta > 0$, the scalar equation* (110) *has at least one solution $\overline{v} \in (\underline{v}_i, \infty)$, hence the system* (107)–(109) *a solution $(P, \overline{v}, S)$ with $P > 0$ and $S > 0$.*

(iv) (Uniqueness and monotonicity.) *Under Assumption 5 the solution is unique, for every $\eta > 0$ including the Cobb–Douglas value $\eta = 1$; write $\overline{v}(t,i,K,z)$, $P(t,i,K,z)$ and $S(t,i,K,z)$ for it. At fixed (K,z) the price and sales are strictly decreasing functions of the threshold, and, wherever $Y_i'(\overline{v}) > 0$, the elasticity of sales with respect to the node's own price along the supply side is*

$$e_S(\overline{v}) := \left.\frac{\mathrm{d}\ln S}{\mathrm{d}\ln P}\right|_{K,z} = 1 + \frac{-\,\mathrm{d}\ln U_i/\mathrm{d}\ln\overline{v}}{\mathrm{d}\ln Y_i/\mathrm{d}\ln\overline{v}} > 1, \tag{111}$$

which tends to the no-capacity value $1/\alpha(i)$ as $\overline{v}\to\infty$ whenever $\overline{v}^{1+\frac{1}{1-\alpha(i)}} g_i(\overline{v}) \to 0$ for the density g_i of G_i (as under the log-normal specification). The threshold is continuously differentiable in (K,z), strictly increasing in K, and strictly increasing in z if $\eta < 1$, strictly decreasing if $\eta > 1$; node utilization $U(t,i,K,z) := U_i(\overline{v}(t,i,K,z))$ moves the opposite way.

(v) (Productivity.) *z enters* (110) *only through the final-demand term, in which $P^{1-\eta} \propto e^{-(1-\eta)z}$ at a given threshold; in the Cobb–Douglas limit $\eta = 1$ the threshold, and with it utilization, the spending share and the revenue split, are independent of z (Appendix B.7).*

Proof. Drop t and the node index on the tail functions, write $\gamma := \frac{1}{1-\alpha(i)}$, $Y := Y_i$,

$U := U_i$, $\mathcal{R} := \mathcal{R}_i$, and set

$$c := \frac{\mathcal{P}}{e^z x(i)^{\alpha(i)}} > 0, \qquad A := \frac{\mathcal{P}K}{x(i)}\,\mu_{i,1/\varepsilon}^{\frac{\varepsilon}{\varepsilon-1}} > 0,$$

so that (107)–(108) read $P = c/\mathrm{Y}(\overline{v})$ and $S = A\,U(\overline{v})/\mathrm{Y}(\overline{v})$. Under Assumption 4 the tail functions are continuously differentiable (Assumption 1), hence so are Y, U and $\mathcal{R}$ on $(\underline{v}_i, \infty)$.

(i) Step 1 (the pricing relation). At the node's symmetric equilibrium the firms post $p_i = P_i\mathcal{R}(\overline{v}_i)$ (Lemma 3(i)) and the threshold is the optimal one of Lemma 2(i) at $k = K$, so $\overline{v} \in (\underline{v}_i, \infty)$. Relation (96) is the first form of (74); dividing both sides by $\Gamma_i(\overline{v})\overline{v}^{-\frac{\alpha(i)}{1-\alpha(i)}}$ gives $P\,\mathcal{R}(\overline{v})\overline{v}^{\frac{\alpha(i)}{1-\alpha(i)}}/\Gamma_i(\overline{v}) = \mathcal{P}/(Z_i x(i)^{\alpha(i)})$, which is (107) with Y of (106); the second form of (106) is the second form of (64). The properties of Y are those of the left-hand side of (81), established in Step 1 of the proof of Proposition 3: continuity, $\mathrm{Y}(\underline{v}_i^+) = 0$ because $(\varepsilon-1)\mathcal{G}_1^- - 1 + G_i$ vanishes at $\underline{v}_i$ while $\mathcal{R}$ and $\mathcal{G}_\gamma^-$ are positive there, and, by Lemma 4(i), $\mathrm{Y}(\overline{v}) = \overline{v}^{\gamma-1}\mathcal{R}(\overline{v})/\Gamma_i(\overline{v}) \sim \overline{v}^{\gamma-1}/\Gamma(i) \to \infty$.

Step 2 (the sales relation). $S = P_i Y_i$ with $Y_i = \overline{y}_i\mathcal{R}(\overline{v})^\varepsilon/\overline{v}$ (Lemma 3(ii)), and by (21)–(22) $\overline{y}_i\mathcal{R}(\overline{v})^\varepsilon/\overline{v} = \overline{Y}\,U(\overline{v})$ with $\overline{Y} = \overline{y}_i\mu_{i,1/\varepsilon}^{\frac{\varepsilon}{\varepsilon-1}}$ and $\overline{y}_i = e^z K x(i)^{\alpha(i)-1}$: this is the first form of (108). Substituting $P = c/\mathrm{Y}(\overline{v})$,

$$S = \frac{\mathcal{P}}{e^z x(i)^{\alpha(i)}} \cdot e^z K x(i)^{\alpha(i)-1}\,\mu_{i,1/\varepsilon}^{\frac{\varepsilon}{\varepsilon-1}}\,\frac{U(\overline{v})}{\mathrm{Y}(\overline{v})} = A\,\frac{U(\overline{v})}{\mathrm{Y}(\overline{v})},$$

e^z cancelling and $x(i)^{\alpha(i)-1}/x(i)^{\alpha(i)} = 1/x(i)$: the second form. This proves (i).

(ii) It is the value-clearing identity (102) for node i, conditioned on the node's own state, with the cross-node terms replaced by their deterministic aggregate versions:

(1) Start from the first identity of (102) at a realisation: $S_i = \beta(i)P_i^{1-\eta}Q + \int_0^1 a(j,i)\,X_j\mathrm{d}j$.

(2) By Lemma 14, the inbound-demand integral is $\mathbb{P}$-a.s. equal to the deterministic $\mathcal{D}(t,i) = (\mathbb{A}^* m_X)(t,i)$, and $Q = Q(t)$ is deterministic. This step presumes what Definition 5 requires of the node maps and no more — joint measurability in (i, K, z) and integrability of X against the ensemble — and the equality it delivers is for a.e. i; where the root of (iii) is not yet known to be unique, "the" node map means a measurable selection, which (iv) then makes superfluous. Node i's own realisation perturbs neither object: i is a single point of a continuum (atomistic), so its own (K_i, z_i) has zero influence on the kernel integrals, which is why the *same* $\mathcal{D}(t,i)$ and $Q(t)$ appear for every realisation of the node state; in particular node i's own utilization moves its own bill X_i but not the bills of its buyers.

(3) What remains random in the identity is only the node-local triple $(P_i, \overline{v}_i, S_i)$ through (K_i, z_i): writing it state by state gives (109), one equation per (K, z), paired with the pricing and sales relations (107)–(108).

(iii) Substitute (i) into (109) and define on $(\underline{v}_i, \infty)$

$$F(\overline{v}) := A\,\frac{U(\overline{v})}{\mathrm{Y}(\overline{v})} - \beta(i)Q\,c^{1-\eta}\,\mathrm{Y}(\overline{v})^{\eta-1} - \mathcal{D}, \tag{112}$$

so that (110) is $F(\overline{v}) = 0$; F is continuous. *Near* $\underline{v}_i$*:* $\mathrm{Y}(\overline{v}) \to 0$ while $U(\overline{v}) \to U(\underline{v}_i) \in (0,1)$ (Definition 1), so the supply term diverges to $+\infty$ like Y^{-1}; the final-demand term is bounded if $\eta \geq 1$ and diverges like $\mathrm{Y}^{-(1-\eta)}$ if $\eta < 1$, more slowly than Y^{-1} because $0 < 1-\eta < 1$; hence $F(\overline{v}) \to +\infty$. *Near* ∞*:* $U(\overline{v}) \to 0$ and $\mathrm{Y}(\overline{v}) \to \infty$, so the supply term vanishes. If $\eta > 1$ the final-demand term tends to $-\infty$ and $F \to -\infty$. If $\eta = 1$, $F \to -\beta(i)Q - \mathcal{D} < 0$. If $\eta < 1$ the final-demand term vanishes and $F \to -\mathcal{D} \leq 0$: when $\mathcal{D} > 0$, $F < 0$ for large $\overline{v}$; when $\mathcal{D} = 0$, compare rates using Lemma 4(i), $U(\overline{v}) = \mathcal{R}(\overline{v})^{\varepsilon}/[\overline{v}\mu_{i,1/\varepsilon}^{\frac{\varepsilon}{\varepsilon-1}}] \sim \mu_{i,1/\varepsilon}^{-\frac{\varepsilon}{\varepsilon-1}}\overline{v}^{-1}$ and $\mathrm{Y}(\overline{v}) \sim \overline{v}^{\gamma-1}/\Gamma(i)$: the supply term decays like $\overline{v}^{-\gamma}$ and the final-demand term like $\overline{v}^{-(1-\eta)(\gamma-1)}$, and $\gamma > (1-\eta)(\gamma-1)$, so $F < 0$ for large $\overline{v}$. In every case F changes sign on $(\underline{v}_i, \infty)$ and has a root by the intermediate value theorem; at the root $P = c/\mathrm{Y} > 0$ and $S = AU/\mathrm{Y} > 0$. This proves (iii).

(iv) Step 1 (log-derivatives). Set

$$\ell_{\mathrm{Y}}(\overline{v}) := \frac{\mathrm{d}\ln \mathrm{Y}}{\mathrm{d}\ln \overline{v}}, \qquad \ell_U(\overline{v}) := -\frac{\mathrm{d}\ln U}{\mathrm{d}\ln \overline{v}}.$$

Under Assumption 5, $\ell_{\mathrm{Y}} \geq 0$ with $\ell_{\mathrm{Y}} > 0$ off a set without interior; and $\ell_U > 0$ everywhere, because $1/U(\overline{v}) = \mu_{i,1/\varepsilon}^{\frac{\varepsilon}{\varepsilon-1}}\,\overline{v}/\mathcal{R}(\overline{v})^{\varepsilon}$ is strictly increasing with a strictly positive derivative by Step 2 of the proof of Lemma 3 ($W' < 0$ there). Hence at fixed (K, z)

$$-\frac{\mathrm{d}\ln P}{\mathrm{d}\ln \overline{v}} = \ell_{\mathrm{Y}}, \qquad -\frac{\mathrm{d}\ln S}{\mathrm{d}\ln \overline{v}} = \ell_U + \ell_{\mathrm{Y}} > 0,$$

so P is nonincreasing and S strictly decreasing in the threshold; P is in fact strictly decreasing, since Y is strictly increasing. Wherever $\ell_{\mathrm{Y}} > 0$, the elasticity of sales with respect to the price along the supply side is $e_S = (\ell_U + \ell_{\mathrm{Y}})/\ell_{\mathrm{Y}} = 1 + \ell_U/\ell_{\mathrm{Y}} > 1$, which is (111). For the limit, write $\ln \mathrm{Y} = \ln \mathcal{R} + (\gamma-1)\ln \overline{v} - \ln \Gamma_i$ and $-\ln U = \ln \overline{v} - \varepsilon \ln \mathcal{R} + \text{const}$. From Step 1 of the proof of Lemma 3,

$$\frac{\mathrm{d}\ln \mathcal{R}^{\varepsilon-1}}{\mathrm{d}\ln \overline{v}} = \frac{(1-\frac{1}{\varepsilon})\,\overline{v}^{1-\frac{1}{\varepsilon}}\int_{\overline{v}}^{\infty} v^{1/\varepsilon}\mathrm{d}G_i(v)}{\mathcal{R}(\overline{v})^{\varepsilon-1}} \leq \frac{(1-\frac{1}{\varepsilon})\int_{\overline{v}}^{\infty} v\,\mathrm{d}G_i(v)}{\mathcal{R}(\overline{v})^{\varepsilon-1}} \to 0,$$

because $v^{1/\varepsilon}\overline{v}^{1-1/\varepsilon} \leq v$ on the upper tail, the tail of the finite first moment vanishes and $\mathcal{R} \to 1$; so $\ell_U \to 1$. By the second form of (64), $\ln \Gamma_i = \text{const} + \ln\left(\overline{v}^{\gamma}\mathcal{G}_{\gamma}^{-}\right) - \ln(\overline{v}D_i)$ with $D_i := (\varepsilon-1)\mathcal{G}_1^{-} - 1 + G_i$, and, writing g_i for the density,

$$\frac{\mathrm{d}\ln\left(\overline{v}^{\gamma}\mathcal{G}_{\gamma}^{-}\right)}{\mathrm{d}\ln \overline{v}} = \frac{\overline{v}^{\gamma+1}g_i(\overline{v})}{\int_0^{\overline{v}} v^{\gamma}\mathrm{d}G_i}, \qquad \frac{\mathrm{d}\ln(\overline{v}D_i)}{\mathrm{d}\ln \overline{v}} = \frac{\varepsilon\,\overline{v}^2 g_i(\overline{v}) - \overline{v}\,[1-G_i(\overline{v})]}{\overline{v}D_i(\overline{v})},$$

whose denominators tend to $\mu_{i,\gamma}$ and $\varepsilon - 1$ (Lemma 4(i)) and whose numerators tend to zero under the hypothesis $\overline{v}^{\gamma+1}g_i(\overline{v}) \to 0$ (which implies $\overline{v}^2 g_i \to 0$ since $\gamma > 1$) together with $\overline{v}[1-G_i] \to 0$. Hence $\mathrm{d}\ln\Gamma_i/\mathrm{d}\ln\overline{v} \to 0$, $\ell_{\mathrm{Y}} \to \gamma - 1 = \frac{\alpha(i)}{1-\alpha(i)}$, and $e_S \to 1 + \frac{1-\alpha(i)}{\alpha(i)} = 1/\alpha(i)$.

Step 2 (uniqueness). F is continuously differentiable and, by Step 1,

$$\overline{v}F'(\overline{v}) = -(\ell_U + \ell_{\mathrm{Y}})\,A\,\frac{U}{\mathrm{Y}} + (1-\eta)\,\ell_{\mathrm{Y}}\,\beta(i)Q\,c^{1-\eta}\mathrm{Y}^{\eta-1}.$$

At a root, $A\,U/\mathrm{Y} = \beta(i)Qc^{1-\eta}\mathrm{Y}^{\eta-1} + \mathcal{D}$, so

$$\overline{v}F'(\overline{v}) = -(\ell_U + \ell_\mathrm{Y})\,\mathcal{D} - \beta(i)Q\,c^{1-\eta}\mathrm{Y}^{\eta-1}[\ell_U + \eta\,\ell_\mathrm{Y}] < 0,$$

because $\ell_U > 0$, $\ell_\mathrm{Y} \geq 0$, $\eta > 0$ and $\beta(i)Qc^{1-\eta}\mathrm{Y}^{\eta-1} > 0$. A continuously differentiable function whose derivative is strictly negative at every zero can vanish only once: between two consecutive zeros it would have to cross from below at the second. Hence the root of (iii) is unique, for every $\eta > 0$, and (P, S) are unique through (i). What was used is only the displayed inequality at the root, namely $\beta(i)Qc^{1-\eta}\mathrm{Y}^{\eta-1}[\ell_U + \eta\ell_\mathrm{Y}] + \mathcal{D}[\ell_U + \ell_\mathrm{Y}] > 0$; at $\mathcal{D} = 0$ this is $\ell_U + \eta\ell_\mathrm{Y} > 0$ alone, and a condition uniform in the aggregates is $\ell_U + \min\{1, \eta\}\ell_\mathrm{Y} > 0$. Assumption 5 delivers all of these at once and, since the lemma is stated for every $\eta > 0$, is close to necessary at that generality; at a fixed η it is stronger than uniqueness requires.

Step 3 (comparative statics). F depends on (K, z) through $A \propto K$ and $c \propto e^{-z}$ only, continuously differentiably, with

$$\frac{\partial F}{\partial K} = \frac{A}{K}\frac{U}{\mathrm{Y}} = \frac{S}{K} > 0, \qquad \frac{\partial F}{\partial z} = (1-\eta)\,\beta(i)Q\,c^{1-\eta}\mathrm{Y}^{\eta-1},$$

the latter because $\partial_z c^{1-\eta} = -(1-\eta)c^{1-\eta}$ and the final-demand term enters F with a minus sign. Since $\partial_{\overline{v}}F < 0$ at the root (Step 2), the implicit function theorem makes $\overline{v}(t, i, K, z)$ continuously differentiable with $\partial_K\overline{v} = -\partial_K F/\partial_{\overline{v}}F > 0$ and $\operatorname{sign}\partial_z\overline{v} = \operatorname{sign}(1-\eta)$. utilization U is strictly decreasing in the threshold (Definition 1), so it moves the opposite way. This proves (iv).

(v) By (i), at a given threshold $P = c/\mathrm{Y} \propto e^{-z}$ and $S = AU/\mathrm{Y}$ does not involve z; in (112) the state variable z therefore appears only in $c^{1-\eta} \propto e^{-(1-\eta)z}$, which is identically one at $\eta = 1$. At $\eta = 1$ the root, and every function of it, is independent of z. ■

[Flag: a shorter route to (iii)–(iv) is available.] Multiplying (110) through by $\mathrm{Y}_i(\overline{v})$ turns it into $H(\overline{v}) := A\,U_i(\overline{v}) - \beta(i)Qc^{1-\eta}\mathrm{Y}_i(\overline{v})^\eta - \mathcal{D}\,\mathrm{Y}_i(\overline{v}) = 0$, in the notation of the proof. H is continuous with $H(\underline{v}_i^+) = A\,U_i(\underline{v}_i) > 0$ and $H(\overline{v}) \to -\infty$, which gives existence with no case split and no rate comparison; and under Assumption 5 H is strictly decreasing, which gives uniqueness for every $\eta > 0$ and every $\mathcal{D} \geq 0$ with no differentiability at all. The proof below is kept because its derivative F' at the root is what the comparative statics of (iv) and the elasticities of Proposition 6 are built from, and those do need the implicit function theorem; but a later session rewriting this block should take (iii) and the uniqueness half of (iv) from H and keep F only for the differentiable part.

Given the solution, node quantities follow mechanically:

$$\begin{aligned} Y_i &= \frac{S_i}{P_i} = \overline{Y}\,U_i(\overline{v}), \qquad X_i = \mathcal{P}(t,i)\,\frac{K}{x(i)}\,\Psi_{i,\frac{1}{1-\alpha(i)}}(\overline{v}) = \xi_i(\overline{v})\,S_i^{\mathrm{f}}, \\ M_i^e &= \frac{X_i}{\mathcal{P}(t,i)}, \qquad m_i(j) = a(i,j)\frac{X_i}{P_j}, \qquad Q_i = \beta(i)P_i^{-\eta}Q, \end{aligned} \tag{113}$$

with firm revenue $S_i^{\mathrm{f}} = S_i \Psi_{i,1}(\overline{v}) / [\mathcal{G}_{i,1}^{-}(\overline{v}) + \mathcal{G}_{i,1/\varepsilon}^{+}(\overline{v})]$ of (97). Operating profit of the node's firms (their revenue net of intermediate purchases; capital is owned, there is no labour), and the rationing rent of the node's bundlers, are

$$\pi_i = S_i^{\mathrm{f}} - X_i = [1 - \xi_i(\overline{v})]\, S^{\mathrm{f}}(t, i, K_i, z_i), \qquad \Pi^{\mathrm{r}}(t, i, K_i, z_i) = S_i - S_i^{\mathrm{f}} \geq 0, \tag{114}$$

so that the node's value added $S_i - X_i$ is the firms' operating profit plus the rent, and node utilization is $U(t, i, K_i, z_i) = U_i(\overline{v}(t, i, K_i, z_i))$. Without capacity constraints $S_i^{\mathrm{f}} = S_i$, the rent vanishes and operating profit is the constant share $1 - \xi(i)$ of sales (Appendix C.4).

Remark 6 (Structure of the coupling) The coupling map of Section 3.2 runs through two moment functions on $[0, 1]$ and one scalar. The two aggregates it builds from the moment functions have a direct reading:

- $m_P(t, i)$ is node i's mean log price across its own states, so $(\mathbb{A} m_P)(t, i)$ is the cost-share-weighted average of the mean log prices of i's *suppliers*: the cross-node part of i's marginal cost.
- $m_X(t, i)$ is node i's mean intermediate bill, so $(\mathbb{A}^* m_X)(t, i)$ is the kernel-weighted spending of i's *buyers*: its inbound order book. Without capacity constraints this is $(\mathbb{A}^* \Xi m_S)(t, i)$, the buyers' mean sales scaled by constant shares; under capacity the buyers' bills carry their utilization, through $\xi_j(\overline{v}) S^{\mathrm{f}} / S$.

Each node therefore needs to know two functions on $[0, 1]$ and one number, not the ensemble of node distributions, which is what makes the computation cheap (Appendix D); what capacity changes is which second function, the spending moment in place of the sales moment. Two ingredients deliver this: the exact law of large numbers, which makes the kernel integrals deterministic, and the Cobb–Douglas value shares of the bundle, which make each buyer's spending on node i a fixed share of a bill that depends on the buyer's state but not on node i's. Constancy of the expenditure share within nodes, which the model without capacity constraints has and the capacity model lacks, is not needed. It is also the static analogue of the coupling in graphon mean field games(Appendix B.9): there the HJB depends on the ensemble through graphon integrals of the node distributions; here through graphon integrals of two low-dimensional moments of $\mathfrak{M}(t, i, \cdot)$.

Remark 7 (GDP identity) Condition (46) is not an additional restriction but an identity implied by the other equilibrium conditions. Explicitly: integrate the node demand relation (109) first against $\mathfrak{M}(t, i, \cdot)$ and then over i. The left-hand side becomes $\int_0^1 m_S(t, i) \mathrm{d}i$, aggregate gross output. On the right, the final-demand term yields

$$Q(t) \int_0^1 \beta(i) \int P(t, i, K, z)^{1-\eta}\, \mathfrak{M}(t, i, \mathrm{d}K, \mathrm{d}z)\, \mathrm{d}i = Q(t),$$

by the numeraire condition (47), while the $\mathcal{D}$ term (already deterministic, so the $\mathfrak{M}$-integration is trivial) yields

$$\int_0^1 (\mathbb{A}^* m_X)(t, i) \mathrm{d}i = \int_0^1 m_X(t, j) \underbrace{\left(\int_0^1 a(j, i) \mathrm{d}i \right)}_{=1} \mathrm{d}j = \int_0^1 m_X(t, j) \mathrm{d}j,$$

by Tonelli and row-stochasticity (42). Collecting terms,

$$\int_0^1 m_S(t,i)\mathrm{d}i = Q(t) + \int_0^1 m_X(t,j)\mathrm{d}j \iff Q(t) = \int_0^1 [m_S(t,j) - m_X(t,j)]\mathrm{d}j,$$

which is (46): aggregate final expenditure (GDP at the numeraire price) equals gross output minus intermediate use, i.e. aggregate value added. Because the derivation consumes the numeraire condition, (46) carries no independent information within Definition 5, but it is worth keeping explicit as a Walras check in computation. The identity also bounds gross output by GDP: under Assumption 3 every node spends less than the share $1 - \frac{1}{\varepsilon}$ of its sales (Remark 4), so $\int m_X \le (1 - \frac{1}{\varepsilon}) \int m_S$ and

$$Q(t) \le \int_0^1 m_S(t,j)\mathrm{d}j \le \varepsilon\, Q(t),$$

the continuum analogue of the bounds on the Domar density in Lemma 9(iii) with $\bar{\xi} = 1 - \frac{1}{\varepsilon}$; without capacity constraints the same argument with the constant shares gives the sharper $\int m_S \le (1 - \bar{\xi})^{-1} Q$ with $\bar{\xi} = \sup_i \xi(i)$ (Appendix C.4).

The individual firm's profit. The object that drives investment in Section 3.3 is the *individual monopolist's* marginal profit. Within a node each monopolist is atomistic in its own sector; its perceived profit as a function of its own capital k, holding the node's symmetric state (K, z) and all aggregates fixed, is derived exactly as in the finite model (Lemma 6), with the node's parameters $(\alpha(i), x(i), G_i)$ in place of the sector's and the node objects $(Y_i, P_i, \mathcal{P}(t,i))$ of the static block in place of the sectoral aggregates:

Lemma 16 (Individual profit function under capacity). *Consider a monopolist at node i with own capital k, when the node is at (K, z) with price $P(t,i,K,z)$, threshold $\overline{v} := \overline{v}(t,i,K,z)$ and sales $S(t,i,K,z)$ of Lemma 15, all its competitors at the symmetric point, and let Assumption 4 hold.*

(i) (Profit at own k.) *Its optimally priced expected operating profit $\pi(t,i,k;K,z)$ is* (85) *evaluated at its own threshold $\overline{v}_i(k)$, the unique root in $(\underline{v}_i, \infty)$ of* (66) *at own capital k, a continuous strictly increasing bijection of $(0,\infty)$ onto $(\underline{v}_i, \infty)$ with $\overline{v}_i(K) = \overline{v}$; it has no closed form, and at the symmetric point $k = K$ it equals the node's operating profit $[1 - \xi_i(\overline{v})]S^{\mathrm{f}}(t,i,K,z)$ of* (114).

(ii) (Marginal profit at the symmetric point.) *At $k = K$,*

$$\left.\frac{\partial \pi}{\partial k}\right|_{k=K} = \zeta_i(\overline{v})[1 - \xi_i(\overline{v})]\,\frac{S^{\mathrm{f}}(t,i,K,z)}{K} = \frac{\varepsilon}{1-\alpha(i)}\,\frac{\mathcal{P}(t,i)}{x(i)}\,\Phi_i(\overline{v}), \tag{115}$$

with the capital share $\zeta_i[1 - \xi_i]$ of (33) *and Φ_i of* (84)*: given $\mathcal{P}(t,i)$, the marginal profit depends on the node's state only through its threshold. In the no-capacity limit it is $\frac{\varepsilon-1}{\varepsilon}\alpha(i)S(t,i,K,z)/K$ (Appendix C.4).*

(iii) (Concavity.) *$\pi(t,i,\cdot;K,z)$ is strictly concave on $(0,\infty)$ if and only if Assumption 3 holds at node i; in that case its derivative decreases strictly from $+\infty$ to 0, so the marginal profit* (115) *is strictly positive.*

Proof. The monopolist's pricing problem at node i is that of Lemma 2 with sector parameters $(\alpha(i), x(i), G_i)$, productivity $Z_i = e^z$ and the node objects $(Y_i, P_i, \mathcal{P}(t,i))$ of Lemma 15 in place of the sectoral aggregates, all of which the atomistic firm takes as given; its competitors are at the symmetric point, so the node's price, output and threshold are $P(t,i,K,z)$, S/P and $\overline{v}(t,i,K,z)$. Nothing in Steps 1–3 of the proof of Lemma 6 uses the finiteness of the set of sectors or the finite-sector form of the aggregates: they treat $(Y_i, P_i, \mathcal{P}_i, Z_i)$ as fixed numbers. Those steps therefore apply verbatim and give (i) and (iii), and the two forms of $k\pi'(k)$ in (86) at every k. At $k = K$: the own threshold is the node's, $\overline{v}_i(K) = \overline{v}(t,i,K,z)$, because both are the unique root of (66) at the symmetric capacity; the price is $p_i(K) = P_i\mathcal{R}_i(\overline{v})$ (Lemma 3(i)); $p\mathbb{E}_v[y] = S^{\mathrm{f}}(t,i,K,z)$ by (97); and $\Pi = S^{\mathrm{f}} - X = [1 - \xi_i(\overline{v})]S^{\mathrm{f}}$ by (98), which is (114). Dividing the first form of (86) by K gives the first form of (115), and the second form is the second form of (86) with $\mathcal{P}_i = \mathcal{P}(t,i)$ and $x_i = x(i)$; since Φ_i is a function of the threshold alone, so is the marginal profit given $\mathcal{P}(t,i)$. The no-capacity limit is Lemma 5(v) together with $S^{\mathrm{f}}/S \to 1$ (Lemma 4(iii)). This proves (ii); the positivity in (iii) is Lemma 6(iii). ∎

Without capacity constraints the individual profit function is exactly a power of own capital, $\pi = [1 - \xi(i)]S(t,i,K,z)(k/K)^{\zeta(i)}$ with the constant elasticity $\zeta(i) \in (0,1)$, and the marginal profit at the symmetric point is the constant fraction $\frac{\varepsilon-1}{\varepsilon}\alpha(i)$ of sales per unit of capital (Lemma 25). Under capacity the elasticity $\zeta_i(\overline{v})$ moves with the node's threshold, from $1 - \frac{1}{\varepsilon}$ where capital is scarce and the node's firms are rationed often to the no-capacity value $\zeta_i(\infty)$ where capacity is slack (Lemma 6), and the flow that prices capital in Section 3.3 is the first form of (115): the capital share of firm revenue, not of sales, per unit of node capital. The second form makes the dependence on the state transparent: by (108) sales per unit of capital are themselves a function of the threshold given $\mathcal{P}(t,i)$, so the entire private return to capital at a node is the function Φ_i of the node's operating position, scaled by the bundle price and the capital intensity; the two channels through which capital raises profit, relaxing the rationing in the constrained states and saving inputs in the unconstrained ones, are those of (76).

Node utilization and its moments. Node utilization $U(t,i,K,z) = U_i(\overline{v}(t,i,K,z)) \in (0, U_i(\underline{v}_i))$ is a function of the node's own state, random across realisations with the node distribution; so are the node's rationed share of firms $1 - G_i(\overline{v}(t,i,K,z))$ and its mean physical utilization $\Psi_{i,1}(\overline{v}(t,i,K,z))$ (Definition 1). Its node mean, and the value-added-weighted aggregate utilization of (23) in moment form, are

$$\begin{aligned} \mathbb{E}_i[U](t) &:= \int U(t,i,K,z)\,\mathfrak{M}(t,i,\mathrm{d}K,\mathrm{d}z), \\ \overline{U}(t) &:= \frac{\int_0^1 \mathbb{E}_i[(S-X)\,U]\,(t)\,\mathrm{d}i}{\int_0^1 [m_S(t,i) - m_X(t,i)]\mathrm{d}i} = \frac{1}{Q(t)}\int_0^1 \mathbb{E}_i[(S-X)\,U]\,(t)\,\mathrm{d}i, \end{aligned} \tag{116}$$

with the value-added weight inside the node expectation, because a node's value added and its utilization move together with its state. By Lemma 13 (the integrand is

bounded by S, which is integrable under Lemma 14), realised aggregate utilization coincides $\mathbb{P}$-a.s. with (116): like every cross-sectional aggregate it is deterministic, although every node's utilization is random. The sales-weighted and capacity-value-weighted alternatives of Definition 1 become

$$\overline{U}^{S}(t) = \frac{\int_0^1 \mathbb{E}_i[S\,U](t)\,\mathrm{d}i}{\int_0^1 m_S(t,i)\,\mathrm{d}i}, \qquad \overline{U}^{\overline{Y}}(t) = \frac{\int_0^1 m_S(t,i)\,\mathrm{d}i}{\int_0^1 \mathbb{E}_i[P\,\overline{Y}](t)\,\mathrm{d}i}.$$

The dispersion of utilization across realisations within a node is a cross-sectional object with no counterpart in the model without capacity constraints, where utilization vanishes in the limit (Lemma 4); the step-kernel model of Section 3.5 carries it block by block.

[Q15 | text | ANSWERED] Wrap the subsection 'sec:degeneracy' and all places where it is cited/referenced in notes box until I decide what to do with it later.

B.7 The Cobb–Douglas Degeneracy under Capacity

[Flag: set aside on 5 November 2026] This subsection is held, unchanged and with its heading and label, inside a maroon box until what to do with it is decided. The question that asks for it is moved above the heading so that it stays outside the box. Every pointer from black text to the subsection or to its results (Theorem 1, Lemma 17, Propositions 4, 5 and 6, Remarks 8 and 9, equations (118), (127), (121) and (131)) is in a maroon box too, and the claim it supported is left black. The pointers sit in the setup and the comparison sections of the continuum model, in the appendices on step kernels and graphon limits, on the primitives, on the model on $[0,1]$, on the static block, on the dynamics and on balanced growth, in Steps 1–6 of the numerical appendix, and in the calibration appendix (the share ϱ_n and the choice of η). The share $\varrho_i(\overline{v}) := S^{\mathrm{f}}/S$ of (118), which black text keeps using, is restated in words where the dynamics subsection first uses it.

Throughout, "Cobb–Douglas" (CD) means the $\eta = 1$ final aggregator, the $\eta \to 1$ limit of (99),

$$Q = \exp\left\{\int_0^1 \beta(i)\,\ln\frac{Q_i}{\beta(i)}\,\mathrm{d}i\right\}, \tag{117}$$

for which nominal final spending on every node is $P_iQ_i = \beta(i)Q$. Two distinct degeneracies follow. The *payoff-side* degeneracy is that no node's cash flow, marginal value or investment responds to its productivity; it is what CES final demand repairs. The *distributional* degeneracy is that node capital is deterministic and the node distribution a Dirac mass in the capital coordinate, so that no fixed point over nondegenerate measures remains. Capacity repairs neither. With every buyer spending a fixed nominal amount on the node, the node's threshold, and with it its utilization, its shares and its cash flow, are free of its productivity (Lemma 15(v)): the operating position of a node responds to its state through its capital alone. What capacity removes is the closed form. The Domar density, log-linear price system and algebraic balanced growth path

of the model without capacity constraints (Appendix C.5) are replaced by a scalar equation per node in the threshold and by a moment fixed point over the cross-section that the cancellation no longer eliminates, so that the Cobb–Douglas model under capacity is degenerate in productivity but not, as the model without capacity constraints is, free of mean-field content in payoffs.

Appendix B.7.1 establishes the cancellation and states what remains of the CD static block; Appendix B.7.2 states investment neutrality and the degenerate balanced growth path; Appendix B.7.3 explains what the shocks still do and why the finite model does not degenerate; Appendix B.7.4 shows that CES restores own-state relevance. The dynamic statements cite the marginal-$\mathfrak{q}$ equation and the detrended system of Section 3.3, which are written for general η and specialise here by substituting the state-independent sales of Lemma 17. Throughout this appendix

$$\varrho_i(\overline{v}) := \frac{S^{\mathrm{f}}}{S} = \frac{\Psi_{i,1}(\overline{v})}{\mathcal{G}^{-}_{i,1}(\overline{v}) + \mathcal{G}^{+}_{i,1/\varepsilon}(\overline{v})} \in (0,1] \tag{118}$$

denotes the share of node sales that accrues to the node's firms (Lemma 5(iv)), the complement $1-\varrho_i$ being the rationing rent's share; ϱ_i is continuous on $(0,\infty)$ and tends to 1 as $\overline{v} \to \infty$ (Lemma 4(i)).

B.7.1 Statics: Cancellation and the Static Block

Lemma 17 (Cancellation lemma under capacity). *Let $\eta = 1$ and let Assumptions 4 and 5 hold. Then in any static equilibrium, at every node i:*

(i) (Sales.) *The sales function is state-independent,*

$$S(t,i,K,z) = \beta(i)\, Q(t) + \mathcal{D}(t,i) =: S(t,i) \qquad \text{for all } (K,z), \tag{119}$$

so that $m_S(t,i) = S(t,i)$ and, by (45), $S = \beta Q + \mathbb{A}^* m_X$. Under Assumption 3, $\beta(i)Q(t) \le S(t,i) \le \left[(\mathrm{Id} - (1-\frac{1}{\varepsilon})\mathbb{A}^*)^{-1}\beta\right](i)\, Q(t)$.

(ii) (Threshold.) *The node's threshold is the unique root in $(\underline{v}_i, \infty)$ of*

$$\frac{\mathcal{P}(t,i)\,K}{x(i)}\, \mu_{i,1/\varepsilon}^{\frac{\varepsilon}{\varepsilon-1}}\, \frac{U_i(\overline{v})}{\mathrm{Y}_i(\overline{v})} = S(t,i), \tag{120}$$

which contains K but not z: $\overline{v} = \overline{v}(t,i,K)$ is continuously differentiable and strictly increasing in K, from $\underline{v}_i$ as $K \downarrow 0$ to ∞ as $K \to \infty$. Consequently node utilization $U_i(\overline{v}(t,i,K))$, strictly decreasing in K, the firms' revenue $S^{\mathrm{f}} = \varrho_i(\overline{v})S(t,i)$, their spending $X = \xi_i(\overline{v})\varrho_i(\overline{v})S(t,i)$ and operating profit $[1-\xi_i(\overline{v})]\varrho_i(\overline{v})S(t,i)$, the rationing rent $[1-\varrho_i(\overline{v})]S(t,i)$ and the marginal profit (115) are functions of (t,i,K) alone: productivity drops out of every node cash flow.

(iii) (Prices absorb productivity.) *The node price is*

$$\ln P(t,i,K,z) = -z + \ln \mathcal{P}(t,i) - \alpha(i) \ln x(i) - \ln \mathrm{Y}_i(\overline{v}(t,i,K)), \tag{121}$$

loading on z with coefficient -1 and strictly decreasing in K; real output $Y = S/P$ and real capacity $\overline{Y}$ both scale with e^z, so that utilization does not.

(iv) (Marginal profit.) *The marginal profit* (115) *equals*

$$\begin{aligned}\pi_k(t,i,K) &= \frac{\varepsilon}{1-\alpha(i)}\,\frac{\mathcal{P}(t,i)}{x(i)}\,\Phi_i(\overline{v}(t,i,K)) \\ &= \zeta_i(\overline{v})[1-\xi_i(\overline{v})]\,\varrho_i(\overline{v})\,\frac{S(t,i)}{K}, \qquad \overline{v}=\overline{v}(t,i,K).\end{aligned} \tag{122}$$

Under Assumption 3 it is strictly decreasing in K, with $K\pi_k \to (1-\frac{1}{\varepsilon})\varrho_i(\underline{v}_i)S(t,i)$ as $K \downarrow 0$ and $K\pi_k \to \frac{\varepsilon-1}{\varepsilon}\alpha(i)S(t,i)$ as $K \to \infty$.

Proof. Drop t where no confusion arises and write $\gamma := \frac{1}{1-\alpha(i)}$.

Step 1: the own price drops out of demand. Set $\eta = 1$ in (109): final spending becomes $\beta(i)P_i^{1-\eta}Q = \beta(i)Q$, so

$$S(t,i,K,z) = \beta(i)\,Q(t) + \mathcal{D}(t,i).$$

Both right-hand terms are deterministic aggregates (Lemma 14) and neither involves (K,z): the node's own price (the only channel through which its state entered its own demand) has vanished. Hence $S(t,i,K,z)$ is the same number at every state; write it $S(t,i)$. This is (119).

Step 2: moments and bounds. Since $S(t,i,\cdot,\cdot)$ is constant in the state, $m_S(t,i) = \int S\,\mathrm{d}\mathfrak{M}(t,i,\cdot) = S(t,i)$, and substituting $\mathcal{D} = \mathbb{A}^* m_X$ from (45) gives $S = \beta Q + \mathbb{A}^* m_X$. The lower bound $S \geq \beta Q$ is $\mathcal{D} \geq 0$. For the upper bound, under Assumption 3 every node spends less than the share $\bar{\xi} := 1 - \frac{1}{\varepsilon}$ of its sales at every state, $X = \xi_j(\overline{v})S^{\mathrm{f}} \leq \xi_j(\overline{v})S < \bar{\xi}S$ (Lemma 6(iii) and $S^{\mathrm{f}} \leq S$), so $m_X \leq \bar{\xi}m_S = \bar{\xi}S$ and, $\mathbb{A}^*$ being a positive operator, $S \leq \beta Q + \bar{\xi}\mathbb{A}^*S$, i.e. $(\mathrm{Id} - \bar{\xi}\mathbb{A}^*)S \leq \beta Q$. By Lemma 9(i)–(ii) with the constant share $\bar{\xi} < 1$, $\mathrm{Id} - \bar{\xi}\mathbb{A}^*$ is invertible on L^1 with the Neumann inverse $\sum_{n\geq 0}\bar{\xi}^n(\mathbb{A}^*)^n$, a positive operator; applying it to the inequality preserves it and gives $S \leq (\mathrm{Id} - \bar{\xi}\mathbb{A}^*)^{-1}\beta\,Q$.

Step 3: the threshold. Substituting Step 1 into the scalar equation (110) at $\eta = 1$ gives (120), in which z does not appear (Lemma 15(v)); in the notation of the proof of Lemma 15, $F(\overline{v}) = A\,U(\overline{v})/\mathrm{Y}(\overline{v}) - S(t,i)$ with $A \propto K$. Under Assumption 5, U/Y is strictly decreasing on $(\underline{v}_i, \infty)$ (both $1/U$ and Y are strictly increasing), from $+\infty$ at $\underline{v}_i$ ($\mathrm{Y} \to 0$, $U \to U(\underline{v}_i) > 0$) to 0 at infinity, so the root is unique; this is Lemma 15(iii)–(iv) specialised to $\eta = 1$, and $\overline{v}$ is continuously differentiable and strictly increasing in K by the implicit function theorem there ($\partial_K F = S/K > 0$, $\partial_{\overline{v}} F < 0$). As $K \downarrow 0$, $A \to 0$ and the root must satisfy $U/\mathrm{Y} = S/A \to \infty$, so $\overline{v} \to \underline{v}_i$; as $K \to \infty$, $U/\mathrm{Y} \to 0$, so $\overline{v} \to \infty$. The remaining objects follow from (113)–(114) and Lemma 5(iv): $S^{\mathrm{f}} = \varrho_i(\overline{v})S$, $X = \xi_i(\overline{v})S^{\mathrm{f}}$, operating profit $S^{\mathrm{f}} - X$, rent $S - S^{\mathrm{f}}$, and the marginal profit (115), each a function of (t,i,K) because $S(t,i)$ and $\overline{v}(t,i,K)$ are. utilization $U_i(\overline{v}(t,i,K))$ is strictly decreasing in K because U_i is strictly decreasing (Definition 1).

Step 4: prices. Taking logarithms in (107) at the root gives (121); the coefficient of z is -1 because z enters (107) only through e^{-z} and the threshold is z-free, and P is strictly decreasing in K because Y_i is strictly increasing and the threshold is. Real

output $Y = S/P \propto e^z$ at fixed K, real capacity $\overline{Y} = e^z K x(i)^{\alpha(i)-1} \mu_{i,1/\varepsilon}^{\varepsilon/(\varepsilon-1)} \propto e^z$, and their ratio U is z-free.

Step 5: the marginal profit. The two forms of (122) are the two forms of (115) with $S^{\mathrm{f}} = \varrho_i(\overline{v})S(t,i)$. Under Assumption 3, Φ_i is strictly decreasing on $(\underline{v}_i, \infty)$ and $\overline{v}(t,i,K)$ is strictly increasing in K, so the first form is strictly decreasing in K. For the limits use the second form, $K\pi_k = \zeta_i(\overline{v})[1-\xi_i(\overline{v})]\varrho_i(\overline{v})S(t,i)$, with $S(t,i)$ fixed. As $K \downarrow 0$, $\overline{v} \to \underline{v}_i$, where $\varepsilon\eta_{i,1}(\underline{v}_i) = 1$ (the definition of $\underline{v}_i$ as the zero of $(\varepsilon-1)\mathcal{G}_{i,1}^- - 1 + G_i = \Psi_{i,1}[\varepsilon\eta_{i,1} - 1]$), so in (33) the second term vanishes and $\zeta_i[1-\xi_i] \to 1 - \eta_{i,1}(\underline{v}_i) = 1 - \frac{1}{\varepsilon}$ (equivalently $\xi_i(\underline{v}_i) = 0$ in (75)), while $\varrho_i(\overline{v}) \to \varrho_i(\underline{v}_i)$ by continuity. As $K \to \infty$, $\overline{v} \to \infty$, $\zeta_i[1-\xi_i] \to \frac{\varepsilon-1}{\varepsilon}\alpha(i)$ (Lemma 5(v)) and $\varrho_i \to 1$ (Lemma 4(i)). This proves (iv). ∎

Without capacity constraints the cancellation goes further: the demand relation iterates to $S = hQ$ with the Domar density $h = (\mathrm{Id} - \mathbb{A}^*\Xi)^{-1}\beta$, cash flows are $[1-\xi(i)]h(i)Q$, constant in K as well as in z, and the marginal profit is that constant divided by the stock (Lemma 26). Under capacity nominal sales are still pinned at the node, but the fixed euro of sales is split between materials, capital income, markup rent and rationing rent in proportions that move with the node's capital: a better-capitalised node has slacker capacity, its firms are rationed less often, they spend a larger share of their revenue on inputs, and the rationing rent shrinks. The marginal profit divides a state-dependent cash flow by the stock and falls with capital for two reasons, the $1/K$ of the model without capacity constraints and the falling capital share; its two ends are fixed by the tail functions. Near $K = 0$ the node's firms are rationed in almost every state, the materials share vanishes, and the whole of firm revenue net of the markup rent accrues to capital, $K\pi_k \to (1-\frac{1}{\varepsilon})\varrho_i(\underline{v}_i)S$; at large K the constant $\frac{\varepsilon-1}{\varepsilon}\alpha(i)S$ of the model without capacity constraints is recovered. Two further differences from the model without capacity constraints matter below. First, sales are pinned by the buyers' *spending* moment, not by primitives: $\mathcal{D} = \mathbb{A}^* m_X$ requires the buyers' utilization, so (119) is a fixed point over the cross-section, bounded above by the Domar density at the uniform share $1 - \frac{1}{\varepsilon}$. Second, the bundle price enters the threshold. At fixed nominal sales, dearer inputs raise the node's price, lower its real output against an unchanged real capacity, and lower its utilization: in (120) the node's operating position is set by the bundle bill at full capacity per euro of sales, $\mathcal{P}(t,i)K/(x(i)S(t,i))$. This is where the price moment, which never touched a payoff without capacity constraints, re-enters the node's problem (Corollary 1).

The rest of the static block is no longer log-linear. We state it in the detrended units (51) of Section 3.3 ($\widehat{s} = S/\overline{K}_t$, $q = Q/\overline{K}_t$, $\kappa = K/\overline{K}_t$, $\widehat{\mathcal{D}} = \mathcal{D}/\overline{K}_t$, $\widehat{m_X} = m_X/\overline{K}_t$; prices, thresholds and utilization need no detrending), in which (119) reads $\widehat{s}(i) = \beta(i)q + \widehat{\mathcal{D}}(i)$ and (120) has κ in place of K, and with $\mathbb{E}_i[\cdot]$ the expectation against the node distribution $\mathfrak{M}(t,i,\cdot)$ as in Section 3.2. The threshold at (i,κ) is written $\overline{v}(i,\kappa)$; it depends on the aggregates through $\mathcal{P}(i)$ and $\widehat{s}(i)$.

Proposition 4 (The Cobb–Douglas static block under capacity). *Let $\eta = 1$, let Assumptions 4, 5 and 3 hold, and suppose $\mathbb{E}_i[\kappa]$, $\mathbb{E}_i[|\ln\kappa|]$ and $\mathbb{E}_i[|z|]$ are finite, uniformly in i.*

Then:

(i) (Sales and inbound demand.) *With the* effective spending share *of node i,*

$$\chi(i) := \mathbb{E}_i[\xi_i(\overline{v}(i,\kappa))\,\varrho_i(\overline{v}(i,\kappa))] \in \Big(0, 1-\tfrac{1}{\varepsilon}\Big), \tag{123}$$

the spending moment is $\widehat{m_X} = \chi\,\widehat{s}$ and sales solve

$$\widehat{s} = \beta\, q + \mathbb{A}^*[\chi\,\widehat{s}] \iff \widehat{s} = (\mathrm{Id} - \mathbb{A}^*\mathbb{D}_\chi)^{-1}\beta\, q, \tag{124}$$

with $\mathbb{D}_\chi$ pointwise multiplication by χ: given the profile of effective shares, sales are Domar-type with χ in place of the constant ξ, and $\beta q \le \widehat{s} \le (\mathrm{Id} - (1-\frac{1}{\varepsilon})\mathbb{A}^*)^{-1}\beta\, q$. The GDP line $q = \int_0^1 [\widehat{s} - \widehat{m_X}]\mathrm{d}i$ holds identically.

(ii) (Price moment.) *The moment $m_P(i) = \mathbb{E}_i[\ln P]$ satisfies*

$$m_P(i) - (\mathbb{A}m_P)(i) = -\mathbb{E}_i[z] - \alpha(i)\ln x(i) - \int_0^1 a(i,j)\ln a(i,j)\mathrm{d}j - \mathbb{E}_i[\ln \mathrm{Y}_i(\overline{v}(i,\kappa))], \tag{125}$$

in which the threshold depends on $\mathcal{P}(i) = \exp\{(\mathbb{A}m_P)(i) - \int a(i,j)\ln a(i,j)\mathrm{d}j\}$ through (120)*. At fixed aggregates and $\widehat{s}(i)$, the node's log price loads on its own log capital with coefficient $-\ell_\mathrm{Y}/(\ell_U + \ell_\mathrm{Y})$, and on the log of its own bundle price with coefficient $\ell_U/(\ell_U + \ell_\mathrm{Y}) \in (0,1]$, where $\ell_\mathrm{Y} \ge 0$ and $\ell_U > 0$ are the log-derivatives of Y_i and $1/U_i$ at the threshold (proof of Lemma 15); wherever $\ell_\mathrm{Y} > 0$ these are $-1/e_S(\overline{v})$ and $1 - 1/e_S(\overline{v})$ in terms of the supply elasticity* (111)*, which are $-\alpha(i)$ and $1-\alpha(i)$ without capacity constraints.*

(iii) (Numeraire.) *The normalisation $P^f \equiv 1$ reads*

$$\int_0^1 \beta(i)\, m_P(i)\,\mathrm{d}i = 0, \tag{126}$$

and is the $\eta \to 1$ limit of the CES numeraire (47)*.*

(iv) (No reduced map.) *Only without capacity constraints, where $e_S \equiv 1/\alpha(i)$ and $\chi \equiv \xi(i)$, is* (125) *linear in m_P with the node distributions entering through $\mathbb{E}_i[\ln\kappa]$ alone and* (124) *the Domar form with constant shares; under capacity the ensemble enters the static block through the two moment functions $\chi(i) = \mathbb{E}_i[\xi_i\varrho_i]$ and $\mathbb{E}_i[\ln \mathrm{Y}_i(\overline{v}(i,\kappa))]$ of the capital distributions, besides $\mathbb{E}_i[z]$, and no log-linear aggregate of the form of the reduced map of the model without capacity constraints exists.*

Proof. Step 1: sales. By Lemma 17(i) in detrended form, $\widehat{s}(i) = \beta(i)q + \widehat{\mathcal{D}}(i)$ is constant in the state, so $\widehat{m_X}(i) = \mathbb{E}_i[X] = \mathbb{E}_i[\xi_i(\overline{v}(i,\kappa))\varrho_i(\overline{v}(i,\kappa))]\,\widehat{s}(i) = \chi(i)\widehat{s}(i)$, with $\chi(i) \in (0, 1-\frac{1}{\varepsilon})$ because $0 < \xi_i < 1 - \frac{1}{\varepsilon}$ (Lemmas 5(i) and 6(iii)) and $0 < \varrho_i \le 1$. Substituting $\widehat{\mathcal{D}} = \mathbb{A}^*\widehat{m_X} = \mathbb{A}^*[\chi\widehat{s}]$ gives the first form of (124). The proof of Lemma 9(i)–(ii) uses only that the share profile is bounded by a constant below one, so it applies verbatim with ξ replaced by χ: $\|\mathbb{A}^*\mathbb{D}_\chi\|_{1\to 1} \le \sup_i \chi(i) < 1 - \frac{1}{\varepsilon}$, $\mathrm{Id} - \mathbb{A}^*\mathbb{D}_\chi$ is invertible on L^1 with a positive Neumann inverse, and the second form follows; the bounds are Lemma 17(i). The GDP line: integrating $\widehat{s} = \beta q + \mathbb{A}^*\widehat{m_X}$ over i and using $\int\beta = 1$ and

$\int_0^1 (\mathbb{A}^* \widehat{m_X})(i)\mathrm{d}i = \int_0^1 \widehat{m_X}(j)\mathrm{d}j$ (Tonelli and row-stochasticity, as in Remark 7) gives $\int \widehat{s} = q + \int \widehat{m_X}$.

Step 2: the price moment. Take $\mathbb{E}_i$ of (121), which is finite by the hypotheses on $\mathbb{E}_i[|z|]$ and $\mathbb{E}_i[|\ln \kappa|]$: the remaining state-dependent term is $\mathbb{E}_i[\ln \mathrm{Y}_i(\overline{v}(i,\kappa))]$, and by Step 4 below $\ln \mathrm{Y}_i(\overline{v}(i,\kappa))$ has slope $1/e_S(\overline{v}) \in (0,1]$ in $\ln \kappa$, so it grows at most linearly in $|\ln \kappa|$. Then substitute the moment form of the bundle price, $\ln \mathcal{P}(i) = (\mathbb{A} m_P)(i) - \int a(i,j) \ln a(i,j)\mathrm{d}j$ from (44); this is (125). For the loadings, fix the aggregates and $\widehat{s}(i)$ and differentiate (121) along (120) in detrended form, $\ln \kappa + \ln \mathcal{P}(i) + \ln U_i(\overline{v}) - \ln \mathrm{Y}_i(\overline{v}) = \text{const}$, using the log-derivatives ℓ_U, ℓ_Y of the proof of Lemma 15(iv): $\mathrm{d}\ln \kappa + \mathrm{d}\ln \mathcal{P} = (\ell_U + \ell_\mathrm{Y})\mathrm{d}\ln \overline{v}$, so

$$\left.\frac{\partial \ln P}{\partial \ln \kappa}\right|_{\mathcal{P},\widehat{s}} = -\ell_\mathrm{Y} \frac{\partial \ln \overline{v}}{\partial \ln \kappa} = \frac{-\ell_\mathrm{Y}}{\ell_U + \ell_\mathrm{Y}}, \qquad \left.\frac{\partial \ln P}{\partial \ln \mathcal{P}}\right|_{\kappa,\widehat{s}} = 1 - \frac{\ell_\mathrm{Y}}{\ell_U + \ell_\mathrm{Y}} = \frac{\ell_U}{\ell_U + \ell_\mathrm{Y}},$$

which are $-1/e_S$ and $1 - 1/e_S$ wherever $\ell_\mathrm{Y} > 0$, by (111); since $\ell_U > 0$ and $\ell_\mathrm{Y} \geq 0$, the capital loading lies in $(-1, 0]$ and the bundle-price loading in $(0, 1]$, both endpoints 0 and 1 being attained exactly where $\ell_\mathrm{Y} = 0$, which Assumption 5 permits. In the no-capacity limit $\ell_U \to 1$ and $\ell_\mathrm{Y} \to \frac{\alpha(i)}{1-\alpha(i)}$ (proof of Lemma 15(iv)), giving $-\alpha(i)$ and $1 - \alpha(i)$.

Step 3: the numeraire. Cost minimisation for (117) (choose Q_i to minimise $\int P_i Q_i \mathrm{d}i$ at given Q) has first-order condition $P_i Q_i = \Lambda \beta(i)$, total cost Λ, and unit cost $P^f = \exp \int_0^1 \beta(i) \ln P_i \mathrm{d}i$: the log price index of a Cobb–Douglas aggregator is the β-weighted mean log price. By the exact law of large numbers (the argument of Lemma 14 with $\psi(j,\omega) = \beta(j) \ln P(t, j, s_j)$, integrable by the first hypothesis of (105)), the realised integral equals $\int \beta\, m_P \mathrm{d}i$ a.s., so $P^f = 1$ is (126). That it is also the $\eta \to 1$ limit of (47) follows from expanding $P^{1-\eta} = 1 + (1-\eta)\ln P + O((1-\eta)^2)$ inside it (details in the derivation after this proof).

Step 4: no reduced map. In (125) the capital distribution enters through $\mathbb{E}_i[\ln \mathrm{Y}_i(\overline{v}(i,\kappa))]$, and by Step 2 the function $\ln \kappa \mapsto \ln \mathrm{Y}_i(\overline{v}(i,\kappa))$ has slope $\ell_\mathrm{Y}/(\ell_U + \ell_\mathrm{Y}) = 1/e_S(\overline{v})$, which varies along the threshold axis unless e_S is constant; its expectation is therefore not a function of $\mathbb{E}_i[\ln \kappa]$ alone. Likewise the threshold, hence $\ln \mathrm{Y}_i$, depends on $\ln \mathcal{P}(i)$ with the state-dependent slope $\ell_U/(\ell_U + \ell_\mathrm{Y})$, so (125) is not linear in m_P; and $\chi(i)$ in (124) depends on the capital distribution through $\mathbb{E}_i[\xi_i \varrho_i]$. Without capacity constraints $e_S \equiv 1/\alpha(i)$ and $\chi \equiv \xi(i)$: $\ln \mathrm{Y}_i(\overline{v}(i,\kappa)) = \alpha(i)[\ln \kappa + \ln \mathcal{P}(i)] + \text{const}$, (125) becomes the linear Fredholm equation (191) in m_P with the capital distributions entering through $\mathbb{E}_i[\ln \kappa]$, and (124) the Domar form $\widehat{s} = hq$; eliminating m_P against the numeraire then gives the log-linear reduced map (193) (Proposition 11). Under capacity the elimination is not available, because (125) is nonlinear in m_P and coupled to the capital distributions through two moment functions rather than one scalar moment per node. ∎

Derivation: the $\eta \to 1$ limit of the numeraire (126).

- Expanding $P^{1-\eta} = \exp[(1-\eta)\ln P] = 1 + (1-\eta)\ln P + O((1-\eta)^2)$ inside (47) gives $1 + (1-\eta)\int_0^1 \beta(i)\, m_P(i)\mathrm{d}i + O((1-\eta)^2) = 1$, hence $\int \beta\, m_P = O(1-\eta) \to$

0: the CES numeraire tends to (126). (The remainder is controlled by a uniform exponential moment of $\ln P$, $\sup_i \mathbb{E}_i[P^c + P^{-c}] < \infty$ for some $c > |1-\eta|$, since $|e^x - 1 - x| \le \frac{1}{2}x^2 e^{|x|}$; at $\eta = 1$ the moment is finite, since $\ln P = -z + f_i(K)$ with z Gaussian and K deterministic (Theorem 1), and for η near 1 it is assumed. The moments m_P themselves depend on η, so the statement is about the family $m_P^{(\eta)}$ and its limit.)

- The normalisation $Q_i/\beta(i)$ inside (117) is what makes the two coincide. The unnormalised aggregator $\exp\int \beta \ln Q_i \mathrm{d}i$ has unit cost $\exp[\int \beta \ln P_i \mathrm{d}i - \int \beta \ln \beta \mathrm{d}i]$ and numeraire $\int \beta\, m_P = \int \beta \ln \beta \mathrm{d}i$, which is not the CES limit; and because the final good is also the investment good, the two closures are different economies (a different g), not a units convention. The CES-limit form (126) is the one to use whenever the Cobb–Douglas block initialises or checks a CES computation.

On the balanced growth path the z-marginal of the node distribution is the stationary OU distribution, so $\mathbb{E}_i[z] = 0$ and the z-term of (125) drops; this includes $\varsigma^Z(i) = 0$, where $z_i(t) = z_i(0)e^{-\theta(i)t}$ decays deterministically to zero. It is kept because (125) also serves off the BGP, where $\mathbb{E}_i[z] = z_i(0)e^{-\theta(i)t} \neq 0$.

[Flag: the Cobb–Douglas price moment is a nonlinear equation whose contraction constant is not that of the model without capacity constraints.] Without capacity constraints the price-moment equation is a linear Fredholm equation of the second kind with operator $\mathbb{D}_{1-\alpha}\mathbb{A}$, a contraction on L^∞ with constant $1 - \underline{\alpha}$ (Proposition 11). Under capacity the analogous local object is the loading $1 - 1/e_S(\overline{v})$ of a node's log price on its bundle price at fixed $(\kappa, \widehat{s})$, which is state-dependent and, on a $9 \times 9 \times 9$ grid of log-normal triples ($\sigma_v \in [0.1, 2]$, $\varepsilon \in [1.2, 15]$, $\alpha \in [0.1, 0.85]$, the ranges of the flag after Proposition 3) scanned over 1501 thresholds, exceeds the no-capacity value $1 - \alpha$ somewhere on the threshold axis in 661 of them (it tends to 0 at $\underline{v}_i$ and to $1 - \alpha(i)$ at infinity, but is not monotone between); its supremum over the grid is 0.924, attained at $\sigma_v = 2$, $\varepsilon = 15$, $\alpha = 0.1$. So the loading stays below one on the grid, but a contraction argument for (125) would need a uniform bound on $1 - 1/e_S$, i.e. on e_S from below, which Assumption 5 does not supply; and the equation is coupled to the capital distributions through $\mathbb{E}_i[\ln Y_i]$ and to the sales fixed point through χ. The existence question is therefore the same as for the CES static block (Section 3.4), not simpler. **[Q16 | text | tag: CD1]** Establish existence, and if possible uniqueness, of the Cobb–Douglas static fixed point $(m_P, \widehat{m_X})$ of Proposition 4 given the capital distributions, e.g. by a contraction argument for (125) under a uniform lower bound $e_S \ge \underline{e} > 1$ (which would follow from a uniform lower bound on $\mathrm{d}\ln Y_i / \mathrm{d}\ln \overline{v}$ relative to $-\mathrm{d}\ln U_i / \mathrm{d}\ln \overline{v}$; see the flag above) combined with the positivity of the Neumann series in (124). If established, state it as an item of Proposition 4 and update the "Established" paragraph of Section 3.4.

Corollary 1 (The moment coupling under Cobb–Douglas). *Let $\eta = 1$ and consider the stationary system of Definition 6 below. Then:*

(i) (Sales and inbound demand.) *The sales moment needs no integration, $\widehat{m_S} = \widehat{s} = \beta q + \widehat{\mathcal{D}}$, but the inbound demand $\widehat{\mathcal{D}} = \mathbb{A}^* \widehat{m_X} = \mathbb{A}^*[\chi \widehat{s}]$ requires the capital distribu-*

tions through the effective shares (123)*; the GDP line of the consistency conditions holds identically.*

(ii) (Price moment.) *The graphon consistency of the price moment is* (125)*, closed by the numeraire* (126)*. Unlike the model without capacity constraints,* $\mathcal{P}(i)$ *enters every node's payoff, through the threshold* (120)*: the price moment is not decoupled from the dynamic block.*

(iii) (What the ensemble pins.) *The ensemble of node distributions enters the equilibrium conditions through the scalar* $\mathbb{E}_i[\kappa]$ *in the adding-up constraint, the moment functions* $\chi(i) = \mathbb{E}_i[\xi_i \varrho_i(\overline{v}(i,\kappa))]$ *in the inbound demand and* $\mathbb{E}_i[\ln Y_i(\overline{v}(i,\kappa))]$ *in the price moment, and* $\mathbb{E}_i[z]$*; the moment* $\mathbb{E}_i[\iota^*\kappa]$ *appears only in the feasibility check* $\widehat{C} > 0$*, which imposes no equation. The coupling is that of Remark 6, two functions on* $[0,1]$ *and one scalar; what the Cobb–Douglas aggregator removes is the dependence of the node's contribution to those functions on* z*, not the functions themselves.*

Proof. $\widehat{s} = \beta q + \widehat{\mathcal{D}}$ is constant in the state, so $\widehat{m_S} = \mathbb{E}_i[\widehat{s}] = \widehat{s}$ needs no integration, whereas $\widehat{m_X} = \chi\widehat{s}$ needs the capital distribution through $\chi(i) = \mathbb{E}_i[\xi_i\varrho_i]$ (Proposition 4(i)); the GDP line is Step 1 of that proof. The price moment satisfies (125) by Proposition 4(ii), and the bundle price $\mathcal{P}(i)$ appears in the threshold equation (120), hence in every function of the threshold: in utilization, in the shares, in the cash flow $[1-\xi_i]\varrho_i S$ and in the marginal profit (122). For (iii), read the stationary conditions of Definition 6 at $\eta = 1$: adding-up (138) involves $\mathbb{E}_i[\kappa]$; the consistency of $\widehat{\mathcal{D}}$ involves $\widehat{m_X} = \chi\widehat{s}$, i.e. $\mathbb{E}_i[\xi_i\varrho_i(\overline{v}(i,\kappa))]$; the consistency of $\mathcal{P}$ and the numeraire involve m_P, i.e. $\mathbb{E}_i[z]$ and $\mathbb{E}_i[\ln Y_i(\overline{v}(i,\kappa))]$ by (125); the growth identity is implied (proof of Proposition 5); and the feasibility check $\widehat{C} > 0$ involves $\mathbb{E}_i[\iota^*\kappa]$ without imposing an equation. No other condition sees a node distribution. The last sentence is Remark 6: the aggregates $(\mathcal{P}, \mathcal{D})$ are the same two kernel-weighted moment functions under either demand system, and at $\eta = 1$ each node's contributions to them, $\ln P$ and X, are functions of (K, z) through $-z + f_i(K)$ and of K alone respectively. ∎

B.7.2 Dynamics: Investment Neutrality

Theorem 1 (Investment neutrality under Cobb–Douglas with capacity). *Let* $\eta = 1$*, let aggregates follow any deterministic path, let* $\varsigma^Z(i) \geq 0$ *be arbitrary, and let Assumptions 4, 5 and 3 hold. Then at every node:*

1. (Payoffs.) *The individual profit function* $\pi(t,i,k;K,z)$ *of Lemma 16 is independent of* z *at every own capital* k*, and the marginal profit* (115) *equals* (122)*: it is independent of* z_i*, depends on the node's state through* K_i *alone, and is strictly decreasing in* K_i*.*
2. (Values and policies.) *The firm's value* w*, the on-diagonal marginal value* $\mathfrak{q}$ *and the investment policy* ι^* *have no* z*-argument.* $\mathfrak{q} = \mathfrak{q}(t,i,K)$ *solves the marginal-*$\mathfrak{q}$ *equation* (134) *with the flow* (122) *and* $\mathcal{L}_z^{(i)}\mathfrak{q} = 0$*: a first-order equation in* K *alone. The value of a firm is unaffected by any realisation of the productivity path.*

3. (Node distribution.) *The closed-loop capital dynamics* (136) *are autonomous in K and deterministic. Hence, for nonrandom $K_i(0)$, the node distribution is the product*

$$\mathfrak{M}(t,i,\mathrm{d}K,\mathrm{d}z) = \delta_{K_i(t)}(\mathrm{d}K)\,\mathfrak{m}^z(t,i,z)\,\mathrm{d}z,$$

with $K_i(t)$ the deterministic solution of the closed-loop capital equation and $\mathfrak{m}^z(t,i,\cdot)$ the Gaussian density of $z_i(t)$ under (93)[28]*. The two coordinates are independent and payoff-decoupled.*

4. (Coupling.) *Each node's problem involves the rest of the economy only through deterministic objects: the scalars $(Q(t), r_t)$ (in detrended form, (q_t, g_t, r_t), with $r = \rho + \gamma g$ on the BGP) and the two node-specific aggregates $(\mathcal{P}(t,i), \mathcal{D}(t,i))$, which are kernel moments of the ensemble through (m_P, m_X). Unlike the model without capacity constraints, the latter enter payoffs (Lemma 17(ii)): the mean-field coupling through payoffs is not removed. It is deterministic, runs through two functions on $[0,1]$ and one scalar exactly as under CES (Remark 6), and, by item 3, involves the node distributions only through the deterministic capital paths $K_i(t)$ together with the mean $\mathbb{E}_i[z]$, which enters the price moment* (125) *and vanishes on the balanced growth path.*

5. (Distributional degeneracy.) *Since $K_i(t)$ is deterministic, the "GMFG" collapses to a continuum of deterministic ODE pairs coupled through scalars and through the static fixed point of Proposition 4, and no fixed point over nondegenerate measures remains.*

Proof. Step 1: the individual profit function is z-free. By Lemma 16(i) and (85), the monopolist's maximised expected profit at own capital k is

$$\pi(t,i,k;K,z) = \max_{\overline{v}>\underline{v}_i} \left\{ P_i\, Y_i^{1/\varepsilon}\, \overline{y}(k)^{1-\frac{1}{\varepsilon}}\, \overline{v}^{1/\varepsilon}\, \Psi_{i,1}(\overline{v}) - \mathcal{P}(t,i)\,\frac{k}{x(i)}\,\Psi_{i,\frac{1}{1-\alpha(i)}}(\overline{v}) \right\},$$

the price at own threshold $\overline{v}$ being $p = P_i[\overline{v}Y_i/\overline{y}(k)]^{1/\varepsilon}$ and the maximiser the root of (66) at own k (Lemma 6(i); the threshold and the price are in bijection at fixed k). The maximisation is written over $\overline{v} > \underline{v}_i$ although the firm may post any $p \in (0,\infty)$, because by Lemma 2(i) the unique maximiser of the full problem already lies in that range. Here $P_i = P(t,i,K,z)$ and $Y_i = S(t,i)/P_i$ are the node's price and output, and $\overline{y}(k) = e^z k x(i)^{\alpha(i)-1}$ the firm's capacity. Write $P_i Y_i^{1/\varepsilon} = S(t,i)^{1/\varepsilon} P_i^{1-1/\varepsilon}$. By Lemma 17, $S(t,i)$ is z-free and $P_i \propto e^{-z}$ at fixed K, so $P_i^{1-1/\varepsilon}\overline{y}(k)^{1-1/\varepsilon} \propto e^{-z(1-1/\varepsilon)}e^{z(1-1/\varepsilon)}$ is z-free; the cost term contains no z. The objective, hence its maximiser and its maximum, do not depend on z: $\pi(t,i,k;K,z) = \pi(t,i,k;K)$, and the expected flow payoff $\pi(t,i,k;K) - \iota k$ contains no z, with $Q(t)$, $\mathcal{P}(t,i)$ and $\mathcal{D}(t,i)$ deterministic by hypothesis. Differentiating at $k = K$ gives (122) by Lemma 17(iv), on the right of which the state enters through K alone; the monotonicity in K is Lemma 17(iv). This is item 1. The remaining items follow in turn, exactly as without capacity constraints.

[28] mean $z_i(0)e^{-\theta(i)t}$, variance $\frac{\varsigma^Z(i)^2}{2\theta(i)}[1-e^{-2\theta(i)t}]$; on the BGP the stationary $\mathcal{N}(0,\varsigma^Z(i)^2/2\theta(i))$, which at $\varsigma^Z(i) = 0$ is the point mass at 0: without its shock, $z_i(t) = z_i(0)e^{-\theta(i)t}$ decays deterministically to zero.

2. In the firm HJB equation (49) with this flow, try a value function without z-argument, $w = w(t, i, k, K)$, together with a z-free symmetric node policy $\bar{\iota}(t, i, K)$ (the on-diagonal maximiser of the same problem, so the guess is internally consistent). Then $\mathcal{L}_z^{(i)} w = 0$, and every remaining coefficient of the equation (the flow, r_t, the drifts $[\Phi_i(\cdot) - \delta(i)]$) is free of z, so the z-free candidate solves the HJB; under the regularity guaranteeing existence and uniqueness for the firm HJB, the value function itself is z-free given the z-free node policy. Then $\mathfrak{q} = w_k\big|_{k=K}$ and $\iota^* = \iota^*(\mathfrak{q})$ from the investment first-order condition are z-free, so the on-diagonal maximiser reproduces the guessed node policy: a z-free symmetric equilibrium exists, and the marginal-$\mathfrak{q}$ equation loses its $\mathcal{L}_z^{(i)}$ term in it; what remains is a first-order equation in (t, K). That it is *the* symmetric equilibrium uses the uniqueness of the node fixed point $(w, \bar{\iota})$ assumedin Section 3.4: a symmetric equilibrium in which the node policy depended on z would make z payoff-relevant through K alone, a sunspot that the fundamentals do not support.

3. With $\iota^* = \iota^*(t, i, K)$, the capital equation of the closed-loop dynamics (136) is a one-dimensional ODE in K with deterministic coefficients, while z_i solves (93) driven by B_i^Z. With K the unique solution of its ODE (the policy being Lipschitz in K under the regularity assumed) and $K_i(0)$ nonrandom, $K_i(t)$ is deterministic, so the joint distribution is the product of the point mass at $K_i(t)$ and the Gaussian OU distribution of $z_i(t)$. Payoff-decoupling is item 1.

4. Sales are pinned by the scalar path $Q(t)$ and the deterministic inbound demand $\mathcal{D}(t, i)$, and the threshold by these and the deterministic bundle price $\mathcal{P}(t, i)$ (Lemma 17(i)–(ii)); by items 1–3 nothing else in the node's problem refers to another node or to a distribution. The two functions are the kernel moments $(\mathbb{A} m_P, \mathbb{A}^* m_X)$ of (44)–(45), and by item 3 and Lemma 17(ii)–(iii) each node's contributions to them, $\ln P = -z + f_i(t, K)$ and $X(t, i, K)$, have expectations that involve the deterministic $K_i(t)$ and $\mathbb{E}_i[z]$ only. In detrended form the scalars are (q_t, g_t, r_t), with $r = \rho + \gamma g$ on the BGP.

5. The capital equation is the ODE $\dot{K} = [\Phi_i(\iota^*(t, i, K)) - \delta(i)]K$ with deterministic coefficients, so $K_i(t)$ is deterministic and the capital marginal is $\delta_{K_i(t)}$; the $\mathfrak{q}$-equation of item 2 is first order in K, and each node is a deterministic ODE pair coupled to the economy only through the scalars and functions of item 4, the functions being determined by the static fixed point of Proposition 4 with Dirac distributions.

■

Items 2, 3 and 5 are those of the model without capacity constraints word for word; item 1 replaces a constant times $1/K$ by a strictly decreasing function of K with the same two ends; item 4 is weaker than its counterpart there, whose node payoffs were pinned by primitives and the scalar Q. The theorem is the precise statement that capacity does not repair the payoff-side degeneracy: as long as a node's buyers spend fixed nominal amounts on it, its productivity is absorbed by its price and never reaches its

cash flow, however often its firms are rationed.

Proposition 5 (The degenerate Cobb–Douglas balanced growth path under capacity). *Let $\eta = 1$, under the assumptions of Theorem 1. Then the stationary equilibrium of Definition 6 is algebraic node by node: the node distribution is $\delta_{\kappa(i)}(\mathrm{d}\kappa)\,\mathfrak{m}^z(i,z)\mathrm{d}z$, and $(\kappa(i), \overline{v}(i), \mathfrak{q}(i), \iota(i))$ solve*

$$\begin{aligned} \Phi_i(\iota(i)) - \delta(i) &= g, \qquad \Phi_i'(\iota(i))\,\mathfrak{q}(i) = 1, \\ [\rho + (\gamma-1)g]\,\mathfrak{q}(i) &= \frac{\varepsilon}{1-\alpha(i)}\,\frac{\mathcal{P}(i)}{x(i)}\,\Phi_i(\overline{v}(i)) - \iota(i), \\ \frac{\mathcal{P}(i)\,\kappa(i)}{x(i)}\,\mu_{i,1/\varepsilon}^{\frac{\varepsilon}{\varepsilon-1}}\,\frac{U_i(\overline{v}(i))}{\mathrm{Y}_i(\overline{v}(i))} &= \widehat{s}(i), \end{aligned} \tag{127}$$

together with $\int_0^1 \kappa(i)\mathrm{d}i = 1$, the static block of Proposition 4 with Dirac distributions ($\widehat{s} = \beta q + \widehat{\mathcal{D}}$, $\widehat{m_X}(i) = \xi_i(\overline{v}(i))\varrho_i(\overline{v}(i))\widehat{s}(i)$, $m_P(i) = \ln P(i)$ from (121) at $z = 0$, and the numeraire (126)) for $(\mathcal{P}, \widehat{\mathcal{D}}, q, g)$, and $r = \rho + \gamma g$. Given g, the first two equations deliver $(\iota(i), \mathfrak{q}(i))$ in closed form; given in addition $(\mathcal{P}(i), \widehat{s}(i))$, the third has a unique root $\overline{v}(i) \in (\underline{v}_i, \infty)$, since Φ_i decreases strictly from $+\infty$ to 0 (Assumption 3), and the fourth delivers

$$\kappa(i) = \frac{x(i)\,\widehat{s}(i)}{\mathcal{P}(i)\,\mu_{i,1/\varepsilon}^{\frac{\varepsilon}{\varepsilon-1}}}\,\frac{\mathrm{Y}_i(\overline{v}(i))}{U_i(\overline{v}(i))}, \tag{128}$$

equivalently, the node's capital is the unique root of $\pi_k(i,\kappa) = [\rho + (\gamma-1)g]\mathfrak{q}(i) + \iota(i)$ with the marginal profit (122), whose left side decreases strictly from $+\infty$ to 0 in κ (the recipe needs $1 + g + \delta(i) > 0$ for the first equation and $[\rho + (\gamma-1)g]\mathfrak{q}(i) + \iota(i) > 0$ for the third, both automatic at $g \geq 0$). Hence $\kappa(i) \propto \widehat{s}(i)$ at given $(g, \mathcal{P}(i))$, and adding-up gives q given g and the profiles $(\mathcal{P}, \widehat{\mathcal{D}}/q)$; but neither profile is available in closed form, so the rest points are coupled through the static fixed point of Proposition 4, and the numeraire is one equation in g once that fixed point is solved. In the no-capacity limit $x(i) \downarrow 0$ the root tends to the closed form $\kappa(i) = \frac{\varepsilon-1}{\varepsilon}\alpha(i)\widehat{s}(i)/\{[\rho + (\gamma-1)g]\mathfrak{q}(i) + \iota(i)\}$ of Proposition 12.

Proof. By Theorem 1, item 5, $\kappa_i(t)$ is deterministic, so on the BGP it is a rest point $\kappa(i)$ of the detrended capital drift (142): $\Phi_i(\iota(i)) - \delta(i) = g$, the first equation. The second is the investment first-order condition (133). The third is the stationary detrended marginal-$\mathfrak{q}$ equation (52) with the flow (122) and $\mathcal{L}_z^{(i)}\mathfrak{q} = 0$, evaluated at the rest point, where the advection term $\mathfrak{q}_\kappa\kappa[\Phi_i - \delta - g]$ vanishes whatever $\mathfrak{q}_\kappa$ is, so that $r - [\Phi_i - \delta] = \rho + \gamma g - g$. The fourth is the threshold equation (120) at the rest point. The node distribution is the product of Theorem 1, item 3, with $\mathfrak{M}^K = \delta_{\kappa(i)}$; adding-up reads $\int \kappa(i)\mathrm{d}i = 1$; the static block of Proposition 4 is evaluated at the Dirac distributions, where $\chi(i) = \xi_i(\overline{v}(i))\varrho_i(\overline{v}(i))$, $\mathbb{E}_i[\ln \mathrm{Y}_i] = \ln \mathrm{Y}_i(\overline{v}(i))$ and $\mathbb{E}_i[z] = 0$ (the stationary OU distribution), so that $m_P(i) = \ln P(i)$ with $P(i)$ from (121) at $z = 0$; and the growth identity (55) holds identically at the rest points, since every node's capital grows at g by the first equation and $\int_0^1 \kappa(i)\mathrm{d}i = 1$. For the solution recipe: given g, the first two equations do not involve the node's capital and

give $(\iota(i), \mathfrak{q}(i))$ as without capacity constraints; given $(\mathcal{P}(i), \widehat{s}(i))$, the third equation is $\Phi_i(\overline{v}) = \frac{(1-\alpha(i))x(i)}{\varepsilon \mathcal{P}(i)}\{[\rho + (\gamma-1)g]\mathfrak{q}(i) + \iota(i)\}$, whose right side is a positive constant and whose left side decreases strictly from $+\infty$ at $\underline{v}_i$ to 0 at infinity under Assumption 3 (Lemma 6(iii)): one root $\overline{v}(i)$. Solving the fourth equation for $\kappa(i)$ at that root gives (128). Equivalently, since $\kappa \mapsto \overline{v}(i,\kappa)$ is a strictly increasing bijection of $(0,\infty)$ onto $(\underline{v}_i, \infty)$ (Lemma 17(ii)), the composite $\kappa \mapsto \pi_k(i,\kappa)$ decreases strictly from $+\infty$ to 0 and the third equation, read as $\pi_k(i,\kappa(i)) = [\rho + (\gamma-1)g]\mathfrak{q}(i) + \iota(i)$, has the same unique root. Reading (128) at fixed $(g, \mathcal{P}(i))$: $\overline{v}(i)$ does not depend on $\widehat{s}(i)$, so $\kappa(i) \propto \widehat{s}(i) = \beta(i)q + \widehat{\mathcal{D}}(i)$, and adding-up gives $q = 1/\int_0^1 [\kappa(i)/q]\mathrm{d}i$ once the profile $\widehat{\mathcal{D}}/q$ is known; but $\widehat{\mathcal{D}} = \mathbb{A}^*[\chi\widehat{s}]$ and $\mathcal{P}$ are themselves determined by the static block at the rest points, so the closed form of the model without capacity constraints (where $\widehat{\mathcal{D}} = (h-\beta)q$ and $\mathcal{P}$ does not enter) is replaced by the fixed point of Proposition 4. The no-capacity limit: as $x(i) \downarrow 0$ at fixed $(\mathcal{P}(i), \widehat{s}(i))$ the root of the third equation diverges ($\Phi_i(\overline{v}(i)) \propto x(i) \to 0$), and along it $\pi_k(i,\kappa) \to \frac{\varepsilon-1}{\varepsilon}\alpha(i)\widehat{s}(i)/\kappa$ (Lemma 17(iv)), so the root in κ tends to $\frac{\varepsilon-1}{\varepsilon}\alpha(i)\widehat{s}(i)/\{[\rho + (\gamma-1)g]\mathfrak{q}(i) + \iota(i)\}$, which is (194) with $\widehat{s} = hq$. ∎

There is nothing mean-field about this object in the distributional sense: a continuum of deterministic q-theory rest points (off the BGP, ODE saddle paths), coupled by two scalars and, through the bundle prices and inbound demands, by a static fixed point of the same kind as the finite static block of Section 2.2, solved once. All distributional content, all payoff-relevant risk, and all use of the FPK equation have vanished. This is the precise sense in which the Cobb–Douglas GMFG is degenerate, and why the CES repair is needed before the graphon limit is a well-posed mean-field problem. The threshold equation of (127) has a direct reading: on the balanced growth path every node's firms operate at the utilization at which the private return to capital, $\frac{\varepsilon}{1-\alpha(i)}\frac{\mathcal{P}(i)}{x(i)}\Phi_i(\overline{v})$, equals the user cost $[\rho + (\gamma-1)g]\mathfrak{q}(i) + \iota(i)$ of a unit of capital, and the node's capital is then whatever stock delivers its pinned sales at that utilization; the common-technology rest point $\kappa(i) \propto \alpha(i)h(i)$ of the model without capacity constraints and its relation to the finite household model's $\kappa_i^* = \omega_i$ are in Appendix C.5. The rest point is also the natural initialiser of the CES computation: Appendix D uses the rest point of the model without capacity constraints, which is the $x(i) \downarrow 0$ limit of this one, as grid centre and warm start, and keeps this one as the refinement.

Mean reversion. The rest point of Proposition 5 is approached from either side if the equilibrium policy is mean-reverting, with node capital growing faster than the economy at small κ and slower at large κ; this is the Cobb–Douglas form of the equilibrium regularity conditionflagged in Section 3.4, and it is what makes the rest point the long-run position of every node. Two features of the node's marginal-q equation, (52) with the flow (122) and $\mathcal{L}_z^{(i)}\mathfrak{q} = 0$, make this the expected case. The only κ-dependence of the node's problem is the marginal dividend $\frac{\varepsilon}{1-\alpha(i)}\frac{\mathcal{P}(i)}{x(i)}\Phi_i(\overline{v}(i,\kappa))$, which is strictly decreasing in κ (Lemma 17(iv)), so $\mathfrak{q}$ should decrease in κ and with it ι^* and the detrended drift $\mu^\kappa(i,\kappa) := \Phi_i(\iota^*(\mathfrak{q}(i,\kappa))) - \delta(i) - g$ of (142). At small κ the dividend is

unbounded, like $(1-\frac{1}{\varepsilon})\varrho_i(\underline{v}_i)\widehat{s}(i)/\kappa$ (the threshold tends to $\underline{v}_i$, where the node's firms are rationed in almost every state and the whole of their revenue net of the markup rent is capital income), and drives ι^*, hence μ^κ, above any bound: the same $1/\kappa$ divergence as in the model without capacity constraints, with a different constant. Unlike its finite-sector counterpart (Section 2.3), the divergence is unconditional, because atomism holds the aggregates fixed while the node's capital vanishes. At large κ the dividend vanishes, like $\frac{\varepsilon-1}{\varepsilon}\alpha(i)\widehat{s}(i)/\kappa$ in the model without capacity constraints, and $\mathfrak{q}$ approaches $\mathfrak{q}_\infty(i)$, the unique root in $(0,1)$ of

$$[\rho+\gamma g+\delta(i)]\,\mathfrak{q}=\max_{\iota}\{\mathfrak{q}\,\Phi_i(\iota)-\iota\}$$

(the right-hand side minus the left is convex in $\mathfrak{q}$, equal to $1/\phi(i)>0$ at $\mathfrak{q}=0$ and to $-[\rho+\gamma g+\delta(i)]<0$ at $\mathfrak{q}=1$, since $\Phi_i(\iota)\le\iota$): the marginal value of a unit of capital that will never earn a dividend and is run down at the optimal disinvestment rate (the limit presumes that $\kappa\mathfrak{q}_\kappa$ vanishes as $\kappa\to\infty$, as it does when $\mathfrak{q}-\mathfrak{q}_\infty=O(1/\kappa)$). Hence $\iota^*(\mathfrak{q}_\infty)<0$ and $\mu^\kappa(i,\infty)=\Phi_i(\iota^*(\mathfrak{q}_\infty))-\delta(i)-g<-\delta(i)-g<0$ for $g>-\delta(i)$, in particular for any $g\ge 0$: capital far above the rest point shrinks relative to the economy. Monotonicity of $\mathfrak{q}$ in κ, which delivers the rest, is a property of the equilibrium policy that we verify numerically rather than prove (Appendix D, Step 7).

B.7.3 Anatomy of the Degeneracy

What the productivity shocks still do. The degeneracy does not delete the productivity shocks from the economy; it deletes them from payoffs and dynamics. By (121), $\ln P(i,\kappa,z)$ loads on z with coefficient -1: prices absorb productivity one-for-one. Hence real node output $Y_i=S_i/P_i$, real capacity $\overline{Y}_i$, real bilateral shipments $m_i(j)=a(i,j)X_i/P_j$, and real final demands $Q_i=\beta(i)Q/P_i$ all fluctuate with the entire cross-section of z's: the *real* input–output economy churns stochastically forever. But nominal sales $S(t,i)$, the threshold, utilization, the rationed share of firms, the split of sales into materials, capital income and the two rents, cash flows, the marginal profit (122), hence $\mathfrak{q}$, ι^* and $K_i(t)$, are untouched by any realisation of any z-path (Theorem 1). The shocks are demoted from state variables to a payoff-irrelevant relabelling of relative prices and real quantities: observable churn with no owner, to which no valuation, no investment decision and no aggregate responds. The removal happens in three places: (i) from the flow payoff (Lemma 17); (ii) therefore from the backward equation, whose $\mathcal{L}_z^{(i)}$ term acts on a constant; (iii) therefore from the forward dynamics, where the OU motion of z is exogenous and decoupled, giving the product distribution of Theorem 1, item 3.

Why the finite model does not degenerate: the two channels. A firm's cash flow can respond to its own (κ,z) through exactly two channels. The *own-price channel* is partial-equilibrium: holding all aggregates fixed, a higher z shifts the node's supply curve out, its relative price falls, and revenue moves by however much quantity and price fail to offset. Its strength is the loading $\partial\ln S/\partial z$ of Proposition 6 below, which

is zero *identically* at $\eta = 1$: with unit-elastic final demand and Cobb–Douglas bundles, every buyer of node i spends a fixed amount on it, so the price falls exactly enough to undo the quantity gain (Lemma 17). The *general-equilibrium channel* is the node moving the aggregates themselves. With $J < \infty$, Q_t is a function of $(\mathbf{K}_t, \mathbf{Z}_t)$ through the static block of Section 2.2, so sector i's cash flow responds to Z_{it} through Q_t and through the prices of its buyers' bundles: sector i moves the aggregates. Without capacity constraints the reduced map (171), $\ln Q = \text{const} + \sum_i \omega_i \ln K_i + \sum_i (\omega_i/\alpha_i) \ln Z_i$, makes this elasticity the explicit ω_i/α_i; under capacity it has no closed form but is nonzero for the same reason. In the continuum every node is atomistic: the exact LLN makes $Q(t)$, $\mathcal{P}(t,i)$ and $\mathcal{D}(t,i)$ deterministic functionals of the ensemble, and Lemma 17 says nothing else in the node's payoff depends on its productivity. The two models divide the channels between them:

	own-price channel	GE channel
Finite CD	0	$\neq 0$ (ω_i/α_i without capacity constraints)
Finite CES	$\neq 0$	$\neq 0$
GMFG CES	$\neq 0$	0 (atomistic)
GMFG CD	0	0

This is why the finite CD and CES models look like near-identical theories: their stationary equilibria differ only in whether the static block is a linear solve or a root-find without capacity constraints, and not even in that under capacity, because *both* live off a nonzero GE channel. In the limit the difference is one of kind: atomism annihilates the GE channel for every specification, so CD, whose own-price channel is identically zero, is left with no own-state relevance in productivity at all, and CES survives precisely because its own-price channel is an $O(1)$ object that needs no general-equilibrium feedback to operate. Capacity opens no third channel, because the threshold is productivity-free at $\eta = 1$: the own-price channel of a node's utilization is exactly the own-price channel of its sales, and vanishes with it.

On a step kernel with N blocks. The CD "GMFG" under capacity is exactly the algebra (127) at N points plus two scalars plus an N-block static fixed point: given g, each block's $(\iota_n, \mathfrak{q}_n)$ is closed-form; given $(\mathcal{P}_n, \widehat{s}_n)$, the block's threshold and capital follow from the third and fourth equations; and $(\mathcal{P}_n, \widehat{\mathcal{D}}_n)_{n \le N}$ solve the N-vector system $\ln \mathcal{P} = \mathbf{A}\, m_P - \text{const}$, $m_{Pn} = \ln P_n$, $\widehat{\mathcal{D}} = \mathbf{A}^{\mathsf{T}} \widehat{m_X}$, $\widehat{m_X}_n = \xi_n \varrho_n \widehat{s}_n$, with (g, q) from the normalisation $\sum_n \kappa_n \Delta_n = 1$ and the numeraire: a nonlinear system of dimension $2N$ in place of the $N \times N$ linear solve of the model without capacity constraints. Placed next to the finite N-*sector* CD model, a genuinely stochastic $(2N-1)$-dimensional problem, the comparison isolates what the within-block continuum does: in the finite model block n's shock is common to the whole block and moves the aggregates; in the step-kernel GMFG the block contains a continuum of independently shocked industries, the exact LLN averages them, block moments are deterministic, and the GE channel is dead *even at the block level*. The CD-GMFG with N blocks is therefore

not a coarse version of the finite N-sector model but its *shock-free skeleton*: the deterministic BGP the finite model fluctuates around, with the entire fluctuation (aggregate risk, the stochastic SDF, all reallocation dynamics) removed; the $O(J^{-1/2})$ aggregate-risk termof Section 3.4 is then not an error term on top of the limit dynamics but the entirety of the finite model's dynamics.

Remark 8 (What a solver reports at $\eta = 1$) Running the numerical pipeline of Appendix D under its no-capacity mapping at $\eta = 1$ shows how the two degeneracies manifest in floating point. The $\mathfrak{q}$-step converges to a $\mathfrak{q}$ constant along the z-grid: the z-dimension is dead weight, which is the payoff-side degeneracy. The FPK step is where the distributional degeneracy surfaces: the κ-dynamics are pure transport into the rest point, no stationary *density* exists, the discrete adjoint $\mathsf{A}^{\mathsf{T}}\mathfrak{m} = 0$ is solvable only in the measure sense, and the implicit-transient iteration piles all mass into the one or two cells straddling κ_n^*, with conditioning deteriorating as it does. Deceptively, the outer loop still converges, and fast, because every moment it consumes is exact however badly $\mathfrak{m}$ is resolved ($\widehat{m_{S_n}} = h_n q$ needs no integration, and $\mathbb{E}_n[\ln \kappa] \to \ln \kappa_n^*$): the fixed point over ensembles has silently collapsed to a fixed point over (g, q) and the Fredholm vector, and the run looks like a successful GMFG computation while containing no mean-field content. The honest use of $\eta = 1$ is as the algebraic initialiser and the $\eta \to 1$ continuation limit of the CES computation (Appendix B.7.4). The same two symptoms are to be expected from a solver for the capacity model, with one difference: the moments the outer loop consumes at $\eta = 1$, $\widehat{m_{X_n}}$ and m_{P_n}, are no longer exact without $\mathfrak{m}$; they are the values at the rest point, which a pile-up in one or two cells reproduces to first order, so the deceptive convergence persists, and the collapsed fixed point is the $2N$-dimensional static system of the preceding paragraph rather than the Fredholm vector.

[Flag: the body of Remark 8 is the no-capacity model's, its last sentence a prediction.] Everything up to the last sentence was observed on the no-capacity pipeline (the no-capacity mapping of Appendix D); the last sentence states the expected behaviour of the capacity pipeline that Appendix D now describes, and is a prediction, not a report. The capacity path of Proposition 5 has since been solved, and the capacity pipeline now returns it directly in this configuration rather than running the loop on a degenerate distribution (Step 5), so the deceptive convergence described here still has not been observed under capacity — it would be, if the loop were run there. What *is* now measured is the object the last sentence names as the collapse's limit: the $\eta \to 1$ continuation under capacity converges to that $2N$-dimensional static system, with $|\Delta g|$ of order $\eta - 1$ (10^{-6} at $\eta = 0.999$). The last sentence's identification of the collapsed fixed point is therefore confirmed; its account of the pile-up in one or two cells remains a prediction, and would be checked by forcing the loop to run at $\eta = 1$ under capacity.

B.7.4 CES Restores Own-State Relevance

Proposition 6 (Own-state relevance under CES with capacity). *Let $\eta \neq 1$ and $\beta(i) > 0$, and let Assumptions 4 and 5 hold. Define the final-demand share of sales $\vartheta(t,i,K,z) := \beta(i)P_i^{1-\eta}Q/S_i \in (0,1]$, and let $\ell_Y := \mathrm{d}\ln Y_i/\mathrm{d}\ln\overline{v} \geq 0$ and $\ell_U := -\,\mathrm{d}\ln U_i/\mathrm{d}\ln\overline{v} > 0$ be evaluated at the equilibrium threshold $\overline{v}(t,i,K,z)$ of Lemma 15, with*

$$\Delta := \ell_U + [1-(1-\eta)\vartheta]\,\ell_Y > 0.$$

Then the equilibrium node maps satisfy, pointwise,

$$\begin{aligned}
\frac{\partial \ln S}{\partial \ln K} &= \frac{-(1-\eta)\,\vartheta\,\ell_Y}{\Delta}, & \frac{\partial \ln S}{\partial z} &= \frac{-(1-\eta)\,\vartheta\,(\ell_U+\ell_Y)}{\Delta}, \\
\frac{\partial \ln \overline{v}}{\partial \ln K} &= \frac{1}{\Delta}, & \frac{\partial \ln \overline{v}}{\partial z} &= \frac{(1-\eta)\,\vartheta}{\Delta}, \\
\frac{\partial \ln P}{\partial \ln K} &= \frac{-\ell_Y}{\Delta}, & \frac{\partial \ln P}{\partial z} &= \frac{-(\ell_U+\ell_Y)}{\Delta}, \\
\frac{\partial \ln U}{\partial \ln K} &= \frac{-\ell_U}{\Delta}, & \frac{\partial \ln U}{\partial z} &= \frac{-(1-\eta)\,\vartheta\,\ell_U}{\Delta},
\end{aligned} \tag{129–130}$$

and, wherever $\ell_Y > 0$, in terms of the supply elasticity $e_S = 1 + \ell_U/\ell_Y$ of (111) alone,

$$\begin{aligned}
\frac{\partial \ln S}{\partial \ln K} &= \frac{-(1-\eta)\,\vartheta}{e_S-(1-\eta)\vartheta}, & \frac{\partial \ln S}{\partial z} &= \frac{-(1-\eta)\,\vartheta\,e_S}{e_S-(1-\eta)\vartheta}, \\
\frac{\partial \ln U}{\partial \ln K} &= \frac{-(e_S-1)}{e_S-(1-\eta)\vartheta}, & \frac{\partial \ln U}{\partial z} &= \frac{-(1-\eta)\,\vartheta\,(e_S-1)}{e_S-(1-\eta)\vartheta},
\end{aligned} \tag{131}$$

with $e_S - (1-\eta)\vartheta > 0$ always. Sales and utilization load on z, and sales on $\ln K$ (wherever $\ell_Y > 0$), with the sign of $\eta - 1$, nonzero whenever $\eta \neq 1$ and $\vartheta > 0$: strictly positive if $\eta > 1$ (elastic final demand: higher capital or productivity lowers the relative price by less than it raises quantity sold) and strictly negative if $\eta < 1$; utilization is strictly decreasing in K and the price strictly decreasing in both. Under Assumption 3 the marginal profit (115) is strictly decreasing in K and strictly monotone in z, increasing if and only if $\eta > 1$. As $\overline{v} \to \infty$ (the no-capacity limit, under the density condition of Lemma 15(iv)) $\ell_U \to 1$, $\ell_Y \to \frac{\alpha(i)}{1-\alpha(i)}$, $e_S \to 1/\alpha(i)$, and (131) reduces to (195) of the model without capacity constraints. In either case cash flows (114) and marginal profits (115) load on (K_i, z_i), and the node problem is a genuine stochastic control problem in (κ, z) (and, if $\varsigma^Z(i) = 0$, a deterministic one-dimensional problem in κ, at $z = 0$ on the BGP).

Proof. Hold the aggregates fixed and use the function F of (112), whose root is the equilibrium threshold, and the log-derivatives $\ell_U > 0$, $\ell_Y \geq 0$ of the proof of Lemma 15(iv). *Step 1: the threshold.* From Steps 2–3 of that proof, at the root

$$\overline{v}F' = -(\ell_U+\ell_Y)\,\mathcal{D} - \beta(i)Q\,P^{1-\eta}[\ell_U + \eta\,\ell_Y], \qquad K\,\partial_K F = S, \qquad \partial_z F = (1-\eta)\,\beta(i)QP^{1-\eta},$$

where $\beta(i)Qc^{1-\eta}Y^{\eta-1} = \beta(i)QP^{1-\eta}$ by (107). Divide by S and use $\beta(i)QP^{1-\eta}/S = \vartheta$ and $\mathcal{D}/S = 1-\vartheta$ (the demand relation (109)):

$$-\frac{\overline{v}F'}{S} = (\ell_U+\ell_Y)(1-\vartheta) + \vartheta(\ell_U+\eta\ell_Y) = \ell_U + [1-(1-\eta)\vartheta]\ell_Y = \Delta,$$

which is strictly positive because $\ell_U > 0$, $\ell_Y \geq 0$ and $1-(1-\eta)\vartheta > 0$ (for $\eta \geq 1$ trivially; for $\eta < 1$ because $(1-\eta)\vartheta < 1$). The implicit function theorem then gives

$$\frac{\partial \ln \overline{v}}{\partial \ln K} = -\frac{K\,\partial_K F}{\overline{v}F'} = \frac{1}{\Delta}, \qquad \frac{\partial \ln \overline{v}}{\partial z} = -\frac{\partial_z F}{\overline{v}F'} = \frac{(1-\eta)\vartheta}{\Delta}.$$

Step 2: price, sales, utilization. At fixed aggregates, (107) gives $\ln P = \text{const} - z - \ln Y_i(\overline{v})$, the second form of (108) gives $\ln S = \text{const} + \ln K + \ln U_i(\overline{v}) - \ln Y_i(\overline{v})$, and $\mathrm{d}\ln U_i = -\ell_U \mathrm{d}\ln \overline{v}$, $\mathrm{d}\ln Y_i = \ell_Y \mathrm{d}\ln \overline{v}$. Hence

$$\begin{aligned}
\frac{\partial \ln P}{\partial \ln K} &= -\frac{\ell_Y}{\Delta}, & \frac{\partial \ln P}{\partial z} &= -1 - \frac{\ell_Y(1-\eta)\vartheta}{\Delta} = -\frac{\ell_U + \ell_Y}{\Delta},\\
\frac{\partial \ln S}{\partial \ln K} &= 1 - \frac{\ell_U + \ell_Y}{\Delta} = -\frac{(1-\eta)\vartheta\,\ell_Y}{\Delta}, & \frac{\partial \ln S}{\partial z} &= -\frac{(\ell_U + \ell_Y)(1-\eta)\vartheta}{\Delta},\\
\frac{\partial \ln U}{\partial \ln K} &= -\frac{\ell_U}{\Delta}, & \frac{\partial \ln U}{\partial z} &= -\frac{\ell_U(1-\eta)\vartheta}{\Delta},
\end{aligned}$$

using $\Delta - (\ell_U + \ell_Y) = -(1-\eta)\vartheta\ell_Y$ in the sales loading; these are (129)–(130). Wherever $\ell_Y > 0$, divide numerators and denominators by ℓ_Y and use $\ell_U + \ell_Y = e_S\ell_Y$ (from (111)) and $\Delta = \ell_Y[e_S - (1-\eta)\vartheta]$ to obtain (131); $e_S - (1-\eta)\vartheta > 0$ because $e_S > 1 > (1-\eta)\vartheta$ for $\eta < 1$ and $-(1-\eta)\vartheta \geq 0$ otherwise.

Step 3: signs and the marginal profit. The signs are read off the formulas: the loadings of S and U on z and of S on $\ln K$ carry the factor $-(1-\eta)\vartheta$, nonzero iff $\eta \neq 1$ and $\vartheta > 0$ (the loading of S on $\ln K$ also needs $\ell_Y > 0$), with the sign of $\eta - 1$; the loadings of U on $\ln K$ and of P on z are strictly negative, being $-\ell_U/\Delta$ and $-(\ell_U + \ell_Y)/\Delta$ with $\ell_U > 0$. The loading of P on $\ln K$ is $-\ell_Y/\Delta \leq 0$ and vanishes wherever $\ell_Y = 0$, which Assumption 5 permits; the price is nonetheless *strictly* decreasing in K, not by the formula but by composition: $P = \mathcal{P}(t,i)e^{-z}x(i)^{-\alpha(i)}/Y_i(\overline{v})$ with Y_i strictly increasing and $\overline{v}$ strictly increasing in K, since $\partial \ln \overline{v}/\partial \ln K = 1/\Delta > 0$ everywhere. The marginal profit (115) is $\frac{\varepsilon}{1-\alpha(i)}\frac{\mathcal{P}(t,i)}{x(i)}\Phi_i(\overline{v}(t,i,K,z))$ at fixed aggregates; under Assumption 3 Φ_i is strictly decreasing, and the threshold is strictly increasing in K and strictly monotone in z with the sign of $1-\eta$ (Step 1, or Lemma 15(iv)), so the marginal profit is strictly decreasing in K and strictly monotone in z, increasing iff $\eta > 1$; its loading on z is $\mathrm{d}\ln\Phi_i/\mathrm{d}\ln\overline{v} \cdot (1-\eta)\vartheta/\Delta$. Cash flows $[1-\xi_i(\overline{v})]\varrho_i(\overline{v})S$ load on (K,z) through S and through the threshold.

Step 4: the no-capacity limit. By Step 1 of the proof of Lemma 15(iv), as $\overline{v} \to \infty$ under the density condition stated there, $\ell_U \to 1$ and $\ell_Y \to \frac{\alpha(i)}{1-\alpha(i)}$, so $e_S \to 1/\alpha(i)$ and $\Delta \to \frac{1-\alpha(i)(1-\eta)\vartheta}{1-\alpha(i)}$; substituting into (129) gives $\partial \ln S/\partial \ln K \to \frac{-(1-\eta)\vartheta\alpha(i)}{1-\alpha(i)(1-\eta)\vartheta}$ and $\partial \ln S/\partial z \to \frac{-(1-\eta)\vartheta}{1-\alpha(i)(1-\eta)\vartheta}$, (195) of the model without capacity constraints, and $\partial \ln P/\partial z \to \frac{-1}{1-\alpha(i)(1-\eta)\vartheta}$ as in its proof. ∎

The structure of (131) is that of the model without capacity constraints, with the constant supply elasticity $1/\alpha(i)$ replaced by the state-dependent $e_S(\overline{v})$: the numerator is the final-demand share of sales times the gap $1-\eta$ between unit elasticity and the elasticity of final spending in the node's price, and the denominator is the gap between

the node's supply elasticity and the demand elasticity of its nominal sales, $(1-\eta)\vartheta$. What capacity adds is the decomposition of the supply response into a price part and a utilization part. The loadings of utilization carry the factor $e_S - 1 = \ell_U/\ell_Y$, the utilization share of the supply response: a node with a steep threshold map (large ℓ_Y, so e_S close to one) meets a productivity or capital shock through its price and barely moves its utilization, whereas a node with a flat threshold map meets it through utilization. Both parts have the same sign structure, so the qualitative statements of the model without capacity constraints survive unchanged: for substitutes a more productive or better-capitalised node sells more in nominal terms and runs at higher utilization, for complements it sells less and, for productivity, runs at lower utilization; capital lowers utilization under either demand system, because it raises capacity faster than it raises sales.

Remark 9 (Nodes with $\beta(i) = 0$ stay degenerate) If $\beta(i) = 0$ on a positive-measure set, those nodes sell only to other nodes' Cobb–Douglas bundles, so $\vartheta = 0$, $S_i = \mathcal{D}(t,i)$ is state-independent, and Theorem 1 applies locally. The CES repair ($\eta \neq 1$ in the final aggregator) repairs the degeneracy exactly on $\{\beta > 0\}$; hence Assumption 4(v) imposes $\beta \geq \underline{\beta} > 0$. Making the *intermediate* bundle CES as well would extend own-state relevance to all nodes regardless of β, at the cost of losing the Cobb–Douglas value shares of the bundle, and hence the scalar node block and the coupling through the spending moment (Remark 6); we do not pursue it here.

What CES repairs, and continuity at $\eta = 1$. Proposition 6 removes the payoff-side degeneracy, the only one CES is needed for: at $\eta \neq 1$ the marginal dividend $\frac{\varepsilon}{1-\alpha(i)}\frac{\mathcal{P}(i)}{x(i)}\Phi_i(\overline{v}(i,\kappa,z))$ loads on z through the threshold, the $\mathcal{L}_z^{(i)}$ term of the detrended marginal-$\mathfrak{q}$ equation (52) acts on a genuine function of z, the policy transmits z into the drift of κ, and the node distribution is no longer a product; the node problem is two-dimensional. The repair is continuous: as $\eta \to 1$ the CES equilibrium objects converge to the Cobb–Douglas system of this appendix, (127) (by Lemma 15, the node root being a continuously differentiable function of η by the implicit function theorem, since $F' < 0$ at the root for every $\eta > 0$), with the z-loadings (131) vanishing linearly in $\eta - 1$ and the K-loading of sales along with them. The degenerate system closed with (126) is therefore both an initial guess and a sanity limit for the CES computation, and the $\eta \to 1$ continuation is a natural correctness test for its numerical implementation (Appendix D, Step 7).

B.8 Dynamics: Households, the SDF and the Firm Problem

This appendix collects the results behind the first four paragraphs of Section 3.3: the deterministic stochastic discount factor, the firm's objective and first-order condition, the marginal-$\mathfrak{q}$ equation with its proof and its Feynman–Kac derivation, the node-level value equation, and the closed-loop node dynamics with their Fokker–Planck–Kolmogorov equation.

Households and the SDF.

Lemma 18 (Deterministic SDF). *Along any equilibrium path, the household SDF is the deterministic process $\Lambda_t = e^{-\rho t} C_t^{-\gamma}$, equivalently the instantaneous risk-free rate is* (48)*. Every shareholder unanimously instructs each firm to maximise $\int_0^\infty e^{-\int_0^t r_s \mathrm{d}s}\, \mathbb{E}[\mathit{dividend}_t]\, \mathrm{d}t$: with a unique deterministic SDF, Fisher separation is immediate, and firms discount expected cash flows at the risk-free rate. There are no risk premia in the limit economy.*

Proof. C_t deterministic makes Λ_t deterministic; (48) is $-\mathrm{d}\ln\Lambda_t/\mathrm{d}t$. Value maximisation with a common deterministic discount is unanimous across shareholders regardless of their exposure, which is the content of Fisher separation; no proof beyond linearity of present values is required, in contrast to the finite model, where the SDF is stochastic and separation must be argued. Note also that idiosyncratic node risk is fully diversified in the household portfolio by the exact LLN: the shock B_i^Z reaches the household only through D_t, whose realised value is deterministic because the family $\{B_i^Z\}$ is e.p.i. across nodes. The rationing rents add nothing to the argument: $\Pi_i^{\mathrm{r}} = [1 - \varrho_i(\overline{v})]S(t, i, K_i, z_i)$, with $\varrho_i(\overline{v}) := S^{\mathrm{f}}/S$ the share of node sales that accrues to the node's firms (Lemma 5(iv)), is a bounded multiple of node sales, a function of the node's own state like π_i, so $\int_0^1 \Pi_i^{\mathrm{r}} \mathrm{d}i$ is deterministic by Lemma 14 (the integrand is dominated by S) and enters D_t as one more diversified flow. ∎

In the language of asset pricing, a deterministic SDF means that the risk-neutral and physical measures coincide: node-level risk exists under $\mathbb{P}$ but is idiosyncratic, fully diversified in the household's portfolio, and therefore carries no premium.

The firm problem. The generic agent is the individual monopolist at node i, with own capital k, the node state (K, z) as a locally common exogenous-to-it pair, and the deterministic aggregate paths $\Theta_t := (\mathcal{P}(t,\cdot), \mathcal{D}(t,\cdot), Q(t), r_t)$. Using Lemma 16, the firm solves

$$w(t, i, k, K, z) = \max_{\{\iota\}} \int_t^\infty e^{-\int_t^u r_s \mathrm{d}s}\, \mathbb{E}[\pi(u, i, k_u; K_u, z_u) - \iota_u k_u] \mathrm{d}u, \tag{132}$$

subject to (103), taking as given the node's symmetric policy $\bar{\iota}(t, i, K, z)$ (so $\mathrm{d}K = [\Phi_i(\bar{\iota}) - \delta(i)]K\mathrm{d}t$), the OU process (93) of z, and Θ. The HJB is (49), with first-order condition

$$w_k\, \Phi_i'(\iota) = 1 \iff \iota = \frac{\mathfrak{q}^{\phi(i)/(\phi(i)-1)} - 1}{\phi(i)}, \qquad \mathfrak{q} := w_k, \tag{133}$$

using $\Phi_i'(\iota) = (1 + \phi\iota)^{-(\phi-1)/\phi}$ exactly as in the finite model's investment FOC (there $\tilde{q}_i = V_{K_i}/u'(C)$; here $\mathfrak{q}$ is directly the marginal value of installed capital in final-good units, since discounting already uses the SDF).

Write

$$\mathbb{G}^{(i)} f := f_K\, K[\Phi_i(\bar{\iota}) - \delta(i)] + \mathcal{L}_z^{(i)} f$$

for the generator of the node state (K, z) under the symmetric policy: pure advection in K and the OU generator in z. Differentiating the HJB in own capital and restricting to the diagonal gives the equation for the on-diagonal marginal value (50).

Proposition 7 (Marginal-$\mathfrak{q}$ equation). *Under regularity sufficient for differentiating* (49) *in k across the diagonal (or, by the Feynman–Kac route in the derivation after the proof, for $\mathfrak{q}$ that is $\mathcal{C}^1$ in (t, K) and $\mathcal{C}^2$ in z alone), the on-diagonal marginal value* (50) *solves the quasi-linear PDE*

$$r_t\,\mathfrak{q} = \underbrace{\zeta_i(\overline{v})[1-\xi_i(\overline{v})]\,\frac{S^{\mathrm{f}}(t,i,K,z)}{K}}_{\pi_k \text{ at } k=K,\ (115)} - \iota^* + \mathfrak{q}[\Phi_i(\iota^*)-\delta(i)] + \mathfrak{q}_K\,K[\Phi_i(\iota^*)-\delta(i)] + \mathcal{L}_z^{(i)}\mathfrak{q} + \partial_t\mathfrak{q}, \quad (134)$$

with $\iota^* = \iota^*(\mathfrak{q})$ given by (133), *$\overline{v} = \overline{v}(t,i,K,z)$ the node's threshold and $S^{\mathrm{f}}(t,i,K,z)$ its firm revenue (Lemma 15,* (113)*). Equation* (134) *is an asset-pricing equation for the marginal unit of installed capital: the required return $r_t\mathfrak{q}$ equals the marginal dividend $\pi_k - \iota^*$, plus the marginal unit's own growth in value $\mathfrak{q}[\Phi-\delta]$, plus capital gains from the drift of the node state and of z, plus the time derivative. It is first order in K and parabolic in z alone (first order throughout when $\varsigma^Z(i)=0$).*

Proof. Differentiate (49) with respect to k. The left side gives $r_t w_k$; the flow gives $\pi_k - \iota^*$ by the envelope theorem, the ι-derivative of the maximised term vanishing; $\mathcal{L}_z^{(i)}$ acts on z only and commutes with ∂_k, as does ∂_t. For the drift terms, on the diagonal a change in the node state moves both slots of $w_k(\cdot,K,K,\cdot)$, so by the chain rule $\mathfrak{q}_K = [w_{kk}+w_{kK}]_{k=K}$, and ($\bar\iota$ not depending on k)

$$\begin{aligned}&\partial_k\{w_k k[\Phi_i(\iota)-\delta(i)] + w_K K[\Phi_i(\bar\iota)-\delta(i)]\}\\ &\quad = w_k[\Phi_i(\iota)-\delta(i)] + w_{kk}k[\Phi_i(\iota)-\delta(i)] + w_{kK}K[\Phi_i(\bar\iota)-\delta(i)],\end{aligned}$$

which on the diagonal, where $\bar\iota = \iota^*$, is $\mathfrak{q}[\Phi_i(\iota^*)-\delta(i)] + \mathfrak{q}_K K[\Phi_i(\iota^*)-\delta(i)]$: the own-growth term and the advection part of $\mathbb{G}^{(i)}\mathfrak{q}$. Finally π_k at $k=K$ is (115), the capital share $\zeta_i(\overline{v})[1-\xi_i(\overline{v})]$ of firm revenue per unit of node capital (Lemma 16(ii)), evaluated at the node's threshold $\overline{v}(t,i,K,z)$ and firm revenue $S^{\mathrm{f}}(t,i,K,z)$. The sanity check $\Phi_i(\iota)=\iota$ (no adjustment costs) forces $\mathfrak{q}\equiv 1$, and (134) collapses to $r_t = \pi_k - \delta$, the frictionless user-cost condition; every derivative term drops out of this check. ∎

Feynman–Kac derivation of (134). Differentiating (49) across the diagonal asks for w_k to be $\mathcal{C}^1$ in (k,K) jointly and $\mathcal{C}^{1,2}$ in (t,z), one derivative more than the HJB itself. The same equation follows from the objective directly, under regularity on $\mathfrak{q}$ alone.

1. Fix the optimal control as an open-loop process $\{\iota_u\}$. By the envelope theorem, w_k is the derivative of the objective (132) at that control. Along it, (103) gives $\partial k_u/\partial k_t = k_u/k_t$, and K_u does not depend on k_t, so

$$w_k(t,k_t,K_t,z_t) = \mathbb{E}_t\int_t^\infty e^{-\int_t^u r_s\mathrm{d}s}\,[\pi_k(u,k_u;K_u,z_u)-\iota_u]\,\frac{k_u}{k_t}\,\mathrm{d}u.$$

2. On the diagonal $k_t = K_t$ the firm follows the node's policy and stays on it, $k_u = K_u$, so
$$K_t\, \mathfrak{q}(t, K_t, z_t) = \mathbb{E}_t \int_t^\infty e^{-\int_t^u r_s \mathrm{d}s} \left[\pi_k - \iota^*\right]_u K_u \, \mathrm{d}u,$$
a Feynman–Kac functional of the Markov state (K, z) with running payoff $K[\pi_k - \iota^*]$. Under the usual growth conditions, if $K\mathfrak{q}$ is $\mathcal{C}^1$ in (t, K) and $\mathcal{C}^2$ in z it solves $r_t(K\mathfrak{q}) = K[\pi_k - \iota^*] + \mathbb{G}^{(i)}(K\mathfrak{q}) + \partial_t(K\mathfrak{q})$.
3. By $\mathbb{G}^{(i)}(K\mathfrak{q}) = K\,\mathbb{G}^{(i)}\mathfrak{q} + \mathfrak{q}\,\mathbb{G}^{(i)}K$ (the generator is first order in K and K does not depend on z, so the product rule has no cross term) and $\mathbb{G}^{(i)}K = K[\Phi_i(\iota^*) - \delta(i)]$, dividing by K gives (134).

The marginal dividend is the capital share of firm revenue per unit of node capital, the first form of (115): given the bundle price $\mathcal{P}(t, i)$ it is the function $\frac{\varepsilon}{1-\alpha(i)} \frac{\mathcal{P}(t,i)}{x(i)} \Phi_i(\overline{v})$ of the node's threshold alone (Lemma 16(ii)), strictly decreasing in the threshold under Assumption 3, so that a node whose capacity is slack earns less on the marginal unit than one whose firms are rationed often. The node state enters (134) through this one term, the threshold and the firm revenue being functions of (K, z) through the static block; without capacity constraints the term is the constant fraction $\frac{\varepsilon-1}{\varepsilon}\alpha(i)$ of sales per unit of capital (Appendix C.6).

[Q17 | tag: hide | ANSWERED] the remark "rem:naive"

[Flag: set aside on 5 November 2026] This remark is held, unchanged and with its title and labels, inside a maroon box until what to do with it is decided. Its black-text pointers, to the remark in Remark 11 and to its value equation (135) in Assumption 6, Remark 11 and the step-kernel subsection of the model without capacity constraints, are in maroon boxes too, with the claims left black.

Remark 10 (Why a naive node-level HJB is wrong) It is tempting to shortcut the above by writing a "representative firm of node i" HJB directly in (K, z) with flow $S^{\mathrm{f}}(t, i, K, z) - X(t, i, K, z) - \iota K$ and to read investment off its K-derivative. This is *not* the equilibrium: the on-path value $W(t, i, K, z) := w(t, i, K, K, z)$ satisfies (given policies) the linear PDE

$$r_t W = S^{\mathrm{f}} - X - \iota^* K + W_K\, K[\Phi_i(\iota^*) - \delta(i)] + \mathcal{L}_z^{(i)} W + \partial_t W, \tag{135}$$

with $S^{\mathrm{f}} - X = [1 - \xi_i(\overline{v})]S^{\mathrm{f}}$ the firms' operating profit of (114) (the rationing rent accrues to the aggregators, not to the firms, and does not enter their value), but its derivative $W_K = [w_k + w_K]_{k=K}$ is *not* the marginal $\mathfrak{q}$ that governs investment: $w_K \le 0$ captures the demand-curve externality (more node capital depresses the node price), which an atomistic firm rightly ignores. Choosing ι to satisfy $W_K \Phi' = 1$ would let the node internalise its own demand curve, a planner-of-the-node distortion. Investment must come from (133) with $\mathfrak{q} = w_k$, i.e. from (134); equation (135) is only needed afterwards, if one wants equity values. In the finite model the analogous wedge is the observation there that the private marginal profit $\pi_i^k = \zeta_i(\overline{v}_i)[1 - \xi_i(\overline{v}_i)]S_i^{\mathrm{f}}/K_i$ of (33) differs from Q_{K_i} (Remark 1).

Node state dynamics and the Fokker–Planck–Kolmogorov equation. In equilibrium the node state follows the closed-loop dynamics

$$\mathrm{d}K_i = [\Phi_i(\iota^*(t,i,K_i,z_i)) - \delta(i)]K_i\,\mathrm{d}t, \qquad \mathrm{d}z_i = -\theta(i)z_i\,\mathrm{d}t + \varsigma^Z(i)\,\mathrm{d}B_i^Z. \tag{136}$$

The K-coordinate carries no Brownian term: randomness enters K only through the policy's dependence on z, so the pair is hypoelliptic whenever the policy depends on z (Hörmander's condition); where it does not, as under Cobb–Douglas demand(Appendix B.7), the K-marginal is a moving point mass. The node distribution $\mathfrak{M}(t,i,\cdot)$, with density $\mathfrak{m}(t,i,K,z)$ when it exists, solves the Fokker–Planck–Kolmogorov (FPK) equation

$$\partial_t \mathfrak{m} = -\frac{\partial}{\partial K}\{[\Phi_i(\iota^*) - \delta(i)]K\,\mathfrak{m}\} + \frac{\partial}{\partial z}\{\theta(i)\,z\,\mathfrak{m}\} + \frac{\varsigma^Z(i)^2}{2}\frac{\partial^2}{\partial z^2}\,\mathfrak{m}, \tag{137}$$

one equation per node, in two state variables, adjoint to the state generator $\mathbb{G}^{(i)}$ defined above. This is the FPK(α) of the Caines–Huang system; the coupling differs from that system as described in Section 3.3.

B.9 Balanced Growth, Detrending and the GMFG System

This appendix collects the results behind the last two paragraphs of Section 3.3: the scale invariance of the economy, the detrended static block, state dynamics, $\mathfrak{q}$-equation and FPK equation with the proposition that collects them, the effective-discounting assumption behind Definition 6, and the recovery of values and other equilibrium objects from a solution. It then gives the full account of what is established, assumed and open(Section 3.4) and sets the equilibrium concept against the graphon mean field games of Caines and Huang (2021).

Homogeneity and detrending.

Lemma 19 (Scale invariance). *Fix the z-profile and the ensemble $\mathfrak{M}_G$, and let Assumptions 4 and 5 hold. Scale every node's capital, and hence $\overline{K}$, by $\Lambda > 0$. Then in the static block all prices $(P, \mathcal{P})$ and all thresholds $\overline{v}$ are invariant, and with the thresholds every function of them: utilization U, the shares ξ_i, ζ_i, ϱ_i, the markup and the rationed share of firms; all nominal quantities $(S, S^{\mathrm{f}}, X, \mathcal{D}, Q, \Pi^{\mathrm{r}}, \pi)$ scale by Λ; consequently π_k in* (115) *is invariant, and $\mathfrak{q}$ in* (50) *is invariant.*

Proof. Let $\Lambda > 0$. The scaled configuration replaces every node state (K,z) by $(\Lambda K, z)$ (each distribution $\mathfrak{M}(t,i,\cdot)$ becomes its pushforward under $K \mapsto \Lambda K$, the z-marginals untouched), and $\overline{K} \mapsto \Lambda\overline{K}$.

Step 1: a candidate static equilibrium for the scaled economy. Guess that prices and thresholds are unchanged and nominal quantities scale: $P'(t,i,\Lambda K,z) = P(t,i,K,z)$, $\overline{v}'(t,i,\Lambda K,z) = \overline{v}(t,i,K,z)$ and $S'(t,i,\Lambda K,z) = \Lambda\, S(t,i,K,z)$, so that the spending map of (113) satisfies $X'(t,i,\Lambda K,z) = \xi_i(\overline{v})\varrho_i(\overline{v})\Lambda S = \Lambda X(t,i,K,z)$. Check the conditions of Definition 5 in turn:

(a) *Moments and aggregates.* Under the guess and the pushforward distributions, $m'_P = m_P$ (log prices unchanged), $m'_X = \Lambda m_X$ and $m'_S = \Lambda m_S$. Then (44) gives $\mathcal{P}' = \mathcal{P}$, and (45)–(46), being linear in (m_X, m_S), give $\mathcal{D}' = \Lambda\mathcal{D}$ and $Q' = \Lambda Q$.

(b) *Pricing/threshold and sales.* The pricing/threshold relation (107) contains neither K nor S: at unchanged $(P, \mathcal{P}, z)$ it holds at the unchanged threshold. The right-hand side of the sales relation (108) is linear in K at fixed $(\mathcal{P}, \overline{v})$: replacing K by ΛK at the unchanged threshold multiplies it by Λ, matching $S' = \Lambda S$.

(c) *Demand.* In (109), the left side is ΛS and the right side is $\beta P^{1-\eta}(\Lambda Q) + \Lambda\mathcal{D} = \Lambda\big[\beta P^{1-\eta} Q + \mathcal{D}\big]$: both sides scale by Λ, so the equation holds at the *same* P. Equivalently, the scalar equation (110) is homogeneous of degree one in $(K, Q, \mathcal{D})$ at fixed $(\mathcal{P}, z)$, so its root $\overline{v}$ is homogeneous of degree zero.

(d) *Numeraire.* (47) involves only prices and the z-marginals, both unchanged.

Step 2: the candidate is the equilibrium. Given the scaled aggregates, the node block has a unique solution at every state (Lemma 15(iv), under Assumption 5); the guess solves it, so it *is* the solution. Hence $(P, \mathcal{P}, \overline{v})$ are invariant, and with the threshold every function of it (U_i, ξ_i, ζ_i, ϱ_i, the markup $\varepsilon_i/(\varepsilon_i - 1)$, the rationed share $1 - G_i$), while $(S, \mathcal{D}, Q)$ and with them $S^{\mathrm{f}} = \varrho_i S$, $X = \xi_i S^{\mathrm{f}}$, $\pi = [1 - \xi_i]S^{\mathrm{f}}$ and $\Pi^{\mathrm{r}} = [1 - \varrho_i]S$ scale by Λ: the static block is homogeneous of degree one in the capital profile.

Step 3: marginal profit is invariant. π_k in (115) equals $\zeta_i(\overline{v})[1 - \xi_i(\overline{v})]S^{\mathrm{f}}/K$, a degree-zero factor times a ratio of two degree-one objects:

$$\pi'_k(t, i, \Lambda K, z) = \zeta_i(\overline{v})[1 - \xi_i(\overline{v})]\,\frac{\Lambda S^{\mathrm{f}}}{\Lambda K} = \pi_k(t, i, K, z);$$

equivalently, in the second form of (115) it is $\frac{\varepsilon}{1-\alpha(i)}\frac{\mathcal{P}(t,i)}{x(i)}\Phi_i(\overline{v})$, a function of the invariant $(\mathcal{P}, \overline{v})$ alone.

Step 4: marginal value is invariant. Define the candidate $\mathfrak{q}'(t, i, \Lambda K, z) := \mathfrak{q}(t, i, K, z)$ and check it solves (134) in the scaled economy: the dividend term is invariant (Step 3); the terms $r_t\mathfrak{q}$, $\iota^*(\mathfrak{q})$ and $\mathfrak{q}[\Phi_i(\iota^*) - \delta]$ involve only $\mathfrak{q}$ itself; the advection term is invariant because, by the chain rule, $\mathfrak{q}'_K(t, i, \Lambda K, z) \cdot \Lambda K = [\mathfrak{q}_K(t, i, K, z)/\Lambda] \cdot \Lambda K = \mathfrak{q}_K K$; and $\mathcal{L}_z^{(i)}\mathfrak{q}$ and $\partial_t\mathfrak{q}$ act on untouched arguments. Under the regularity guaranteeing uniqueness for (134), $\mathfrak{q}'$ is the scaled marginal value, i.e. $\mathfrak{q}$ is invariant. Equivalently, at the level of the value function the objective (132) is linear in the common scale of (k, K, Θ) (the flow payoff $\pi - \iota k$ scales by Λ along scaled paths, and scaled paths are feasible because (103) is linear in k), so $w(t, i, \Lambda k, \Lambda K, z; \Lambda\Theta) = \Lambda\, w(t, i, k, K, z; \Theta)$, with $\Lambda\Theta$ meaning $(\mathcal{P}, \Lambda\mathcal{D}, \Lambda Q, r)$; differentiating in k shows w_k is homogeneous of degree zero, the same statement. ∎

With $\overline{K}_t$, g_t and κ as in (51), the adding-up constraint reads

$$\int_0^1 \mathbb{E}[\kappa_i(t)]\,\mathrm{d}i = 1 \qquad \text{for all } t. \tag{138}$$

Detrended (hatted) nominal objects per unit of aggregate capital:

$$\widehat{s} := S/\overline{K}_t, \qquad \widehat{s}^{\mathrm{f}} := S^{\mathrm{f}}/\overline{K}_t, \qquad \widehat{x} := X/\overline{K}_t, \qquad \widehat{\mathcal{D}} := \mathcal{D}/\overline{K}_t, \qquad q_t := Q_t/\overline{K}_t,$$
$$\widehat{C}_t := C_t/\overline{K}_t, \qquad \widehat{m_S} := m_S/\overline{K}_t, \qquad \widehat{m_X} := m_X/\overline{K}_t.$$

Prices, thresholds and utilization need no detrending by Lemma 19. The per-node static block (107)–(109) becomes, verbatim,

$$\text{(pricing/threshold)} \qquad P\,\mathrm{Y}_i(\overline{v}) = \frac{\mathcal{P}(t,i)}{e^z\, x(i)^{\alpha(i)}}, \tag{139}$$

$$\text{(sales)} \qquad \widehat{s} = \frac{\mathcal{P}(t,i)\,\kappa}{x(i)}\, \mu_{i,1/\varepsilon}^{\frac{\varepsilon}{\varepsilon-1}}\, \frac{U_i(\overline{v})}{\mathrm{Y}_i(\overline{v})}, \tag{140}$$

$$\text{(demand)} \qquad \widehat{s} = \beta(i)\, P^{1-\eta}\, q_t + \widehat{\mathcal{D}}(t,i), \tag{141}$$

with $P = P(t,i,\kappa,z)$, $\overline{v} = \overline{v}(t,i,\kappa,z)$, $\widehat{s} = \widehat{s}(t,i,\kappa,z)$, the detrended firm revenue and spending

$$\widehat{s}^{\mathrm{f}} = \varrho_i(\overline{v})\,\widehat{s}, \qquad \widehat{x} = \xi_i(\overline{v})\,\widehat{s}^{\mathrm{f}}$$

from (113)and (118), and Lemma 15 applying unchanged: the scalar equation (110) with $(\kappa, q_t, \widehat{\mathcal{D}})$ in place of $(K, Q, \mathcal{D})$ has a root in $(\underline{v}_i, \infty)$ at every (κ, z), unique under Assumption 5. Without capacity constraints the threshold drops out and the triple collapses to the pair (196) (Appendix C.6).

Detrended state dynamics. The quotient rule on $\kappa = K_i/\overline{K}_t$, both factors being differentiable in t, gives

$$\dot{\kappa} = \kappa\,[\Phi_i(\iota^*) - \delta(i) - g_t] \qquad \text{i.e.} \qquad \mu^{\kappa}(t,i,\kappa,z) = \Phi_i(\iota^*) - \delta(i) - g_t, \tag{142}$$

the continuum analogue of the finite share dynamics (38): subtracting g_t is exactly the μ^K correction, with the cross-sectional growth rate (55) (with g_t in place of g off the balanced growth path) playing the role of μ^K. This is what makes (138) propagate: $\frac{\mathrm{d}}{\mathrm{d}t}\int \mathbb{E}[\kappa]\mathrm{d}i = \int \mathbb{E}[\kappa(\Phi-\delta)]\mathrm{d}i - g_t \int \mathbb{E}[\kappa]\mathrm{d}i = 0$.

Detrended $\mathfrak{q}$-equation. On a path where $\mathfrak{q}(t,i,K,z) = \mathfrak{q}(t,i,\kappa,z)$ (justified by Lemma 19), holding (K,z) fixed we have $\partial_t \mathfrak{q}\big|_K = \partial_t \mathfrak{q}\big|_\kappa - g_t\,\kappa\,\mathfrak{q}_\kappa$ and $\mathfrak{q}_K K = \mathfrak{q}_\kappa \kappa$. Substituting into (134):

$$\begin{aligned} r_t\,\mathfrak{q} = {} & \zeta_i(\overline{v})[1-\xi_i(\overline{v})]\,\frac{\widehat{s}^{\mathrm{f}}(t,i,\kappa,z)}{\kappa} - \iota^*(\mathfrak{q}) + \mathfrak{q}[\Phi_i(\iota^*) - \delta(i)] \\ & + \mathfrak{q}_\kappa\,\kappa[\Phi_i(\iota^*) - \delta(i) - g_t] + \mathcal{L}_z^{(i)}\mathfrak{q} + \partial_t \mathfrak{q}, \end{aligned} \tag{143}$$

with $\iota^*(\mathfrak{q})$ from (133) and, from the detrended static block (139)–(141), $\overline{v} = \overline{v}(t,i,\kappa,z)$ and $\widehat{s}^{\mathrm{f}} = \varrho_i(\overline{v})\widehat{s}(t,i,\kappa,z)$; equivalently the flow is $\frac{\varepsilon}{1-\alpha(i)}\frac{\mathcal{P}(t,i)}{x(i)}\Phi_i(\overline{v}(t,i,\kappa,z))$, the second form of (115). The detrending is transparent: the marginal unit's dividend and own-growth terms are unchanged (they are per-unit objects: the dividend is the ratio of two degree-one quantities, or equivalently a function of the degree-zero threshold), the capital-gains drift acquires the $-g_t$ of (142), and the explicit ∂_t is now with respect to the stationary coordinates. On the balanced growth path, with $\partial_t \equiv 0$ and $r = \rho + \gamma g$, (143) is the backward equation (52) of Definition 6.

Detrended FPK equation. The distribution of (κ, z) at node i, with density $\mathfrak{m}(t, i, \kappa, z)$ when it exists, solves

$$\partial_t \mathfrak{m} = -\frac{\partial}{\partial \kappa}\{\kappa[\Phi_i(\iota^*(\mathfrak{q})) - \delta(i) - g_t]\,\mathfrak{m}\} + \frac{\partial}{\partial z}\{\theta(i)\, z\, \mathfrak{m}\} + \frac{\varsigma^Z(i)^2}{2}\frac{\partial^2}{\partial z^2}\,\mathfrak{m}, \tag{144}$$

whose stationary form is the forward equation (53) of Definition 6.

Household in detrended form. $\widehat{C}_t = q_t - \int_0^1 \mathbb{E}[\iota^*\kappa]\mathrm{d}i$ by final-good clearing $Q = C + I$,[29] and (48) becomes

$$r_t = \rho + \gamma\left(g_t + \frac{\mathrm{d}}{\mathrm{d}t}\ln \widehat{C}_t\right). \tag{145}$$

Proposition 8 (Detrended system). *Along any equilibrium path, under Assumptions 4 and 5:*

(i) the per-node static block in detrended variables is (139)–(141)*, prices, thresholds and utilization needing no detrending, and the detrended revenue, spending and rent are* $\widehat{s}^{\mathrm{f}} = \varrho_i(\overline{v})\widehat{s}$, $\widehat{x} = \xi_i(\overline{v})\widehat{s}^{\mathrm{f}}$ *and* $\widehat{\Pi}^{\mathrm{r}} = [1 - \varrho_i(\overline{v})]\widehat{s}$*;*

(ii) the capital density follows (142)*;*

(iii) the growth rate satisfies (55)*, and the adding-up constraint* (138) *propagates;*

(iv) the detrended marginal value solves (143)*;*

(v) the detrended node distribution solves (144)*;*

(vi) the household Euler equation reads (145)*.*

Proof. (i) is Lemma 19: the static block is homogeneous of degree one in the capital profile and prices and thresholds are invariant, so dividing the sales relation (108) and the demand relation (109) by $\overline{K}_t$ gives (140)–(141), while the pricing/threshold relation (107), which contains no nominal quantity, is (139) verbatim; the detrended revenue, spending and rent follow from (113)–(114)and (118) by the same division, since ϱ_i and ξ_i are functions of the invariant threshold.

(ii)–(iii) Fix a node i and write $\kappa_t = K_i(t)/\overline{K}_t$. $\overline{K}_t$ is deterministic (Lemma 13 applied to $\{K_i(t)\}_i$, which is e.p.i. across nodes because the driving B_i^Z are) and $\mathcal{C}^1$, with $g_t = \dot{\overline{K}}_t/\overline{K}_t$ by (51); $K_i(t)$ is differentiable, its growth rate being the continuous $\Phi_i(\iota^*) - \delta(i)$. The quotient rule therefore gives

$$\dot{\kappa} = \frac{\dot{K}_i}{\overline{K}_t} - \frac{K_i\,\dot{\overline{K}}_t}{\overline{K}_t^2} = \underbrace{\frac{K_i}{\overline{K}_t}}_{=\kappa}[\Phi_i(\iota^*) - \delta(i)] - \kappa\, g_t,$$

which is (142). The consistency condition pinning g_t is (55) (with g_t in place of g off the balanced growth path): differentiate $\overline{K}_t = \int_0^1 \mathbb{E}[K_i(t)]\mathrm{d}i$ under the integral and the

[29] Walras' law with respect to (46) and dividends, the rationing rents included, is immediate: $D_t = \int_0^1 \mathbb{E}[\pi_i - \iota_i K_i + \Pi_i^{\mathrm{r}}]\mathrm{d}i = \int_0^1 \mathbb{E}[S_i - X_i]\mathrm{d}i - I_t = Q_t - I_t$ by (114) and (46).

expectation (justified by dominated convergence, $(\Phi_i(\iota^*) - \delta(i))K_i$ being integrable against the ensemble locally uniformly in t),

$$\dot{\overline{K}}_t = \int_0^1 \frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[K_i(t)]\mathrm{d}i = \int_0^1 \mathbb{E}[(\Phi_i(\iota^*) - \delta(i))K_i]\mathrm{d}i,$$
$$g_t = \frac{\dot{\overline{K}}_t}{\overline{K}_t} = \int_0^1 \mathbb{E}[(\Phi_i(\iota^*) - \delta(i))\,\kappa]\mathrm{d}i,$$

dividing by the deterministic $\overline{K}_t$ inside the expectation. Subtracting g_t in (142) is thus exactly the finite model's $-\mu^K$ correction: a node's capital *share* drifts up or down according to whether its own growth rate beats the cross-sectional average.

(iv) By Lemma 19, $\mathfrak{q}(t,i,K,z) = \mathfrak{q}(t,i,\kappa,z)$ along the path; holding (K,z) fixed, $\partial_t\mathfrak{q}\big|_K = \partial_t\mathfrak{q}\big|_\kappa - g_t\,\kappa\,\mathfrak{q}_\kappa$ and $\mathfrak{q}_K K = \mathfrak{q}_\kappa\kappa$, and substituting into (134) gives (143).

(v) (144) is the FPK equation adjoint to the generator of the detrended closed-loop dynamics (142), $\mathrm{d}z = -\theta(i)z\mathrm{d}t + \varsigma^Z(i)\mathrm{d}B_i^Z$, exactly as (137) is adjoint to (136) (the adjoint identity uses vanishing boundary fluxes at $\kappa \in \{0,\infty\}$ and $z \to \pm\infty$; $\kappa = 0$ is unattainable, the drift of (142) being proportional to κ, so no condition is imposed there).

(vi) With $C_t = \widehat{C}_t\overline{K}_t$, $\dot{C}_t/C_t = g_t + \frac{\mathrm{d}}{\mathrm{d}t}\ln\widehat{C}_t$ in (48). ■

The GMFG system. Section 3.3 states the equilibrium as a Caines–Huang-type system on the stationary balanced growth path (Definition 6), where $\partial_t \equiv 0$, $g_t \equiv g$ and $r = \rho + \gamma g$. Well-posedness of values requires:

Assumption 6 (Effective discounting). $\rho + (\gamma - 1)g > 0$ at the equilibrium g (verified ex post). This makes $r - g > 0$, so detrended firm values(135) and household utility are finite.

Remark 11 (Recovering values and other equilibrium objects) Given a solution, detrended equity value $\widehat{w}(i,\kappa,z) := W/\overline{K}$ solves the *linear* PDE(from (135))

$$(r-g)\,\widehat{w} = [1 - \xi_i(\overline{v})]\widehat{s}^{\mathrm{f}} - \iota^*\kappa + \widehat{w}_\kappa\,\kappa[\Phi_i(\iota^*) - \delta(i) - g] + \mathcal{L}_z^{(i)}\widehat{w}, \tag{146}$$

with $\overline{v} = \overline{v}(i,\kappa,z)$ and $\widehat{s}^{\mathrm{f}} = \varrho_i(\overline{v})\widehat{s}$, and with the policies fixed from (52); recall $\widehat{w}_\kappa \neq \mathfrak{q}$(Remark 10). Aggregate wealth is $\int \mathbb{E}_i[\widehat{w}]\mathrm{d}i \cdot \overline{K}_t$ plus the capitalised rationing rents, GDP is $q\overline{K}_t$, node Domar weights are $\widehat{m_S}(i)/q$, node utilization is $\mathbb{E}_i[U]$ and aggregate utilization is (116) in detrended form, $\overline{U} = \frac{1}{q}\int_0^1 \mathbb{E}_i[(\widehat{s} - \widehat{x})U]\mathrm{d}i$, the aggregate rationing rent is $\int_0^1 \mathbb{E}_i[(1 - \varrho_i(\overline{v}))\widehat{s}]\mathrm{d}i \cdot \overline{K}_t$, and so on.

[Q18 | tag: hide | ANSWERED] from the start of paragraph "Status: what is established, what is assumed, what is open." until and including paragraph "Two clarifications."

[Flag: set aside on 5 November 2026] This block, from the paragraph on the status of the equilibrium concept to the paragraph of two clarifications, is held, unchanged

and with its headings and labels, inside a maroon box until what to do with it is decided. It crosses a subsection boundary: it holds the whole subsection on the three models side by side, with the step-kernel system (156)–(159), Section B.10, Lemma 20 and the relation to graphon mean field games (Appendices B.9–B.9.4). Every pointer to it from black text is in a maroon box too, with the claim left black; the following black claims now cite a boxed source:

- the opening of the continuum-model section (the relation to graphon mean field games);
- the step-kernel paragraph of the section on the dynamic equilibrium (the finite system the paper computes);
- the comparison of the finite and continuum models (the three systems stated in full; the goods-units discounting of Lemma 20);
- the appendix on step kernels and graphon limits (section-wise rather than cut convergence; the bases of the failures table);
- the model on $[0, 1]$ (the six differences from the cluster-based construction) and the static block (the static analogue of graphon mean-field coupling);
- the step-kernel subsection of the model without capacity constraints (the two structural contrasts; what the numerical appendix solves);
- the numerical appendix: its scope paragraph, the operator conventions, Step 0 (block weights and entropy constants), Step 1 (the transport operator), Step 2 (the static block), Step 3 (the backward equation), Step 4 (the forward equation) and Step 5 (the coupling and the scale family);
- the overview of the calibration appendix (the calibrated object).

Status: what is established, what is assumed, what is open. This paragraph gives the full account, result by result, of the status of the equilibrium concept of Definition 6, which Section 3.4 summarises.

Established. The static layer is characterised up to the shape assumptions on the demand-shock distribution: cross-sectional aggregates are deterministic functionals of the ensemble (Lemma 14, on the exact law of large numbers of Lemma 13, which is cited from Sun (2006) and Bayraktar et al. (2023) rather than proved here); the per-node static block has a solution at every state, unique under Assumption 5, with the threshold increasing in the node's capital and moving with its productivity according to the sign of $1 - \eta$ (Lemma 15); node-level sales are state-independent under Cobb–Douglas final demand and state-relevant under CES (Lemma 17, Theorem 1, Proposition 6), and under Cobb–Douglas the threshold, hence utilization and every revenue share, is free of the node's productivity (Lemma 17(ii)), so that the Cobb–Douglas degeneracy in productivity survives capacity while the closed forms of the model without capacity constraints do not (Proposition 4); under Cobb–Douglas the node distribution is a point mass in capital, the deterministic rest point of Proposition 5, and the

marginal dividend diverges as the node's capital density vanishes, unconditionally, because atomism holds the aggregates fixed (Lemma 17(iv) and the mean-reversion paragraph of Section B.7.2); and the static block is homogeneous of degree one in the capital profile, with prices, thresholds and utilization of degree zero (Lemma 19). The dynamic layer is characterised conditionally on equilibrium: the household SDF is deterministic (Lemma 18), the individual firm's marginal value solves the marginal-q equation (Proposition 7), and the detrended system is Proposition 8. The step-kernel embedding is exact (Lemma 10): for a finite input–output table the continuum system *is* the finite system of Appendix B.10, with no kernel approximation.

Assumed where flagged. Three shape restrictions on the demand-shock distribution carry the capacity block and are verified numerically for the log-normal specification rather than proved: single-peakedness of the pricing problem (Assumption 2), concavity of profit in own capital (Assumption 3) and monotonicity of the threshold map (Assumption 5); the first is used for the firm's pricing, the second for the sign and monotonicity of the marginal dividend and for the uniqueness of the balanced-growth thresholds, the third for the uniqueness of the per-node static root. Regularity sufficient for existence and uniqueness of solutions of the firm HJB, and, at $\eta = 1$, of the symmetric node fixed point $(w, \bar{\iota})$ (which excludes sunspot node policies that depend on z), is assumed where it is used (Theorem 1, Lemma 19); $w \in \mathcal{C}^{1,2}$ is assumed in Proposition 7; and the integrability hypotheses of Lemma 14 (finite first moments of $\ln P$, S, X and $P^{1-\eta}$ against the node distributions, uniformly in i) are treated as equilibrium regularity conditions rather than derived from primitives. The last point is mild without capital shocks. $\mathbb{E}_i[\kappa] < \infty$, which (138) needs, holds once the stationary κ-marginal has bounded support, and it has bounded support whenever there are $0 < \underline{\kappa} < \overline{\kappa} < \infty$ such that the geometric drift $\Phi_i(\iota^*) - \delta(i) - g$ of κ in (142) is positive for all $\kappa < \underline{\kappa}$ and negative for all $\kappa > \overline{\kappa}$, uniformly in $z \in \mathbb{R}$: outside $[\underline{\kappa}, \overline{\kappa}]$ the node's capital then moves monotonically towards the band, so no stationary mass sits there. This is a mean-reversion property of the equilibrium policy, verified ex post rather than proved, and the uniformity in z over the unbounded Gaussian stationary distribution of z is what a check on a truncated grid cannot verify. $\mathbb{E}_i[P^{1-\eta}] < \infty$, which the numeraire needs, is not a consequence of the bounded κ-support alone, because the threshold moves with z when $\eta \neq 1$ (Lemma 15); it follows instead from the sales moment, since $\beta(i)P^{1-\eta}q \leq \widehat{s}$ by (141), so that $\mathbb{E}_i[P^{1-\eta}] \leq \widehat{m_S}(i)/(\beta(i)q)$, and $\widehat{m_S}(i) < \infty$ is the operative hypothesis (empty at $\eta = 1$, where $P^{1-\eta} = 1$). Under Cobb–Douglas final demand the κ-marginal is a point mass at the rest point of Proposition 5, and the mean reversion that makes the rest point the long-run position of every node is again verified numerically (Appendix B.7.2).

Open. Two statements are conjectured, not proved, and two questions are left open.

Conjecture: existence. Under Assumptions 4, 5 and 6, a stationary GMFG equilibrium in the sense of Definition 6 exists. For a step kernel with N blocks the statement reduces to a finite-dimensional fixed point: the map from the coupling objects $(\Theta, g) = (\mathcal{P}_n, \widehat{\mathcal{D}}_n, q, g)$ to the objects regenerated by (159), each evaluation of which

solves N decoupled two-dimensional PDE pairs.

The natural route mirrors the graphon mean field game literature: define the best-response map $\mathfrak{M}_G \mapsto$ (statics) $\mapsto$ (HJB(i) solutions) $\mapsto$ (FPK(i) distributions) $= \Phi(\mathfrak{M}_G)$ and apply a Schauder-type fixed-point argument on a set of ensembles with uniform moment and continuity-in-i bounds. Continuity of the node problem in $(\mathcal{P}, \widehat{\mathcal{D}}, q)$ follows from Lemma 15 under Assumption 5 (the root is $\mathcal{C}^1$ in the aggregates by the implicit function theorem, since the derivative of the scalar equation (110) at the root is nonzero), and continuity in i from continuity of the primitives; what is missing is the compactness and the regularity of the per-node marginal-$\mathfrak{q}$ problem under unbounded payoffs, the step that the boundedness hypotheses of Theorem 3.12 of Caines and Huang (2021) supply in their setting and that fails here (Appendix B.9.4). For the step kernel the moment map is finite-dimensional, so a degree-theoretic or Brouwer argument on (Θ, g) with standard HJB/FPK regularity for each block is the most tractable path; this is the theorem the paper most needs.

Conjecture: ε-optimality in the finite economy. Let the finite J-sector economy of Section 2.1 have primitives given by the step functions of Section 3.1, and let every sector's firms use the step-kernel GMFG policy $\iota^*(n, \kappa, z)$ evaluated at their own block. Then no firm can improve its value by more than ε_J by deviating, where $\varepsilon_J \to 0$ as $J \to \infty$ and ε_J contains two terms: the within-block granularity term of the mean-field literature and an $O(J^{-1/2})$ aggregate-fluctuation term, because the finite economy carries aggregate risk (the J sector productivity shocks do not wash out) that the limit policies ignore.

The deviation metric should be stated in value terms, following the structure of Theorem 4.4 of Caines and Huang (2021); numerical evidence would come from finite-J solutions on the same input–output table, which this paper does not compute, set against the step-kernel solution (Appendix D, Step 7).

Static fixed point under Cobb–Douglas. The first open question is new to the capacity model. Without capacity constraints the static block at $\eta = 1$ is solved in closed form (a linear Fredholm equation for the price moment and the Domar density for sales, Appendix C.5), so that the existence conjecture concerns the dynamic layer only. Under capacity even the Cobb–Douglas static block is a fixed point over the two functions $(m_P, \widehat{m_X})$ on $[0, 1]$ given the capital distributions (Proposition 4), because the bundle price enters every node's threshold and the inbound demand needs the buyers' spending moment; its existence and uniqueness are not established (the contraction constant of the Fredholm equation of the model without capacity constraints does not carry over, Appendix B.7.1), and the same fixed point sits inside the CES system at every η.

Uniqueness. Uniqueness should not be expected in general: monotonicity in the sense of Lasry and Lions (2007) has no obvious analogue here, because the mean field

enters through general-equilibrium prices rather than through an additively separable cost. Multiplicity is therefore probed numerically (Appendix D, Step 7) and reported rather than assumed away.

Relation to graphon mean field games. The equilibrium concept of Definition 6 is a graphon mean field game (GMFG) in the sense of Caines and Huang (2021),[30] but the model is not an instance of their theorems. The rest of this appendix makes the relation precise for a reader who has not met GMFGs: Appendix B.9.1 recalls the classical mean field game in the form macroeconomists know, Appendix B.9.2 summarises the Caines–Huang construction and results, Appendix B.9.3 lists exactly how the present model differs, and Appendix B.9.4 records what can be borrowed from their theory and what has to be developed in-house. Throughout, the notation is that of this paper: nodes are $i, j \in [0,1]$, the generic graphon of their theory is written $g(i,j)$ to keep it apart from the input–output kernel $a(i,j)$, and $\mathfrak{M}_G = \{\mathfrak{M}_i\}_{i\in[0,1]}$ is the ensemble of local mean fields.[31]

B.9.1 The Classical Mean Field Game

The construction. A continuous-time heterogeneous-agent economy of the Bewley–Huggett–Aiyagari type is a mean field game (Achdou et al., 2022). A continuum of ex-ante identical agents, each with an idiosyncratic state hit by an independent shock, solves a Hamilton–Jacobi–Bellman (HJB) equation taking as given a small set of aggregates (prices) that are functionals of the cross-sectional distribution m_t; the distribution evolves according to the Fokker–Planck–Kolmogorov equation driven by the agents' policies; and a market-clearing or consistency condition requires the aggregates generated by m_t to be the ones the agents took as given.

Why it is tractable. Two facts do the work. First, with a continuum of agents carrying independent shocks the cross-sectional distribution is deterministic, by a law of large numbers, so there is no aggregate risk and agents need not forecast anything stochastic. Second, because no single agent can move the aggregates, the equilibrium is Walrasian in the aggregates and the fixed point is over the deterministic path of m_t — in a stationary equilibrium, over one constant per aggregate — rather than over strategies against named opponents. The distribution is thus an *unknown of the problem* but not an *argument of any value function*, so the HJB keeps the dimension of the individual state, while the general-equilibrium coupling survives the limit as the consistency fixed point. This is why an Aiyagari equilibrium is simultaneously a Nash equilibrium of the continuum game. The limit pays for itself through ε-Nash theorems (Huang et al., 2006; Lasry & Lions, 2006b, 2006a): if all N agents of the finite economy use the mean-field policy, no agent can gain more than $\varepsilon_N \to 0$ by deviating. The

[30] Theorem and hypothesis numbers refer to the arXiv version 2008.10216v4 of that paper.

[31] Dictionary: our $i, j, g(i,j), \mathfrak{M}_i, \mathfrak{M}_G$ are their $\alpha, \beta, g(\alpha,\beta), \mu_\alpha, \mu_G$; their clusters $\mathcal{C}_l$, node count M_k and hypotheses (H1)–(H11) are cited under their own names.

graphon mean field game generalises the construction from an unstructured population to a population living on a network whose size also goes to infinity.

B.9.2 Caines–Huang in Brief

The finite primitive. Caines and Huang start from a sequence of finite weighted graphs G_k with M_k nodes and weights $g^k_{ij} \in [0,1]$. A node is not an agent: node l hosts a *cluster* $\mathcal{C}_l$ of agents, $\mathcal{C}(i)$ denotes the cluster of agent i, and the total population is $N = \sum_l |\mathcal{C}_l|$. Agent i's scalar state x_i follows a controlled diffusion whose drift has two averaged interaction terms, one within the cluster and one across the graph,

$$f_0(x_i,u_i,\mathcal{C}(i)) = \frac{1}{|\mathcal{C}(i)|}\sum_{j\in\mathcal{C}(i)} f_0(x_i,u_i,x_j), \qquad f_{G_k}\Big(x_i,u_i,g^k_{\mathcal{C}(i)}\Big) = \frac{1}{M_k}\sum_{l=1}^{M_k} g^k_{\mathcal{C}(i)\mathcal{C}_l}\,\frac{1}{|\mathcal{C}_l|}\sum_{j\in\mathcal{C}_l} f(x_i,u_i,x_j), \tag{147}$$

combined additively, $\mathrm{d}x_i = \big[f_0 + f_{G_k}\big]\mathrm{d}t + \sigma\,\mathrm{d}w_i$, with independent Brownian motions w_i and i.i.d. initial states; the running cost $l_0 + l_{G_k}$ has the same two-term anatomy, and agent i minimises $\mathbb{E}\int_0^T (l_0 + l_{G_k})\mathrm{d}t$ over a finite horizon. The graph therefore acts between clusters, not between individuals: cluster l affects agent i only through the average of $f(x_i,u_i,\cdot)$ over its members' states, weighted by the graph entry $g^k_{\mathcal{C}(i)\mathcal{C}_l}/M_k$, and an agent's problem depends on the graph only through its own row, its *section*.

The limit objects. As $M_k \to \infty$ and every cluster grows, node l maps to a point $i \in [0,1]$, each cluster's empirical distribution converges to a deterministic *local mean field* $\mathfrak{M}_i(t)$, and the ensemble $\mathfrak{M}_G(t) = \{\mathfrak{M}_i(t)\}_{i\in[0,1]}$ replaces the single mean field of the classical theory. The two averages become integrals against the own local mean field and against the graphon-weighted ensemble,

$$f_0[x,u,\mathfrak{M}_i] := \int f_0(x,u,y)\,\mathfrak{M}_i(\mathrm{d}y), \qquad f[x,u,\mathfrak{M}_G;g_i] := \int_0^1\!\!\int f(x,u,y)\,g(i,j)\,\mu_j(\mathrm{d}y)\,\mathrm{d}j, \tag{148}$$

with $g_i := g(i,\cdot)$ the section of the graphon at i, and identically for the cost. Only the agent's own section enters: two nodes with the same section face literally the same problem.

The GMFG equations. For each $i \in [0,1]$ on $[0,T]$, with terminal condition $V^i(T,x) = 0$,

$$\text{[HJB]}(i) \qquad -\partial_t V^i = \inf_{u\in U}\Big\{\tilde f[x,u,\mathfrak{M}_G;g_i]\,\partial_x V^i + \tilde l[x,u,\mathfrak{M}_G;g_i]\Big\} + \tfrac{\sigma^2}{2}\partial_{xx}V^i, \tag{149}$$

$$\text{[FPK]}(i) \qquad \partial_t p_i = -\partial_x\Big\{\tilde f[x,u^0,\mathfrak{M}_G;g_i]\,p_i\Big\} + \tfrac{\sigma^2}{2}\partial_{xx}p_i, \tag{150}$$

$$\text{[BR]}(i) \qquad u^0 = \varphi(t,x\mid \mathfrak{M}_G;g_i),\ \text{the minimiser in [HJB]}(i), \tag{151}$$

where $\tilde f = f_0[\cdot] + f[\cdot]$, $\tilde l = l_0[\cdot] + l[\cdot]$, and $p_i(t,\cdot)$ is the density of $\mathfrak{M}_i(t)$; the forward equation may equivalently be replaced by the closed-loop McKean–Vlasov equation $\mathrm{d}x_i = \tilde f[x_i,\varphi,\mathfrak{M}_G;g_i]\mathrm{d}t + \sigma\mathrm{d}w_i$, with $\mathfrak{M}_i(t)$ defined as the distribution of $x_i(t)$.

Structurally the system is a continuum of classical mean field games, one per node, each consistent with its own local mean field and coupled to the others only through graphon-weighted integrals of the ensemble; the classical case is $g \equiv 1$ (complete graph, one mean field), and $g \equiv 0$ gives a continuum of disconnected games.

Existence and uniqueness (their Theorem 3.12). Admissible ensembles are continuous in i under the Wasserstein metric W_1 (their condition (C1)) and uniformly Hölder in t (C2). The hypotheses are: (H1) the control set U is compact; (H2) f_0, f, l_0, l are continuous, bounded, and Lipschitz in the state arguments uniformly in u; (H3) f_0, f are Lipschitz in u; (H4) the minimiser of the Hamiltonian is a singleton and Lipschitz in $(x, \mathfrak{q})$, uniformly in the ensemble and in i; (H5) $i \mapsto \int_0^1 g(i,j)h(j)\mathrm{d}j$ is continuous for every bounded measurable h; (H6) the minimiser is continuous in $(i, x, \mathfrak{q})$; (H7) a sensitivity condition, $\sup_{t,x,i} |\phi_i(t, x \mid \mathfrak{M}_G) - \bar{\phi}_i(t, x \mid \bar{\mathfrak{M}}_G)| \leq c_1\, d(\mathfrak{M}_G, \bar{\mathfrak{M}}_G)$, where d is the *sup-over-nodes* Wasserstein distance between path-space ensembles and ϕ_i the best response to an arbitrary admissible ensemble; and (H8) the contraction condition $c_1 c_2 < 1$, where c_2 (their Lemma 3.9) bounds the response of the closed-loop distributions to a change in the control laws. Under (H1)–(H8) the map "best-respond to an ensemble, then regenerate the ensemble from the closed-loop distributions" is a contraction in d, so the GMFG equations have a unique solution (Theorem 3.12); the resulting feedback $\varphi(t, x \mid \mathfrak{M}_G; g_i)$ depends only on own state, own section and the ensemble, and it constitutes a Nash equilibrium of the continuum game. The economically relevant feature is that consistency is required *uniformly across network positions at once*, a continuum of coupled fixed points rather than one scalar market-clearing condition; weak coupling or a short horizon makes c_2 small (their Remark 3.10).

The ε-Nash theorem (their Theorem 4.4). Three further hypotheses connect the limit to the finite graphs: (H9) the double limit $M_k \to \infty$ and $\min_l |\mathcal{C}_l| \to \infty$; (H10) i.i.d. initial states with finite mean; and (H11) *sectional convergence* of the graphs to the graphon,

$$\lim_{k\to\infty} \max_i \sum_{j=1}^{M_k} \left| \frac{g^k_{\mathcal{C}_i \mathcal{C}_j}}{M_k} - \int_{I_j} g(I_i^*, j')\mathrm{d}j' \right| = 0, \tag{152}$$

where I_i^* is the midpoint of the interval of length $1/M_k$ that node i occupies: every node's connection profile must converge to the corresponding row of the graphon, uniformly over nodes. Cut-metric convergence is explicitly insufficient (their Section 2 and Lemma A.2): the cut norm controls rectangle integrals, whereas an agent's problem depends on a row, a measure-zero slice. Theorem 4.4 states that if every agent of the finite-graph game plays the GMFG best response evaluated at its node's position, the profile is an ε-Nash equilibrium with $\varepsilon \to 0$ along the sequence, even against a deviator who may use a centralised feedback on the entire finite cross-section. The proof compares four systems (the finite closed loop; a decoupled system of independent copies with cluster distributions; the infinite-population dynamics at the node midpoints; and the finite system with one deviator), and the error has two separable sources:

within-cluster granularity, of order $O(\min_l |\mathcal{C}_l|^{-1/2})$ (their Lemma 4.6), and the graph-to-graphon sectional error controlled by (H11) together with the continuity of the ensemble in (t, i) (their Lemma 4.7 and Lemma A.1). The paper closes with a linear-quadratic-Gaussian example in which existence reduces to a contraction condition on an explicit linear operator.

Mean-field-dependent performance functions. Caines et al. (2022) extend the GMFG equations to agents with affine individual dynamics whose performance functions depend on the mean field. A model in which the network enters *payoffs* rather than drifts, as here, is in that respect closer to their setting than to the drift-coupled primitive above. Both papers, however, work with additively separable interaction functionals on a finite horizon, whereas here the mean field enters through a nonlinear static general-equilibrium map (Definition 5).

[Warning: verify before citing Caines et al. (2022) more specifically.] The paragraph above is written at the level of that paper's title and abstract. Their hypotheses were *not* checked when it was drafted, so the claim "closer to their setting" is a structural impression, not a verified statement about their assumptions. Before this sentence goes into a submitted version, read their hypothesis list and either sharpen the claim or drop it to a bare citation.

B.9.3 Six Differences

Each item gives the mathematical statement and its economic reading.

1. *Where randomness lives.* Caines and Huang have idiosyncratic shocks within clusters and no node-level common noise, so the cluster limit makes each local mean field deterministic *within* the node. Here the within-node population (the unit measure of monopolists) is degenerate by construction: all monopolists at node i share the common node shock z_i, choose identical policies in the symmetric within-node equilibrium, and sit at the single point (K_i, z_i), the firm-specific demand shocks v being transient and integrated out inside expected profit. A within-node law of large numbers would have nothing to average. Randomness survives at the node level, $\mathfrak{M}(t, i, \cdot)$ being the distribution of the node state across realisations of z_i, and determinism of economy-wide aggregates comes instead from essential pairwise independence *across* the continuum of nodes, via Sun's exact law of large numbers (Lemma 13). Economically, node-level uncertainty, the nondegenerate stationary joint distribution of (κ, z) and the $\mathcal{L}_z^{(i)}\mathfrak{q}$ term all exist only because the shocks sit at the node.
2. *Which limit does the work.* Caines and Huang take the double limit (H9); here there is a single construction, a continuum of nodes on a Fubini extension (Sun, 2006; Sun & Zhang, 2009; Bayraktar et al., 2023), with no within-node limit at all. Consequently the step-kernel case of this paper is *not* their (H9) regime (finite graph, infinite clusters); the pre-limit object it approximates is the finite J-sector

economy of Section 2.1 with one common shock per sector, and the gap between the two is the granularity term of Section 3.5.

3. *Where the kernel acts.* In (147)–(148) the graphon enters the drift and the running cost additively. Here the kernel enters only the *payoff*, through the nonlinear static general-equilibrium fixed point in the aggregates $(\mathcal{P}, \mathcal{D}, Q)$ of Definition 5: the node's flow profit depends on the ensemble through $\mathbb{A}m_P$, $\mathbb{A}^* m_X$ and Q, and the state dynamics (136) contain no graphon term. The model is a GMFG with mean-field-dependent performance functions, with general-equilibrium structure rather than an additively separable cost.

4. *Directedness.* Caines and Huang's graphons are symmetric; the input–output kernel is a digraphon (Section 3.1). Symmetry is nowhere load-bearing in their interaction functional or in the structure of their lemmas, whose regularity hypotheses ((H5), (H11)) are conditions on sections, but their *theorems* are stated for symmetric graphons. The extension to directed kernels is therefore available by inspection of their arguments rather than by citation of their results. The compactness of the graphon space in the cut metric, developed for symmetric kernels, is never used here, because finite networks are represented exactly by step kernels.

5. *Horizon and stationarity.* Caines and Huang work on a finite horizon $[0, T]$ with a terminal condition; the present model is infinite-horizon with a balanced growth path and detrending (Section 3.3), and the object of interest is the stationary system of Definition 6. None of their existence machinery applies as stated to a stationary forward–backward system.

6. *Boundedness and compactness.* (H1)–(H2) require a compact control set and bounded Lipschitz coefficients. Here payoffs are unbounded (CRRA utility; profits homogeneous of degree one in scale), the control is bounded only below ($\iota \geq -1/\phi$), the coupling is a nonlinear solve, and the per-node problem is a two-dimensional quasi-linear marginal-$\mathfrak{q}$ equation (Proposition 7) rather than a max-form HJB in one state. No analogue of their contraction constants c_1, c_2 is available here.

[The Caines–Huang architecture applied to this model: what the alternative would look like.] Difference 1 is worth stating as a counterfactual specification, since it delimits the model that is *not* being solved here. Suppose productivity were firm-specific, z_{in} i.i.d. across the monopolists of a node, rather than common to the node. Then:

- the within-node distribution $\mathfrak{M}(t, i, \mathrm{d}k, \mathrm{d}z)$ would become a nondegenerate local mean field, in the exact sense of Caines and Huang's cluster limit;
- same-node firms would post different prices, so the node price index would be the CES aggregate $P_i = \left[\int p(k,z)^{1-\varepsilon} \mathfrak{M}(t, i, \mathrm{d}k, \mathrm{d}z)\right]^{1/(1-\varepsilon)}$ over that distribution, in place of the scalar equation per node state that the symmetric equilibrium delivers here;
- every aggregate, including each node's own K_i and P_i, would be deterministic,

because the within-node law of large numbers would already have removed node-level randomness.

That is the "genuinely three-layer" model of Remark 5, which we do not pursue.

B.9.4 What Can Be Borrowed and What Must Be Developed

Table 8 separates what the paper takes from Caines and Huang as *architecture* (citable for structure and language) from what it cannot take as theorems, and lists what has to be developed in-house; Section 3.4 states the resulting status of the equilibrium concept in one place.

B.10 The Three Models Side by Side

This appendix states in full the three stationary systems that Section 3.5 compares, the finite J-sector system of Section 2, the general-kernel system of Definition 6 and its step-kernel form, reads their backward equations line by line, gives the full account of what the limit removes and of the exactness of the step-kernel representation, and tabulates the correspondence between the three models object by object. The matched pair is that of Section 3.5: the finite CES economy on a row-stochastic matrix $\mathbf{A}$ and the CES GMFG on the step kernel built from the same matrix.

The three systems. Both economies are solved by the *same* generic agent, an atomistic monopolist choosing ι under adjustment costs, discounting profits with the household's SDF, facing environments that differ only in what the LLN has made deterministic.

Finite side. The stationary equilibrium is the system (40)–(41): J backward equations on one shared $(2J-1)$-dimensional state, bound pointwise by the scalar c and the shared drifts, with the single ergodic stationary density $\mathfrak{m}_\infty(\boldsymbol{\kappa}, \mathbf{Z})$ solved downstream.

The finite system in goods units. To put the finite and continuum displays in the same units (the final good), we use the finite-sector system in its *Tobin's-$\tilde{q}$* form.

Lemma 20 (Goods-units form of the finite envelope system). *Substituting $\widehat{\mathfrak{q}}_n = \tilde{q}_n\, c^{-\gamma}$ into the utility-weighted envelope system* (40) *and dividing by $c^{-\gamma}$ gives, for sector $n = 1, \dots, J$, on the shared $(2J-1)$-dimensional state $(\boldsymbol{\kappa}, \mathbf{Z}) \in \Delta^{J-1} \times \mathbb{R}^J_{++}$,*

$$\begin{aligned} r(\boldsymbol{\kappa}, \mathbf{Z})\, \tilde{q}_n = \zeta_n(\overline{v}_n)[1 - \xi_n(\overline{v}_n)]\, \frac{\widehat{s}^{\mathrm{f}}_n(\boldsymbol{\kappa}, \mathbf{Z})}{\kappa_n} - \iota_n + \tilde{q}_n\, \mu^{K_n}(\iota_n) + \sum_{j=1}^{J} \frac{\partial \tilde{q}_n}{\partial \kappa_j}\, \mu_{\kappa_j} \\ + \mathcal{L}_Z\, \tilde{q}_n - \gamma \sum_{l=1}^{J} \sigma_{Z_l}(Z_l)^2\, \partial_{Z_l} \tilde{q}_n\, \partial_{Z_l} \ln c, \\ \Phi_n'(\iota_n)\, \tilde{q}_n = 1, \end{aligned} \tag{153}$$

with $\overline{v}_n = \overline{v}_n(\boldsymbol{\kappa}, \mathbf{Z})$ *and* $\widehat{s}^{\mathrm{f}}_n = \widehat{s}^{\mathrm{f}}_n(\boldsymbol{\kappa}, \mathbf{Z})$ *the threshold and detrended firm revenue of the static block as in* (40)*, closed pointwise by the resource constraint and shared drifts* (41)*, where*

$$r(\boldsymbol{\kappa}, \mathbf{Z}) := \rho + \gamma\, \mu^K - c^{\gamma} \left[\sum_{j=1}^{J} \mu_{\kappa_j}\, \partial_{\kappa_j} c^{-\gamma} + \mathcal{L}_Z\, c^{-\gamma} \right] \tag{154}$$

is the state-dependent risk-free rate*: minus the drift rate of the SDF* $\Lambda_t = e^{-\rho t} C_t^{-\gamma}$ *along the equilibrium dynamics.*

Proof. Write $\widehat{\mathfrak{q}}_n = \tilde{q}_n\, c^{-\gamma}$ with $\tilde{q}_n, c$ functions of $(\boldsymbol{\kappa}, \mathbf{Z})$ and substitute into (40) term by term.

1. *Flow.* The flow term $c^{-\gamma}[\zeta_n(\overline{v}_n)[1 - \xi_n(\overline{v}_n)]\widehat{s}^{\mathrm{f}}_n/\kappa_n - \iota_n]$ divides through by $c^{-\gamma}$ unchanged: the threshold and the detrended firm revenue are objects of the static block, functions of $(\boldsymbol{\kappa}, \mathbf{Z})$ that do not involve $\widehat{\mathfrak{q}}_n$ or c.
2. *Advection.* $\partial_{\kappa_j} \widehat{\mathfrak{q}}_n = c^{-\gamma}\, \partial_{\kappa_j} \tilde{q}_n + \tilde{q}_n\, \partial_{\kappa_j} c^{-\gamma}$.
3. *Generator.* By the product rule inside $\mathcal{L}_Z$ (which carries second derivatives),

$$\mathcal{L}_Z\left(\tilde{q}_n c^{-\gamma}\right) = c^{-\gamma} \mathcal{L}_Z \tilde{q}_n + \tilde{q}_n\, \mathcal{L}_Z c^{-\gamma} + \sum_{l=1}^{J} \sigma_{Z_l}(Z_l)^2\, \partial_{Z_l} \tilde{q}_n\, \partial_{Z_l} c^{-\gamma}.$$

4. *Divide by* $c^{-\gamma} > 0$ *and collect the* $\tilde{q}_n$ *terms.* All terms proportional to $\tilde{q}_n$ join the discount:
$$\left[\rho + \gamma \mu^K - c^{\gamma} \left(\sum_j \mu_{\kappa_j} \partial_{\kappa_j} c^{-\gamma} + \mathcal{L}_Z c^{-\gamma} \right) \right] \tilde{q}_n,$$
and $c^{\gamma}\, \partial_{Z_l} c^{-\gamma} = -\gamma\, \partial_{Z_l} \ln c$ turns the cross term of step 3 into the risk-premium term of the display.
5. *The coefficient is the risk-free rate.* With $\Lambda_t = e^{-\rho t} K^{-\gamma} c^{-\gamma}$, where K carries no diffusion and no cross-variation because capital has no shock, minus the drift rate of Λ is $\rho + \gamma \mu^K - c^{\gamma} \left[\sum_j \mu_{\kappa_j} \partial_{\kappa_j} c^{-\gamma} + \mathcal{L}_Z c^{-\gamma} \right] = r(\boldsymbol{\kappa}, \mathbf{Z})$, the instantaneous risk-free rate of the finite economy. The FOC transforms as $c^{-\gamma} = \widehat{\mathfrak{q}}_n \Phi_n'(\iota_n) \iff \Phi_n'(\iota_n) \tilde{q}_n = 1$.

Sanity checks: (i) on a deterministic BGP (c constant, $\mu^K = g$) the rate collapses to $\rho + \gamma g$ and the covariance term vanishes, reproducing the GMFG display's structure exactly; (ii) the sign is standard: for a procyclical firm ($\partial_{Z_l} \tilde{q}_n$ and $\partial_{Z_l} \ln c$ of the same sign) the covariance term lowers value, i.e. procyclical assets are discounted more heavily. ∎

The last term of (153) is an explicit *risk premium*: the covariance of $\tilde{q}_n$'s productivity exposure with marginal utility's. Read in utility units, (40) itself carries the same pricing implicitly: the stochastic marginal utility $c^{-\gamma}(\boldsymbol{\kappa}, \mathbf{Z})$ multiplies the flow and the effective discount $\rho + \gamma \mu^K$ moves with the state; together they *are* the stochastic SDF,

and their covariance with the flow is how risk premia get priced without any explicit covariance term appearing. The goods-units display makes that implicit pricing explicit; the utility-weighted form is the one a numerical solution would use, because it needs no derivatives of c (the finite model is not solved in this paper).

The two forward equations. The backward comparison is drift for drift; the forward objects are where the dimensionality actually shows. The finite model has *one* stationary FPK equation, adjoint to the closed-loop generator on the shared state: its ergodic density $\mathfrak{m}_\infty(\boldsymbol{\kappa}, \mathbf{Z})$ on $\Delta^{J-1} \times \mathbb{R}^J_{++}$ solves

$$0 = -\sum_{j=1}^{J} \partial_{\kappa_j}\Big[\mu_{\kappa_j}\,\mathfrak{m}_\infty\Big] - \sum_{l=1}^{J} \partial_{Z_l}\big[\mu_{Z_l}(Z_l)\,\mathfrak{m}_\infty\big] + \frac{1}{2}\sum_{l=1}^{J} \partial^2_{Z_l Z_l}\Big[\sigma_{Z_l}(Z_l)^2\,\mathfrak{m}_\infty\Big], \qquad \int \mathfrak{m}_\infty = 1,$$

one $(2J-1)$-dimensional joint density, carrying all realised cross-sector correlations.

General-kernel side. For a general kernel the stationary system is Definition 6: for a.e. node $i \in [0,1]$ the static block (139)–(141) given $(\mathcal{P}(i), \widehat{\mathcal{D}}(i), q)$, a scalar root in the threshold at each (κ, z), the backward equation (the stationary (52)) and the forward equation (the stationary (53)),

$$\begin{aligned}(\rho+\gamma g)\,\mathfrak{q} &= \zeta_i(\overline{v})[1-\xi_i(\overline{v})]\,\frac{\widehat{s}^{\mathrm{f}}(i,\kappa,z)}{\kappa} - \iota^*(\mathfrak{q}) + \mathfrak{q}[\Phi_i(\iota^*)-\delta(i)] + \mathfrak{q}_\kappa\,\kappa[\Phi_i(\iota^*)-\delta(i)-g] \\ &\quad + \mathcal{L}_z^{(i)}\mathfrak{q}, \\ 0 &= -\frac{\partial}{\partial\kappa}\{\kappa[\Phi_i(\iota^*)-\delta(i)-g]\,\mathfrak{m}\} + \frac{\partial}{\partial z}\{\theta(i)\,z\,\mathfrak{m}\} + \frac{\varsigma^Z(i)^2}{2}\frac{\partial^2}{\partial z^2}\mathfrak{m}, \qquad \int \mathfrak{m}(i,\kappa,z)\mathrm{d}\kappa\mathrm{d}z = 1,\end{aligned} \tag{155}$$

with $\overline{v} = \overline{v}(i,\kappa,z)$ and $\widehat{s}^{\mathrm{f}} = \varrho_i(\overline{v})\widehat{s}(i,\kappa,z)$ from the static block, coupled across nodes only through the consistency map (54), $\ln\mathcal{P} = \mathbb{A}m_P + H$ and $\widehat{\mathcal{D}} = \mathbb{A}^*\widehat{m_X}$ with $(m_P, \widehat{m_X})$ the node moments of log prices and spending (the sales moment $\widehat{m_S}$ enters through $q = \int_0^1[\widehat{m_S} - \widehat{m_X}]\mathrm{d}j$), and closed by the numeraire, capital adding-up, $g = \int_0^1 \mathbb{E}_i[(\Phi_i(\iota^*)-\delta(i))\kappa]\mathrm{d}i$ and $r = \rho + \gamma g$. The general kernel differs from the step kernel below only in the index set ($[0,1]$ versus N blocks), in the operators ($\mathbb{A}, \mathbb{A}^*$ versus the block matrix A and its transpose) and in the unknown coupling objects (functions on $[0,1]$ versus N-vectors); the per-node pair is identical.

Step-kernel side. Take a step kernel with blocks $I_1, \ldots, I_N$ of lengths Δ_n, block input–output matrix and entropy constants

$$\begin{aligned}A_{nm} &:= \int_{I_m} a(i,j)\,\mathrm{d}j \quad (i \in I_n), \\ \sum_{m=1}^{N} A_{nm} &= 1, \\ e_n &:= \int_0^1 a(i,j)\ln a(i,j)\,\mathrm{d}j = \sum_{m=1}^{N} A_{nm}\ln\frac{A_{nm}}{\Delta_m},\end{aligned}$$

and block-constant primitives $(\alpha_n, x_n, G_n, \beta_n, \phi_n, \delta_n, \theta_n, \varsigma_n^Z)$, so that the tail functions of G_n and with them the threshold map Υ_n, utilization U_n and the shares ξ_n, ζ_n, ϱ_n are block-constant functions of the threshold, with β_n the density value on I_n ($\sum_n \beta_n \Delta_n = 1$). The stationary GMFG is then *exactly* the following finite system. For each block $n = 1, \dots, N$, given the coupling objects $\left(\mathcal{P}_n, \widehat{\mathcal{D}}_n, q, g\right)$ and $r = \rho + \gamma g$: the per-block statics, pointwise in (κ, z),

$$P_n\, \Upsilon_n(\overline{v}) = \frac{\mathcal{P}_n}{e^z x_n^{\alpha_n}}, \qquad \widehat{s}_n = \frac{\mathcal{P}_n\, \kappa}{x_n}\, \mu_{n,1/\varepsilon}^{\frac{\varepsilon}{\varepsilon-1}}\, \frac{U_n(\overline{v})}{\Upsilon_n(\overline{v})}, \qquad \widehat{s}_n = \beta_n\, P_n^{1-\eta}\, q + \widehat{\mathcal{D}}_n, \tag{156}$$

one scalar root $\overline{v}_n(\kappa, z) \in (\underline{v}_n, \infty)$ in the threshold at each state (Lemma 15), with detrended firm revenue $\widehat{s}_n^{\mathrm{f}} = \varrho_n(\overline{v})\widehat{s}_n$ and spending $\widehat{x}_n = \xi_n(\overline{v})\widehat{s}_n^{\mathrm{f}}$; the backward equation (marginal $\mathfrak{q}$) with its FOC, whose flow is the same private capital-income flow as in (40), evaluated at the block's own threshold,

$$\begin{aligned} r\,\mathfrak{q}_n &= \zeta_n(\overline{v})[1-\xi_n(\overline{v})]\,\varrho_n(\overline{v})\,\frac{\widehat{s}_n(\kappa,z)}{\kappa} - \iota^* + \mathfrak{q}_n[\Phi_n(\iota^*) - \delta_n] + \mathfrak{q}_{n,\kappa}\,\kappa[\Phi_n(\iota^*) - \delta_n - g] \\ &\quad + \mathcal{L}_z^{(n)}\mathfrak{q}_n, \\ \Phi_n'(\iota^*)\,\mathfrak{q}_n &= 1; \end{aligned} \tag{157}$$

and the forward equation

$$0 = -\frac{\partial}{\partial\kappa}\{\kappa[\Phi_n(\iota^*) - \delta_n - g]\,\mathfrak{m}_n\} + \frac{\partial}{\partial z}\{\theta_n\, z\, \mathfrak{m}_n\} + \frac{(\varsigma_n^Z)^2}{2}\frac{\partial^2}{\partial z^2}\mathfrak{m}_n, \qquad \int \mathfrak{m}_n(\kappa, z)\,\mathrm{d}\kappa\,\mathrm{d}z = 1. \tag{158}$$

Given $\left(\mathcal{P}, \widehat{\mathcal{D}}, q, g\right)$, the N pairs (157)–(158) are mutually decoupled two-dimensional problems. The input–output network re-enters *only* through the consistency map. With block moments

$$\mathbb{E}_n[\cdot] := \int \cdot\ \mathfrak{m}_n(\kappa, z)\,\mathrm{d}\kappa\,\mathrm{d}z, \qquad m_{P,n} := \mathbb{E}_n[\ln P_n],$$
$$\widehat{m}_{X,n} := \mathbb{E}_n[\widehat{x}_n], \qquad \widehat{m}_{S,n} := \mathbb{E}_n[\widehat{s}_n], \qquad \mathcal{N}_n := \mathbb{E}_n\left[P_n^{1-\eta}\right],$$

the fixed point closes via

$$\begin{aligned} \ln \mathcal{P}_n &= \sum_{m=1}^N A_{nm}\, m_{P,m} - e_n, && \text{(rows of } A\text{: what } n \text{ buys; this is } \mathbb{A}) \\ \widehat{\mathcal{D}}_n &= \frac{1}{\Delta_n}\sum_{m=1}^N A_{mn}\, \widehat{m}_{X,m}\, \Delta_m, && \text{(columns of } A\text{: what is sold to } n\text{; this is } \mathbb{A}^*) \\ q &= \sum_{m=1}^N [\widehat{m}_{S,m} - \widehat{m}_{X,m}]\,\Delta_m, && \sum_{n=1}^N \beta_n\, \mathcal{N}_n\, \Delta_n = 1 \quad \text{(numeraire)}, \\ g &= \sum_{n=1}^N \mathbb{E}_n[(\Phi_n(\iota^*) - \delta_n)\,\kappa]\,\Delta_n, && \sum_{n=1}^N \mathbb{E}_n[\kappa]\,\Delta_n = 1, \qquad r = \rho + \gamma g. \end{aligned} \tag{159}$$

The loop is: guess $\left(\mathcal{P}, \widehat{\mathcal{D}}, q, g\right) \to$ solve N independent 2D pairs (each with the scalar static root (156) at every grid point) $\to$ integrate the moments $\to$ regenerate the coupling objects via (159) $\to$ iterate. Two structural contrasts with the finite system are

now explicit. First, the coupling is *in distribution, not in state*: the distributions $\mathfrak{m}_n$ feed back into the backward equations through (159), so the stationary problem is genuinely forward–backward, whereas (40) never sees $\mathfrak{m}_\infty(\boldsymbol{\kappa}, \mathbf{Z})$. Second, the same matrix A that coupled realised states in the finite system (through the static block of Section 2.2 and (41); without capacity constraints through $\mathbf{h}$ and $\boldsymbol{\omega}$) here acts only on N deterministic moments per object: the block forms of $\mathbb{A}$ and $\mathbb{A}^*$ in the first two lines of (159).

Reading the systems line by line. Set the finite backward equation (153) against the step-kernel backward equation (157), and against its general-kernel twin, the first line of (155): the flow is the *same* private capital-income flow, the discount, drift and coupling terms are not. Likewise set the single finite FPK equation against the N two-dimensional marginal FPK equations (158): no joint distribution across blocks exists or is needed; the ensemble of marginals is the entire forward content of the limit, and the correlations that the finite $\mathfrak{m}_\infty$ carries are exactly the granular content the limit discards. Reading the backward equations line by line:

- *Flow.* The marginal-profit flow $\zeta_n(\overline{v}_n)[1 - \xi_n(\overline{v}_n)]\widehat{s}^{\mathrm{f}}_n/\kappa_n$ is *identical*: both firms are priced at the private rental flow (33), the capital share of firm revenue per unit of capital, so the GMFG is the limit of exactly this firm system: the comparison is pure granularity, with no change in what is internalised. What differs is the argument of the threshold and of the revenue: the finite $(\overline{v}_n, \widehat{s}^{\mathrm{f}}_n)(\boldsymbol{\kappa}, \mathbf{Z})$ come from the static system of Section 2.2 solved on the full realised state (every sector's price, and hence every sector's rationing margin, responds to every sector's state), the GMFG $(\overline{v}_n, \widehat{s}^{\mathrm{f}}_n)(\kappa, z)$ from the scalar node block (156) solved against the deterministic aggregates $(\mathcal{P}_n, \widehat{\mathcal{D}}_n, q)$. Without capacity constraints the flow is the constant coefficient $\frac{\varepsilon-1}{\varepsilon}\alpha_n$ times $\widehat{s}_n/\kappa_n$ on both sides (Appendix C.6).
- *Discounting.* Both displays are now in final-good units, so the comparison is term for term: the finite discount is the *state-dependent* risk-free rate $r(\boldsymbol{\kappa}, \mathbf{Z})$, fluctuating with the SDF, and the equation carries an explicit risk-premium term (the covariance of $\tilde{q}_n$'s productivity exposure with marginal utility's); the GMFG discount is the *constant* $r = \rho + \gamma g$, and the premium term is identically zero because the SDF is deterministic (Lemma 18). What the limit removes from the backward equation is exactly this pair: rate fluctuations and priced covariances.
- *Drift.* The finite advection term carries the shared vector field $\mu_{\kappa_j} = \kappa_j[\mu^{K_j}(\iota_j) - \mu^K]$: every sector's realised policy enters every sector's equation. The GMFG carries only the own drift $\kappa[\Phi_n(\iota^*) - \delta_n - g]$, with the single deterministic scalar g summarising everyone else's accumulation.
- *Coupling.* The finite system is coupled *pointwise in the state* (at each $(\boldsymbol{\kappa}, \mathbf{Z})$ the scalar c and the pair $(\mu^K, \mu_{\boldsymbol{\kappa}})$ tie the J equations together), and no distribution appears anywhere; the ergodic stationary density $\mathfrak{m}_\infty(\boldsymbol{\kappa}, \mathbf{Z})$ is computed downstream, from one FPK equation on the shared state. The GMFG is coupled *in distribution*: the N per-block distributions $\mathfrak{m}_n(\kappa, z)$ feed back into the backward

equations through the moments (159), making the stationary problem genuinely forward–backward; but given those moments the N two-dimensional pairs are completely decoupled.

The general kernel adds nothing to this reading: relative to the step kernel it changes only the index set, the operators and the coupling unknowns (Section B.10), so every line above applies to it verbatim with the block index n replaced by $i \in [0,1]$ and the block matrix by $\mathbb{A}, \mathbb{A}^*$.

What the limit removes.

(i) *Aggregate risk and risk premia.* At finite J the sector shocks do not wash out: $(\widehat{s}_n, c, \mu^K)$ and hence the SDF are stochastic, firm values covary with consumption, and those covariances are priced (explicitly, in the risk-premium term of the goods-units display above; implicitly, through the utility-weighting of $\widehat{q}_n$ in (40)). In the limit the exact LLN (Lemma 13) makes every aggregate deterministic: no aggregate fluctuations, no consumption-smoothing dynamics, no risk premia, discounting at the risk-free rate. The approximation error this induces against a finite economy is of order $O(J^{-1/2})$ (the aggregate-fluctuation term) on top of any graphon-approximation term(Section 3.4). MIT shocks (Appendix D, Step 6) restore *perfect-foresight* aggregate transitions, but not priced aggregate risk.

(ii) *Own-shock general-equilibrium feedback.* A finite sector has mass $1/J$: its own Z_n moves aggregate output, consumption, the SDF, and, through the static block, every other sector's prices, and part of the state-relevance of its flow runs through exactly this channel. In the GMFG a node is atomistic: its own (κ, z) moves no moment and no aggregate. The starkest consequence is the Cobb–Douglas degeneracy (Appendix B.7): at $\eta = 1$ the finite model remains a nondegenerate stochastic system *only* because of the $1/J$ own-shock feedback, and in the limit that feedback dies, so the node's *stochastic* state-relevance dies with it: z becomes payoff-irrelevant and the node distributions degenerate to point masses in κ (Theorem 1). Capacity does not change this: at $\eta = 1$ the node's threshold, and with it its utilization and its revenue shares, is free of z (Lemma 17), so the rationing margin opens no channel through which productivity could re-enter the payoff. At $\eta \neq 1$ the CES relative-price channel $\beta_n P_n^{1-\eta} q$ survives as the *sole* source of own-state relevance, which is precisely why this paper insists on CES final demand.

(iii) *State-contingent cross-sector comovement.* Finite policies are functions $\iota_n(\boldsymbol{\kappa}, \mathbf{Z})$: sector n's investment responds to the *realised* state of every other sector, propagation of a given shock through the network is resolved state by state, and the joint distribution $\mathfrak{m}_\infty(\boldsymbol{\kappa}, \mathbf{Z})$ encodes realised cross-sector correlations (comovement, granular concentration dynamics). GMFG policies are functions $\iota^*(n, \kappa, z)$ of own states alone; other nodes enter only through deterministic moments, and across nodes the equilibrium randomness is essentially pairwise independent by construction. First moments of the cross-section survive exactly; realised cross-node

second moments (comovement and its asset-pricing and volatility implications) are zero in the limit by fiat.

(iv) *Not lost: strategic interaction and internalisation.* Neither economy has them. In both, the decision unit is an atomistic monopolist that takes its sector's price and quantity, all aggregates, and (at finite J) the SDF as given, and internalises the effect of its investment on nothing: capital is priced at the private rental flow in both systems, not at the internalised derivative Q_{K_n} (Remark 10). Sectors as units take no decisions in either model, so there is no oligopolistic margin to lose. What changes between the two models is *statistical*, not strategic: which objects in the generic firm's environment are random, and which states its policy is allowed to condition on.

Exactness and the two error terms. Two distinct approximations are easily conflated, and the step kernel makes exactly one of them. *Representing a finite input–output table by a kernel is exact*: by Section 3.1, Lemma 10 and (90), the step kernel a^J reproduces the table's block weights, its Leontief inverse and its entropy up to the bookkeeping shift, and both the sectional and the cut distance between the table and its kernel are zero (Appendix B.2.1). In the language of the graphon mean field game literature, the graph-to-graphon approximation error controlled by their sectional-convergence hypothesis (H11) is identically zero here: the continuum formulation introduces no discretisation of the network, and the only discretisation in Appendix D is of the (κ, z) grids.

What separates the step-kernel GMFG from the finite J-sector economy on the same table is therefore a single thing: *granularity*. In the finite economy each sector is one industry carrying one common shock, so aggregates are stochastic and each sector's shock moves them with weight of order $1/J$; in the step-kernel GMFG each block contains a continuum of independently shocked industries, so block moments are deterministic (Section 3.5) and a block's state moves no aggregate. The gap is the aggregate-fluctuation term of the ε-optimality conjecture (Appendix B.9): the $O(J^{-1/2})$ sampling noise of J independent sector shocks that the limit policies ignore, together with the within-block granularity term familiar from the mean-field literature. This also fixes which pre-limit object the step-kernel GMFG approximates: the finite J-sector economy of Section 2.1 on the same table, not the finite-graph, infinite-cluster regime of the double limit of Caines and Huang (2021) (Appendix B.9.3, item 2), whose clusters would correspond to firm-level idiosyncratic states this model does not have. Under Cobb–Douglas final demand the gap is not a correction term at all but the entirety of the finite model's dynamics: the step-kernel Cobb–Douglas GMFG is the deterministic skeleton that the finite Cobb–Douglas economy fluctuates around (Proposition 5).

The general-kernel GMFG is one step further removed from any finite economy. It requires no approximation of the model either, but if it is to be read as the limit of a growing sequence of finite networks, the convergence of the networks to the kernel has to be sectional, not merely in the cut metric, and the kernel has to satisfy the conditions

of Appendix B.2.2; existence and ε-optimality are conjectured, not proved, for either kernel(Section 3.4).

The correspondence table. Table 9 tabulates the correspondence between the three models object by object.

Two clarifications. First, the backward object of the per-block pair is not the value function of a representative firm: it is the *on-diagonal marginal value* $\mathfrak{q}_m(\kappa, z) = w_k\big|_{k=K}$ of the individual monopolist at a representative industry of block m, solving the quasi-linear marginal-q PDE; the equity value W is recovered afterwards from a separate linear PDE. Second, "the representative firm of each sector" should read "the individual firm at a representative industry of each block": it is the double atomism (firm within industry, industry within economy) that reduces the state to own (κ, z), not any representativeness at block level.

C The Model without Capacity Constraints

This appendix collects the model without capacity constraints: the limit $x_i \downarrow 0$ of the model of the main text, in which capacity $\overline{y}_i = Z_i K_i x_i^{\alpha_i - 1}$ becomes infinite, no firm is ever rationed, the threshold diverges and every capacity object converges to a constant (Lemma 4). It serves two purposes: it is the counterfactual against which the effects of capacity are measured, and it is the model whose static block admits a closed form, so it supplies the initialisers and the tests of the numerical work (Appendix D). Three conventions govern the appendix. First, every result of the main text whose statement does not involve the capacity block holds verbatim without capacity constraints and is not restated; the list of such results is given at the head of each subsection. Second, an equation of this model that follows from a main-text equation by the capacity limit is stated as a limit, with a pointer to Lemma 4 and no proof. Third, a result whose content the capacity model lacks (a closed form, or a uniqueness argument that uses the power form of the profit function without capacity constraints) is stated in full, with its proof beside it. All labels of this appendix carry the prefix `nc-`.

C.1 Finite Sectors: Pricing, Supply and Shares

Throughout this subsection the sector is at the symmetric point of Lemma 3 and capacity is slack, so that the threshold is infinite, $\varphi_{in} \equiv 0$, the shadow index coincides with the posted price and the aggregator's demand (3) reduces to

$$y_{in} = v_{in} Y_i \left(\frac{p_{in}}{P_i} \right)^{-\varepsilon}, \qquad P_i = \left[\int_0^1 v_{in} p_{in}^{1-\varepsilon} \mathrm{d}n \right]^{1/(1-\varepsilon)}. \tag{160}$$

Expected output and bundle use. By Lemma 4(i), without capacity constraints the expectations (6)–(12) become

$$\mathbb{E}_v[y_{in}] = Y_i \left(\frac{p_{in}}{P_i} \right)^{-\varepsilon} \tag{161}$$

$$\mathbb{E}_v[M_{in}] = \mu_{i,\frac{1}{1-\alpha_i}} \left[\frac{Y_i (p_{in}/P_i)^{-\varepsilon}}{Z_i k_{in}^{\alpha_i}} \right]^{\frac{1}{1-\alpha_i}} \tag{162}$$

Because there is no capacity constraint, the upper tail of v is never truncated, so the demand distribution enters only through the moment $\mu_{i,\frac{1}{1-\alpha_i}}$. Using the sectoral demand function (160), the bundle requirement implied by the technology (8) can be written as

$$M_{in} = \left(\frac{v_{in} Y_i (p_{in}/P_i)^{-\varepsilon}}{Z_i k_{in}^{\alpha_i}} \right)^{\frac{1}{1-\alpha_i}}$$

Using (161), (162) can also be written as

$$\mathbb{E}_v[M_{in}] = \mu_{i,\frac{1}{1-\alpha_i}} \left[\frac{\mathbb{E}_v[y_{in}]}{Z_i k_{in}^{\alpha_i}} \right]^{\frac{1}{1-\alpha_i}}$$

Firm pricing. Before observing v_{in} the firm chooses p_{in} to maximise expected operating profit

$$\max_{p_{in}} p_{in} \mathbb{E}_v[y_{in}] - \mathcal{P}_i(\mathbf{P}) \mathbb{E}_v[M_{in}] \ \text{ s.t. } (160)$$

taking $Y_i, P_i, \mathcal{P}_i$ as given. The optimum is the constant CES markup over an effective expected marginal cost. The result is the version of Lemma 2 without capacity constraints; it is stated and proved in full because its uniqueness argument uses the power form of the profit function, which the capacity model does not have.

Lemma 21 (Pre-shock pricing: a constant markup over effective expected marginal cost). *Let $\varepsilon > 1$ and $\alpha_i \in (0,1)$, and fix $(Y_i, P_i, \mathcal{P}_i(\mathbf{P}), Z_i)$ and own capital $k_{in} > 0$. Then:*

(i) The pricing problem above has a unique maximiser $p_{in} \in (0, \infty)$, and it satisfies

$$p_{in} = \frac{\varepsilon}{\varepsilon - 1} \widetilde{mc}_{in}, \qquad \widetilde{mc}_{in} := \frac{1}{1-\alpha_i} \mathcal{P}_i(\mathbf{P}) \frac{\mathbb{E}_v[M_{in}]}{\mathbb{E}_v[y_{in}]}. \tag{163}$$

(ii) Evaluated with (161)–(162)*, the effective marginal cost is*

$$\widetilde{mc}_{in} = \frac{1}{1-\alpha_i} \mathcal{P}_i(\mathbf{P}) \, \mu_{i,\frac{1}{1-\alpha_i}} \left(\frac{[Y_i (p_{in}/P_i)^{-\varepsilon}]^{\alpha_i}}{Z_i k_{in}^{\alpha_i}} \right)^{\frac{1}{1-\alpha_i}} \tag{164}$$

so (163) *pins down p_{in} implicitly; explicitly, $p_{in}^{\Theta_i}$ is proportional to $k_{in}^{-\frac{\alpha_i}{1-\alpha_i}}$ with $\Theta_i := 1 + \frac{\varepsilon \alpha_i}{1-\alpha_i} > 0$.*

(iii) The markup $\frac{\varepsilon}{\varepsilon-1}$ is a constant: it depends on neither the state $(k_{in}, \mathbf{K}, \mathbf{Z})$ nor the demand-shock distribution G_i. The latter enters the price only through the single moment $\mu_{i,\frac{1}{1-\alpha_i}} \geq 1$, and only inside $\widetilde{mc}_{in}$.

Proof. Write $p := p_{in}$ and $k := k_{in}$, and let $\Pi(p) := p\,\mathbb{E}_v[y_{in}] - \mathcal{P}_i(\mathbf{P})\,\mathbb{E}_v[M_{in}]$ be expected operating profit: capital is fixed within the instant, so the bundle is the only cost.

Step 1 (the objective as a function of p alone). Substituting (161)–(162) at $p_{in} = p$, $k_{in} = k$,

$$\Pi(p) = \mathcal{A}p^{1-\varepsilon} - \mathcal{C}p^{-\frac{\varepsilon}{1-\alpha_i}}, \qquad \mathcal{A} := Y_i P_i^{\varepsilon} > 0, \qquad \mathcal{C} := \mathcal{P}_i(\mathbf{P})\,\mu_{i,\frac{1}{1-\alpha_i}} \frac{\mathcal{A}^{\frac{1}{1-\alpha_i}}}{Z_i^{\frac{1}{1-\alpha_i}} k^{\frac{\alpha_i}{1-\alpha_i}}} > 0.$$

Both constants collect exactly the objects the firm takes as given. The two exponents are negative and ordered,

$$-\frac{\varepsilon}{1-\alpha_i} < 1-\varepsilon < 0, \qquad \text{since} \qquad (1-\varepsilon) - \left(-\frac{\varepsilon}{1-\alpha_i}\right) = 1 + \frac{\varepsilon\alpha_i}{1-\alpha_i} = \Theta_i > 0,$$

i.e. the cost term is the steeper power of p. This ordering is what drives Step 4.

Step 2 (the two elasticities). Demand (160) has constant price elasticity ε, and the technology (8) makes $\mathbb{E}_v[M_{in}]$ proportional to $\mathbb{E}_v[y_{in}]^{\frac{1}{1-\alpha_i}}$ at fixed k, so

$$\frac{\partial \mathbb{E}_v[y_{in}]}{\partial p} = -\varepsilon\,\frac{\mathbb{E}_v[y_{in}]}{p}, \qquad \frac{\partial \mathbb{E}_v[M_{in}]}{\partial p} = -\frac{\varepsilon}{1-\alpha_i}\,\frac{\mathbb{E}_v[M_{in}]}{p}:$$

the bundle's elasticity is the demand elasticity scaled by $\frac{1}{1-\alpha_i} > 1$, the inverse output elasticity of the bundle, because serving one per cent less output at fixed capital frees $\frac{1}{1-\alpha_i}$ per cent of the bundle.

Step 3 (first-order condition). Differentiating Π and inserting Step 2,

$$\Pi'(p) = \mathbb{E}_v[y_{in}] + p\,\frac{\partial \mathbb{E}_v[y_{in}]}{\partial p} - \mathcal{P}_i\,\frac{\partial \mathbb{E}_v[M_{in}]}{\partial p} = (1-\varepsilon)\mathbb{E}_v[y_{in}] + \frac{\varepsilon}{1-\alpha_i}\,\mathcal{P}_i\,\frac{\mathbb{E}_v[M_{in}]}{p}.$$

Setting $\Pi'(p) = 0$ and dividing by $(\varepsilon-1)\mathbb{E}_v[y_{in}] > 0$ rearranges to (163).

Step 4 (the critical point is unique and is the maximiser). In the power form of Step 1 the same condition reads

$$\Pi'(p) = (1-\varepsilon)\mathcal{A}p^{-\varepsilon} + \frac{\varepsilon}{1-\alpha_i}\mathcal{C}\,p^{-\frac{\varepsilon}{1-\alpha_i}-1} = 0 \iff p^{-\Theta_i} = \frac{(\varepsilon-1)(1-\alpha_i)}{\varepsilon}\frac{\mathcal{A}}{\mathcal{C}},$$

because $-\frac{\varepsilon}{1-\alpha_i} - 1 + \varepsilon = -\Theta_i$. The right-hand side is a strictly positive constant and $p \mapsto p^{-\Theta_i}$ is a decreasing bijection of $(0,\infty)$ onto itself, so there is exactly one critical point $p^\star > 0$. That it is the maximum follows from the ordering of Step 1: as $p \downarrow 0$ the steeper cost term dominates and $\Pi(p) \to -\infty$; $\Pi(p) > 0$ exactly for $p^{\Theta_i} > \mathcal{C}/\mathcal{A}$; and $\Pi(p) \to 0^+$ as $p \to \infty$. Hence Π is continuous, positive somewhere and vanishing at

infinity, so its supremum is attained at an interior point, which must be a critical point, and $p^\star$ is the only one.

Step 5 (the closed form). Dividing (162) by (161) and using $\frac{1}{1-\alpha_i} - 1 = \frac{\alpha_i}{1-\alpha_i}$,

$$\frac{\mathbb{E}_v[M_{in}]}{\mathbb{E}_v[y_{in}]} = \mu_{i,\frac{1}{1-\alpha_i}} \frac{\left[Y_i (p_{in}/P_i)^{-\varepsilon}\right]^{\frac{\alpha_i}{1-\alpha_i}}}{Z_i^{\frac{1}{1-\alpha_i}} k_{in}^{\frac{\alpha_i}{1-\alpha_i}}},$$

which is (164); the displayed proportionality of $p_{in}^{\Theta_i}$ to $k_{in}^{-\frac{\alpha_i}{1-\alpha_i}}$ is the solved form of Step 4 with $\mathcal{C} \propto k^{-\frac{\alpha_i}{1-\alpha_i}}$.

Step 6 (part (iii)). The markup in (163) is $\frac{\varepsilon}{\varepsilon-1}$, which came from Step 2's constant demand elasticity alone; nothing else in Steps 3–4 touches it. Since there is no capacity constraint, the upper tail of v is never truncated, so ε stays the elasticity of expected demand and G_i enters Π only through the single moment $\mu_{i,\frac{1}{1-\alpha_i}}$ inside $\mathcal{C}$, hence only through $\widetilde{mc}_{in}$. ■

Reading $\widetilde{mc}_{in}$. At given capital the only flexible input is the bundle with output elasticity $1-\alpha_i$, so the marginal cost of an extra unit of output is the bundle price divided by the bundle's marginal product,

$$\frac{\mathcal{P}_i}{\partial y_{in}/\partial M_{in}} = \frac{\mathcal{P}_i M_{in}}{(1-\alpha_i) y_{in}}$$

that is, average bundle cost per unit of output, scaled up by $\frac{1}{1-\alpha_i} > 1$ because of diminishing returns to the bundle. $\widetilde{mc}_{in}$ is this object with the realised (M_{in}, y_{in}) replaced by the expectations the firm prices against, $\mathbb{E}_v[M_{in}]/\mathbb{E}_v[y_{in}]$ as the price is set before the realisation of v_{in}. Since $\mathbb{E}_v[M_{in}]$ carries the convexity moment $\mu_{i,\frac{1}{1-\alpha_i}} \geq 1$, demand risk raises this effective marginal cost while leaving the markup untouched. Under capacity the markup is no longer untouched (Lemma 2(iii)).

Symmetric sectoral equilibrium. Since price and investment are set before v_{in} is realised, there is a symmetric equilibrium $p_{in} = p_i$, $k_{in} = K_i$ for all n. By the law of large numbers and $\mu_{i,1} = 1$, $P_i = p_i$ and $p_i/P_i = 1$; realised firm output is $y_i(v) = vY_i$ and $\mathbb{E}_v[y_i] = Y_i$, so sales at the index and firm revenue coincide, $S_i = S_i^{\mathrm{f}} = P_i Y_i$, and the rationing rent is zero. By Lemma 4(iii), average (hence aggregate) conditional bundle demand (18) becomes

$$M_i^e = \mu_{i,\frac{1}{1-\alpha_i}} \left(\frac{Y_i}{Z_i K_i^{\alpha_i}} \right)^{\frac{1}{1-\alpha_i}} \tag{165}$$

and the supply relation (74) becomes the sectoral supply curve

$$P_i = \Gamma_i \mathcal{P}_i(\mathbf{P}) \left(\frac{Y_i^{\alpha_i}}{Z_i K_i^{\alpha_i}} \right)^{\frac{1}{1-\alpha_i}}, \qquad \Gamma_i := \frac{\varepsilon}{\varepsilon-1} \frac{1}{1-\alpha_i} \mu_{i,\frac{1}{1-\alpha_i}}, \tag{166}$$

The sectoral price is the constant wedge Γ_i (markup $\times$ diminishing returns $\times$ demand risk) times a marginal-cost index that rises with the input-bundle price, falls with productivity, and rises with output per unit of installed capital with elasticity $\frac{\alpha_i}{1-\alpha_i}$, the decreasing returns to the flexible input at fixed capital. Directly from the pricing rule (163):

1. In the symmetric equilibrium the pricing rule (163) reads $P_i = \frac{\varepsilon}{\varepsilon-1}\frac{1}{1-\alpha_i}\,\mathcal{P}_i(\mathbf{P})\,M_i^e/Y_i$.
2. By (165), $M_i^e/Y_i = \mu_{i,\frac{1}{1-\alpha_i}}\,Y_i^{\frac{\alpha_i}{1-\alpha_i}}Z_i^{-\frac{1}{1-\alpha_i}}K_i^{-\frac{\alpha_i}{1-\alpha_i}}$; substituting and collecting the constants into Γ_i gives the price form of (166).
3. Solving for Y_i/K_i: $(Y_i/K_i)^{\frac{\alpha_i}{1-\alpha_i}} = P_i Z_i^{\frac{1}{1-\alpha_i}}/(\Gamma_i\mathcal{P}_i)$, and raising both sides to $\frac{1-\alpha_i}{\alpha_i}$ gives the output form.

Expenditure shares and the sales split. Multiplying nominal gross output $S_i := P_iY_i$ by the pricing FOC gives the expenditure-share identity: intermediate spending is a constant share of nominal gross output,

$$X_i \equiv \mathcal{P}_i(\mathbf{P})M_i^e = \xi_i S_i, \qquad \xi_i := \frac{\varepsilon-1}{\varepsilon}(1-\alpha_i). \tag{167}$$

The input–output network therefore pins down nominal expenditure shares independently of current capital and productivity. This is the $\overline{v}\to\infty$ limit of Lemma 5(i), and the constant ξ_i is the limit of the share function $\xi_i(\overline{v})$ of (75).

The share ξ_i is a sales-split identity: $1-\alpha_i$ is the Cobb–Douglas cost share of the intermediate bundle, and with market power only the fraction $\frac{\varepsilon-1}{\varepsilon}$ of the sales price reflects cost, so every cost share of revenue is scaled by the inverse markup. Every euro of sales thus splits into three constant pieces,

$$\underbrace{\frac{1}{\varepsilon}}_{\text{pure rents}} + \underbrace{\frac{\varepsilon-1}{\varepsilon}\alpha_i}_{\text{capital income}} + \underbrace{\frac{\varepsilon-1}{\varepsilon}(1-\alpha_i)}_{\text{materials }(=\xi_i)} = 1 \tag{168}$$

the markup share $1/\varepsilon$ removed first and the remainder split by the technology. The capital piece is the capital income of (177) below, the rent piece the dividend S_i/ε. Under capacity the split acquires a fourth piece, the rationing rent, and the first three become functions of the threshold, with the rent piece alone keeping its value $1/\varepsilon$ of firm revenue (Lemma 5(iii)–(iv)).

Remark 12 (Scope of the constant-share result) Equation (167) is a firm-level pricing identity that needs only a constant demand elasticity (a constant markup) and a Cobb–Douglas requirement of the flexibly chosen bundle (a constant output elasticity of materials); the final-demand system, the network, the levels of $\mathbf{K}$ and $\mathbf{Z}$, the demand-risk distribution and constant returns play no role in it. It fails once the demand elasticity varies with the state (e.g. a binding capacity constraint truncating v) or the technology over the flexible inputs is not Cobb–Douglas.

Scope of the constant-share result, detailed version. The constant-share result is a firm-level pricing identity, and it depends on exactly two ingredients:

1. A constant price elasticity of demand for the variety (CES within the sector, with no capacity truncation of the demand shock, so the markup is constant)
2. A power-function (Cobb–Douglas) requirement of the flexibly chosen input bundle, so that the output elasticity of the bundle is constant.

What is *not* needed:

1. The final-demand system (Cobb–Douglas or CES never enters (167))
2. The network $\mathbf{A}$
3. The levels of $\mathbf{K}$ and $\mathbf{Z}$,
4. The demand-risk distribution ($\mu_{i,\frac{1}{1-\alpha_i}}$ scales expected costs and expected revenue equally, so it cancels from the ratio)
5. Constant returns is *not* needed: for $y = Z\,f(k)\,M^{\theta}$ with arbitrary f and $\theta \in (0,1)$ the same steps give $X_i/S_i = \frac{\varepsilon-1}{\varepsilon}\theta$.
6. Nor is the absence of other factors: additional flexibly chosen inputs entering the Cobb–Douglas nest with constant exponents would leave every input's revenue share constant. What would break constancy is a non-constant demand elasticity (e.g. a binding capacity constraint truncating v) or a non-Cobb–Douglas (e.g. CES) technology over the flexible inputs, either of which makes the cost share move with prices or with the state.

C.2 Finite Sectors: The Cobb–Douglas Closed Form

Throughout this subsection final demand is Cobb–Douglas and capacity is slack. The definition of the static equilibrium (Section 2.2) applies with item 1 replaced by the supply curve (166), item 2 by the constant-share identity (167), and the thresholds removed from the list of unknowns; Proposition 3(ii)–(iii) hold verbatim without capacity constraints, and the closed form below sharpens them. The aggregate capital K, the shares κ_i and the simplex Δ^{J-1} are those of (30)–(32).

Nominal shares. Let $\Xi := \operatorname{diag}(\xi_1,\dots,\xi_J)$, $\mathbf{S} := (S_1,\dots,S_J)'$, $\boldsymbol{\beta} := (\beta_1,\dots,\beta_J)'$, and let $\mathbf{A}$ be the input–output matrix with entry a_{ij} (row i buys from column j). Market clearing (27) in value form reads

$$S_j = \beta_j Q + \sum_{i=1}^{J} a_{ij} X_i, \qquad Q = \sum_{j=1}^{J} S_j - \sum_{j=1}^{J} X_j. \tag{169}$$

and using (167) this becomes, in vector form, $\mathbf{S} = \boldsymbol{\beta} Q + \mathbf{A}'\Xi\mathbf{S}$, hence

$$\mathbf{S} = \mathbf{h}\,Q, \qquad \mathbf{h} := \left[\mathbf{I} - \mathbf{A}'\Xi\right]^{-1}\boldsymbol{\beta}.$$

The inverse exists because $\xi_i < 1$ under $\varepsilon > 1$, $\alpha_i \in (0,1)$, and row-stochastic $\mathbf{A}$, so the spectral radius of $\mathbf{A}'\boldsymbol{\Xi}$ is below one. Thus $S_i = h_i Q$, which implies that nominal gross-output shares are fixed by the network, independent of $\mathbf{K}$ and $\mathbf{Z}$.

Closed-form prices and output. Combining $S_i = P_i Y_i = h_i Q$ with the supply curve (166) gives one log-linear equation per sector,

$$\frac{1}{\alpha_i}\ln P_i - \frac{1-\alpha_i}{\alpha_i}\sum_{j=1}^{J} a_{ij}\ln P_j - \ln Q = r_i - \ln K_i - \frac{1}{\alpha_i}\ln Z_i, \tag{170}$$

where $r_i := \ln h_i + \frac{1-\alpha_i}{\alpha_i}\ln\Gamma_i - \frac{1-\alpha_i}{\alpha_i}\sum_{j=1}^{J} a_{ij}\ln a_{ij}$ collects primitives.

With the numeraire $\sum_j \beta_j \ln P_j = 0$, equations (170) form a $(J{+}1)\times(J{+}1)$ linear system in $(\ln P_1, \ldots, \ln P_J, \ln Q)$. Given $\mathbf{P}$ and Q, all quantities follow:

$$\begin{aligned} &S_i = h_i Q, \quad Y_i = S_i/P_i, \quad X_i = \xi_i S_i \\ &M_i^e = X_i/\mathcal{P}_i(\mathbf{P}), \quad m_{ij} = a_{ij}X_i/P_j, \quad Q_j = \beta_j Q/P_j \end{aligned}$$

Static factorisation. Collect $\mathbf{p} := (\ln P_i)$, $\mathbf{k} := (\ln K_i)$, $\mathbf{z} := (\ln Z_i)$, and define $\mathbf{D}_\alpha := \mathrm{diag}(\alpha_i)$, $\mathbf{D}_{1-\alpha} := \mathrm{diag}(1-\alpha_i)$, and

$$\mathbf{M} := \mathbf{D}_\alpha^{-1}[\mathbf{I} - \mathbf{D}_{1-\alpha}\mathbf{A}]$$

Then (170) reads $\mathbf{M}\mathbf{p} - \mathbf{1}\ln Q = \mathbf{r} - \mathbf{k} - \mathbf{D}_\alpha^{-1}\mathbf{z}$. Eliminating prices with the numeraire yields the reduced final-output map

$$\ln Q = \ln \mathcal{B} + \sum_{i=1}^{J}\omega_i \ln K_i + \sum_{i=1}^{J}\frac{\omega_i}{\alpha_i}\ln Z_i, \qquad \boldsymbol{\omega}' = \boldsymbol{\beta}'\mathbf{M}^{-1} \tag{171}$$

where $\mathbf{M}$ is nonsingular and $\sum_i \omega_i = 1$ (Proposition 9 below), and $\mathcal{B}$ is a positive constant depending only on primitives, given explicitly in (173). Equivalently

$$Q(\mathbf{K},\mathbf{Z}) = \mathcal{B}\left(\prod_{i=1}^{J} K_i^{\omega_i}\right)\left(\prod_{i=1}^{J} Z_i^{\omega_i/\alpha_i}\right),$$

AK **form and capital composition.** This factorisation is the central property of the model: Cobb–Douglas demand and bundles fix expenditure shares, and constant markups fix nominal shares, so sectoral productivity shocks move the level of output and prices but not the relative weights on capital.

With aggregate capital K and the capital shares $\boldsymbol{\kappa} \in \Delta^{J-1}$ of (30)–(32), Then we have the factorisation

$$Q(\mathbf{K},\mathbf{Z}) = K\, q(\boldsymbol{\kappa},\mathbf{Z}), \qquad q(\boldsymbol{\kappa},\mathbf{Z}) = \mathcal{Z}(\mathbf{Z})\,\mathcal{K}(\boldsymbol{\kappa}), \tag{172}$$

with productivity shifter and composition term

$$\mathcal{Z}(\mathbf{Z}) := \prod_{i=1}^{J} Z_i^{\omega_i/\alpha_i}, \qquad \mathcal{K}(\boldsymbol{\kappa}) := \mathcal{B}\prod_{i=1}^{J}\kappa_i^{\omega_i}.$$

Proposition 9 (Static equilibrium under Cobb–Douglas final demand). *Let $\varepsilon > 1$, $\alpha_i \in (0,1)$, $\beta_i > 0$ with $\sum_i \beta_i = 1$, and $\mathbf{A}$ row-stochastic. For every $(\mathbf{K}, \mathbf{Z}) \in \mathbb{R}^J_{++} \times \mathbb{R}^J_{++}$ the Cobb–Douglas static equilibrium exists and is unique, and:*

(i) (Nominal shares.) $\mathbf{S} = \mathbf{h}Q$ *with* $\mathbf{h} = [\mathbf{I} - \mathbf{A}'\Xi]^{-1}\boldsymbol{\beta} > 0$*; the inverse exists because the spectral radius of $\mathbf{A}'\Xi$ is below one.*

(ii) (Prices.) $(\ln \mathbf{P}, \ln Q)$ *is the unique solution of the linear system* (170) *together with the numeraire $\sum_j \beta_j \ln P_j = 0$; the matrix $\mathbf{M} = \mathbf{D}_\alpha^{-1}[\mathbf{I} - \mathbf{D}_{1-\alpha}\mathbf{A}]$ is nonsingular and $\mathbf{M}\mathbf{1} = \mathbf{1}$.*

(iii) (Reduced output map.) (171) *holds with $\boldsymbol{\omega}' = \boldsymbol{\beta}'\mathbf{M}^{-1}$, $\omega_i > 0$ and $\sum_i \omega_i = 1$, and*

$$\ln \mathcal{B} = -\underbrace{\sum_{i=1}^{J} \omega_i \ln h_i}_{\text{Domar weights}} - \underbrace{\sum_{i=1}^{J} \omega_i \frac{1-\alpha_i}{\alpha_i}[\ln \Gamma_i + \mathcal{H}_i]}_{\text{wedges and sourcing mix}}, \tag{173}$$
$$\mathcal{H}_i := -\sum_{j=1}^{J} a_{ij} \ln a_{ij} \geq 0,$$

(iv) (AK factorisation.) $Q(\mathbf{K}, \mathbf{Z}) = K\, q(\boldsymbol{\kappa}, \mathbf{Z})$ *with $q = \mathcal{Z}(\mathbf{Z})\mathcal{K}(\boldsymbol{\kappa})$ as in* (172)*; in particular Q is homogeneous of degree one in $\mathbf{K}$ and $\sum_i K_i Q_{K_i} = Q$.*

(v) (Marginal products.) $Q_{K_i} = \omega_i Q/K_i = \omega_i\, q/\kappa_i$*, so relative marginal products $Q_{K_i}/Q_{K_j} = (\omega_i/\kappa_i)/(\omega_j/\kappa_j)$ do not depend on $\mathbf{Z}$; and $\mathcal{K}$ is maximised on the simplex at $\boldsymbol{\kappa}^* = \boldsymbol{\omega}$, where $Q_{K_i} = Q/K$ for every i.*

Proof. (i) Sectoral clearing was (27); multiply by P_j to get

$$Y_j P_j \equiv S_j = P_j Q_j + \sum_{i=1}^{J} m_{ij} P_j$$

use $m_{ij} = a_{ij} X_i / P_j$ and (80) to get

$$S_j = \beta_j Q + \sum_{i=1}^{J} a_{ij} X_i \tag{$\star$}$$

and using (167), $X_i = \xi_i S_i$,

$$S_j = \beta_j Q + \sum_{i=1}^{J} a_{ij}\xi_i S_i \implies \mathbf{S} = \boldsymbol{\beta} Q + \mathbf{A}'\Xi\mathbf{S}$$

The second equation in (169) comes from summing over j in equation ($\star$):

$$\sum_j S_j = \sum_j \beta_j Q + \sum_j \sum_i a_{ij} X_i = Q + \sum_i X_i$$

Since $\mathbf{A}$ is row-stochastic and $\xi_i \leq \bar{\xi} < 1$, the maximum row sum of $\Xi\mathbf{A}$ is $\max_i \xi_i < 1$, so $\rho(\mathbf{A}'\Xi) = \rho(\Xi\mathbf{A}) < 1$, the Neumann series converges, and $\mathbf{h} = \sum_{n\geq 0}(\mathbf{A}'\Xi)^n \boldsymbol{\beta} \geq \boldsymbol{\beta} > 0$.

(ii)

1. *Log supply.* Taking logs of (166),

$$\ln Y_i = \ln K_i + \frac{1-\alpha_i}{\alpha_i}\left[\ln P_i + \frac{1}{1-\alpha_i}\ln Z_i - \ln \Gamma_i - \ln \mathcal{P}_i(\mathbf{P})\right].$$

2. *Log demand.* Nominal shares are fixed by the network, $S_i = P_i Y_i = h_i Q$, so

$$\ln P_i + \ln Y_i = \ln h_i + \ln Q.$$

3. *Combine.* Substituting Step 1 into Step 2 and collecting $\ln P_i$,

$$\frac{1}{\alpha_i}\ln P_i - \frac{1-\alpha_i}{\alpha_i}\ln \mathcal{P}_i(\mathbf{P}) + \frac{1}{\alpha_i}\ln Z_i + \ln K_i - \frac{1-\alpha_i}{\alpha_i}\ln \Gamma_i = \ln h_i + \ln Q.$$

4. *Expand the bundle price.* By (13), $\ln \mathcal{P}_i = \sum_j a_{ij} \ln P_j - \sum_j a_{ij} \ln a_{ij}$. Substituting and moving all primitives to the right-hand side,

$$\frac{1}{\alpha_i}\ln P_i - \frac{1-\alpha_i}{\alpha_i}\sum_{j=1}^{J} a_{ij}\ln P_j - \ln Q$$
$$= \underbrace{\ln h_i + \frac{1-\alpha_i}{\alpha_i}\ln \Gamma_i - \frac{1-\alpha_i}{\alpha_i}\sum_{j=1}^{J} a_{ij}\ln a_{ij}}_{=r_i} - \ln K_i - \frac{1}{\alpha_i}\ln Z_i,$$

   which is (170).

Row i of $\mathbf{M}$ sums to $1/\alpha_i - (1-\alpha_i)/\alpha_i = 1$, so $\mathbf{M1} = \mathbf{1}$; and $\mathbf{M} = \mathbf{D}_\alpha^{-1}[\mathbf{I} - \mathbf{N}]$ with $\mathbf{N} := \mathbf{D}_{1-\alpha}\mathbf{A}$, whose row sums $1-\alpha_i$ are below one, so $\rho(\mathbf{N}) < 1$ and $\mathbf{M}$ is nonsingular. Given $\ln Q$, the prices are therefore $\mathbf{p} = \mathbf{M}^{-1}\left[\mathbf{1}\ln Q + \mathbf{r} - \mathbf{k} - \mathbf{D}_\alpha^{-1}\mathbf{z}\right]$, and the numeraire pins $\ln Q$ uniquely because $\boldsymbol{\beta}'\mathbf{M}^{-1}\mathbf{1} = 1$ (Step 4 of (iii) below); all remaining equilibrium quantities then follow mechanically from $(\mathbf{P}, Q)$ as displayed above, which proves existence and uniqueness.

(iii) Stack (170) across sectors as

$$\mathbf{Mp} - \mathbf{1}\ln Q = \mathbf{r} - \mathbf{k} - \mathbf{D}_\alpha^{-1}\mathbf{z}, \qquad \mathbf{M} = \mathbf{D}_\alpha^{-1}[\mathbf{I} - \mathbf{D}_{1-\alpha}\mathbf{A}]$$

1. *Prices given Q.* Solving the above for nonsingular $\mathbf{M}$,

$$\mathbf{p} = \mathbf{M}^{-1}\left[\mathbf{1}\ln Q + \mathbf{r} - \mathbf{k} - \mathbf{D}_\alpha^{-1}\mathbf{z}\right]$$

2. *Numeraire.* $P = 1$ in (80) reads $\boldsymbol{\beta}'\mathbf{p} = 0$. Premultiplying Step 1 by $\boldsymbol{\beta}'$,

$$\boldsymbol{\beta}'\mathbf{M}^{-1}\mathbf{1}\,\ln Q + \boldsymbol{\beta}'\mathbf{M}^{-1}\left[\mathbf{r} - \mathbf{k} - \mathbf{D}_\alpha^{-1}\mathbf{z}\right] = 0.$$

3. *Isolating* $\ln Q$. Solving,

$$\ln Q = \underbrace{-\frac{\boldsymbol{\beta}'\mathbf{M}^{-1}\mathbf{r}}{\boldsymbol{\beta}'\mathbf{M}^{-1}\mathbf{1}}}_{=\ln \mathcal{B}} + \boldsymbol{\omega}'\mathbf{k} + \boldsymbol{\omega}'\mathbf{D}_\alpha^{-1}\mathbf{z}, \qquad \boldsymbol{\omega}' := \frac{\boldsymbol{\beta}'\mathbf{M}^{-1}}{\boldsymbol{\beta}'\mathbf{M}^{-1}\mathbf{1}},$$

   and since $\mathbf{D}_\alpha^{-1}$ is diagonal the coefficient on $\ln Z_i$ is ω_i/α_i, which is (171).

4. *The denominator is one.* Row i of $\mathbf{M}$ sums to $1/\alpha_i - \frac{1-\alpha_i}{\alpha_i} = 1$ because $\mathbf{A}$ is row-stochastic, so $\mathbf{M1} = \mathbf{1}$, hence $\mathbf{M}^{-1}\mathbf{1} = \mathbf{1}$ and $\boldsymbol{\beta}'\mathbf{M}^{-1}\mathbf{1} = 1$: the normalisation in (171) is automatic, $\boldsymbol{\omega}' = \boldsymbol{\beta}'\mathbf{M}^{-1}$ exactly, and $\sum_i \omega_i = \boldsymbol{\omega}'\mathbf{1} = 1$.

5. *Expression for the constant* $\ln \mathcal{B}$. With the denominator equal to one, Step 3 reads $\ln \mathcal{B} = -\boldsymbol{\omega}'\mathbf{r}$. Substituting $r_i = \ln h_i + \frac{1-\alpha_i}{\alpha_i} \ln \Gamma_i - \frac{1-\alpha_i}{\alpha_i} \sum_{j=1}^{J} a_{ij} \ln a_{ij}$ and collecting the two primitive terms of r_i that involve sector i's cost side gives (173), where $\mathcal{H}_i$ is the entropy of sector i's sourcing weights: a more dispersed input mix raises the unbundling term in (13) and so lowers the level of output.

Positivity: $\mathbf{M}^{-1} = [\mathbf{I} - \mathbf{N}]^{-1}\mathbf{D}_\alpha = \sum_{n \geq 0} \mathbf{N}^n \mathbf{D}_\alpha$, so $\boldsymbol{\omega}' = \boldsymbol{\beta}'\mathbf{M}^{-1} \geq \boldsymbol{\beta}'\mathbf{D}_\alpha > 0$ componentwise.

(iv) substitute $K_i = \kappa_i K$ into the closed form and use $\sum_i \omega_i = 1$:

$$Q = \mathcal{B} \prod_{i=1}^{J} (\kappa_i K)^{\omega_i} \prod_{i=1}^{J} Z_i^{\omega_i/\alpha_i} = K^{\sum_i \omega_i}\, \mathcal{B} \prod_{i=1}^{J} \kappa_i^{\omega_i} \prod_{i=1}^{J} Z_i^{\omega_i/\alpha_i} = K\, \mathcal{Z}(\mathbf{Z})\, \mathcal{K}(\boldsymbol{\kappa}).$$

Degree-one homogeneity in $\mathbf{K}$ is thus equivalent to $\sum_i \omega_i = 1$, which holds automatically because $\mathbf{M1} = \mathbf{1}$. Finally, Euler's theorem gives

$$\sum_i K_i Q_{K_i} = \sum_i \omega_i Q = Q$$

so Q is homogeneous of degree one in $\mathbf{K}$.

In detail: differentiating the closed form $Q = \mathcal{B} \prod_j K_j^{\omega_j} \prod_j Z_j^{\omega_j/\alpha_j}$ in K_i gives

$$Q_{K_i} = \frac{\omega_i}{K_i}\, \mathcal{B} \prod_j K_j^{\omega_j} \prod_j Z_j^{\omega_j/\alpha_j} = \omega_i \frac{Q}{K_i}, \qquad \text{hence} \qquad \sum_{i=1}^{J} K_i Q_{K_i} = \sum_{i=1}^{J} \omega_i\, Q = Q,$$

using $\sum_i \omega_i = 1$ from Step 4. This is Euler's identity for a function homogeneous of degree one; equivalently, scaling $\mathbf{K} \mapsto \lambda \mathbf{K}$ multiplies Q by $\lambda^{\sum_i \omega_i} = \lambda$. Degree-one homogeneity of Q in $\mathbf{K}$ is thus the same statement as $\sum_i \omega_i = 1$, i.e. as $\mathbf{M1} = \mathbf{1}$.

(v) The first claim is Euler's identity in (iv): $Q_{K_i} = \omega_i Q/K_i$, and $Q/K_i = q/\kappa_i$. Since $\boldsymbol{\omega}$ is constant, the ratio $Q_{K_i}/Q_{K_j} = (\omega_i/\kappa_i)/(\omega_j/\kappa_j)$ involves no $\mathbf{Z}$. Finally, $\ln \mathcal{K}(\boldsymbol{\kappa}) = \ln \mathcal{B} + \sum_i \omega_i \ln \kappa_i$ is strictly concave on the simplex with first-order conditions $\omega_i/\kappa_i = \lambda$ for a common multiplier λ; summing over i gives $\lambda = 1$, hence $\kappa_i^* = \omega_i$, and there $Q_{K_i} = \omega_i Q/K_i = Q/K$. ■

Reading the weights $\boldsymbol{\omega}$ and constant $\mathcal{B}$. By Proposition 9(v), relative marginal products of capital are independent of $\mathbf{Z}$: productivity shocks scale all marginal products through $\mathcal{Z}(\mathbf{Z})$ but do not rotate them. $\mathcal{K}$ is maximised on the simplex at $\kappa_i^* = \omega_i$, where $Q_{K_i} = Q/K$ for every i, so $\boldsymbol{\omega}$ is the output-maximising *composition* of a given aggregate stock. Since $Q_{K_i} = (\omega_i/\kappa_i)(Q/K)$, the ratio ω_i/κ_i is sector i's marginal product of capital relative to the economy-wide average product. As such, a sector with $\kappa_i < \omega_i$ has too little capital relative to its importance for final output, directly (final demand) and

indirectly (through the network), so reallocating a marginal unit towards it raises Q; at $\boldsymbol{\kappa}^* = \boldsymbol{\omega}$ all marginal products are equalised at Q/K and no reallocation gain remains. The gaps $\omega_i - \kappa_i$ are therefore the natural reallocation signals of the dynamic model; the decentralised economy prices capital at the private flow instead (Section 2.3). Because $\boldsymbol{\omega}$ is constant, the target composition is invariant to $\mathbf{Z}$: productivity shocks do not change the direction of reallocation, in contrast to the CES case below.

All level effects of the distortions are in $\mathcal{B}$: markups and demand risk through $\Gamma_i = \frac{\varepsilon}{\varepsilon-1}\frac{1}{1-\alpha_i}\mu_{i,\frac{1}{1-\alpha_i}}$ (and ε additionally through $\mathbf{h}$ via ξ_i), giving $\partial \ln \mathcal{B}/\partial \ln \Gamma_i = -\omega_i \frac{1-\alpha_i}{\alpha_i} < 0$.

Two influence vectors: $\mathbf{h}$ and $\boldsymbol{\vartheta}$. The weights $\boldsymbol{\omega}$ set the *composition* (the exponents of the AK map and the target $\boldsymbol{\kappa}^* = \boldsymbol{\omega}$), the constant $\mathcal{B}$ the *level*. For the weights write

$$\boldsymbol{\omega}' = \boldsymbol{\vartheta}' \mathbf{D}_\alpha, \qquad \boldsymbol{\vartheta}' := \boldsymbol{\beta}' [\mathbf{I} - \mathbf{D}_{1-\alpha}\mathbf{A}]^{-1} = \boldsymbol{\beta}' \sum_{n=0}^{\infty} [\mathbf{D}_{1-\alpha}\mathbf{A}]^n,$$

i.e. $\omega_i = \alpha_i \vartheta_i$. This places a second upstream-propagation vector next to the Domar weights of Proposition 9(i), which by $\boldsymbol{\Xi} = \frac{\varepsilon-1}{\varepsilon}\mathbf{D}_{1-\alpha}$ read

$$\mathbf{h}' = \boldsymbol{\beta}' [\mathbf{I} - \boldsymbol{\Xi}\mathbf{A}]^{-1} = \boldsymbol{\beta}' \sum_{n=0}^{\infty} \left(\frac{\varepsilon - 1}{\varepsilon}\right)^n [\mathbf{D}_{1-\alpha}\mathbf{A}]^n.$$

Both answer "how important is sector i for the aggregate?" by passing final demand $\boldsymbol{\beta}$ upstream through the same supply chains, but they count different things and carry different per-link factors.

1. $\mathbf{h}$ *counts euros.* By the sales split (168), of each euro of sales the rent share $1/\varepsilon$ is paid out as dividends and only the cost share $\frac{\varepsilon-1}{\varepsilon}$ reaches the factors, of which the technology sends $1 - \alpha_i$ to suppliers. Spending therefore reaches the next link at the rate $\xi_i = \frac{\varepsilon-1}{\varepsilon}(1-\alpha_i)$, and $h_i = S_i/Q$ is an *accounting* object (Domar weight): sector i's equilibrium sales relative to GDP ($\sum_i h_i > 1$, since gross output double-counts intermediates).

2. $\boldsymbol{\vartheta}$ *counts elasticities.* In the log price system (170) a change in supplier j's log price enters buyer i's log price with coefficient $(1-\alpha_i)a_{ij}$, the *full* cost share: the markup multiplies marginal cost by the constant $\frac{\varepsilon}{\varepsilon-1}$ and so shifts the intercept r_i, not the pass-through. Summing over all chains to final demand, $\vartheta_j = \partial \ln Q/\partial \ln Z_j$ is the TFP elasticity of GDP and $\omega_j = \alpha_j \vartheta_j = \partial \ln Q/\partial \ln K_j$ its capital elasticity (171). As such ϑ_i is a *technological* objects. Explicitly, as a sum over walks through the network,

$$\vartheta_j = \sum_{n \geq 0} \sum_{i_0, \ldots, i_{n-1}} \beta_{i_0} (1 - \alpha_{i_0}) a_{i_0 i_1} \cdots (1 - \alpha_{i_{n-1}}) a_{i_{n-1} j},$$

each term being final demand β_{i_0} passed upstream link by link, shrunk by the materials share $1-\alpha$ at every node it crosses (the $n = 0$ term is β_j itself).

Why they disagree. In a competitive economy ($\varepsilon \to \infty$) $\xi_i = 1 - \alpha_i$ and $\mathbf{h} = \boldsymbol{\vartheta}$: the sales share *is* the TFP elasticity (Hulten's theorem). At finite ε every link of the sales recursion falls short of the technology's by the factor $\frac{\varepsilon-1}{\varepsilon}$, and the shortfall compounds with chain depth: $h_i < \vartheta_i$ whenever sector i has any customers, with the gap largest deep in supply chains, and $h_i \to \vartheta_i$ only as $\varepsilon \to \infty$. The equilibrium spending that reaches a sector, $h_j Q$, thus understates its output elasticity $\omega_j = \alpha_j \vartheta_j$ by exactly the compounded markup skim; the same ratio prices the private–social wedge of the dynamic model, $R_i/Q_{K_i} = \frac{\varepsilon-1}{\varepsilon} h_i/\vartheta_i$ (Proposition 10). Under CES final demand $\mathbf{h}(\mathbf{P})$ becomes price-dependent and $\boldsymbol{\vartheta}$ survives only as a local, state-dependent object; the comparison here is the Cobb–Douglas one.

[Q19 | tag: hide | ANSWERED] paragraph'nc-par:cd-comparative-statics'

Comparative statics. **[Flag: set aside on 5 November 2026]** This paragraph is held, unchanged and with its heading and label, inside a maroon box until what to do with it is decided. No black text refers to it. By (171) and the walk expansion, $\boldsymbol{\omega}$ is built from $(\boldsymbol{\beta}, \mathbf{A}, \boldsymbol{\alpha})$ alone; everything else enters through $\mathcal{B}$.

Final-demand weights $\boldsymbol{\beta}$ and the network $\mathbf{A}$. A single channel: $\boldsymbol{\vartheta}'$ is linear in $\boldsymbol{\beta}'$, so

$$\frac{\partial \vartheta_j}{\partial \beta_i} = \left[(\mathbf{I} - \mathbf{D}_{1-\alpha}\mathbf{A})^{-1}\right]_{ij} \geq 0,$$

the total-requirements coefficient from buyer i to supplier j, strictly positive exactly when $j = i$ or j supplies i directly or indirectly.

1. *Mechanism.* A euro of final spending shifted towards sector i makes final output more exposed to i's marginal cost and, through i's cost function at the share $1 - \alpha_i$, to its suppliers' costs, and so on up the chain; at every node the share α_j of that exposure is attributable to capital, so $\omega_j = \alpha_j \vartheta_j$ rises for i and all of its direct and indirect suppliers.
2. *Renormalisation.* $\sum_j \omega_j = 1$ holds identically and $\sum_j \beta_j = 1$ forces the rise in β_i to be paid for elsewhere, so the chain's gain is the rest of the economy's loss: $\boldsymbol{\kappa}^* = \boldsymbol{\omega}$ tilts towards i's supply chain and away from the sectors off it.

Rewiring a row of $\mathbf{A}$ towards a supplier (raising a_{ij} at the expense of $a_{ij'}$) is the same mechanism one link upstream: it redirects the exposure that passes through i, $(1 - \alpha_i)\vartheta_i$, from j' and its chain to j and its chain. $\boldsymbol{\beta}$ and $\mathbf{A}$ also move the level $\mathcal{B}$, through the Domar and entropy terms, but their first-order role is compositional.

Capital intensity α_i. Two channels, both visible in the walk expansion: α_i enters once as the claim $\omega_i = \alpha_i \vartheta_i$ and once through the factor $1 - \alpha_i$ carried by every walk that crosses sector i.

1. *Claim channel.* Holding the exposure ϑ_i reaching the sector fixed, a higher α_i attributes more of i's activity to its own capital: ω_i rises, and κ_i^* with it.
2. *Sink channel.* The complement $1 - \alpha_i$ is i's pass-through to its inputs. A higher α_i makes i more of a *sink*: exposure reaching it stays with its own capital instead of

being transmitted upstream, so ϑ_j falls for every sector that supplies i directly or indirectly.

Net effect: if no walk to i passes through i itself (no cycles through i), ϑ_i is unchanged (in the walk expansion α_i multiplies only *non-terminal* nodes), so ω_i strictly rises while every upstream ω_j falls, the gains and losses netting to zero by $\sum_j \omega_j = 1$; with cycles the sink channel also feeds back onto ϑ_i. In the one-sector limit ($J = 1$, $a_{11} = 1$) the two channels cancel exactly: $\vartheta_1 = 1/\alpha_1$ and $\omega_1 \equiv 1$. α_i also moves $\mathcal{B}$, through Γ_i and the weights on the entropy and wedge terms; we do not consider that here.

Markups ε and demand risk $\mu_{i,1/(1-\alpha_i)}$. These are *level-only* parameters: the slope coefficients of (170) are built from $(\boldsymbol{\alpha}, \mathbf{A})$ alone, and ε and $\mu_{i,\frac{1}{1-\alpha_i}}$ enter only its intercepts r_i (through Γ_i and h_i), hence only $\mathcal{B}$:

$$\frac{\partial \ln \mathcal{B}}{\partial \ln \Gamma_i} = -\omega_i \frac{1-\alpha_i}{\alpha_i} < 0, \qquad \Gamma_i = \underbrace{\frac{\varepsilon}{\varepsilon - 1}}_{\text{markup}} \times \underbrace{\frac{1}{1-\alpha_i}}_{\text{returns}} \times \underbrace{\mu_{i,\frac{1}{1-\alpha_i}}}_{\text{demand risk}}.$$

1. *Markup: a tax at every link.* At every link of every supply chain $1/\varepsilon$ of sales is diverted to rents before factors are paid, so each round of intermediate purchases is priced above marginal cost and input use (with it value added) is compressed cumulatively with chain depth (double marginalisation). Raising ε shrinks the per-link tax and lifts the level of activity.
2. *Demand risk: a convexity cost.* The input requirement is convex in the demand shock ($\frac{1}{1-\alpha_i} > 1$), so serving volatile demand takes $\mu_{i,\frac{1}{1-\alpha_i}} \geq 1$ times the inputs of serving its average: a higher effective marginal cost at unchanged expected output, i.e. a negative productivity shock with exposure $\omega_i \frac{1-\alpha_i}{\alpha_i}$.
3. *Why the composition is untouched.* With unit elasticities everywhere (Cobb–Douglas demand and bundles) all expenditure shares are price-invariant, so both distortions scale every sector's marginal product by the same factor and rotate none: $\boldsymbol{\omega}$, the exponents of $\mathcal{Z}$, and $\boldsymbol{\kappa}^*$ are unchanged.

Net effect: higher $\mu_{i,\frac{1}{1-\alpha_i}}$ lowers $\mathcal{B}$ unambiguously. For ε two channels oppose: Γ_i falls (the per-link tax shrinks, raising $\mathcal{B}$) while h_i rises (more of each euro is re-spent, and $\mathcal{B}$ loads on $-\sum_i \omega_i \ln h_i$); the net is positive, so $\mathcal{B}$ is increasing in ε. The markup's cross-sectional footprint shows up not in $\boldsymbol{\omega}$ but in the gap between $\mathbf{h}$ and $\boldsymbol{\vartheta}$ above.

CES final demand. Under CES final demand, since η enters only the final-good block, all within-sector objects of the model are unchanged (Section A.10). The intermediate block is intact, so $X_i = \xi_i S_i$, (167) and the supply relation (166) still hold; the nominal-output system (169) keeps its constant Leontief inverse but acquires a price-dependent final-demand loading,

$$\mathbf{S} = \mathbf{b}(\mathbf{P})Q + \mathbf{A}'\boldsymbol{\Xi}\mathbf{S} \implies \mathbf{S} = \left[\mathbf{I} - \mathbf{A}'\boldsymbol{\Xi}\right]^{-1}\mathbf{b}(\mathbf{P})Q \equiv \mathbf{h}(\mathbf{P})Q, \qquad \mathbf{b}(\mathbf{P}) := \left(\beta_j P_j^{1-\eta}\right)_{j=1}^{J}. \tag{174}$$

Because $\ln h_i(\mathbf{P})$ is the log of a price-weighted sum of the $\beta_j P_j^{1-\eta}$, it is not affine in $\ln \mathbf{P}$. Hence the demand side $\ln S_i = \ln h_i(\mathbf{P}) + \ln Q$ is no longer log-linear: the constant-coefficient system (170) no longer applies, and the factorisation (171) fails. Consequently:

- $Q(\mathbf{K}, \mathbf{Z})$ is not Cobb–Douglas in $\mathbf{K}$; the weights $\boldsymbol{\omega}$ in (171) are not constant, and $\kappa_i^* = \omega_i$ no longer holds.
- The separability $q(\boldsymbol{\kappa}, \mathbf{Z}) = \mathcal{Z}(\mathbf{Z})\mathcal{K}(\boldsymbol{\kappa})$ fails.
- Relative marginal products Q_{K_i}/Q_{K_j} now depend on $\mathbf{Z}$: sectoral TFP shocks rotate relative MPKs rather than merely scaling them.

Degree-one homogeneity survives: $Q = Kq(\boldsymbol{\kappa}, \mathbf{Z})$ still holds, so the AK/dynamic reduction is intact, but q is now defined implicitly by the static solve rather than in closed form.

C.3 Finite Sectors: Dynamics

The following results of Section 2.3 hold verbatim without capacity constraints and are not restated: the state generator (87) and Proposition 1 (the capital-pricing system, with the flow π_i^k of (176) below), Lemma 7, and Proposition 2 with the flow $\frac{\varepsilon-1}{\varepsilon}\alpha_i \widehat{s}_i/\kappa_i$ in (40). This subsection collects what the model without capacity constraints adds: the power form of the profit function and its proof, the cost-share route to the flow that prices capital, the closed-form reading of the reduced system and of the boundary behaviour, and the markup-wedge results, all of which use the power form of the profit function or the Cobb–Douglas closed form of Appendix C.2.

Profit as a function of own capital. The dynamic problem requires firm profit at an arbitrary own stock k, not only at the symmetric point.

Lemma 22 (Profit and marginal profit of a firm with own capital k). *Consider a sector-i monopolist with own capital k that prices optimally against* (160)*, taking $(Y_i, P_i, \mathcal{P}_i(\mathbf{P}), Z_i)$ as given, while its competitors are at the symmetric point.*

(i) Its expected operating profit is

$$\pi_i(k; \mathbf{K}, \mathbf{Z}) = [1 - \xi_i]\, S_i \left(\frac{k}{K_i}\right)^{\zeta_i}, \qquad \zeta_i := \frac{(\varepsilon-1)\alpha_i}{1 + (\varepsilon-1)\alpha_i} \in (0,1), \tag{175}$$

where S_i is equilibrium sectoral revenue.

(ii) At every k its input cost is the fraction ξ_i of its revenue, so profit inherits the own-capital elasticity ζ_i of revenue.

(iii) Its marginal profit at the symmetric point $k = K_i$ is, using the identity $\zeta_i[1-\xi_i] = \frac{\varepsilon-1}{\varepsilon}\alpha_i$,

$$\pi_i^k := \left.\frac{\partial \pi_i}{\partial k}\right|_{k=K_i} = \frac{\varepsilon - 1}{\varepsilon}\,\alpha_i\,\frac{S_i}{K_i}, \tag{176}$$

(iv) Items (i)–(iii) hold verbatim under Cobb–Douglas and CES final demand: η enters only through the equilibrium value of S_i, with $S_i = h_i Q$ in the former case and $S_i = h_i(\mathbf{P})Q$ in the latter.

Proof. Fix a sector-i monopolist holding own capital k. Being atomistic *within* its sector, it takes $(Y_i, P_i, \mathcal{P}_i(\mathbf{P}), Z_i)$ as given and chooses only its price p; quantities then follow from demand (160) and the technology (8). Every object below is conditional on those four.

1. *Expected sales and expected bundle at an arbitrary (p,k).* Evaluating (161)–(162) at $p_{in} = p, k_{in} = k$,

$$\mathbb{E}_v[y] = Y_i \left(\frac{p}{P_i}\right)^{-\varepsilon},$$

$$\mathbb{E}_v[M] = \mu_{i,\frac{1}{1-\alpha_i}} \left[Y_i \left(\frac{p}{P_i}\right)^{-\varepsilon}\right]^{\frac{1}{1-\alpha_i}} Z_i^{-\frac{1}{1-\alpha_i}} k^{-\frac{\alpha_i}{1-\alpha_i}}.$$

   Own capital enters *only* through $k^{-\frac{\alpha_i}{1-\alpha_i}}$ in the input requirement: more capital means a smaller bundle for the same output.

2. *The optimal price as a function of k.* Dividing,

$$\frac{\mathbb{E}_v[M]}{\mathbb{E}_v[y]} = \mu_{i,\frac{1}{1-\alpha_i}} Z_i^{-\frac{1}{1-\alpha_i}} [Y_i P_i^{\varepsilon}]^{\frac{\alpha_i}{1-\alpha_i}} p^{-\frac{\varepsilon\alpha_i}{1-\alpha_i}} k^{-\frac{\alpha_i}{1-\alpha_i}},$$

   and substituting into the pricing rule (163) gives an equation in p alone,

$$p^{\Theta_i} = \mathcal{C}_i\, k^{-\frac{\alpha_i}{1-\alpha_i}}, \quad \Theta_i := 1 + \frac{\varepsilon\alpha_i}{1-\alpha_i} > 0, \quad \mathcal{C}_i := \frac{\varepsilon}{\varepsilon-1}\frac{1}{1-\alpha_i}\mathcal{P}_i \mu_{i,\frac{1}{1-\alpha_i}} Z_i^{-\frac{1}{1-\alpha_i}} [Y_i P_i^{\varepsilon}]^{\frac{\alpha_i}{1-\alpha_i}}.$$

   We write it this way so that $\mathcal{C}_i$ collects *exactly* the objects the firm takes as given, hence does not depend on k. Therefore $p(k) \propto k^{-\frac{\alpha_i}{1-\alpha_i}/\Theta_i}$: a better-capitalised firm has lower expected marginal cost and, at a constant markup, charges a lower price.

3. *Optimised own revenue.* Own revenue is $s(k) := p(k)\mathbb{E}_v[y] = Y_i P_i^{\varepsilon}\, p(k)^{1-\varepsilon}$, so

$$\frac{\mathrm{d}\ln s}{\mathrm{d}\ln k} = (1-\varepsilon)\frac{\mathrm{d}\ln p}{\mathrm{d}\ln k} = \frac{(\varepsilon-1)\frac{\alpha_i}{1-\alpha_i}}{\Theta_i} =: \zeta_i \implies s(k) \propto k^{\zeta_i}.$$

4. *Costs are a fixed fraction of revenue at every k.* Rearranging (163) itself,

$$\mathcal{P}_i \mathbb{E}_v[M] = \frac{\varepsilon-1}{\varepsilon\frac{1}{1-\alpha_i}}\, p\, \mathbb{E}_v[y] = \xi_i\, s(k) \implies \pi_i(k) = s(k) - \mathcal{P}_i\mathbb{E}_v[M] = [1-\xi_i]s(k).$$

   This is (167) *off* the symmetric point: it is the pricing FOC evaluated at the firm's own k, not an equilibrium restriction, which is why it holds for every k and not only at $k = K_i$.

5. *Pinning the constant.* Since C_i is independent of k, the proportionality of Step 3 may be normalised anywhere. At the symmetric point $k = K_i$ we have $p = P_i$ and $s = P_i Y_i = S_i$, so $s(k) = S_i(k/K_i)^{\zeta_i}$ and therefore $\pi_i(k) = [1 - \xi_i]S_i(k/K_i)^{\zeta_i}$, which is (175).

6. *Marginal profit.* Differentiating and evaluating at $k = K_i$ gives $\pi_i^k = \zeta_i[1 - \xi_i]S_i/K_i$. Since $1 - \xi_i = [1 + (\varepsilon - 1)\alpha_i]/\varepsilon$ is exactly the denominator of the second form of ζ_i, the two cancel,

$$\zeta_i[1 - \xi_i] = \frac{(\varepsilon - 1)\alpha_i}{1 + (\varepsilon - 1)\alpha_i} \cdot \frac{1 + (\varepsilon - 1)\alpha_i}{\varepsilon} = \frac{\varepsilon - 1}{\varepsilon}\,\alpha_i,$$

which is (176).

■

Reading the profit function. The three pieces of the profit function (175) are

- $[1 - \xi_i]S_i$ is *sector-level operating profit*: sales minus the materials bill, i.e. the value-added share of the sector's gross output. By (177) below it comprises capital income $\frac{\varepsilon-1}{\varepsilon}\alpha_i S_i$ plus pure markup rents S_i/ε; since the sector has a unit measure of firms, it is also *average* firm profit at the symmetric point.
- $(k/K_i)^{\zeta_i}$ scales a firm's slice of that profit by its own capital relative to the sector's. The channel is entirely a price/demand one: a better-capitalised firm needs a smaller input bundle per unit of output, has lower expected marginal cost, and at the constant markup passes this into a lower price, gaining demand along the sectoral CES curve.
- $\zeta_i \in (0, 1)$ measures how strongly this works. It is increasing in ε because flatter demand converts a given cost advantage into more revenue with less price-cutting ($\zeta_i \to 1$ in the competitive limit: the concavity is purely demand-side, as the technology is CRS in (k, M))

  ζ_i is increasing in α_i, because a more capital-intensive technology turns own capital into a larger cost advantage (the elasticity of marginal cost with respect to k is $-\frac{\alpha_i}{1-\alpha_i}$).

Since $\zeta_i < 1$, profit is strictly concave in own capital (the curvature is inherited from the CES demand curve), so the firm problem of Section 2.3 is well posed; the same result holds node by node in the continuum (Lemma 25).

By Lemma 4(ii) the price of a firm with own capital k converges to the price without capacity constraints as $x_i \downarrow 0$ at fixed aggregates, and with it the profit function (85) to (175); by Lemma 5(v) the flow (33) converges to (176). Without capacity constraints Assumption 3 holds automatically: the function Φ_i of (84) reduces to the strictly decreasing power $\frac{\alpha_i}{\varepsilon}\mu_{i,\frac{1}{1-\alpha_i}}\overline{v}^{-\frac{1}{1-\alpha_i}}$, which is the statement $\zeta_i < 1$ of the lemma.

Rental rates and dividends. By Lemma 5(v), without capacity constraints the flow that prices capital, (33), and the firms' dividend become

$$R_i(\mathbf{K},\mathbf{Z}) = \pi_i^k = \frac{\varepsilon-1}{\varepsilon}\,\alpha_i\,\frac{S_i}{K_i}, \qquad \pi_i(K_i) - R_iK_i = \left[1-\xi_i-\frac{\varepsilon-1}{\varepsilon}\alpha_i\right]S_i = \frac{S_i}{\varepsilon}, \quad (177)$$

i.e. markup rents are exactly the revenue share $1/\varepsilon$, sector by sector, and $S_i = S_i^{\mathrm{f}}$ since there is no rationing rent. Together with (167) this completes the constant three-way split of sales: $1/\varepsilon$ to rents, $\frac{\varepsilon-1}{\varepsilon}\alpha_i$ to capital, $\xi_i = \frac{\varepsilon-1}{\varepsilon}(1-\alpha_i)$ to materials: the factor share $1-1/\varepsilon$ of each euro of sales is divided by the technology, α_i to capital and $1-\alpha_i$ to inputs. Total household income is rental income plus dividends,

$$\sum_i R_iK_i + \Pi = \sum_i [1-\xi_i]S_i = Q \quad (178)$$

with $\Pi := \sum_i S_i/\varepsilon$: value added exhausts final output; under capacity the household also receives the rationing rent Π^{r}.

Without capacity constraints (177) can also be reached from static cost minimisation alone, without the profit function (175). Suppose the firm hired capital at a rental rate R_i and bought the bundle at $\mathcal{P}_i$, both flexibly. Two steps:

1. *Cobb–Douglas cost shares.* Cost minimisation over (k, M) in the technology (8) makes factor outlays proportional to the exponents,

$$\frac{R_ik}{\mathcal{P}_iM} = \frac{\alpha_i}{1-\alpha_i} \implies R_iK_i = \frac{\alpha_i}{1-\alpha_i}\,X_i,$$

   aggregating over the sector ($X_i = \mathcal{P}_iM_i^e$).

2. *The expenditure-share identity.* By (167), $X_i = \xi_iS_i = \frac{\varepsilon-1}{\varepsilon}(1-\alpha_i)S_i$, so

$$R_iK_i = \frac{\varepsilon-1}{\varepsilon}\,\alpha_i\,S_i,$$

   which is K_i times (177).

Equivalently in one step: with both inputs flexible the technology is CRS, so total factor cost is the inverse-markup fraction $\frac{\varepsilon-1}{\varepsilon}$ of sales, split $\alpha_i : (1-\alpha_i)$ between capital and materials by the Cobb–Douglas exponents. The static rental payment equals K_i times the marginal profit flow π_i^k of the fixed-capital problem (176). Under capacity this route is closed: capital is committed before v_{in} is observed, and at capacity the realised input requirement is not on the Cobb–Douglas expansion path, so the flow that prices capital is the first-order condition of the pre-shock problem, (33) (Lemma 5(v)).

The reduced system under Cobb–Douglas final demand. Without capacity constraints the static block that delivers $\widehat{s}_i$ and q to (40), whose flow is then $\frac{\varepsilon-1}{\varepsilon}\alpha_i\widehat{s}_i/\kappa_i$, is explicit under Cobb–Douglas final demand: $\widehat{s}_i$ and q are constants computed once from $\mathbf{A}$ plus a closed form,

$$\widehat{s}_i = h_iq, \qquad \mathbf{h} = \left[\mathbf{I}-\mathbf{A}'\Xi\right]^{-1}\boldsymbol{\beta}, \qquad \boldsymbol{\omega}' = \frac{\boldsymbol{\beta}'\mathbf{M}^{-1}}{\boldsymbol{\beta}'\mathbf{M}^{-1}\mathbf{1}'}, \quad q = \mathcal{B}\prod_{i=1}^{J}\kappa_i^{\omega_i}\prod_{i=1}^{J}Z_i^{\omega_i/\alpha_i};$$

under CES the same objects are state-dependent, $\widehat{s}_i = h_i(\mathbf{P})\, q$ with the price-dependent loadings of (174) and q implicit in the static solve of Section 2.2 at each $(\boldsymbol{\kappa}, \mathbf{Z})$. Either way $\mathbf{A}$ acts in three distinct places: through the sales weights inside each sector's private flow $\widehat{s}_i$; through the shape of the common q (the weights $\boldsymbol{\omega}$ under Cobb–Douglas, the static solve under CES); and indirectly through (41), where the flows shape every sector's policy and the policies feed back into the shared aggregates $(c, \mu^K, \mu_{\boldsymbol{\kappa}})$ entering every other sector's equation.

Boundary behaviour on the simplex. Under Cobb–Douglas the private flow $\frac{\varepsilon-1}{\varepsilon}\alpha_i h_i\, q/\kappa_i \propto \kappa_i^{\omega_i - 1} \to \infty$ as $\kappa_i \downarrow 0$: the marginal value of a vanishing sector's capital diverges, investment there becomes arbitrarily attractive, and the resulting drift pushes the state away from the faces of Δ^{J-1}. So no sector's capital share is ever pushed to zero, and the faces need no imposed boundary data provided the scheme upwinds on $\mu_{\boldsymbol{\kappa}}$. Under CES the conclusion is η-dependent. For substitutes ($\eta > 1$) the same logic goes through: as $\kappa_i \downarrow 0$ the vanishing sector's price rises without bound and $\widehat{s}_i/\kappa_i \to \infty$, so the faces remain repelling. For complements ($\eta < 1$) the numeraire (26) caps the bottleneck sector's price at $P_i \le \beta_i^{-1/(1-\eta)}$ and aggregate output collapses together with the scarce sector, so the flow $\widehat{s}_i/\kappa_i$ diverges at most weakly and repulsion from the faces is not guaranteed a priori.

Underaccumulation and misallocation. Throughout this part final demand is Cobb–Douglas, so both flow vectors are constant. Write

$$\psi_i := \frac{\varepsilon - 1}{\varepsilon}\, \alpha_i h_i,$$

and recall from (171) that $\boldsymbol{\vartheta}' = \boldsymbol{\beta}'[\mathbf{I} - \mathbf{D}_{1-\alpha}\mathbf{A}]^{-1}$ and $\omega_i = \alpha_i \vartheta_i$. The two flows that price sector-i capital are

$$R_i = \psi_i \frac{q}{\kappa_i} \quad \text{(private, by (177))}, \qquad Q_{K_i} = \omega_i \frac{q}{\kappa_i} \quad \text{(internalised)},$$

so $\psi_i = R_i K_i / Q$ is sector i's capital income as a share of GDP and $\omega_i = Q_{K_i} K_i / Q$ its internalised counterpart. Their ratio factors into two wedges,

$$\frac{R_i}{Q_{K_i}} = \frac{\psi_i}{\omega_i} = \underbrace{\frac{\varepsilon - 1}{\varepsilon}}_{\text{own markup}} \times \underbrace{\frac{h_i}{\vartheta_i}}_{\substack{\text{sales-based vs cost-based} \\ \text{network influence}}}, \tag{179}$$

the first common to all sectors, the second (the ratio of the two influence vectors of the Cobb–Douglas comparative statics) sector-specific. Proposition 10 establishes the sign and the cross-sectional behaviour of (179); Lemma 23 then characterises exactly when its second factor is uniform across sectors.

Proposition 10 (Markup wedges: level and cross-section). *Assume $\varepsilon \in (1, \infty)$, $\alpha_i \in (0,1)$, $\beta_i > 0$ with $\sum_i \beta_i = 1$, and $\mathbf{A}$ row-stochastic, and write $\Theta \subset \mathbb{R}^{J^2+J}$ for the set of such primitives $\theta = (\boldsymbol{\beta}, \boldsymbol{\alpha}, \mathbf{A}, \varepsilon)$ in the free coordinates $(\beta_1, \ldots, \beta_{J-1}, \boldsymbol{\alpha}, (a_{ij})_{j<J}, \varepsilon)$. Then:*

1. (Uniform underpricing) *For every sector i,*

$$\omega_i - \psi_i = \alpha_i \sum_{n\geq 0} \left[1 - \left(\tfrac{\varepsilon-1}{\varepsilon}\right)^{n+1}\right] \left[\boldsymbol{\beta}'[\mathbf{D}_{1-\alpha}\mathbf{A}]^n\right]_i \;\geq\; \frac{\alpha_i\beta_i}{\varepsilon} \;>\; 0,$$

so $R_i < Q_{K_i}$ at every interior state ($\kappa_i > 0$). Each gap $\omega_i - \psi_i$ is strictly decreasing in ε and vanishes as $\varepsilon \to \infty$, where $h_i \to \vartheta_i$ and $\psi_i \to \omega_i$.

2. (Aggregate wedge.) $\sum_i \psi_i < 1 = \sum_i \omega_i$*: total capital income falls short of the internalised capital share of GDP.*

3. (Cross-sectional wedge.) $\boldsymbol{\psi} \propto \boldsymbol{\omega}$ *if and only if* $\mathbf{h} \propto \boldsymbol{\vartheta}$*, i.e. if and only if h_i/ϑ_i is the same for all i. For $J \geq 2$ the exceptional set*

$$\mathcal{N} := \{\theta \in \Theta : \; \mathbf{h} \propto \boldsymbol{\vartheta}\}$$

is closed in Θ and has Lebesgue measure zero, and for every $\theta \in \Theta \setminus \mathcal{N}$ the ratio $R_i/Q_{K_i} = \frac{\varepsilon-1}{\varepsilon}\, h_i/\vartheta_i$ takes at least two distinct values across sectors. (Lemma 23 below characterises $\mathcal{N}$ exactly.)

Proof. Throughout this proof and the proof of Lemma 23 write $\mathbf{N} := \mathbf{D}_{1-\alpha}\mathbf{A}$ and $m := \frac{\varepsilon-1}{\varepsilon} \in (0,1)$ for the inverse markup.

Step 0 (series representations). The row sums of $\mathbf{N}$ are $1 - \alpha_i \leq 1 - \min_i \alpha_i < 1$, so the spectral radius is $\rho(\mathbf{N}) < 1$ and the Neumann series for $[\mathbf{I} - \mathbf{N}]^{-1}$ and $[\mathbf{I} - m\mathbf{N}]^{-1}$ converge. Since $\boldsymbol{\Xi} = m\mathbf{D}_{1-\alpha}$, transposing $\mathbf{h} = [\mathbf{I} - \mathbf{A}'\boldsymbol{\Xi}]^{-1}\boldsymbol{\beta}$ gives

$$\mathbf{h}' = \boldsymbol{\beta}'[\mathbf{I} - \boldsymbol{\Xi}\mathbf{A}]^{-1} = \boldsymbol{\beta}' \sum_{n\geq 0} m^n \mathbf{N}^n, \qquad \boldsymbol{\vartheta}' = \boldsymbol{\beta}' \sum_{n\geq 0} \mathbf{N}^n.$$

Part (i).

(a) *The gap as a series.* Using $\omega_i = \alpha_i\vartheta_i$, $\psi_i = m\alpha_i h_i$ and Step 0, and subtracting the two series termwise,

$$\omega_i - \psi_i = \alpha_i(\vartheta_i - m\, h_i) = \alpha_i \sum_{n\geq 0} \left(1 - m^{n+1}\right) \left[\boldsymbol{\beta}'\mathbf{N}^n\right]_i.$$

(b) *Sign and lower bound.* Every term is nonnegative ($\boldsymbol{\beta}, \mathbf{N} \geq 0$ and $m < 1$), and the $n = 0$ term equals $(1-m)\beta_i > 0$; hence $\omega_i - \psi_i \geq (1-m)\alpha_i\beta_i = \alpha_i\beta_i/\varepsilon > 0$, and $R_i < Q_{K_i}$ wherever $\kappa_i > 0$.

(c) *Monotonicity and limit in ε.* m is increasing in ε, so each coefficient $1 - m^{n+1}$ is strictly decreasing in ε and so is the gap. The series is dominated by $\sum_n [\boldsymbol{\beta}'\mathbf{N}^n]_i = \vartheta_i < \infty$, so dominated convergence gives $\omega_i - \psi_i \to 0$, i.e. $h_i \to \vartheta_i$ and $\psi_i \to \omega_i$, as $m \uparrow 1$.

Part (ii). Sum (i)(a)–(b) over i and use $\sum_i \omega_i = 1$: $\sum_i \psi_i = 1 - \sum_i (\omega_i - \psi_i) < 1$.

Part (iii).

(a) *Equivalence.* Since $\alpha_i > 0$, $\psi_i = c\,\omega_i$ for all i iff $m\, h_i = c\,\vartheta_i$ for all i, i.e. iff $\mathbf{h} \propto \boldsymbol{\vartheta}$.

(b) *Genericity.* Each ratio difference $h_i/\vartheta_i - h_j/\vartheta_j$ is a continuous function of $\theta \in \Theta$, and a real-analytic function on the interior Θ° (where every $a_{ij} > 0$), which is open and connected (entries of inverses of matrices depending analytically on the primitives, with $\vartheta_j > 0$); and it is not identically zero: by Lemma 23 (whose proof uses only part (i), so there is no circularity) the differences all vanish at a parameter point exactly when $\boldsymbol{\beta}'$ is a left eigenvector of $\mathbf{N}$, which fails, for instance, at $\mathbf{A} = \mathbf{1}\mathbf{a}'$ with a strictly positive row $\mathbf{a}' \neq \boldsymbol{\beta}'$, where $\boldsymbol{\beta}'\mathbf{N} = (\sum_i \beta_i(1-\alpha_i))\mathbf{a}'$. The zero set of a real-analytic function that is not identically zero on a connected open domain is Lebesgue-null, the boundary $\Theta \setminus \Theta^\circ$ (where some $a_{ij} = 0$) is itself Lebesgue-null, and proportionality requires all $\binom{J}{2}$ such differences to vanish; hence $\mathcal{N}$ is Lebesgue-null, and it is closed in Θ as the common zero set of continuous functions. Off $\mathcal{N}$ at least one difference is nonzero, i.e. h_i/ϑ_i takes at least two values.

■

Consequently in this model, capital's private return is taxed twice, and (179) is the tax schedule.

- *Own markup (common level effect).* The firm's capital is paid its share of *after-markup* revenue: of each euro of sales, $1/\varepsilon$ is diverted to rents before factors are paid, which shrinks every sector's private flow by the same factor $\frac{\varepsilon-1}{\varepsilon}$.
- *Network undercounting (cross-sectional effect).* More capital in sector i lowers P_i, which lowers the costs of i's customers, their customers, and ultimately the price of the final basket. This pecuniary externality is counted by $\omega_i = \alpha_i\vartheta_i$ through the cost-based Leontief inverse, but firm i's own revenue records only i's own sales, and sales themselves undercount network influence. In the recursion behind $\mathbf{h}$, each euro of sales at a customer j passes on to j's suppliers only the re-spent fraction $\xi_j = \frac{\varepsilon-1}{\varepsilon}(1-\alpha_j)$: the rent share $1/\varepsilon$ is skimmed off before factors are paid, and only the materials part of the factor share flows upstream. In the recursion behind $\boldsymbol{\vartheta}$, by contrast, costs pass through at the full technological share $1-\alpha_j$. Every link therefore shrinks h_i relative to ϑ_i by an extra factor $\frac{\varepsilon-1}{\varepsilon}$, the undercounting compounds with network depth, and $h_i/\vartheta_i \leq 1$ is smallest for upstream sectors.

Underaccumulation (parts (i)–(ii)) is the product of the two taxes, both of which scale down the flow that prices capital, so private investment incentives are uniformly too weak. Capital is by this underpaid in *every* sector at *any* finite ε. *Misallocation* (part (iii)) is the second tax alone, and it is present for almost every configuration of primitives: proportionality of the two flow vectors is a knife edge, a closed Lebesgue-null set of primitives, and off it the second factor of (179) genuinely differs across sectors. It is not uniform and is heaviest for sectors whose influence runs through many intermediate links (upstream sectors), so the wedge $R_i/Q_{K_i} = \frac{\varepsilon-1}{\varepsilon}\, h_i/\vartheta_i$ varies with network position and distorts the cross-section of marginal values $\widehat{\mathsf{q}}_i$, hence the long-run composition of capital, and not merely its level.

The knife edge, exactly. Part (iii) of Proposition 10 reduces the cross-sectional wedge to whether h_i/ϑ_i is uniform across sectors. The following lemma characterises that configuration exactly: it is the case in which final demand is a left eigenvector of the transmission matrix.

Lemma 23 (The knife edge). *Under the assumptions of Proposition 10,* $\mathbf{h} \propto \boldsymbol{\vartheta}$ *if and only if* $\boldsymbol{\beta}'$ *is a left eigenvector of* $\mathbf{D}_{1-\alpha}\mathbf{A}$*. In that case the eigenvalue is necessarily the Perron root* $\nu = \rho(\mathbf{D}_{1-\alpha}\mathbf{A}) \in (0,1)$*, and*

$$\mathbf{h}' = \frac{\boldsymbol{\beta}'}{1-\frac{\varepsilon-1}{\varepsilon}\nu}, \qquad \boldsymbol{\vartheta}' = \frac{\boldsymbol{\beta}'}{1-\nu}, \qquad \frac{h_i}{\vartheta_i} = \frac{1-\nu}{1-\frac{\varepsilon-1}{\varepsilon}\nu}, \qquad \frac{R_i}{Q_{K_i}} = \frac{\varepsilon-1}{\varepsilon}\cdot\frac{1-\nu}{1-\frac{\varepsilon-1}{\varepsilon}\nu} \qquad \forall i.$$

Proof. Step 1 (sufficiency). If $\boldsymbol{\beta}'\mathbf{N} = \nu\boldsymbol{\beta}'$ then $\boldsymbol{\beta}'\mathbf{N}^n = \nu^n\boldsymbol{\beta}'$ for all n, and the series of Step 0 in the proof of Proposition 10 give $\mathbf{h}' = \boldsymbol{\beta}'\sum_n (m\nu)^n = \boldsymbol{\beta}'/(1-m\nu)$ and $\boldsymbol{\vartheta}' = \boldsymbol{\beta}'/(1-\nu)$ (geometric series, convergent because $\nu < 1$ by Step 3 below); dividing gives the constant ratio, and (179) gives the last display.

Step 2 (necessity). Suppose $\mathbf{h} = c\,\boldsymbol{\vartheta}$ for a scalar c. The two defining equations are $\mathbf{h}'[\mathbf{I} - m\mathbf{N}] = \boldsymbol{\beta}' = \boldsymbol{\vartheta}'[\mathbf{I} - \mathbf{N}]$; substituting $\mathbf{h}' = c\,\boldsymbol{\vartheta}'$ and subtracting,

$$c\,\boldsymbol{\vartheta}' - cm\,\boldsymbol{\vartheta}'\mathbf{N} = \boldsymbol{\vartheta}' - \boldsymbol{\vartheta}'\mathbf{N} \implies (c-1)\,\boldsymbol{\vartheta}' = (cm-1)\,\boldsymbol{\vartheta}'\mathbf{N}.$$

By Proposition 10(i), $h_i < \vartheta_i$ for all i, so $c < 1$ and the coefficient $cm - 1 < 0$ is nonzero; hence $\boldsymbol{\vartheta}'\mathbf{N} = \nu\,\boldsymbol{\vartheta}'$ with $\nu := (c-1)/(cm-1)$, i.e. $\boldsymbol{\vartheta}'$ is a left eigenvector of $\mathbf{N}$, and so is $\boldsymbol{\beta}' = \boldsymbol{\vartheta}'[\mathbf{I} - \mathbf{N}] = (1-\nu)\boldsymbol{\vartheta}'$, with the same eigenvalue.

Step 3 (the eigenvalue is the Perron root, and lies in $(0,1)$*).* $\boldsymbol{\beta} > 0$ is a strictly positive left eigenvector of the nonnegative matrix $\mathbf{N}$. Pair it with a Perron right eigenvector $\mathbf{y} \geq 0$, $\mathbf{y} \neq 0$, $\mathbf{N}\mathbf{y} = \rho(\mathbf{N})\mathbf{y}$: then $\boldsymbol{\beta}'\mathbf{N}\mathbf{y} = \nu\,\boldsymbol{\beta}'\mathbf{y} = \rho(\mathbf{N})\,\boldsymbol{\beta}'\mathbf{y}$ and $\boldsymbol{\beta}'\mathbf{y} > 0$, so $\nu = \rho(\mathbf{N})$. Finally $\rho(\mathbf{N}) < 1$ by Step 0 of the proposition, and $\rho(\mathbf{N}) \geq \min_i(1-\alpha_i) > 0$ because the spectral radius of a nonnegative matrix is at least its minimum row sum. ∎

Intuition. $\boldsymbol{\beta}'\mathbf{D}_{1-\alpha}\mathbf{A} = \nu\boldsymbol{\beta}'$ says that final demand, re-spent one round through the network, reproduces itself scaled by ν, i.e. every sector is equally central. The markup's per-link skim then compounds *uniformly* (it scales $\mathbf{h}$ down relative to $\boldsymbol{\vartheta}$ without twisting the cross-section), so the network-undercounting factor of (179) collapses to the common number $(1-\nu)/\left(1-\frac{\varepsilon-1}{\varepsilon}\nu\right)$, only the level wedge survives, and the cross-sectional distortion of part (iii) vanishes while the aggregate underaccumulation of parts (i)–(ii) remains.

Example. Common capital intensity $\alpha_i = \alpha$ and $\boldsymbol{\beta}$ the stationary distribution of the chain $\mathbf{A}$ ($\boldsymbol{\beta}'\mathbf{A} = \boldsymbol{\beta}'$): then $\boldsymbol{\beta}'\mathbf{D}_{1-\alpha}\mathbf{A} = (1-\alpha)\boldsymbol{\beta}'\mathbf{A} = (1-\alpha)\boldsymbol{\beta}'$, so $\nu = 1-\alpha$ and $h_i/\vartheta_i = \frac{\varepsilon\alpha}{1+(\varepsilon-1)\alpha}$; the private–social ratio is $\frac{\varepsilon-1}{\varepsilon}\,h_i/\vartheta_i = \frac{(\varepsilon-1)\alpha}{1+(\varepsilon-1)\alpha}$, the same value as in the one-sector economy ($J = 1$, $a_{11} = 1$, $\omega = 1$), of which this example is the network generalisation.

C.4 The Continuum: Statics

The following results of Sections 3–3.2 do not involve the capacity block and hold verbatim without capacity constraints: the graphon primitives of Section 3 in full (Definitions 2 and 3, Lemmas 8, 9, 10 and 11, and Assumption 4 without its item on the capacity technology); the bundle price index (95) and Lemma 12(i); the final CES demand (99)–(100) and Remark 5; the law of motion of capital (103); the exact law of large numbers (Lemma 13) and the measure ensemble (Definition 4); the bundle-price aggregate (44), the numeraire condition (47), Lemma 14 and Remark 7, with the spending moment specialised as below. What the absence of capacity constraints changes is the within-node block: the pricing relation becomes the constant-wedge supply curve, spending the constant share $\xi(i)$ of sales, and the node block a pair of equations in (P, S) whose unique solution follows from the ordering of two exponents. Throughout this subsection

$$\Gamma(i) := \frac{\varepsilon}{\varepsilon-1}\frac{1}{1-\alpha(i)}\mu_{i,\frac{1}{1-\alpha(i)}}, \qquad \xi(i) := \frac{\varepsilon-1}{\varepsilon}(1-\alpha(i)) \le \bar{\xi} := \frac{\varepsilon-1}{\varepsilon}(1-\underline{\alpha}) < 1,$$

are the markup constant and the expenditure share of the model without capacity constraints, node by node (Appendix C.1), the bound $\bar{\xi}$ and $\Gamma(i) \in [\frac{\varepsilon}{\varepsilon-1}\frac{1}{1-\underline{\alpha}}, \frac{\varepsilon}{\varepsilon-1}\frac{\bar{\mu}}{1-\bar{\alpha}}]$ following from Assumption 4.

Within-node block. By Lemma 4(iii) applied node by node, without capacity constraints the pricing/threshold relation (96) and the spending relation (98) become

$$\text{(pricing/supply)} \qquad P_i = \Gamma(i)\mathcal{P}_i Z_i^{-\frac{1}{1-\alpha(i)}}\left(\frac{Y_i}{K_i}\right)^{\frac{\alpha(i)}{1-\alpha(i)}} \tag{180}$$

$$\text{(expenditure share)} \qquad X_i = \mathcal{P}_i M_i^e = \xi(i)\, S_i, \qquad S_i := P_i Y_i, \tag{181}$$

with $S_i^{\mathrm{f}} = S_i$ and $\Pi_i^{\mathrm{r}} = 0$. Lemma 12(ii) then states that (180)–(181) hold at each node's symmetric equilibrium, and (iii) that value clearing reads

$$S_i = \underbrace{\beta(i)P_i^{1-\eta}Q}_{\text{final}} + \underbrace{\int_0^1 a(j,i)\,\xi(j)\,S_j\,\mathrm{d}j}_{=(\mathbb{A}^*\Xi S)(i)}, \qquad Q = \int_0^1 [1-\xi(j)]S_j\,\mathrm{d}j : \tag{182}$$

node j spends $a(j,i)\xi(j)S_j$ on node-i goods, a fixed value share of a fixed share of its sales, and every euro of node i's sales is either final spending (price-sensitive, elasticity η) or a fixed value share $a(j,i)$ of some buyer's intermediate budget $\xi(j)S_j$. Remark 5 then holds in its original form: expenditure *shares* and the markup are constant, and the static block is a scalar equation per node in the price.

Aggregate functionals. Since $X = \xi(i)S$ node by node, the spending moment is $m_X(t,i) = \xi(i)m_S(t,i)$ and Definition 5 reduces to a fixed point in the two moments

(m_P, m_S), with

$$\mathcal{D}(t,i) = (\mathbb{A}^*\Xi\, m_S)(t,i) = \int_0^1 a(j,i)\,\xi(j)\,m_S(t,j)\,\mathrm{d}j, \qquad Q(t) = \int_0^1 [1-\xi(j)]\,m_S(t,j)\,\mathrm{d}j. \tag{183}$$

Lemma 14 applies with the integrability of X implied by that of S (bounded nonrandom factor $\bar{a}\bar{\xi}$ in the inbound-demand integral), and the GDP identity of Remark 7 reads $Q = \int_0^1 [1-\xi(j)] m_S(t,j)\mathrm{d}j$, with the bound $\int_0^1 m_S \mathrm{d}j \le (1-\bar{\xi})^{-1} Q$; under Cobb–Douglas final demand the demand relation iterates to the Domar density of Lemma 9(iii), $m_S = hQ$(Appendix C.5). The coupling remark (Remark 6) holds with (m_P, m_S) as the two functions each node needs to know: $(\mathbb{A}^*\Xi m_S)(t,i)$ is the kernel-weighted intermediate spending of i's buyers, their mean sales scaled by their constant shares.

The per-node static system. Fix aggregates $(\mathcal{P}(t,i), \mathcal{D}(t,i), Q(t))$. Combining (180) with $S_i = P_i Y_i$

$$S_i = P_i Y_i = K_i\,[\Gamma(i)\,\mathcal{P}(t,i)]^{-\frac{1-\alpha(i)}{\alpha(i)}}\,(Z_i P_i)^{1/\alpha(i)}. \tag{184}$$

Demand-side clearing at the node, from (182), is

$$S_i = \beta(i)\,P_i^{1-\eta}\,Q(t) + \mathcal{D}(t,i). \tag{185}$$

Equations (184)–(185) are two equations in the two node-level unknowns (P_i, S_i) given (K_i, z_i) and aggregates: the threshold has dropped out, and with it the third equation of the capacity block. They are the capacity limit of (107)–(109): as $\overline{v} \to \infty$, $\mathrm{Y}_i(\overline{v}) \sim \overline{v}^{\frac{\alpha(i)}{1-\alpha(i)}}/\Gamma(i)$ and $U_i(\overline{v}) \sim \mu_{i,1/\varepsilon}^{-\frac{\varepsilon}{\varepsilon-1}}/\overline{v}$ (Lemma 4(i)), so (108) becomes $S = \mathcal{P}K\,\Gamma(i)\,\overline{v}^{-\frac{1}{1-\alpha(i)}}/x(i)$ and (107) gives $\overline{v}^{-\frac{1}{1-\alpha(i)}} = [e^z x(i)^{\alpha(i)} P/(\Gamma(i)\mathcal{P})]^{\frac{1}{\alpha(i)}}$, which is (184) once the powers of $x(i)$ are collected.

Lemma 24 (The per-node static block). *Fix i and aggregates with $\mathcal{P}(t,i) > 0$, $Q(t) > 0$, $\mathcal{D}(t,i) \ge 0$. Then: (i) node sales satisfy the supply relation* (184)*; (ii) node sales satisfy the demand relation* (185)*; (iii) for every $(K,z) \in (0,\infty)\times\mathbb{R}$ the system* (184)–(185) *has a unique solution $P(t,i,K,z) > 0$, $S(t,i,K,z) > 0$.*

Proof. (i) Invert (180): $Y_i = K_i\left[P_i Z_i^{\frac{1}{1-\alpha}}/(\Gamma\mathcal{P})\right]^{\frac{1-\alpha}{\alpha}}$, exactly as in the finite model (the exponent $\frac{1-\alpha}{\alpha}$ is the reciprocal of the supply exponent $\frac{\alpha}{1-\alpha}$ in (180)). Then $S_i = P_i Y_i = K_i(\Gamma\mathcal{P})^{-\frac{1-\alpha}{\alpha}} Z_i^{\frac{1}{1-\alpha}\frac{1-\alpha}{\alpha}} P_i^{1+\frac{1-\alpha}{\alpha}}$, and both Z- and P-exponents equal $1/\alpha$.

(ii) It is the value-clearing identity (182) for node i, conditioned on the node's own state, with the cross-node terms replaced by their deterministic aggregate versions:

(1) Start from the first identity of (182) at a realisation: $S_i = \beta(i)P_i^{1-\eta}Q + \int_0^1 a(j,i)\,\xi(j)\,S_j\mathrm{d}j$.

(2) By Lemma 14, the inbound-demand integral is $\mathbb{P}$-a.s. equal to the deterministic $\mathcal{D}(t,i) = (\mathbb{A}^*\Xi m_S)(t,i)$, and $Q = Q(t)$ is deterministic. Node i's own realisation

perturbs neither object: i is a single point of a continuum (atomistic), so its own (K_i, z_i) has zero influence on the kernel integrals, which is why the *same* $\mathcal{D}(t,i)$ and $Q(t)$ appear for every realisation of the node state.

(3) What remains random in the identity is only the node-local pair (P_i, S_i) through (K_i, z_i): writing it state by state gives (185), one equation per (K, z), paired with the supply relation (184).

(iii) Substitute (184) into (185) and define

$$F(P) := c\,P^{1/\alpha} - \beta Q\,P^{1-\eta} - \mathcal{D}, \qquad c := K(\Gamma\mathcal{P})^{-\frac{1-\alpha}{\alpha}} e^{z/\alpha} > 0.$$

For $\eta < 1$, $F(0^+) = -\mathcal{D} \le 0$ (and $F(0^+) < 0$ if $\mathcal{D} > 0$; if $\mathcal{D} = 0$ both terms vanish at 0 and the argument below still applies on $(0,\infty)$); for $\eta > 1$ the demand term $-\beta Q P^{1-\eta} \to -\infty$ as $P \downarrow 0$, so $F(0^+) = -\infty < 0$. In either case $F < 0$ near 0, and $F(\infty) = +\infty$ since $1/\alpha > 1 > \max\{0, 1-\eta\}$. At any root $\bar{P}$, $c\bar{P}^{1/\alpha} = \beta Q \bar{P}^{1-\eta} + \mathcal{D} \ge \beta Q \bar{P}^{1-\eta}$, so

$$\bar{P}F'(\bar{P}) = \tfrac{1}{\alpha}\,c\bar{P}^{1/\alpha} - (1-\eta)\,\beta Q\bar{P}^{1-\eta} \ge \Big(\tfrac{1}{\alpha} - (1-\eta)\Big)\beta Q\bar{P}^{1-\eta} + \tfrac{1}{\alpha}\mathcal{D} > 0,$$

because $1/\alpha > 1 > 1-\eta$. F can therefore cross zero only from below, so the root is unique. $S > 0$ follows from (184). ∎

The argument is the proof of Lemma 15(iv) without capacity constraints: the supply elasticity in the price is the constant $1/\alpha(i)$, which exceeds the demand elasticity $1-\eta$ of the final-demand term for every $\eta > 0$; under capacity the supply elasticity is the state-dependent $e_S(\overline{v}) > 1$ of (111), and Assumption 5 is what keeps it above one. Given the solution, node quantities follow mechanically: $Y_i = S_i/P_i$, $X_i = \xi(i)S_i$, $M_i^e = X_i/\mathcal{P}(t,i)$, $m_i(j) = a(i,j)X_i/P_j$, $Q_i = \beta(i)P_i^{-\eta}Q$. Operating profit (sales net of intermediate purchases; capital is owned, there is no labour) is

$$\pi_i = S_i - X_i = [1 - \xi(i)]\,S(t,i,K_i,z_i), \tag{186}$$

which is the whole of the node's value added, there being no rationing rent.

The individual firm's profit. The individual monopolist's perceived profit as a function of its own capital k, holding the node's symmetric state (K,z) and all aggregates fixed, is a power of k, derived exactly as for finite sectors (Lemma 22):

Lemma 25 (Individual profit function). *Consider a monopolist at node i with own capital k, when the node is at (K,z) with sales $S(t,i,K,z)$ and all its competitors at the symmetric point. Its optimally-priced expected operating profit is*

$$\pi(t,i,k;K,z) = [1-\xi(i)]\,S(t,i,K,z)\left(\frac{k}{K}\right)^{\zeta(i)}, \qquad \zeta(i) := \frac{(\varepsilon-1)\alpha(i)}{1+(\varepsilon-1)\alpha(i)} \in (0,1), \tag{187}$$

so its marginal profit at the symmetric point $k = K$ is, using the identity $\zeta(i)[1-\xi(i)] = \frac{\varepsilon-1}{\varepsilon}\alpha(i)$ of (176),

$$\left.\frac{\partial\pi}{\partial k}\right|_{k=K} = \zeta(i)\,[1-\xi(i)]\,\frac{S(t,i,K,z)}{K} = \frac{\varepsilon-1}{\varepsilon}\,\alpha(i)\,\frac{S(t,i,K,z)}{K}. \tag{188}$$

Proof. Step 1: demand and input requirement. A monopolist posting price p faces expected demand $\mathbb{E}_v[y] = Y_i\,(p/P_i)^{-\varepsilon}$ (the node CES demand curve, with the firm-specific shock v integrated out; $Y_i, P_i, \mathcal{P}$ are node objects, exactly as in the finite model). Inverting the technology $y = Zk^{\alpha}M^{1-\alpha}$ for the bundle,

$$M = \left(y\,Z^{-1}k^{-\alpha}\right)^{\frac{1}{1-\alpha}} = y^{\frac{1}{1-\alpha}}\,Z^{-\frac{1}{1-\alpha}}\,k^{-\frac{\alpha}{1-\alpha}},$$

so expected input spending is $\mathcal{P}\,\mathbb{E}_v[M] = \mathcal{P}\,\mu_{i,\frac{1}{1-\alpha}}[Y_i(p/P_i)^{-\varepsilon}]^{\frac{1}{1-\alpha}}Z^{-\frac{1}{1-\alpha}}k^{-\frac{\alpha}{1-\alpha}}$, the moment $\mu_{i,\frac{1}{1-\alpha}}$ arising because M is the convex power $\frac{1}{1-\alpha} > 1$ of realised demand.

Step 2: profit and first-order condition. Write

$$\pi(p) = R(p) - C(p), \qquad R(p) = P_i^{\varepsilon}Y_i\,p^{1-\varepsilon}, \qquad C(p) = \mathcal{P}\,\mu_{i,\frac{1}{1-\alpha}}\,(Y_iP_i^{\varepsilon})^{\frac{1}{1-\alpha}}\,Z^{-\frac{1}{1-\alpha}}k^{-\frac{\alpha}{1-\alpha}}\,p^{-\frac{\varepsilon}{1-\alpha}}.$$

Setting $\pi'(p) = 0$ gives $(\varepsilon-1)R(p) = \frac{\varepsilon}{1-\alpha}\,C(p)$. The critical point is the unique maximum: as $p \downarrow 0$, $C/R \propto p^{-\frac{\varepsilon}{1-\alpha}+\varepsilon-1} \to \infty$ (cost dominates, $\pi \to -\infty$), while as $p \to \infty$ both terms vanish.

Step 3: constant cost share. The FOC rearranges to

$$\frac{C(p)}{R(p)} = \frac{\varepsilon-1}{\varepsilon}\,(1-\alpha) = \xi(i),$$

i.e. at the optimum input cost is the fraction ξ of revenue regardless of (k, K, z), the constant-markup logic. Hence $\pi(k) = [1-\xi]\,R(p(k)) = [1-\xi]\,p(k)^{1-\varepsilon}P_i^{\varepsilon}Y_i$.

Step 4: the posted price. Substituting R and C into the FOC and collecting powers of p,

$$p(k)^{\Theta(i)} = \underbrace{\frac{\varepsilon}{\varepsilon-1}\frac{1}{1-\alpha}\mu_{i,\frac{1}{1-\alpha}}}_{=\Gamma(i)}\,\mathcal{P}\,Y_i^{\frac{\alpha}{1-\alpha}}P_i^{\frac{\varepsilon\alpha}{1-\alpha}}\,Z^{-\frac{1}{1-\alpha}}\,k^{-\frac{\alpha}{1-\alpha}}, \qquad \Theta(i) := 1-\varepsilon+\frac{\varepsilon}{1-\alpha} = 1+\frac{\varepsilon\alpha}{1-\alpha} > 0,$$

the same exponent Θ as in the finite model.

Step 5: constant elasticity. Taking logs in Step 4, $\frac{\partial \ln p}{\partial \ln k} = -\frac{\alpha/(1-\alpha)}{\Theta(i)}$, and from Step 3,

$$\frac{\partial \ln \pi}{\partial \ln k} = (1-\varepsilon)\,\frac{\partial \ln p}{\partial \ln k} = \frac{(\varepsilon-1)\frac{\alpha}{1-\alpha}}{\Theta(i)} = \frac{(\varepsilon-1)\alpha}{1+(\varepsilon-1)\alpha} = \zeta(i),$$

a constant, the second equality using $\Theta(i) = [1+(\varepsilon-1)\alpha]/(1-\alpha)$. Integrating in $\ln k$ between K and k: $\pi(k) = \pi(K)(k/K)^{\zeta}$.

Step 6: the symmetric point. At $k = K$, dividing Step 4 by $P_i^{\frac{\varepsilon\alpha}{1-\alpha}}$ shows $p = P_i$ solves it iff $P_i = \Gamma\mathcal{P}\,Y_i^{\frac{\alpha}{1-\alpha}}/\left(Z^{\frac{1}{1-\alpha}}K^{\frac{\alpha}{1-\alpha}}\right)$, which is precisely the node pricing relation (180): the symmetric price is the equilibrium node price. Hence $p(K) = P_i$, revenue $R = P_iY_i = S(t, i, K, z)$, and $\pi(K) = [1-\xi]S$, giving (187).

Step 7: marginal profit and the range of ζ. Differentiating (187) in k and evaluating at $k = K$ gives (188), where the second form follows from $1-\xi(i) = [1+(\varepsilon-1)\alpha(i)]/\varepsilon$

cancelling the denominator of $\zeta(i)$. Finally $\zeta > 0$ since $\varepsilon > 1$ and $\alpha > 0$, and $\zeta < 1$ since $1 + (\varepsilon - 1)\alpha - (\varepsilon - 1)\alpha = 1 > 0$; moreover $\zeta \uparrow 1$ as $\varepsilon \to \infty$ (competitive limit: marginal approaches average profit) and $\zeta \downarrow 0$ as $\alpha \to 0$ (capital drops out of production). ■

Since $\zeta(i) < 1$, profit is strictly concave in own capital, the curvature inherited from the CES demand curve; this is Lemma 16(iii) without capacity constraints, where the same conclusion requires Assumption 3. The power form is the large-capital asymptote of the capacity profit function (Lemma 6(iv)), and the constant flow (188) is the limit of (115) as the node's threshold diverges (Lemma 5(v)).

[Q20 | tag: hide | ANSWERED] subsection "app:nc-cd"

C.5 The Continuum: The Cobb–Douglas Closed Forms

[Flag: set aside on 5 November 2026] This subsection is held, unchanged and with its heading and labels, inside a maroon box until what to do with it is decided. Its black-text pointers, in the continuum statics of this appendix (the Domar density), in the subsection on the step-kernel system that follows, in the appendix on graphon limits (the Fredholm equation (191)) and in Steps 0, 1, 2, 3 and 5 of the numerical appendix (the closed forms and the balanced growth path of Proposition 12 as initialiser), are in maroon boxes too, with the claims left black.

The following results of Section B.7 do not involve the capacity block and hold verbatim without capacity constraints: the definition of the Cobb–Douglas aggregator (117) and the two degeneracies; Theorem 1, items 2, 3 and 5; the two-channel argument of Section B.7.3; Remark 9; and the continuity of the CES repair at $\eta = 1$. What the absence of capacity constraints changes is that the cancellation goes all the way to a closed form: sales are pinned by primitives and the scalar Q through the Domar density, cash flows are constant in K as well as in z, the price moment solves a linear Fredholm equation, the balanced growth path is algebraic, and the mean-field coupling through payoffs is removed entirely. Throughout this subsection $\xi(i)$, $\Gamma(i)$ and $\bar{\xi}$ are the constants of Appendix C.4, Ξ is pointwise multiplication by ξ, and $h = (\mathrm{Id} - \mathbb{A}^*\Xi)^{-1}\beta$ is the Domar density of Lemma 9(iii).

Cancellation. By Lemma 4, without capacity constraints $\varrho_i \equiv 1$, $\xi_i(\overline{v}) \to \xi(i)$ and the threshold drops out of Lemma 17; its items (i)–(iv) then collapse to the following statement, whose proof is given beside it.

Lemma 26 (Cancellation lemma). *Let $\eta = 1$. Then in any static equilibrium the sales function is state-independent:*

$$S(t, i, K, z) = h(i)\, Q(t) \qquad \text{for all } (K, z), \qquad h := (\mathrm{Id} - \mathbb{A}^*\Xi)^{-1}\beta, \tag{189}$$

with h the Domar density of Lemma 9, satisfying $\int_0^1 [1 - \xi(i)]h(i)\mathrm{d}i = 1$. Consequently node operating profits $\pi_i = [1 - \xi(i)]h(i)Q(t)$ contain neither K_i nor z_i: productivity and capital drop out of node cash flows entirely. Prices absorb everything: $P(t, i, K, z)^{1/\alpha(i)} \propto e^{-z/\alpha(i)}K^{-1}$.

Proof. Step 1: the own price drops out of demand. Set $\eta = 1$ in (109): final spending becomes $\beta(i)P_i^{1-\eta}Q = \beta(i)Q$, so

$$S(t,i,K,z) = \beta(i)\,Q(t) + \mathcal{D}(t,i).$$

Both right-hand terms are deterministic aggregates (Lemma 14) and neither involves (K,z): the node's own price (the only channel through which its state entered its own demand) has vanished. Hence $S(t,i,K,z)$ is the same number at every state; write it $S(t,i)$.

Step 2: moments collapse. Since $S(t,i,\cdot,\cdot)$ is constant in the state, $m_S(t,i) = \int S\,\mathrm{d}\mathfrak{M}(t,i,\cdot) = S(t,i)$.

Step 3: the linear fixed point and its Neumann solution. Substituting $\mathcal{D} = \mathbb{A}^*\Xi\, m_S = \mathbb{A}^*\Xi\, S$ from (183) into Step 1 gives a linear equation on L^1,

$$S = \beta\,Q + \mathbb{A}^*\Xi\, S \iff (\mathrm{Id} - \mathbb{A}^*\Xi)\,S = \beta\,Q.$$

By Lemma 9(i), $\|\mathbb{A}^*\Xi\|_{1\to 1} \le \bar{\xi} < 1$, so $\mathrm{Id} - \mathbb{A}^*\Xi$ is invertible with Neumann inverse (89), and the unique solution is

$$S(t,\cdot) = (\mathrm{Id} - \mathbb{A}^*\Xi)^{-1}\beta\; Q(t) = h\,Q(t),$$

pulling the scalar $Q(t)$ through by linearity. This is (189).

Step 4: normalisation. Apply $\mathrm{Id} - \mathbb{A}^*\Xi$ to h and integrate: $h - \mathbb{A}^*\Xi h = \beta$ gives

$$\int_0^1 h\mathrm{d}i - \int_0^1 (\mathbb{A}^*\Xi h)(i)\mathrm{d}i = \int_0^1 \beta\mathrm{d}i = 1, \qquad \int_0^1 (\mathbb{A}^*\Xi h)(i)\mathrm{d}i = \int_0^1 \xi(j)\,h(j)\mathrm{d}j,$$

the second equality by Tonelli and row-stochasticity, as in Remark 7; hence $\int_0^1[1-\xi]h\mathrm{d}i = 1$. (Consistently, the GDP identity of (183) then holds identically: $\int[1-\xi]S = Q\int[1-\xi]h = Q$.)

Step 5: cash flows. By (186), $\pi_i = [1-\xi(i)]S(t,i) = [1-\xi(i)]h(i)Q(t)$: neither K_i nor z_i appears.

Step 6: prices absorb the state. Solving the supply relation (184) for the price at the (state-independent) sales level,

$$P(t,i,K,z)^{1/\alpha(i)} = \frac{S(t,i)\,[\Gamma(i)\,\mathcal{P}(t,i)]^{\frac{1-\alpha(i)}{\alpha(i)}}}{K}\,e^{-z/\alpha(i)} \;\propto\; e^{-z/\alpha(i)}\,K^{-1},$$

with a deterministic (t,i) factor of proportionality. A more productive or better-capitalised node sells a proportionally larger quantity at a proportionally lower price, leaving revenue (a fixed value share of aggregate spending) exactly unchanged: relative prices, not revenues, respond to (K,z). ∎

The rest of the static block is then log-linear and solvable in closed form. We state it in the detrended units of Section 3.3, (51) ($\widehat{s} = S/\overline{K}_t$, $q = Q/\overline{K}_t$, $\kappa = K/\overline{K}_t$; prices need no detrending), in which (189) reads $\widehat{s}(i) = h(i)q$, and with $\mathbb{E}_i[\cdot]$ the expectation against the node distribution $\mathfrak{M}(t,i,\cdot)$ as in Definition 5.

Proposition 11 (The Cobb–Douglas static block in closed form). *Let $\eta = 1$, and suppose $\mathbb{E}_i[|\ln \kappa|]$ and $\mathbb{E}_i[|z|]$ are finite, uniformly in i. Then:*

(i) (Prices.) *The node price function is log-linear in the state,*

$$\ln P(i,\kappa,z) = \alpha(i)[\ln(h(i)\,q) - \ln\kappa] - z + [1-\alpha(i)][\ln\Gamma(i) + \ln\mathcal{P}(i)]. \tag{190}$$

(ii) (Fredholm system for the log-price moment.) *The moment $m_P(i) = \mathbb{E}_i[\ln P]$ is the unique bounded solution of the linear Fredholm equation*

$$\begin{aligned} &m_P(i) - [1-\alpha(i)]\,(\mathbb{A}m_P)(i) \\ &= \alpha(i)[\ln h(i) + \ln q - \mathbb{E}_i[\ln\kappa]] - \mathbb{E}_i[z] + [1-\alpha(i)]\left[\ln\Gamma(i) - \int_0^1 a(i,j)\ln a(i,j)\mathrm{d}j\right], \end{aligned} \tag{191}$$

the exact continuum analogue of the finite log-linear system (170).

(iii) (Numeraire.) *The normalisation $P^f \equiv 1$ reads*

$$\int_0^1 \beta(i)\, m_P(i)\,\mathrm{d}i = 0, \tag{192}$$

and is the $\eta \to 1$ limit of the CES numeraire (47).

(iv) (Reduced map.) *Let $\mathbb{D}_\alpha$ and $\mathbb{D}_{1-\alpha}$ denote pointwise multiplication by $\alpha(i)$ and $1-\alpha(i)$, and define*

$$\mathbb{M} := \mathbb{D}_\alpha^{-1}[\mathrm{Id} - \mathbb{D}_{1-\alpha}\mathbb{A}], \qquad \omega := (\mathbb{M}^*)^{-1}\beta = \mathbb{D}_\alpha(\mathrm{Id} - \mathbb{A}^*\mathbb{D}_{1-\alpha})^{-1}\beta,$$

the operator analogues of the finite $\mathbb{M}$ and $\boldsymbol{\omega}$ of Appendix C.2. Then $\omega \geq 0$, $\int_0^1 \omega(i)\mathrm{d}i = 1$, and eliminating m_P between (ii) and (iii) gives

$$\begin{aligned} \ln q &= const + \int_0^1 \omega(i)\,\mathbb{E}_i[\ln\kappa]\,\mathrm{d}i + \int_0^1 \frac{\omega(i)}{\alpha(i)}\,\mathbb{E}_i[z]\,\mathrm{d}i, \\ const &:= -\int_0^1 \omega(i)\left\{\ln h(i) + \frac{1-\alpha(i)}{\alpha(i)}\left[\ln\Gamma(i) - \int_0^1 a(i,j)\ln a(i,j)\mathrm{d}j\right]\right\}\mathrm{d}i, \end{aligned} \tag{193}$$

a log-linear aggregate with the Domar-type density ω: the continuum analogue of the reduced map (171), *and an* integral *over the continuum, so that no individual node's (κ, z) moves it.*

Proof. Step 1: prices. The supply relation (184) in detrended units, $\widehat{s} = \kappa[\Gamma\mathcal{P}]^{-(1-\alpha)/\alpha}(e^z P)^{1/\alpha}$, solved for $\ln P$ at $\widehat{s} = hq$ (multiply the logarithm through by α), is (190).

Step 2: the Fredholm system. Take $\mathbb{E}_i$ of (190) and substitute the moment form of the bundle price, $\ln\mathcal{P}(i) = (\mathbb{A}m_P)(i) - \int a(i,j)\ln a(i,j)\mathrm{d}j$ from (44); this is (191). The right-hand side is bounded: h is bounded above ($h = \beta + \mathbb{A}^*\Xi h$ with $(\mathbb{A}^*\Xi h)(i) \leq \bar{a}\bar{\xi}\|h\|_1$) and below by $\underline{\beta} > 0$, $\Gamma(i) \in [\frac{\varepsilon}{\varepsilon-1}, \bar{\Gamma}]$ by its definition in Appendix C.4, the entropy term is bounded by Lemma 11, and the moments are bounded by hypothesis.

The operator $\mathbb{A}$ is an average against a probability density, so $\|\mathbb{A}\|_{\infty\to\infty} \le 1$ and $\|\mathbb{D}_{1-\alpha}\mathbb{A}\|_{\infty\to\infty} \le 1-\underline{\alpha} < 1$: the Neumann series of $(\mathrm{Id} - \mathbb{D}_{1-\alpha}\mathbb{A})^{-1}$ converges in L^∞, and the bounded solution is unique; m_P is bounded a priori by the first hypothesis of Lemma 14, so it is that solution.

Step 3: the numeraire. Cost minimisation for (117) (choose Q_i to minimise $\int P_i Q_i \mathrm{d}i$ at given Q) has first-order condition $P_i Q_i = \Lambda\beta(i)$, total cost Λ, and unit cost $P^f = \exp\int_0^1 \beta(i)\ln P_i \mathrm{d}i$: the log price index of a Cobb–Douglas aggregator is the β-weighted mean log price. By the exact law of large numbers (the argument of Lemma 14 with $\psi(j,\omega) = \beta(j)\ln P(t,j,s_j)$, integrable by the first hypothesis of (105)), the realised integral equals $\int \beta\, m_P \mathrm{d}i$ a.s., so $P^f = 1$ is (192). That it is also the $\eta \to 1$ limit of (47) follows from expanding $P^{1-\eta} = 1 + (1-\eta)\ln P + O((1-\eta)^2)$ inside it (details in Appendix B.7.1).

Step 4: the reduced map. Dividing (191) by $\alpha(i)$ puts it in the form $\mathbb{M}m_P = R$ with $R(i) = \ln h(i) + \ln q - \mathbb{E}_i[\ln\kappa] - \mathbb{E}_i[z]/\alpha(i) + \frac{1-\alpha(i)}{\alpha(i)}[\ln\Gamma(i) - \int a(i,j)\ln a(i,j)\mathrm{d}j]$, so $m_P = \mathbb{M}^{-1}R$ and the numeraire is $\langle\beta, \mathbb{M}^{-1}R\rangle = \langle(\mathbb{M}^*)^{-1}\beta, R\rangle = \langle\omega, R\rangle = 0$, the pairing being that of L^1 against L^∞ (the adjoint of $\mathbb{A}$ on L^∞ is $\mathbb{A}^*$ on L^1, Definition 3). Three properties of ω: (a) the two expressions for ω agree, since $\mathbb{M}^* = [\mathrm{Id} - \mathbb{A}^*\mathbb{D}_{1-\alpha}]\mathbb{D}_\alpha^{-1}$; (b) $\omega > 0$, because $\|\mathbb{A}^*\mathbb{D}_{1-\alpha}\|_{1\to1} \le 1-\underline{\alpha} < 1$ ($\mathbb{A}^*$ conserves mass, as in the proof of Lemma 9), so the inverse is the L^1-convergent Neumann series $\sum_{n\ge0}(\mathbb{A}^*\mathbb{D}_{1-\alpha})^n\beta$, a sum of nonnegative terms whose first term gives $\omega \ge \alpha\beta \ge \underline{\alpha}\underline{\beta} > 0$; (c) $\int_0^1 \omega\mathrm{d}i = \langle 1, (\mathbb{M}^*)^{-1}\beta\rangle = \langle\mathbb{M}^{-1}1, \beta\rangle = \langle 1, \beta\rangle = 1$, because $\mathbb{M}1 = \mathbb{D}_\alpha^{-1}[1-(1-\alpha)\mathbb{A}1] = \mathbb{D}_\alpha^{-1}\alpha = 1$ by row-stochasticity ($\mathbb{A}1 = 1$). Expanding $\langle\omega, R\rangle = 0$ and using (c) to pull $\ln q$ out gives (193). ∎

[Background: Fredholm integral equations: what is needed here.]

- *The object.* A linear Fredholm equation of the *second kind* on $[0,1]$ is

$$u(i) - \int_0^1 k(i,j)\,u(j)\,\mathrm{d}j = f(i), \qquad \text{i.e.} \qquad (\mathrm{Id} - K)u = f,$$

  with a given kernel k and right-hand side f; the unknown u appears both outside and inside the integral. (An equation of the *first kind*, $Ku = f$ alone, is an ill-posed inverse problem; nothing of that kind appears here.) Equation (191) is exactly this with $K = \mathbb{D}_{1-\alpha}\mathbb{A}$, i.e. $k(i,j) = [1-\alpha(i)]\,a(i,j)$, and (193) is its adjoint form with $\mathbb{A}^*\mathbb{D}_{1-\alpha}$.

- *Neumann series.* If $\|K\| < 1$ in some operator norm, then $(\mathrm{Id} - K)^{-1} = \sum_{n\ge0} K^n$ converges in that norm and $u = \sum_n K^n f$ is the unique solution in the corresponding space. This is the only tool the proof of Proposition 11 uses: $\mathbb{A}$ averages against a probability density, so $\|\mathbb{A}\|_{\infty\to\infty} \le 1$ and $\|\mathbb{D}_{1-\alpha}\mathbb{A}\|_{\infty\to\infty} \le 1-\underline{\alpha} < 1$. Economically the series is the round-by-round propagation of costs upstream, the continuum Leontief inverse: the same argument as for $(\mathrm{Id} - \mathbb{A}^*\Xi)^{-1}$ in Lemma 9, and on a step kernel it is literally the $J \times J$ linear system $[\mathbf{I} - \mathbf{D}_{1-\alpha}\mathbf{A}]\mathbf{m} = \mathbf{r}$ of the finite model.

- *Why "unique bounded solution".* Uniqueness is relative to a function space. On L^∞ the contraction gives a unique bounded solution; in a larger space one could add unbounded solutions of the homogeneous equation $u = Ku$, which is why the proposition specifies boundedness, and why m_P is shown to be bounded a priori before it is identified with the Neumann-series solution.
- *Fredholm theory proper* (background only, not used). For a compact operator K, which any bounded kernel on $[0,1]$ defines on L^p, the Fredholm alternative says: either $(\mathrm{Id} - K)u = f$ has a unique solution for every f, or the homogeneous equation has finitely many independent solutions and f must be orthogonal to the corresponding solutions of the adjoint equation. The contraction case above is the first branch, with the sharper information that the solution is a convergent series. The alternative is the infinite-dimensional version of "a square matrix is invertible if and only if its kernel is trivial"; the finite-dimensional intuition is the right one here.
- *Adjoint pairing.* The two forms of the numeraire in Step 4, $\langle \beta, \mathbb{M}^{-1} R\rangle = \langle (\mathbb{M}^*)^{-1}\beta, R\rangle$, are the adjoint identity $\langle g, (\mathrm{Id} - K)^{-1} f\rangle = \langle (\mathrm{Id} - K^*)^{-1} g, f\rangle$ with K^* the transposed kernel $k(j,i)$: costs propagate upstream through K, influence weights propagate downstream through K^*, and pairing with β turns the one into the other.
- (Standard material, e.g. Kress, *Linear Integral Equations*; not re-derived.)

Derivation in Appendix B.7.1.

On the balanced growth path the z-marginal of the node distribution is the stationary OU distribution, so $\mathbb{E}_i[z] = 0$ and the z-terms of (191) and (193) drop; this includes $\varsigma^Z(i) = 0$, where $z_i(t) = z_i(0)e^{-\theta(i)t}$ decays deterministically to zero. They are kept because (191) also serves off the BGP, where $\mathbb{E}_i[z] = z_i(0)e^{-\theta(i)t} \neq 0$.

Corollary 2 (Collapse of the moment coupling). *Let $\eta = 1$ and consider the stationary system of Definition 6 with the consistency conditions* (198)*. Then:*

(i) (Sales and inbound demand.) *The sales moment and the inbound demand are pinned by primitives and the scalar q, $\widehat{m_S} = h\,q$ and $\widehat{\mathcal{D}} = (h - \beta)\,q$, and the GDP line of* (198) *holds identically.*

(ii) (Price moment.) *The graphon consistency of item 4 reduces to the Fredholm system* (191) *for m_P, closed by the numeraire* (192)*; neither m_P nor $\mathcal{P}$ enters any node's payoff.*

(iii) (What the ensemble pins.) *The ensemble of node laws enters the equilibrium conditions only through the scalar moments $\mathbb{E}_i[\ln \kappa]$ and $\mathbb{E}_i[z]$ in* (191) *(equivalently* (193)*) and $\mathbb{E}_i[\kappa]$ in the adding-up constraint* (138)*; the moment $\mathbb{E}_i[\iota^*\kappa]$ appears only in the feasibility check $\widehat{C} > 0$, which imposes no equation.*

Proof. $\widehat{s} = hq$ is constant in the state, so $\widehat{m_S} = \mathbb{E}_i[\widehat{s}] = hq$ needs no integration; $\widehat{\mathcal{D}} = \mathbb{A}^*\Xi\widehat{m_S} = q\,\mathbb{A}^*\Xi h = q\,(h - \beta)$ by $h = \beta + \mathbb{A}^*\Xi h$; and $\int [1 - \xi]\widehat{m_S}\mathrm{d}i = q\int [1 - \xi] h \mathrm{d}i = q$ by the normalisation in Lemma 26. The price moment m_P and the bundle price $\mathcal{P}$ appear in (190) but not in $\widehat{s}$, hence neither in the cash flow (186) nor in the marginal profit (188). ∎

Proposition 12 (The degenerate Cobb–Douglas balanced growth path). *Let $\eta = 1$. Then the stationary equilibrium of Definition 6 (with* (196), (197) *and* (198)*) is algebraic node by node: the node distribution is $\delta_{\kappa(i)}(\mathrm{d}\kappa)\,\mathfrak{m}^z(i,z)\mathrm{d}z$, and $(\kappa(i), \mathfrak{q}(i), \iota(i))$ solve*

$$\Phi_i(\iota(i)) - \delta(i) = g, \qquad \Phi_i'(\iota(i))\,\mathfrak{q}(i) = 1, \qquad [\rho + (\gamma-1)g]\,\mathfrak{q}(i) = \frac{\varepsilon-1}{\varepsilon}\,\alpha(i)\,\frac{h(i)\,q}{\kappa(i)} - \iota(i), \tag{194}$$

together with $\int_0^1 \kappa(i)\mathrm{d}i = 1$, the Fredholm–numeraire block (191)–(192) *(equivalently* (193) *with $\mathbb{E}_i[\ln\kappa] = \ln\kappa(i)$) for (q, g), and $r = \rho + \gamma g$. Given g, the first two equations deliver $(\iota(i), \mathfrak{q}(i))$ in closed form and the third delivers*

$$\kappa(i) = \frac{\frac{\varepsilon-1}{\varepsilon}\,\alpha(i)\,h(i)\,q}{[\rho + (\gamma-1)g]\,\mathfrak{q}(i) + \iota(i)},$$

so that adding-up gives $q(g)$ in closed form and the numeraire is one equation in g (the recipe needs $1 + g + \delta(i) > 0$ for the first equation and $[\rho + (\gamma-1)g]\mathfrak{q}(i) + \iota(i) > 0$ for the third, both automatic at $g \geq 0$). If the accumulation technology (ϕ, δ) is common across nodes, $\mathfrak{q}$ and ι are common and $\kappa(i) \propto \alpha(i)h(i)$.

Proof. By Theorem 1, item 5, $\kappa_i(t)$ is deterministic, so on the BGP it is a rest point $\kappa(i)$ of the detrended capital drift (142): $\Phi_i(\iota(i)) - \delta(i) = g$, the first equation. The second is the investment FOC (133). The third is the stationary detrended marginal-$\mathfrak{q}$ equation (197) with $\widehat{s} = hq$ and $\mathcal{L}_z^{(i)}\mathfrak{q} = 0$, evaluated at the rest point, where the advection term $\mathfrak{q}_\kappa\kappa[\Phi_i - \delta - g]$ vanishes whatever $\mathfrak{q}_\kappa$ is, so that $r - [\Phi_i - \delta] = \rho + \gamma g - g$. The node distribution is the product of Theorem 1, item 3, with $\mu^K = \delta_{\kappa(i)}$; adding-up (138) reads $\int \kappa(i)\mathrm{d}i = 1$; the numeraire is (192) with $\mathbb{E}_i[\ln\kappa] = \ln\kappa(i)$ in (191); and the growth identity (55) holds identically at the rest points. The closed forms follow by reading the three equations in order, $\kappa(i) \propto q$ giving $q(g) = 1/\int_0^1 [\kappa(i)/q]\mathrm{d}i$ from adding-up. Under a common technology the first two equations do not depend on i. ∎

There is nothing mean-field about this object: a continuum of deterministic $\mathfrak{q}$-theory rest points (off the BGP, ODE saddle paths), coupled by two scalars. All distributional content, all risk, and all use of the FPK equation have vanished. This is the precise sense in which the Cobb–Douglas GMFG is degenerate, and why the CES repair is needed before the graphon limit is a well-posed mean-field problem. Two readings of the common-technology rest point $\kappa(i) \propto \alpha(i)h(i)$ are useful. It is the firm-side, private-rental counterpart of the finite household model's $\kappa_i^* = \omega_i$: since $\omega = \mathbb{D}_\alpha(\mathrm{Id} - \mathbb{A}^*\mathbb{D}_{1-\alpha})^{-1}\beta$ is α times the cost-based influence density while h is the sales-based one, the two differ by the sales-versus-cost factor of the finite markup wedge (179). And Proposition 12 provides the grid centres and the warm start of Appendix D. Under capacity the common-technology proportionality is lost, because the rest-point threshold of Proposition 5 depends on $x(i)/\mathcal{P}(i)$ and on G_i even when (ϕ, δ) are common, and the sales-based density h is replaced by the fixed point (124).

What a solver reports at $\eta = 1$ in this model is described in Remark 8.

Own-state relevance. The form of Proposition 6 without capacity constraints follows from (131) with $e_S \equiv 1/\alpha(i)$ (the supply elasticity of (184)); it is stated in full because its proof is the two-line power-form argument.

Proposition 13 (Own-state relevance under CES). *Let $\eta \neq 1$ and $\beta(i) > 0$. Define the final-demand share of sales $\vartheta(t,i,K,z) := \beta(i)P_i^{1-\eta}Q/S_i \in (0,1]$. Then the equilibrium node sales function satisfies, pointwise,*

$$\frac{\partial \ln S}{\partial \ln K} = \frac{-(1-\eta)\,\vartheta\,\alpha(i)}{1-\alpha(i)(1-\eta)\vartheta}, \qquad \frac{\partial \ln S}{\partial z} = \frac{-(1-\eta)\,\vartheta}{1-\alpha(i)(1-\eta)\vartheta}, \tag{195}$$

with $1-\alpha(1-\eta)\vartheta > 0$ always. Both derivatives are nonzero whenever $\eta \neq 1$ and $\vartheta > 0$: strictly positive if $\eta > 1$ (elastic final demand: higher capital or productivity lowers the relative price by less than it raises quantity sold) and strictly negative if $\eta < 1$. In either case cash flows (186) *and marginal profits* (188) *load on (K_i, z_i), and the node problem is a genuine stochastic control problem in (κ, z) (one-dimensional in κ, at $z = 0$ on the BGP, if $\varsigma^Z(i) = 0$).*

Proof. Total-differentiate the node block at the solution of Lemma 24, holding aggregates fixed. From (184), $\mathrm{d}\ln S = \mathrm{d}\ln K + \alpha^{-1}\mathrm{d}z + \alpha^{-1}\mathrm{d}\ln P$; from (109), $\mathrm{d}S = (1-\eta)\beta P^{1-\eta}Q\,\mathrm{d}\ln P$, i.e. $\mathrm{d}\ln S = (1-\eta)\vartheta\,\mathrm{d}\ln P$. Eliminating $\mathrm{d}\ln P$:

$$\frac{\partial \ln P}{\partial \ln K} = \frac{-\alpha}{1-\alpha(1-\eta)\vartheta}, \qquad \frac{\partial \ln P}{\partial z} = \frac{-1}{1-\alpha(1-\eta)\vartheta},$$

and (195) follows. Positivity of the denominator: if $\eta \geq 1$ it is ≥ 1; if $\eta < 1$ then $\alpha(1-\eta)\vartheta < 1$ since each factor is in $(0,1]$ with $\alpha < 1$. ∎

C.6 The GMFG System and Its Step-Kernel Form

The dynamic layer of Sections 3.3–3.3 holds verbatim without capacity constraints: the household block and the deterministic SDF (Lemma 18), the firm problem (132)–(133), the generator of the node state, the marginal-q equation (134) and its Feynman–Kac derivation, the closed-loop dynamics (136) and the FPK equation (137), the detrending (51)–(138), the detrended state dynamics (142)–(55), the detrended FPK equation (53), the household Euler equation (145), items (ii)–(vi) of Proposition 8, Assumption 6, and items 2, 3, 5, 6 and 7 of Definition 6, none of which involves the capacity block. What changes is the static block inside them and the flow that prices capital: by Lemma 4, as $x(i) \downarrow 0$ the threshold diverges, firm revenue and sales coincide ($\varrho_i \to 1$, no rationing rent), the spending share is the constant $\xi(i)$, and the marginal profit (115) tends to $\frac{\varepsilon-1}{\varepsilon}\alpha(i)S/K$ ((188)). The household budget then carries no rent, and the on-path value equation(135) and the detrended value equation (146) have the flow $[1-\xi(i)]S - \iota^*K$ and $[1-\xi(i)]\widehat{s} - \iota^*\kappa$ respectively. Lemma 19 holds with the statement for this model (prices invariant, $(S, \mathcal{D}, Q, X, \pi)$ of degree one, no thresholds to speak of), by the same proof with Step 1(b) replaced by the linearity of (184) in K at fixed prices. This subsection records the forms without capacity constraints of the objects that the numerical appendix (Appendix D) maps onto at the end of each of its steps.

The detrended static block and $\mathfrak{q}$-equation. Without capacity constraints the per-node static block (184)–(185), divided by $\overline{K}_t$, is the pair

$$\widehat{s} = \kappa \left[\Gamma(i)\,\mathcal{P}(t,i)\right]^{-\frac{1-\alpha(i)}{\alpha(i)}} (e^z P)^{1/\alpha(i)}, \qquad \widehat{s} = \beta(i)\, P^{1-\eta}\, q_t + \widehat{\mathcal{D}}(t,i), \tag{196}$$

with $P = P(t,i,\kappa,z)$, $\widehat{s} = \widehat{s}(t,i,\kappa,z)$, and Lemma 24 applying unchanged; it is the capacity limit of the triple (139)–(141), obtained as in Appendix C.4 for the undetrended block. The detrended marginal-$\mathfrak{q}$ equation (52) becomes, with the limit flow,

$$\boxed{r_t\,\mathfrak{q} = \frac{\varepsilon-1}{\varepsilon}\,\alpha(i)\,\frac{\widehat{s}(t,i,\kappa,z)}{\kappa} - \iota^*(\mathfrak{q}) + \mathfrak{q}[\Phi_i(\iota^*) - \delta(i)] + \mathfrak{q}_\kappa\,\kappa[\Phi_i(\iota^*) - \delta(i) - g_t] + \mathcal{L}_z^{(i)}\mathfrak{q} + \partial_t\mathfrak{q},} \tag{197}$$

with $\iota^*(\mathfrak{q})$ from (133): the marginal dividend is the constant fraction $\frac{\varepsilon-1}{\varepsilon}\alpha(i)$ of detrended sales per unit of capital density, and the node state enters the equation through $\widehat{s}(t,i,\kappa,z)$ alone.

The consistency conditions. With $X = \xi(i)S$ the spending moment is $\widehat{m_X} = \xi\widehat{m_S}$ and item 4 of Definition 6 reduces to two moment functions: with $\mathbb{E}_i[\cdot] := \int \cdot\, \mathfrak{m}(i,\kappa,z)\mathrm{d}\kappa\mathrm{d}z$,

$$\begin{aligned} m_P(i) &= \mathbb{E}_i[\ln P(i,\kappa,z)], & \widehat{m_S}(i) &= \mathbb{E}_i[\widehat{s}(i,\kappa,z)], \\ \ln\mathcal{P}(i) &= (\mathbb{A}m_P)(i) - \textstyle\int a(i,j)\ln a(i,j)\mathrm{d}j, & \widehat{\mathcal{D}}(i) &= (\mathbb{A}^*\Xi\widehat{m_S})(i), \\ q &= \textstyle\int_0^1 [1-\xi(j)]\widehat{m_S}(j)\mathrm{d}j. \end{aligned} \tag{198}$$

The stationary equilibrium without capacity constraints is Definition 6 with (196) in item 1, (197) in item 2 and (198) in item 4. Its coupling runs through exactly two kernel-weighted moment functions $(\mathbb{A}m_P, \mathbb{A}^*\Xi\widehat{m_S})$ and one scalar, delivered jointly by the exact LLN and by the constancy of expenditure shares within nodes; the capacity model keeps the count and replaces the second function by the spending moment (Definition 6).

The general-kernel pair. For a general kernel the stationary system is, for a.e. node $i \in [0,1]$, the static block (196) given $(\mathcal{P}(i), \widehat{\mathcal{D}}(i), q)$, the backward equation (the stationary (197)),

$$(\rho+\gamma g)\,\mathfrak{q} = \frac{\varepsilon-1}{\varepsilon}\,\alpha(i)\,\frac{\widehat{s}(i,\kappa,z)}{\kappa} - \iota^*(\mathfrak{q}) + \mathfrak{q}[\Phi_i(\iota^*) - \delta(i)] + \mathfrak{q}_\kappa\,\kappa[\Phi_i(\iota^*) - \delta(i) - g] + \mathcal{L}_z^{(i)}\mathfrak{q}, \tag{199}$$

and the forward equation (53), coupled across nodes only through the consistency map (198), $\ln\mathcal{P} = \mathbb{A}m_P + H$ and $\widehat{\mathcal{D}} = \mathbb{A}^*\Xi\widehat{m_S}$ with $(m_P, \widehat{m_S})$ the node moments, and closed by the numeraire, capital adding-up, $g = \int_0^1 \mathbb{E}_i[(\Phi_i(\iota^*) - \delta(i))\kappa]\mathrm{d}i$ and $r = \rho + \gamma g$. The general kernel differs from the step kernel below only in the index set ($[0,1]$ versus N blocks), in the operators ($\mathbb{A}, \mathbb{A}^*$ versus the block matrix A and its transpose) and in the unknown coupling objects (functions on $[0,1]$ versus N-vectors); the per-node pair is identical.

The step-kernel system. Take a step kernel with blocks $I_1, \dots, I_N$ of lengths Δ_n, block input–output matrix and entropy constants

$$A_{nm} := \int_{I_m} a(i,j)\,\mathrm{d}j \quad (i \in I_n),$$
$$\sum_{m=1}^{N} A_{nm} = 1,$$
$$e_n := \int_0^1 a(i,j)\ln a(i,j)\,\mathrm{d}j = \sum_{m=1}^{N} A_{nm}\ln\frac{A_{nm}}{\Delta_m},$$

and block-constant primitives $(\alpha_n, \xi_n, \zeta_n, \beta_n, \phi_n, \delta_n, \theta_n, \varsigma_n^Z)$, with β_n the density value on I_n ($\sum_n \beta_n\Delta_n = 1$) and Γ_n the constant wedge of (180). The stationary GMFG without capacity constraints is then *exactly* the following finite system. For each block $n = 1, \dots, N$, given the coupling objects $\left(\mathcal{P}_n, \widehat{\mathcal{D}}_n, q, g\right)$ and $r = \rho + \gamma g$: the per-block statics, pointwise in (κ, z),

$$\widehat{s}_n = \kappa\,[\Gamma_n\,\mathcal{P}_n]^{-\frac{1-\alpha_n}{\alpha_n}}\,(e^z P_n)^{1/\alpha_n}, \qquad \widehat{s}_n = \beta_n\,P_n^{1-\eta}\,q + \widehat{\mathcal{D}}_n; \tag{200}$$

the backward equation (marginal $\mathfrak{q}$) with its FOC, whose flow coefficient $\frac{\varepsilon-1}{\varepsilon}\alpha_n$ is the same private-rental coefficient as in (176),

$$r\,\mathfrak{q}_n = \frac{\varepsilon-1}{\varepsilon}\,\alpha_n\,\frac{\widehat{s}_n(\kappa,z)}{\kappa} - \iota^* + \mathfrak{q}_n[\Phi_n(\iota^*) - \delta_n] + \mathfrak{q}_{n,\kappa}\,\kappa[\Phi_n(\iota^*) - \delta_n - g] + \mathcal{L}_z^{(n)}\mathfrak{q}_n,$$
$$\Phi_n'(\iota^*)\,\mathfrak{q}_n = 1; \tag{201}$$

and the forward equation

$$0 = -\frac{\partial}{\partial\kappa}\{\kappa[\Phi_n(\iota^*) - \delta_n - g]\,\mathfrak{m}_n\} + \frac{\partial}{\partial z}\{\theta_n\,z\,\mathfrak{m}_n\} + \frac{(\varsigma_n^Z)^2}{2}\frac{\partial^2}{\partial z^2}\,\mathfrak{m}_n, \qquad \int \mathfrak{m}_n(\kappa,z)\,\mathrm{d}\kappa\,\mathrm{d}z = 1. \tag{202}$$

Given $\left(\mathcal{P}, \widehat{\mathcal{D}}, q, g\right)$, the N pairs (201)–(202) are mutually decoupled two-dimensional problems. The input–output network re-enters *only* through the consistency map. With block moments

$$\mathbb{E}_n[\cdot] := \int \cdot\;\mathfrak{m}_n(\kappa,z)\,\mathrm{d}\kappa\,\mathrm{d}z, \qquad m_{P,n} := \mathbb{E}_n[\ln P_n], \qquad \widehat{m}_{S,n} := \mathbb{E}_n[\widehat{s}_n], \qquad \mathcal{N}_n := \mathbb{E}_n\left[P_n^{1-\eta}\right],$$

the fixed point closes via

$$\ln\mathcal{P}_n = \sum_{m=1}^{N} A_{nm}\,m_{P,m} - e_n, \qquad \text{(rows of } A\text{: what } n \text{ buys; this is } \mathbb{A}\text{)}$$
$$\widehat{\mathcal{D}}_n = \frac{1}{\Delta_n}\sum_{m=1}^{N} A_{mn}\,\xi_m\,\widehat{m}_{S,m}\,\Delta_m, \qquad \text{(columns of } A\text{: what is sold to } n\text{; this is } \mathbb{A}^*\Xi\text{)}$$
$$q = \sum_{m=1}^{N}[1-\xi_m]\,\widehat{m}_{S,m}\,\Delta_m, \qquad \sum_{n=1}^{N}\beta_n\,\mathcal{N}_n\,\Delta_n = 1 \quad \text{(numeraire)},$$
$$g = \sum_{n=1}^{N}\mathbb{E}_n[(\Phi_n(\iota^*) - \delta_n)\,\kappa]\,\Delta_n, \qquad \sum_{n=1}^{N}\mathbb{E}_n[\kappa]\,\Delta_n = 1, \qquad r = \rho + \gamma g. \tag{203}$$

The loop is: guess $\left(\mathcal{P}, \widehat{\mathcal{D}}, q, g\right) \to$ solve N independent 2D pairs $\to$ integrate the moments $\to$ regenerate the coupling objects via (203) $\to$ iterate; the two structural contrasts with the finite systemdrawn in Section B.10 apply word for word.

What the numerical appendix solves. Appendix D discretises and solves the *capacity* model's step-kernel system(156)–(159), and obtains the system without capacity constraints (200)–(203) from it by a mapping stated at the end of each step it changes (Steps 2, 5 and 6). The two systems differ in exactly three places, and the mapping is the $x_n \downarrow 0$ limit of Lemma 4 applied to each: the scalar threshold root inside the static step becomes the power-form root in the price of Lemma 24; the capital-income flow $\zeta_n[1-\xi_n]\varrho_n\widehat{s}_n/\kappa$ in the backward equation becomes $\frac{\varepsilon-1}{\varepsilon}\alpha_n\widehat{s}_n/\kappa$ (188); and the spending moment $\widehat{m}_{X,n}$ in the coupling becomes the product $\xi_n\widehat{m}_{S,n}$, so that the three regenerated block moments collapse to the two $(m_{P,n}, \widehat{m}_{S,n})$ of (203). The Cobb–Douglas closed formsof Appendix C.5 and the balanced growth pathof Proposition 12 are the initialiser and warm start of *both* solves, because they are the $x_n \downarrow 0$ limit of the capacity model as well as this model's own solution, and the pair (196)–(197) is this model's general-kernel form.

D Numerical Solution of the Graphon Mean-Field Game

[Scope of this appendix.] This appendix is written for the capacity model in the case the paper uses, $\varsigma^K \equiv 0$: the only idiosyncratic risk is the firm's productivity z, the capital coordinate is pure transport, and the backward and forward operators of a node coincide. The reference implementation also solves the general case with node-level capital shocks $\varsigma_n^K > 0$, in which the capital coordinate is a diffusion, the backward and forward operators differ by a first-order term, and the stationary distribution has power-law tails that the grid must cover; that case is documented in the general-risk version of this paper, is not treated here, and nothing below depends on it. The system discretised and solved is the step-kernel system (156)–(159) of Section B.10, with the detrended triple (139)–(141) and the pair (52)–(53) as its general-kernel form; the model without capacity constraints of Appendix C is obtained by the mapping stated at the end of Steps 2, 5 and 6, and the cross-references inside those paragraphs point at the `nc-` statements of Appendix C. Four things separate the two procedures, and all four sit in Steps 2 and 5: the per-node static solver of Step 2 is a scalar root in the *threshold* rather than in the price ((156), Lemma 15); the flow array of Step 2 is the capacity flow (115), $\zeta_n[1-\xi_n]\varrho_n\widehat{s}_n/\kappa$, in place of $\frac{\varepsilon-1}{\varepsilon}\alpha_n\widehat{s}_n/\kappa$; the coupling of Step 5 regenerates *three* block moments $(m_{P,n}, \widehat{m}_{X,n}, \widehat{m}_{S,n})$ instead of two, the spending moment $\widehat{m}_{X,n}$ of (159) replacing the product $\xi_n\widehat{m}_{S,n}$, which is exact only without capacity constraints (these three are the three places in which the two systems differ, Appendix C.6); and the initialiser of Step 5 is the algebraic path of the model without capacity constraints, which under capacity is the $x_n \to 0$ limit rather than the model's own Cobb–Douglas path (Proposition 5, no longer algebraic). Steps 0, 1, 3 and 4 are common to the two models;

Step 6 differs only through the two objects Step 2 hands it, the static block it calls at each date and the supply elasticity in its numeraire step; and Step 7 gains the capacity checks. Notation carried over from the main text: the state generator is written $\mathbb{G}$ (its discretisation stays G); the capacity block's tail functions, threshold map Y_n, utilization U_n and shares $\xi_n, \zeta_n, \varrho_n$ are the block-constant objects of Section B.10, evaluated at the node's own threshold. Notation deliberately left: the moment fixed point keeps the letter Θ (the capacity constant $\mu_{n,1/\varepsilon}^{\varepsilon/(\varepsilon-1)}$ took no symbol); the low-rank kernel factors $\varphi_r(i)$ of Step 0 are unrelated to the capacity multiplier φ_{in}; the CES elasticity η is unrelated to the lower-tail share $\eta_{n,\theta}(\overline{v})$, which always carries its subscripts and argument; the marginal-profit function $\Phi_n(\overline{v})$ of (84) and the installation function $\Phi_n(\iota)$ are told apart by their argument, as in the main text. The letter x is not used for a coordinate anywhere here: the log-capital coordinate of Steps 1, 3 and 4 is $\varkappa := \ln \kappa$, leaving x_n to the capital intensity, which the capacity model needs in every static display of Step 2.

The object computed is the step-kernel systemof Section B.10, (156)–(159): given the coupling objects $(\mathcal{P}_n, \widehat{\mathcal{D}}_n, q, g)$, N decoupled problems in (κ, z) (statics, marginal-$\mathfrak{q}$ equation, FPK equation) are solved in parallel, their moments are integrated against the block distributions, the coupling objects are regenerated through the block forms of $\mathbb{A}$ and $\mathbb{A}^*$, and the loop is iterated to a fixed point, with the two scalar closures (capital adding-up and the numeraire) assigned to the inner and outer loops as explained in Step 5. The static within-node algebra is that of the finite model and is not repeated; every step where the graphon appears and every step where the capacity block replaces the power form of the model without capacity constraints is spelled out. Throughout, "node loop" means an embarrassingly parallel loop over the blocks $n = 1, \ldots, N$, and the block index is suppressed where no confusion arises.

Conventions. Four conventions hold throughout the appendix.

- *Model.* Every step is written for the capacity model. Steps 2, 5 and 6 end with a paragraph "The model without capacity constraints" giving the mapping that turns them into the procedure for that model; the other steps are common to the two models. The mapping is always the $x_n \to 0$ limit of Lemma 4 evaluated in advance: the threshold is sent to infinity, so $\mathrm{Y}_n(\overline{v})$ becomes the power $\overline{v}^{\alpha_n/(1-\alpha_n)}/\Gamma_n$ and the state-dependent objects $\xi_n(\overline{v}), \zeta_n(\overline{v})[1-\xi_n(\overline{v})], \varrho_n(\overline{v}), e_S(\overline{v})$ collapse to the constants ξ_n, $\frac{\varepsilon-1}{\varepsilon}\alpha_n$, 1 and $1/\alpha_n$. Runs without capacity constraints never evaluate a tail function.
- *Shocks.* Every display is written for the case the paper uses, $\varsigma_n^K \equiv 0$ and $\varsigma_n^Z > 0$ (Assumption 4): the node's capital follows the deterministic drift (142) and its productivity the Ornstein–Uhlenbeck process, so each node problem is two-dimensional in (κ, z), first order in κ and second order in z.
- *Implementation.* The algorithm is fixed by Definition 6 and Proposition 8. Every procedure below is implemented in the reference implementation and validated by the checksof Step 7, under both demand systems, except where an item is marked *proposed*. Practical settings (grid sizes, step sizes, tolerances) and meas-

ured magnitudes are not part of the algorithm and are confined to the private notes.

- *Operators.* Sans-serif letters denote *discretised* operators (matrices), blackboard-bold letters the continuum operators of Definition 3 ($\mathbb{A}$, $\mathbb{A}^*$, $\mathbb{D}_\alpha$) and the state generator $\mathbb{G}$ of Section B.8. Discrete operators act on vectors indexed by the (κ, z) grid of a node, in column-major order with κ varying fastest. A denotes the discretised operator of the backward equation(157), G the discretised state generator, L_z the discretised $\mathcal{L}_z^{(n)}$, and I the identity. Without capital shocks the first two coincide, $\mathsf{A} = \mathsf{G}$; both letters are kept because the two operators play different roles, and the implementation builds them separately (Step 4).

D.1 Step 0: Discretising the Graphon

Choose N nodes $i_1 < \cdots < i_N$ with quadrature weights $\Delta_n > 0$, $\sum_n \Delta_n = 1$ (uniform $\Delta_n = 1/N$ unless the kernel warrants adaptivity). The kernel and densities are represented by

$$A_{nm} := a(i_n, i_m), \qquad \beta_n := \beta(i_n), \qquad \text{row check: } \sum_m A_{nm}\Delta_m = 1. \tag{204}$$

The graphon operators become matrix–vector products with the weights attached:

$$(\mathbb{A}f)_n \approx \sum_m A_{nm}\, f_m\, \Delta_m, \qquad (\mathbb{A}^* g)_n \approx \sum_m A_{mn}\, g_m\, \Delta_m, \tag{205}$$

i.e. $\mathbb{A} \leftrightarrow A\,\mathrm{diag}(\Delta)$ and $\mathbb{A}^* \leftrightarrow A^{\mathsf{T}}\mathrm{diag}(\Delta)$. (Convention: here A_{nm} denotes kernel *values*, with $\sum_m A_{nm}\Delta_m = 1$, whereas the step-kernel systemof Section B.10 uses the integrated block weights $\int_{I_m} a(i,j)\mathrm{d}j = A_{nm}\Delta_m$, which sum to one; the two are interchangeable and each is internally consistent.) The entropy constants $e_n := \int a(i_n, j)\ln a(i_n, j)\mathrm{d}j$of Section B.10 are precomputed once. If the data are J observed sectors, take $N = J$ and $A_{nm} = J\, a^{\text{data}}_{nm}$, $\Delta_n = 1/J$: the step-kernel economy is then *exactly* the data economy, and the continuum method is exact, not an approximation (Section 3.1). For a smooth estimated kernel, refine N until aggregates stabilise. The step carries no approximation. Nothing here depends on the model: the block constants ξ_n, Γ_n and the Domar weights h_n of the model without capacity constraints(Corollary 2) are precomputed alongside the kernel in both models, because they are the inputs of the initialiser of Step 5 whether or not capacity binds.

[Warning: the $-\ln J$ shift makes *levels* incomparable across J]

- At a *fixed* J the shift is pure bookkeeping: e_n enters the bundle price only as the constant $\mathcal{P}_n = \exp((\mathbb{A}m_P)_n - e_n)$, and the numeraire absorbs it. Nothing in a calibration to one IO table is affected.
- It does bite when two economies with *different* J are put side by side: they sit on different price scales, so g, welfare and price levels are not comparable across J without an explicit normalisation. Only unit-free objects are.

- Verified numerically: block-replicating the two-sector reference economy into $J/2$ identical copies — which ought to leave the economy unchanged — moves g by several percentage points. Since $A_{nm} = a_{nm}/\Delta_m$, halving Δ doubles the kernel *density*, e_n rises by $\ln 2$, bundle prices fall and measured productivity rises.
- *Replication is not refinement.* Replicating blocks concentrates a onto ever-thinner diagonal blocks, so a^J approaches a Dirac-like kernel rather than a fixed a: it violates the section-wise convergence of Appendix B.2. A genuine refinement test holds a fixed and sets $a^J_{nm} = \Delta_m^{-1} \iint_{I_n \times I_m} a$, and then e_n converges instead of drifting.
- Practical fix when a cross-J comparison is really wanted: pin a common sector-TFP level $\bar{Z}$ so that the Cobb–Douglas balanced-growth root is the same at every J. Without capacity constraints this is cheap: the Fredholm right-hand side of (191) is affine in $\ln \bar{Z}$ with coefficient $-1/\alpha$, so the required $\ln \bar{Z}$ solves a scalar *linear* equation. Under capacity the same scalar condition has no closed form (Proposition 4(iv)) and is solved numerically, one static fixed point per candidate $\bar{Z}$.
- In the code the entropy constant is stored with this sign, $e_n = +\int a \ln a$, i.e. as *minus* the row entropy H of Section 3; hence $\mathcal{P} = \exp(\mathbb{A} m_P - e)$ there.

Remark 13 (Complexity) The cost is $N \times$ (one PDE pair in (κ, z)) plus $O(N^2)$ for the moment products (205): *linear* in the number of sectors for the PDE part, versus exponential ($(2J-1)$-dimensional) for the finite-model HJB. If the kernel is low rank, $a(i,j) = \sum_{r=1}^{R} \varphi_r(i)\psi_r(j)$, the coupling collapses further to $2R+1$ scalars per iteration (R inner products for $\mathbb{A} m_P$, R for $\mathbb{A}^* \widehat{m_X}$, plus q; without capacity constraints $\widehat{m_X} = \Xi \widehat{m_S}$ and the count is the same), making even very large N cheap: the third block moment the capacity coupling regenerates, $\widehat{m_S}$, enters only through the scalar q and costs no kernel product.

D.2 Step 1: State Coordinates, Grids and Boundary Conditions

The capital coordinate. Write $\mu^\kappa(\kappa, z) = \Phi_n(\iota^*(\mathfrak{q}(\kappa, z))) - \delta_n - g$ for the geometric drift of κ in (142). Without capital shocks the κ-parts of the state generator and of the backward operator, read off (158) and (157), coincide and are the first-order transport operator

$$\mathbb{G}_\kappa f = \mathcal{A}_\kappa f = \mu^\kappa \, \kappa \, f_\kappa. \tag{206}$$

[Q21 | here | ANSWERED] Notation changed in the main text: x_i, $x(i)$ and x_n now denote the capital intensity of the capacity technology (Section 2.1, Assumption 4), while this appendix uses $x := \ln \kappa$ for the log-capital coordinate, here and in Steps 3–4 ("in the coordinate $x = \ln \kappa$"). Check and fix this passage: either rename the coordinate (for instance $\ell := \ln \kappa$, with F_ℓ, $\mathbb{G}_\ell$, $\mathcal{A}_\ell$ and the label of the display below), or keep it with a one-line disclaimer here. *Answered by the rename:* the log-capital coordinate is now $\varkappa := \ln \kappa$ throughout Steps 1, 3 and 4, with $F_\varkappa$, $\mathbb{G}_\varkappa$, $\mathcal{A}_\varkappa$, $\Delta\varkappa$ and the display relabelled (207); x_n is the capital intensity everywhere in this appendix. The capacity

model needs the letter: x_n appears in every static display of Step 2. In the logarithmic coordinate $\varkappa := \ln \kappa$ (the letter x_n is the capital intensity of the capacity technology throughout), with $F(\varkappa) = f(e^{\varkappa})$ and $\kappa f_\kappa = F_\varkappa$,

$$\mathbb{G}_\varkappa F = \mathcal{A}_\varkappa F = \mu^\kappa F_\varkappa, \tag{207}$$

with a drift that is the drift of $\ln \kappa$ itself (Section 3.3) and the density transforming as $\mathfrak{m}^\varkappa(\varkappa) = e^\varkappa \mathfrak{m}(e^\varkappa)$. The two formulations are equivalent; either may be discretised. In both, the natural grid is uniform in $\varkappa$ (log-spaced in κ): the scale family of Remark 14 is then a translation of the grid. The grid is centred on the Cobb–Douglas rest point κ_n of the path without capacity constraints(194) (Proposition 12), which is the $x_n \to 0$ limit of the capacity rest point(127) (Proposition 5) and a valid centre in both models. Under capacity the converged node distributions sit away from that centre — capacity moves the rest point, and the shift grows with x_n — so the grid may be recentred on the converged $\mathbb{E}_n[\ln \kappa]$ and re-solved, a manual step taken only when the edge mass of (a) below warrants it, never inside a solve. The computation discretises in the coordinate κ on the log-spaced grid, where the stencil below reproduces the identity function exactly and the stationarity diagnostic of Step 4 is exact.

Truncation and boundary conditions in κ. The state space is $(0, \infty)$. The point $\kappa = 0$ is unattainable, and without capital shocks the κ-direction is pure transport: at each level of z the state moves towards the zero of $\mu^\kappa(\cdot, z)$, which mean reversion delivers(Section 3.4; for the Cobb–Douglas case Section B.7.2), so the stationary κ-marginal is concentrated on the band of rest points $\{\kappa : \mu^\kappa(\kappa, z) = 0 \text{ for some } z \in [-\bar{z}, \bar{z}]\}$, which is narrow (the box below), and has no tails beyond it. The grid is nevertheless a finite interval $[\underline{\kappa}, \bar{\kappa}]$, closed at both ends by one-sided rows on both operators, and two facts govern the error this can introduce.

(a) *Forward equation.* Where the drift points inward at an end no boundary condition is needed: the upwind stencil of Step 3 uses interior points only, and the one-sided row at that end is exact. When the grid contains the band, the stationary distribution of the discrete equation is then the stationary distribution of the untruncated process restricted to the grid, and the diagnostic is the mass in the edge cells, which is zero to rounding. Mass in an edge cell appears exactly when the drift at that edge points outward, in which case the one-sided row discards the outward component (the process is stopped at the edge) and the grid is too narrow on that side; the implementation warns when the drift at an edge points outward.

(b) *Backward equation.* The same one-sided rows close the backward operator. The marginal value at a point depends on the values along the path the state follows from it; where the drift points inward at both ends every such path stays on the grid, so the truncation introduces no error in $\mathfrak{q}$ on the band the distribution occupies. Where the drift points outward at an end, the paths from the cells near that end leave the grid and the one-sided row is a modelling choice, whose effect

is confined to those cells, outside the band. Both facts are checked by widening the grid(Step 7).

Under Cobb–Douglas demand the drift does not depend on z(Theorem 1), the band is a single point and the stationary distribution a point mass; that case is not solved on the grid (Remark 8; Step 5).

The productivity coordinate. z is discretised on a uniform grid $[-\bar{z}, \bar{z}]$ with $\bar{z}$ a fixed number of stationary standard deviations $\varsigma_n^Z / \sqrt{2\theta_n}$, and $\mathcal{L}_z^{(n)}$ by the upwind stencil for the drift $-\theta_n z$ (or the central stencil wherever the cell-Péclet condition $(\varsigma_n^Z)^2 \geq \theta_n |z| \Delta z$ holds, which keeps the operator an M-matrix and is second order), central differences for the diffusion, and reflecting rows at $\pm\bar{z}$. The discrete z-process is a reflected Ornstein–Uhlenbeck process whose stationary distribution is known in closed form; since z is exogenous, the z-marginal of the stationary node distribution is exactly this distribution, which Step 7 checks.

Quadrature. Densities are carried as mass vectors: with w_i the trapezoid weights of the κ-grid and Δz the z-spacing, the mass at grid point (i, j) is $p_{ij} = \mathfrak{m}(\kappa_i, z_j)\, w_i \Delta z$ (halved at the z-ends), $\sum_{ij} p_{ij} = 1$, and every moment of Step 5 is a sum $\sum_{ij} f_{ij} p_{ij}$.

[Grid settings and cautions, Step 1.]

- *Centre and parity.* The κ-grid is log-spaced around the Cobb–Douglas rest point κ_n of the model without capacity constraints with an *odd* number of points, so that κ_n is itself a grid node: the degenerate Cobb–Douglas branch of Step 5 places a point mass there, and the $\eta \to 1$ check of Step 7 reads the solution on that node. Grid-refinement checks must preserve the parity (refine to $2N_\kappa - 1$, not $2N_\kappa$).
- *Recentring under capacity.* The converged node distributions sit further from that centre under capacity than without capacity constraints, and on the side capacity pushes them: at the two-node calibration used here ($N = 2$, capital intensities chosen so that $\mathbb{E}_n[U] = 0.80$ at both nodes) $\mathbb{E}_n[\ln \kappa] - \ln \kappa_n$ was -0.157 and $+0.100$ under capacity against -0.081 and $+0.046$ without capacity constraints, roughly a doubling: the node with the larger capital intensity gains capital. With a grid four log units wide on either side this is harmless — the edge mass stayed at 10^{-17} throughout — so no recentring was needed, and the recentring described in the text was left as a manual option. It becomes worth running when x_n is pushed far above its calibrated value, or when a coarse grid is used deliberately (a calibration or a transition sweep), where the natural schedule is to recentre and re-widen once between stages, never inside one.
- *Edge rows drop the outward drift.* The reflecting rows are one-sided zero-flux rows built from the inward spacing only: an edge drift that points *outward* is silently discarded by the operator. That is what the edge-drift warning guards; a warning means the grid is too narrow on that side, not that the operator is wrong.

- *Stencils.* The z-operator uses the central drift row wherever the M-matrix condition allows (second-order z-marginal at no cost). The κ-drift is upwinded.
- Refinement in N_κ, in N_z and in the κ-width is a standing check of Step 7; the z-refinement is the least sensitive by far.
- *The window, not the point count, is the binding grid parameter.* The stationary distribution is about a hundred times narrower than the default window — $\mathrm{sd}_n(\ln\kappa) \approx 0.02$ against ± 2 log units — so it spans roughly one cell whatever N_κ is, and refining the point count on that window buys resolution at the wrong price. Narrowing to ± 0.5 at $N_\kappa = 101$ resolves the same distribution over 2.8 cells at the wall time at which $N_\kappa = 151$ on ± 2 resolves it over 1.3, and it improves the standing refinement check by an order of magnitude: g moves by 2×10^{-6} at $2N_\kappa - 1$, 2×10^{-8} at $2N_z - 1$ and 1×10^{-5} for one more log unit of width, against 7×10^{-5}, 6×10^{-9} and 3×10^{-5} on the wide default. Once the window holds the distribution it is the *cell width* and not the window that the cross-sectional dispersions see: 101 points on ± 0.25 reproduce 201 points on ± 0.5 digit for digit.
- *The window has a second job, and it floors how narrow it may be.* Fitting the window to the distribution is right for resolving it and says nothing about the other thing the κ-grid does — host the $\mathfrak{q}$-equation, whose Newton method needs the state drift to point inward at both reflecting ends, while the outer loop evaluates the equilibrium at growth rates away from the solution, where the zero-drift rest point has moved. Measured on a three-node capacity economy at the centres the distribution itself suggests: the solve succeeds at half-widths 0.50, 0.30 and 0.20 and *fails* at 0.15 and at 0.12 — at every one of $N_\kappa = 61, 81, 101$. The failure is a property of the window alone (g agrees to 5×10^{-6} across every grid that solves), so any automatic narrowing rule needs a floor, a quarter of a log unit here.
- *A continuation in a structural parameter recentres on the model's own algebraic rest points, not on the previous point's converged distribution.* The Cobb–Douglas rest points of Proposition 5 move *with* the parameter; the previous continuation step's distribution is one step stale. Measured over a 240-fold sweep in the capacity intensities: 138 solves, no edge warning, edge mass at most 4×10^{-18}.
- *What is grid-converged and what is not.* At four times the κ-points, node utilization $\mathbb{E}_n[U]$ and the within-node dispersion of *firm* utilization move by at most 4×10^{-5} relative — the second because it is a function of the threshold and the shape parameter alone and sees no grid — while the dispersion of *sector* utilization across a node moves by 8–16 % and the dispersion of the capital distribution itself by 20–38 %, neither having settled. A cross-sectional dispersion of $\ln\kappa$, or any statistic inheriting it, is not a reportable number on any grid this method affords; a moment of the threshold is.

D.3 Step 2: Per-Node Static Solver

Given $(\mathcal{P}_n, \widehat{\mathcal{D}}_n, q)$, solve at every grid point (κ, z) the per-node static block(156), the detrended form of Lemma 15. Substituting the pricing/threshold line and the second form of the sales line into the demand line eliminates $(P, \widehat{s})$ and leaves one scalar equation in the threshold; multiplying it through by $\mathrm{Y}_n(\overline{v})$ clears the pole that the price has at $\underline{v}_n$ and leaves the residual

$$\begin{aligned} H(\overline{v}) &:= A\, U_n(\overline{v}) - b\, \mathrm{Y}_n(\overline{v})^{\eta} - \widehat{\mathcal{D}}_n\, \mathrm{Y}_n(\overline{v}) = 0, \\ A := \frac{\mathcal{P}_n\, \kappa}{x_n}\, \mu_{n,1/\varepsilon}^{\frac{\varepsilon}{\varepsilon-1}}, &\qquad b := \beta_n\, q\, c^{1-\eta}, \qquad c := \frac{\mathcal{P}_n}{e^z x_n^{\alpha_n}}, \end{aligned} \tag{208}$$

solved by a safeguarded Newton method in the coordinate $\ln \overline{v}$, in which the derivative is available in closed form from the two log-derivativesof Proposition 6,

$$\begin{aligned} \frac{\mathrm{d}H}{\mathrm{d}\ln \overline{v}} &= -A\, U_n\, \ell_U - \eta\, b\, \mathrm{Y}_n^{\eta}\, \ell_{\mathrm{Y}} - \widehat{\mathcal{D}}_n\, \mathrm{Y}_n\, \ell_{\mathrm{Y}} < 0, \\ \ell_{\mathrm{Y}} := \frac{\mathrm{d}\ln \mathrm{Y}_n}{\mathrm{d}\ln \overline{v}} \geq 0, &\qquad \ell_U := -\frac{\mathrm{d}\ln U_n}{\mathrm{d}\ln \overline{v}} > 0, \end{aligned} \tag{209}$$

every factor of which the same tail-function evaluation already delivers. Capital enters only A and productivity only b. The residual is positive at the lower end, $H(\underline{v}_n^{+}) = A\, U_n(\underline{v}_n) > 0$, tends to $-\infty$, and is strictly decreasing under Assumption 5: the sign of H therefore identifies the side of the root, exactly as the sign of the price residual does without capacity constraints, and the method is Newton safeguarded by a bracket, a step being accepted only inside the current bracket and replaced by the bisection midpoint otherwise.

The bracket is one-sided by construction. Its lower end is the pole $\underline{v}_n$ plus a margin, the pole being the zero of the elastic-demand margin $(\varepsilon-1)\mathcal{G}_{n,1}^{-} - 1 + G_n$ in the numerator of (106) (Assumption 2) and a datum of the node, computed once when the node's tail functions are built. No upper end is searched for: the sign of H at the iterates supplies one as soon as an iterate overshoots, and until then an upward Newton step is capped at a fixed number of log units. Convergence is tested on the Newton increment at the evaluated point rather than on the bracket width, because at a root reached to rounding the increment is the only quantity still meaningful and a strict interior test would reject the converged iterate. Every grid point is warm-started from its own threshold at the previous aggregate iteration, and on a cold start from the neighbouring grid point: the root moves smoothly and monotonically with (κ, z) (Lemma 15(iv)), so a warm-started sweep costs one evaluation of the tail functions per grid point.

At the root, seven arrays are stored on the node's grid:

$$\begin{aligned} &\ln P = \ln c - \ln \mathrm{Y}_n(\overline{v}), \qquad \widehat{s} = A\, \frac{U_n(\overline{v})}{\mathrm{Y}_n(\overline{v})}, \qquad P^{1-\eta}, \qquad \widehat{x} = \xi_n(\overline{v})\, \varrho_n(\overline{v})\, \widehat{s}, \\ &\pi_k = \zeta_n(\overline{v})[1-\xi_n(\overline{v})]\, \varrho_n(\overline{v})\, \frac{\widehat{s}}{\kappa} = \frac{\varepsilon}{1-\alpha_n}\, \frac{\mathcal{P}_n}{x_n}\, \Phi_n(\overline{v}), \qquad \ln \overline{v}, \qquad e_S(\overline{v}) = 1 + \frac{\ell_U}{\ell_{\mathrm{Y}}}, \end{aligned} \tag{210}$$

with $\varrho_n$of (118), the share functions of Lemma 5, the two forms of the flow (115) and the supply elasticity (111). The first three and the flow array π_k are what the model without capacity constraints also stores, π_k being what the backward equation of Step 3 consumes in either model; the spending array $\widehat{x}$ is new, and is what the coupling of Step 5 integrates in place of the product $\xi_n \widehat{m}_{S,n}$ of that model; $\ln \overline{v}$ is the warm start of the next sweep; and e_S, a pure function of the threshold, is what the numeraire step of Step 6 consumes. Only one elasticity is needed downstream, and it is recorded here once: at the root, by Proposition 6,

$$\frac{\partial \ln P}{\partial \ln \kappa} = \frac{-1}{e_S(\overline{v}) - (1-\eta)\vartheta}, \qquad \vartheta := \frac{\beta_n \, P^{1-\eta} \, q}{\widehat{s}} \in (0,1], \tag{211}$$

the final-demand share of the node's sales ϑ being read off the two terms of the demand line.

Cobb–Douglas demand. At $\eta = 1$ node sales are state-independentby Lemma 17(i), so (208) loses its z-dependence and the price is filled along z afterwards, from (121):

$$\widehat{s}_n = \beta_n q + \widehat{\mathcal{D}}_n, \qquad A\, U_n(\overline{v}) = \widehat{s}_n \, \Upsilon_n(\overline{v}), \qquad \ln P = \ln \mathcal{P}_n - z - \alpha_n \ln x_n - \ln \Upsilon_n(\overline{v}(\kappa)).$$

The root-find therefore runs once per κ-row rather than once per grid point, and the threshold, utilization, the shares, the spending array and the flow are all free of z(Lemma 17(ii)); only the price and $P^{1-\eta}$ carry it. Unlike the model without capacity constraints, the root-find is still run: capacity removes the closed form even where it leaves the cancellation intact.

Shape checks. The block rests on three shape restrictions on the node's demand-shock distribution — Assumptions 2, 3 and 5 — each of which is a condition on $(G_n, \varepsilon, \alpha_n)$ alone and none of which is proved. Each is checked numerically on the node's own parameters when its tail functions are built, before any solve, so that a parameter draw that leaves the verified region is reported rather than silently producing a non-monotone Υ_n and a spurious root.

The model without capacity constraints. As $x_n \to 0$ the threshold diverges and every tail function tends to its constant (Lemma 4), so (208) is replaced by the power equation in the price of Lemma 24, in detrended form (200),

$$c\, P^{1/\alpha_n} = \beta_n \, q \, P^{1-\eta} + \widehat{\mathcal{D}}_n, \qquad c = \kappa(\Gamma_n \mathcal{P}_n)^{-\frac{1-\alpha_n}{\alpha_n}} e^{z/\alpha_n}, \tag{212}$$

solved by Newton on $\ln P$; the left-hand side minus the right-hand side crosses zero exactly once and from below (proof of Lemma 24), and the bracket is two-sided rather than one-sided: the two one-term roots of (212) (supply against final demand alone, supply against inbound demand alone) are lower bounds for the root, and an upper bound is found by geometric expansion. The seven arrays of (210) are the same with the tail functions at their limits: $\varrho_n \equiv 1$, so $\widehat{x} = \xi_n \widehat{s}$ with the constant share ξ_n; $\pi_k = \frac{\varepsilon-1}{\varepsilon}\alpha_n \widehat{s}/\kappa$ (188); $\ln \overline{v} = \infty$; and $e_S \equiv 1/\alpha_n$, which turns (211) into $-\alpha_n/[1 - \alpha_n(1-\eta)\vartheta]$. Under Cobb–Douglas demand the node block without capacity constraints has

the closed formof Appendix C.5 ($\ln P$ affine in z, $\ln\kappa$ and $\ln\mathcal{P}_n$, and $\widehat{s}_n = h_n q$ state-independent) and no root-find is run at all. The solved block reproduces the closed-form elasticities(195) (Step 7).

[Static block, Step 2: closed forms, hygiene, tolerances.]

- *What one evaluation costs.* Under the log-normal specification every tail function is a normal distribution function (60), and the whole of (208)–(210) needs G_n, $\mathcal{G}^-_{n,1}$, $\mathcal{G}^+_{n,1/\varepsilon}$, $\mathcal{G}^-_{n,\frac{1}{1-\alpha_n}}$ and the log-density $\overline{v}g_n(\overline{v})$: four normal distribution functions and one density per evaluation, from which $\mathcal{R}_n$, U_n, Γ_n, Υ_n, ϱ_n, ξ_n, $\zeta_n[1-\xi_n]$, ℓ_U and ℓ_Υ all follow by ratios. That is why the capacity sweep is not the expensive step it looks: warm-started, it was *cheaper* than the price Newton of the model without capacity constraints, 4.0 against 4.5 ms per node for a 301×31 grid (re-measured alone on an otherwise idle machine, 3.9 against 5.0), because that Newton takes two or three steps per point and this one takes a single evaluation.
- *Everything in logs, and the upper tail through the complementary normal cdf.* $\mathcal{G}^-$ and $\mathcal{G}^+$ are formed as $\exp[\ln\mu_{n,\theta} - \theta\ln\overline{v} + \ln\Phi_{\mathcal{N}}]$ and the same with the complementary distribution function; never as $1-\Phi_{\mathcal{N}}$. This is not cosmetic: evaluating the upper tail by subtraction breaks the monotonicity of Φ_n at large α_n and manufactures spurious violations of the three shape assumptions in the far tail. Beyond roughly eight standard deviations above the pole G_n rounds to exactly 1, so a check on the derivative of G_n has to be run as a check on the derivative of $1-G_n$.
- *The one unavoidable cancelling difference.* The margin $(\varepsilon-1)\mathcal{G}^-_{n,1} - 1 + G_n = [\varepsilon_n(\overline{v}) - 1]\Psi_{n,1}(\overline{v})$ in the numerator of (106) is a difference of $O(1)$ terms whose zero is the pole, so near the pole its relative precision is 10^{-16} divided by its own value, and Υ_n and Φ_n inherit that. Nothing removes it; the bracket's margin above the pole is what keeps it away from the cliff, and the capital share is computed from the lower-tail shares $\eta_{n,\theta}$ rather than by subtracting near-equal products.
- *Tolerances and the two constants.* The Newton tolerance is on the increment in $\ln\overline{v}$ and has to be tight: at 10^{-12} the demand relation was satisfied only to 2×10^{-10} at roots close to the pole, where Υ_n is steep; at 10^{-14} it holds to 10^{-12}, and quadratic convergence makes that at most one extra evaluation. The bracket's lower end is the pole plus 10^{-8} in $\ln\overline{v}$, and an upward Newton step before an upper end exists is capped at four log units. A *strict* interior test on the bracket is the thing to avoid: it rejected an iterate that had landed on the root to rounding and sent the solver into forty bisections from a far upper end.
- *The Cobb–Douglas branch carries a $-\ln\bar{Z}_n$ term,* $\ln P = \ln\mathcal{P}_n - \ln\bar{Z}_n - z - \alpha_n\ln x_n - \ln\Upsilon_n(\overline{v}(\kappa))$, which the displays of Section B.7 drop because they set $\bar{Z}\equiv 1$; the closed form without capacity constraints carries it in the same place, $\ln P = \alpha_n[\ln(h_nq) - \ln\kappa] - z - \ln\bar{Z}_n + (1-\alpha_n)[\ln\Gamma_n + \ln\mathcal{P}_n]$. It matters for the $\bar{Z}$ sweeps and for the MIT shock of Step 6, both of which run through this branch under Cobb–Douglas demand.

- *The shape checks warn, they do not refuse.* They are run on the node's own $(\sigma_{v,n}, \varepsilon, \alpha_n)$ at build time; a warning means the parameters have left the region over which the assumptions were verified, not that the solve is wrong. At the two-node calibration used here none warned.
- *Measured at that calibration* ($N = 2$, x_n chosen so that $\mathbb{E}_n[U] = 0.80$ at both nodes, CES, 301×31 grids): the node equation's own residual $\sup |H|/(AU_n)$ and the demand relation (141) both at most 7×10^{-13}; the two forms of the flow in (210) agreeing to 3×10^{-15}; the rent identity (77) to 2×10^{-16}; the closed-form tail functions against quadrature in $\ln v$ to 2×10^{-14}; and ℓ_U, ℓ_Y against central differences of $\ln U_n$ and $\ln Y_n$ to 8×10^{-10}.
- *The same magnitudes over the whole configuration matrix.* Rerun over all of {with, without capacity constraints} $\times$ {CES, Cobb–Douglas} at $N = 2$ and at $N = 3$, the node equation's residual and the demand relation stay at most 7×10^{-13}, the two forms of the flow at most 6×10^{-14}, the rent identity and the four-way split at 2×10^{-16}, and Cobb–Douglas z-freeness at 2×10^{-15}: the block is not configuration-sensitive. The one quantity that moves is the elasticity check, and it moves only with the grid — 1.5 to 5.7×10^{-4} at 301 κ-points against up to 5×10^{-3} at 101, which is the $O(h^2)$ of the central differences and not of the solver (Step 7).

D.4 Step 3: Per-Node Backward Equation

The backward equation(157) is quasi-linear rather than a max-form HJB: the optimisation is already encoded in $\iota^*(\mathfrak{q})$ of (133). The upwind finite-difference machinery of the Achdou–Han–Lasry–Lions–Moll toolkit therefore applies, and the discrete equation is solved by Newton's method with the false transient of the toolkit retained as a fallback. This step is common to the two models: it reads the flow array π_k of (210) as data, and under capacity that array is the capacity flow (115), a function of the node's threshold given $\mathcal{P}_n$. What matters here is that the threshold, and hence π_k, is set by the static block alone and does not depend on $\mathfrak{q}$: the flow is exogenous to this equation in both models, so the semismooth Newton method and its Jacobian (215) below are unchanged by capacity, and only the numbers in π_k differ. Stationary case, per node:

1. Initialise $\mathfrak{q}^0$: the Cobb–Douglas closed form $\mathfrak{q}_n$ of the model without capacity constraints, (194), constant on the grid — the initialiser of Step 5 in both models — or the previous outer iterate.
2. Given $\mathfrak{q}^\ell$: compute $\iota^\ell = \iota^*(\mathfrak{q}^\ell)$ from (133), the geometric drift $\mu^{\kappa,\ell} = \Phi_n(\iota^\ell) - \delta_n - g$, the effective discount rate $R^\ell_{\text{eff}} := r - \Phi_n(\iota^\ell) + \delta_n$ and the flow $\pi_k - \iota^\ell$, all arrays on the grid.
3. Build the sparse operator $\mathsf{A}^\ell = \mathsf{D}^\ell_\kappa + \mathsf{L}_z$, where D^ℓ_κ discretises $\mathcal{A}_\kappa$ of (206) at each z-level: the drift $\mu := \mu^{\kappa,\ell}\kappa$ by first-order one-sided differences in the direction of its own sign. With $h_i^+ = \kappa_{i+1} - \kappa_i$, $h_i^- = \kappa_i - \kappa_{i-1}$ and $\mu^\pm = \max\{\pm\mu, 0\}$, the

interior row is

$$(\mathsf{D}_\kappa f)_i = \mu_i^+ \frac{f_{i+1} - f_i}{h_i^+} - \mu_i^- \frac{f_i - f_{i-1}}{h_i^-}, \tag{213}$$

and the two boundary rows carry the inward drift component only (one-sided closure, Step 1). Every row of A^ℓ sums to zero and every off-diagonal entry is nonnegative on any grid, because the drift is upwinded; the same holds for L_z (Step 1), whose diffusion weights are positive. In the coordinate $\varkappa = \ln \kappa$ the stencil is (213) with the constant spacing $\Delta\varkappa$ and drift $\mu^{\kappa,\ell}$, by (207).

4. *Newton step.* With the term $\mathfrak{q}[\Phi_n - \delta_n]$of (157) absorbed into the discount, the discrete equation is

$$F(\mathfrak{q}) := \mathsf{A}(\mathfrak{q})\,\mathfrak{q} + \pi_k - \iota^*(\mathfrak{q}) - R_{\text{eff}}(\mathfrak{q})\,\mathfrak{q} = 0, \tag{214}$$

where $\mathsf{A}(\mathfrak{q})$, $\iota^*(\mathfrak{q})$ and $R_{\text{eff}}(\mathfrak{q})$ are the arrays of items 2–3 evaluated at $\mathfrak{q}$. Its Jacobian is

$$J(\mathfrak{q}) = \mathsf{A}(\mathfrak{q}) - \operatorname{diag}(R_{\text{eff}}(\mathfrak{q})) + \operatorname{diag}\left(\frac{\kappa\,\iota^{*\prime}(\mathfrak{q})}{\mathfrak{q}}\,(\mathsf{D}\mathfrak{q})\right), \tag{215}$$

where $(\mathsf{D}\mathfrak{q})_i$ is the one-sided difference of $\mathfrak{q}$ that the upwind stencil applies in row i: the derivatives of the flow and of the discount cancel by the envelope property, $-\iota^{*\prime} + \Phi_n'(\iota^*)\iota^{*\prime}\mathfrak{q} = 0$ by (133), and only the derivative of the drift, $\partial\mu_i/\partial\mathfrak{q}_i = \kappa_i\iota^{*\prime}(\mathfrak{q}_i)/\mathfrak{q}_i$, survives, applied to the difference of $\mathfrak{q}$ it multiplies. Because the upwind direction switches with the sign of the drift, F is only piecewise smooth and the method is a semismooth Newton method. Solve $J(\mathfrak{q}^\ell)\Delta\mathfrak{q} = -F(\mathfrak{q}^\ell)$, backtrack on $\|F\|_\infty$ along $\Delta\mathfrak{q}$, and set $\mathfrak{q}^{\ell+1} = \mathfrak{q}^\ell + t\Delta\mathfrak{q}$. J has the sparsity pattern of A plus the diagonal, fixed across iterations once both upwind directions are included in the pattern, so one numerical factorisation per step with a reused symbolic analysis is all that changes.

5. *Fallback: capped false transient.* Whenever no backtracked Newton step decreases $\|F\|_\infty$, take instead one implicit false-transient step with step $\Delta\tau_{\text{eff}}$,

$$\underbrace{\left[\left(\frac{1}{\Delta\tau_{\text{eff}}} + R_{\text{eff}}^\ell\right)\mathsf{I} - \mathsf{A}^\ell\right]}_{=:\,\mathsf{M}^\ell}\mathfrak{q}^{\ell+1} = \pi_k - \iota^\ell + \frac{\mathfrak{q}^\ell}{\Delta\tau_{\text{eff}}}. \tag{216}$$

M^ℓ has nonpositive off-diagonals, and because the rows of A^ℓ sum to zero its i-th diagonal entry exceeds the sum of the moduli of its off-diagonal entries by exactly $1/\Delta\tau_{\text{eff}} + R_{\text{eff},i}^\ell$; hence M^ℓ is a nonsingular M-matrix, and the step monotone, if

$$\frac{1}{\Delta\tau_{\text{eff}}} + R_{\text{eff}}^\ell(\kappa, z) > 0 \qquad \text{at every grid point.} \tag{217}$$

R_{eff} is positive at the rest point (there it is $\rho + (\gamma - 1)g > 0$ by Assumption 6) but *not* on the whole grid: at low-κ, high-π_k states the policy has $\Phi_n(\iota) - \delta_n > r$, so $R_{\text{eff}} < 0$ there, and a fixed large step violates (217). The step is therefore capped

adaptively: with $R^{\ell}_{\min} := \min_{(\kappa,z)} R^{\ell}_{\text{eff}}$ and a safety factor $c_{\text{cap}} \in (0,1)$,

$$\Delta\tau_{\text{eff}} = \begin{cases} \Delta\tau_{\max}, & R^{\ell}_{\min} \geq 0, \\ \min\left\{\Delta\tau_{\max},\ c_{\text{cap}}/(-R^{\ell}_{\min})\right\}, & R^{\ell}_{\min} < 0. \end{cases} \tag{218}$$

Iterating (216) alone, in the manner of policy iteration, is a valid solver of (214); the two solvers converge to the same discrete solution.

6. Iterate 2–5 until $\left\|\mathfrak{q}^{\ell+1} - \mathfrak{q}^{\ell}\right\|_{\infty}$ is below tolerance, the test being evaluated after the policy and operator have been rebuilt at the new iterate, so that the three are mutually consistent. Each linear solve is a sparse banded system of size $N_\kappa N_z$; the node loop is parallel.

Configurations. Without capital shocks A^{ℓ} is the pure upwind operator in κ and coincides with G^{ℓ} of Step 4. Under Cobb–Douglas the converged $\mathfrak{q}$ is constant along z(Theorem 1).

[The step cap and the line search, Step 3.]

- *Do not replace the global cap by a fixed step or by a row-wise cap.* A fixed step of the size the toolkit recommends leaves (217) violated at the bottom of any grid wide enough to contain low-κ states; the implicit matrix loses the M-matrix property and the iteration diverges. A row-by-row cap, which satisfies (217) at every grid point with the largest admissible local step, *also diverges*: the condition is sufficient for a monotone linear step, not for convergence of the quasi-linear iteration, and the global cap is what damps that iteration. The safety factor c_{cap} is a setting, not a modelling choice, and a value well inside $(0,1)$ is what makes the damping reliable.
- *Line search.* A Newton step whose sup-norm is already below the solver tolerance is accepted without backtracking: at a converged point $\|F\|_{\infty}$ sits at rounding and cannot fall. Monotonicity is not guaranteed for a Newton step and is not needed; the converged $\mathfrak{q}$ satisfies the same discrete equation as the transient's.

D.5 Step 4: Per-Node Forward Equation

The forward equation(158) is the adjoint of the state generator $\mathbb{G}^{(n)}$ (Section B.8; Proposition 8, (v)), which without capital shocks is the operator $\mathcal{A}^{(n)}$ of the backward equation. Nothing in this step depends on the model: the forward equation(158) contains no capacity object, and the two models differ only through the converged policy at which the drift is evaluated. Discretely:

1. Build the discrete state generator $\mathsf{G} = \mathsf{D}_\kappa[\mu^\kappa \kappa] + \mathsf{L}_z$ with the same routine and the same stencil (213) as in Step 3, the drift $\mu^\kappa \kappa$ upwinded by the sign of μ^κ and evaluated at the converged policy. Without capital shocks $\mathsf{G} = \mathsf{A}$, and the converged operator of Step 3 could be reused as it stands; the reference implementation nevertheless builds G separately, because the backward operator of the valuation and the generator of the state are different objects that coincide only in this case.

2. Solve for the stationary mass vector

$$\mathsf{G}^{\mathsf{T}}\, p = 0, \qquad \sum_{ij} p_{ij} = 1, \qquad m_{ij} = \frac{p_{ij}}{w_i \Delta z}, \tag{219}$$

with the quadrature weights of Step 1. Mass is conserved exactly: the rows of G sum to zero, so the columns of G^{T} do, and $\sum_{ij}(\mathsf{G}^{\mathsf{T}} p)_{ij} = 0$ for every p. This is the transpose-consistency of the standard toolkit, attached to the state generator.

3. *Solution method.* Replace one equation of (219) by the normalisation and solve the linear system directly. The sparsity pattern of G^{T} is fixed across iterations once both upwind directions are included in the pattern, so the symbolic factorisation is reused and only the numerical factorisation is redone. If the pinned system is singular, fall back to the general stationary solver (the same pinned solve on a fresh factorisation, then the leading eigenvector of G^{T}). The pinned system is singular to rounding when the stationary distribution is narrower than a cell of the grid, which is a statement about the grid and not about the solver (the box below); the limiting case, a stationary distribution singular in κ, arises at η near 1 (pure transport into a rest point, Remark 8), and the pipeline declines to run on the grid at $\eta = 1$ (Step 5).

4. *Diagnostic.* Node-wise stationarity of $\mathbb{E}_n[\kappa]$,

$$\mathbb{E}_n[\kappa\, \mu^{\kappa}(\kappa, z)] = 0 \quad \Longleftrightarrow \quad \mathbb{E}_n[(\Phi_n(\iota^*) - \delta_n)\kappa] = g\, \mathbb{E}_n[\kappa], \tag{220}$$

holds for the distribution of (219) up to the boundary-flux terms of the two one-sided end rows, which are of the order of the mass in the edge cells and absent where the edge drift points inward. In the κ-coordinate the stencil (213) reproduces the identity function exactly, $(\mathsf{D}_\kappa \kappa)_i = \mu_i$ at interior rows, so $\sum_{ij} p_{ij}(\mathsf{G}\kappa)_{ij} = 0$ *is* (220) to rounding; in the coordinate $\varkappa = \ln \kappa$ the stencil applied to $e^{\varkappa}$ is exact only to first order and the identity holds to $O(\Delta\varkappa)$.

[KFE cautions, Step 4.]

- *Sign-safe normalisation.* The solution of the pinned linear system is determined up to an overall sign that depends on the pinned row and on rounding. Divide by the (possibly negative) sum *before* clipping tiny negative entries; clipping first can annihilate a globally negative but correct solution. This step is load-bearing.
- *When the eigenvector fallback fires, it is a diagnosis of resolution.* It is triggered exactly where the κ-distribution is narrower than a cell of the grid it is being solved on, so that the pinned linear system is singular to machine precision: inside the $\eta \to 1$ continuation of Step 7 on a coarse κ-grid, and generally whenever the grid is cut for cost — at $N_\kappa = 61$ on a half-width of 0.5 the narrower node's distribution has $\mathrm{sd}(\ln \kappa) = 0.006$ against a cell of 0.017, and the fallback then fires once per solve. It is the documented alternative route and not a wrong answer: on the transitions where it fired, the reported peak responses reproduced the uncut grid's to 0.002 %. But it is a statement about the grid, not a routine event, and a run whose object is the *shape* of the distribution rather than an aggregate response should be refined until it stops firing.

D.6 Step 5: Moments, Graphon Products, and the Outer Fixed Point

Every moment in this step is a sum against the mass vector of Step 4. This is where the two models differ most visibly. The couplingof (159) regenerates *three* block moments: the log-price moment $m_{P,n}$, the spending moment $\widehat{m}_{X,n}$ and the sales moment $\widehat{m}_{S,n}$, the first two feeding the bundle price and the inbound demand and the third entering only through GDP. Without capacity constraints the spending moment is not an independent object — the expenditure share is the constant ξ_n, so $\widehat{m}_{X,n} = \xi_n \widehat{m}_{S,n}$ — and two moments suffice; under capacity the share $\xi_n(\overline{v})\varrho_n(\overline{v})$ moves with the node's state, the product of the means is not the mean of the product, and the spending array of (210) must be integrated in its own right.

Algorithm 1 (Stationary GMFG solver). Initialisation. *Solve the degenerate Cobb–Douglas system of the model without capacity constraints, (194) (algebraic; Proposition 12), closed with the numeraire(192), for $(g^0, q^0, \kappa_n, \mathfrak{q}_n)$, and set $\mathcal{P}^0, \widehat{\mathcal{D}}^0, m_P^0$from (191) with $\mathbb{E}_n[\ln \kappa] = \ln \kappa_n$ and $\mathbb{E}_n[z] = 0$. This is a valid initialiser at any x_n: it is the $x_n \to 0$ limit of the model being solved (Lemma 4), so it solves a nearby economy exactly rather than approximately, and its whole content — the profile of rest points, the growth rate and the price profile — survives the limit as the leading term. The capacity model's own Cobb–Douglas path, Proposition 5, is the optional refinement: it is node-wise algebra once $(\mathcal{P}_n, \widehat{s}_n)$ are known, but those are no longer given by a Fredholm solve, so it carries a $2N$-dimensional static fixed point(Proposition 4) with it and is worth solving only where the path without capacity constraints is too far away to converge from.*

Outer loop on g (safeguarded secant). *Given g, $r = \rho + \gamma g$:*

1. Inner loop on aggregates $\Theta = (\mathcal{P}_n, \widehat{\mathcal{D}}_n, q)$:
   (a) *Node loop: Step 2 (statics), Step 3 ($\mathfrak{q}$), Step 4 (distribution).*
   (b) *Moments per node:* $m_{P,n} = \mathbb{E}_n[\ln P]$, $\widehat{m}_{X,n} = \mathbb{E}_n[\widehat{x}]$, $\widehat{m}_{S,n} = \mathbb{E}_n[\widehat{s}]$, $\mathcal{N}_n = \mathbb{E}_n[P^{1-\eta}]$, $\mathbb{E}_n[\kappa]$, $\mathbb{E}_n[\iota^*\kappa]$ *and* $\mathbb{E}_n[(\Phi_n(\iota^*) - \delta_n)\kappa]$, *as sums against p_n.*
   (c) *Graphon products* (205), *the three lines of* (159): $\ln \mathcal{P}_n^{new} = \sum_m A_{nm} m_{P,m} \Delta_m - e_n$, $\widehat{\mathcal{D}}_n^{new} = \sum_m A_{mn} \widehat{m}_{X,m} \Delta_m$, $q^{new} = \sum_m [\widehat{m}_{S,m} - \widehat{m}_{X,m}] \Delta_m$.
   (d) Capital adding-up (scale normalisation). *Compute $u = [\sum_n \mathbb{E}_n[\kappa]\, \Delta_n]^{-1}$ and rescale the nominal aggregates, $\widehat{\mathcal{D}}^{new} \mapsto u\, \widehat{\mathcal{D}}^{new}$, $q^{new} \mapsto u\, q^{new}$, leaving $\mathcal{P}^{new}$ untouched. The inner system at fixed g has an exact one-parameter scale-invariance family along which* all prices, including the final price index P^f, are invariant *(Remark 14), so no numeraire rescaling can pin the scale; only capital adding-up can, and it must therefore be enforced inside the inner loop.*
   (e) *Fixed-point update of Θ towards Θ^{new}, on $(\ln \mathcal{P}, \widehat{\mathcal{D}}, q)$: the damped map $\Theta \leftarrow (1 - \varpi)\Theta + \varpi\, \Theta^{new}$, $\varpi \in (0, 1]$, accelerated by Anderson mixing of the same residual, with the damped step as the fallback whenever the Anderson extrapolate is unusable. Repeat (a)–(e) until $\|\Theta^{new} - \Theta\|$ is below tolerance, the test being taken on the undamped residual before the update so that the returned objects are mutually consistent.*

2. Numeraire residual. *At the converged inner fixed point (at which* $u\,(P^f)^{1-\eta} = 1$ *by the GDP identity, so that* P^f *is a function of* g *alone and adding-up and the numeraire are one condition) compute the final price index* $P^f(g) = [\sum_n \beta_n \mathcal{N}_n \Delta_n]^{1/(1-\eta)}$ *and the residual* $N(g) := P^f(g)^{1-\eta} - 1$*, which has the same root as* $P^f(g) - 1$*, avoids the power* $1/(1-\eta)$ *near* $\eta = 1$ *and equals* $1/u - 1$ *for the scale factor* u *of step (d) (under Cobb–Douglas demand the residual is the log form* $\sum_n \beta_n \mathbb{E}_n[\ln P]\Delta_n$ *of the numeraire). Update* g *by a safeguarded secant on* N*: the secant step through the two most recent iterates, confined to the tightest sign-change bracket among all evaluations once one exists, with bisection as the fallback, and every candidate kept inside the region* $\rho + (\gamma - 1)g > 0$ *of Assumption 6. The first step uses the slope in* g *of the Cobb–Douglas block's numeraire residual, available in closed form from the initialiser, as a Newton-like step of bounded size. The outer scalar root-find plays the role of the interest-rate update on the asset-market residual in Aiyagari-type algorithms; monotonicity of* N *in* g *is not guaranteed, so a sign change must be verified, not assumed. At the root, adding-up and the numeraire* (47) *hold simultaneously.*

3. *Ex-post checks:* $\widehat{C} = q - \sum_n \mathbb{E}_n[\iota^* \kappa]\Delta_n > 0$*;* $\rho + (\gamma - 1)g > 0$*; adding-up* $\sum_n \mathbb{E}_n[\kappa]\Delta_n - 1 \approx 0$ *(holds at the outer root, not at every inner fixed point: step (d) rescales the nominal aggregates, so an inner fixed point at the wrong* g *has* $u \neq 1$*, and the adding-up residual* $1/u - 1$ *then coincides with the numeraire residual); the node-wise stationarity diagnostic* (220)*, which must hold node by node at a stationary distribution; and the edge mass of Step 1.*

Remark 14 (Scale invariance of the inner system) The inner system at fixed g has an exact one-parameter family of solutions. For any $u > 0$, given an inner solution, the maps $\overline{v}'(\kappa, z) = \overline{v}(\kappa/u, z)$, $P'(\kappa, z) = P(\kappa/u, z)$, $\widehat{s}'(\kappa, z) = u\,\widehat{s}(\kappa/u, z)$, $\widehat{x}'(\kappa, z) = u\,\widehat{x}(\kappa/u, z)$, $\mathfrak{q}'(\kappa, z) = \mathfrak{q}(\kappa/u, z)$, $\mathfrak{m}'(\kappa, z) = \mathfrak{m}(\kappa/u, z)/u$, together with $(\widehat{\mathcal{D}}, q) \mapsto (u\widehat{\mathcal{D}}, uq)$ and $\mathcal{P}$ unchanged, solve the statics, the $\mathfrak{q}$-equation and the FPK equation exactly (up to κ-grid truncation). The static part is Lemma 19: the node equation (208) is homogeneous of degree zero in $(\kappa, q, \widehat{\mathcal{D}})$ at fixed prices, so the *threshold* is invariant along the family and with it every function of the threshold — utilization, the shares ξ_n, ζ_n, ϱ_n, the markup and the supply elasticity e_S — while the euro amounts $(\widehat{s}, \widehat{s}^{\mathrm{f}}, \widehat{x})$ scale by u; hence the flow $\pi_k = \zeta_n[1 - \xi_n]\varrho_n \widehat{s}/\kappa$ is scale-free, as is $\frac{\varepsilon-1}{\varepsilon}\alpha_n \widehat{s}/\kappa$ without capacity constraints, and the spending moment scales exactly like the sales moment, so the three graphon productsof (159) regenerate the scaled aggregates. Capacity therefore leaves the family untouched. The dynamic part is unchanged: the advection term $\mathfrak{q}_\kappa \kappa[\Phi - \delta - g]$of (157) and the term $-\partial_\kappa\{\kappa[\Phi - \delta - g]\mathfrak{m}\}$of (158) are both invariant under translation in $\ln \kappa$ (each is homogeneous of degree zero in κ once $\mathfrak{m}$ is rescaled as a density), so policies, $\mathfrak{q}$ and the distribution simply translate along the log-κ grid. *All prices (*P*,* $\mathcal{P}$*, and hence* P^f*) are invariant along this family*, so no rescaling by any power of P^f can pin the scale; only adding-up can ($\mathbb{E}'_n[\kappa] = u\,\mathbb{E}_n[\kappa]$ moves along the family, so it selects a unique member), and it must sit *inside* the inner loop. Conversely, at the adding-up-normalised inner fixed point P^f is a function of g alone, so the numeraire is the natural outer residual; the Cobb–Douglas blockof Section B.7 already has this structure ($q(g)$ from adding-up in closed form, g from the numeraire). Both closures are equilibrium

conditions of Definition 6; only their algorithmic roles are assigned here. A rescaling by the price level, by contrast, has no bite for $\eta \neq 1$: supply (140) is degree-one in the price level in either model (a common factor on $(P, \mathcal{P})$ leaves the threshold, hence U_n/Y_n, unchanged), but the CES demand term $\beta P^{1-\eta} q$ scales as (price level)$^{2-\eta}$, so dividing nominal aggregates by P^f is not a consistent renormalisation.

Under Cobb–Douglas demand the loop is not run(Remark 8): there is no stationary density to solve for, only a point mass at the node's rest point. Without capacity constraints the algebraic blockof Appendix C.5 is returned directly. Under capacity the corresponding object is the node-wise algebraof Proposition 5 closed by the $2N$-dimensional static fixed pointof Proposition 4 — a small nonlinear solve rather than a closed form — which is solved and returned directly, the loop of Algorithm 1 not being run on a degenerate distribution.

The model without capacity constraints. The mapping is one substitution: $\widehat{m}_{X,n} = \xi_n \widehat{m}_{S,n}$ with the constant expenditure share, which turns the two products of step (c) into that model's $\widehat{\mathcal{D}}_n^{\text{new}} = \sum_m A_{mn} \xi_m \widehat{m}_{S,m} \Delta_m$ and $q^{\text{new}} = \sum_m [1 - \xi_m] \widehat{m}_{S,m} \Delta_m$ and removes the spending moment from the list of integrals in step (b). Everything else in Algorithm 1 — the adding-up normalisation inside, the numeraire residual outside, the safeguarded secant, the ex-post checks — is common to the two models, because Remark 14 holds in both.

[Inner/outer loop cautions, Step 5.]

- *Do not rescale nominal aggregates by a power of P^f inside the inner loop, and do not use adding-up as the outer residual.* The reversed assignment looks natural ("all nominal objects are homogeneous of degree one in the price level") and fails twice: price-level homogeneity holds only at $\eta = 1$ (the CES demand term scales as (price level)$^{2-\eta}$), so a $1/P^f$ rescaling locks P^f away from one while q drains along a neutral ray; and that ray is the exact scale family of Remark 14, along which P^f is flat, so no power of P^f can pin the scale. Adding-up inside, numeraire outside, as in Algorithm 1; the document's own Cobb–Douglas block has the same structure. Adding-up is not an inner convergence criterion: at any inner fixed point u $(P^f)^{1-\eta} = 1$ holds identically, so the raw adding-up residual *is* the outer residual.
- *The inner tolerance must be strictly tighter than the outer one.* The inner loop tests convergence on the residual *before* applying the update, so the accepted inner fixed point carries up to one inner tolerance of error, and the outer numeraire residual inherits it — amplified about threefold under the Cobb–Douglas log numeraire. With the two tolerances equal at 10^{-7} the Cobb–Douglas capacity solves stalled: the outer residual sat at 2.2×10^{-7} for fifty growth iterations while the inner loop accepted its unchanged first iterate every time. The model without capacity constraints escapes the same trap only because Anderson's last step happens to overshoot to 10^{-9}. At 10^{-8} inside against 10^{-7} outside every config-

uration converged in four growth iterations at the cost of about one extra inner iteration each. One order of magnitude is enough; the ordering is what matters.

- *The initialiser without capacity constraints was never a binding constraint.* At the two-node calibration both demand systems, CES and Cobb–Douglas, converged from the algebraic path of the model without capacity constraints in four growth iterations (six, six, five and three inner iterations), with no fallback and no need for the capacity path of Proposition 5. Warm-starting from the corresponding solve without capacity constraints instead saves one inner iteration per outer step and lands on the same root to 10^{-9} in g. Expect the capacity path to be worth implementing only at capacity intensities far above the calibrated ones, where the $x_n \to 0$ limit is a poor guess.
- *Anderson mixing as implemented.* Type-II Anderson acceleration on $x = (\ln \mathcal{P}, \widehat{\mathcal{D}}, q)$ with a short history (depth four at the two-sector reference economy), mixing parameter one, and the damped step as the first step and as the restart: the history is discarded and a damped step taken whenever the extrapolate is non-finite or leaves the positive orthant in $(\widehat{\mathcal{D}}, q)$. It cuts the inner iteration count by an order of magnitude relative to the damped map. There is no regularisation of the least-squares fit; the system is overdetermined by construction ($2N + 1$ rows against at most the depth in columns), which leaves a margin of one row at $N = 2$; the restart is what makes the fit safe there.
- *First outer step.* The bounded Newton-like step on the slope of the Cobb–Douglas block's numeraire residual replaces a fixed first step of prescribed size and sign, which overshoots the root by an order of magnitude at the two-sector reference economy and costs a long inner solve at a bad g. The prescribed size survives only as the bound on the first step.
- All objects ($\mathfrak{q}$, the distributions, the static cache, now seven arrays per node rather than four) are warm-started across g-iterations and across inner iterations; the node contexts, including the cached factorisations, are reused across successive solves that share the grids (the $\bar{Z}$ sweeps). The threshold array is the one that matters most: it is what makes the capacity static sweep one evaluation per grid point (Step 2).

[Q22 | mynotes | ANSWERED] Record what the capacity Cobb–Douglas path of Proposition 5 costs and buys once it has been solved: how many iterations its $2N$-dimensional static fixed point takes from the path without capacity constraints, how far its rest points sit from that path's at the calibrated capacity intensities, and whether initialising the loop from it saves growth iterations at intensities where that path is a poor guess. Until then the text's claim that it is an optional refinement rests on the path without capacity constraints having converged everywhere it was tried. **[The capacity Cobb–Douglas path: what it costs and what it buys.]** The path of Proposition 5, closed by the static fixed point of Proposition 4, has since been solved. It is cheap; it is *not* needed as the initialiser of Algorithm 1 at calibrated capacity intensities; and it is nevertheless the object two other parts of the pipeline want.

- *How it solves, and what it costs.* Given g, the first two lines of (127) give $(\iota, \mathfrak{q})$ in closed form and the third is a scalar root in the threshold whose left-hand side carries the *bundle* price and not the node's own sales — so the operating positions are pinned by $(g, \mathcal{P})$ alone. The log prices then close as a fixed point in the price moments (damped Picard; the loading of a node's log price on its bundle price at a fixed user cost is $1 - \ell_Y/\ell_\Phi$, which vanishes in the limit without capacity constraints), sales by the Domar-type linear system of Proposition 4 with the effective spending shares $\xi_n(\overline{v})\varrho_n(\overline{v})$, capital by (128), q by adding-up and g by the numeraire. The whole path costs 0.03 s — a scan in g and a bracketed root, each evaluation a handful of Picard passes over N scalar root-finds — against 0.7 s for the same configuration through the general degenerate branch.
- *Accuracy.* Its own equations hold at the returned solution to 4×10^{-16} (user cost, sales, price moment, $\Phi_n(\iota^*) - \delta_n - g$), adding-up exactly, the numeraire to 7×10^{-14} and the GDP identity to 6×10^{-17}; the mean-field pipeline reproduces every node moment of it to 10^{-15}. Its own capacity limit against the algebraic path of the model without capacity constraints is $|\Delta g| = 6 \times 10^{-4}$, 4×10^{-12}, 1×10^{-15} at $x_n \times \{1, 0.1, 0.01\}$.
- *Distance from the rest points without capacity constraints.* At the calibrated intensities ($\mathbb{E}_n[U] = 0.80$ at both nodes) the capacity rest points differ from those by 0.05 in κ, one node below and one above, with $g = 0.0227$ against 0.0233; the thresholds sit about 0.3 log units above their poles, at $\mathbb{E}_n[U] = 0.805$, 0.808, rationed shares 0.144, 0.155 and $\varrho_n = 0.993$. The gap is small, which is exactly why the path without capacity constraints is a good initialiser and the box above found no configuration that needed a fallback.
- *What it is actually for.* Not the initialiser of the CES loop, which never needed it, but (i) the *answer* in the degenerate Cobb–Douglas configuration, where the loop is not run at all; (ii) the provider of grid centres in a continuation, where the rest points of the model without capacity constraints do not move with x_n and its own do — a 240-fold sweep in the capacity intensities ran on its centres with no edge warning at any of 138 solves; and (iii) the target the $\eta \to 1$ check of Step 7 converges to, which is what took that check out of *proposed* under capacity.

D.7 Step 6: Transition Dynamics and MIT Shocks

Aggregate shocks are reintroduced as unanticipated deterministic paths (MIT shocks). The implemented experiment is a temporary shift in a sector productivity level, $\ln \bar{Z}_n(t) = s\, e^{-\nu t}$ on one node; shifts in the kernel $a \to a'$, in β, in the productivity process (θ, ς^Z), and permanent shifts, enter the same code path as time-varying coefficients and are *proposed*. Fix a horizon $[0, T]$ with T large, initial ensemble $\mathfrak{m}(0, n, \cdot)$ from the pre-shock BGP and terminal condition $\mathfrak{q}(T, \cdot)$ at the post-shock BGP; for a temporary shock the two BGPs coincide, for a permanent one the new BGP is solved first by Algorithm 1. The unknowns iterated on are the deterministic paths $\left\{\mathcal{P}(t, \cdot), \widehat{\mathcal{D}}(t, \cdot), q_t, r_t\right\}$; the growth

path g_t is determined inside the two sweeps and consumption $\widehat{C}_t$ by the Euler equation, as follows. The scheme is the standard MFG forward–backward sweep with five closure choices, each of which replaces a more obvious alternative that does not converge (the box below records them).

1. Initialise the aggregate paths at their BGP values (or, for a permanent shock, at an interpolation between the two BGPs).
2. **Backward.** Solve (52) node by node from $\mathfrak{q}(T,\cdot)$ backwards, *fully* implicit in time: at each step the discrete equation (214) with the additional term $-(\mathfrak{q}^k - \mathfrak{q}^{k+1})/\Delta t_k$ (and $-1/\Delta t_k$ added to the diagonal of (215)) is solved by the Newton method of Step 3, warm-started at $\mathfrak{q}^{k+1}$, with the static block of Step 2 solved at the aggregates of date t_k. The growth rate g_k of the interval $[t_k, t_{k+1}]$ is an aggregate that the firms take as given: it enters (52) through the drift of κ, $\Phi_n(\iota^*) - \delta_n - g_k$. It is solved *together with* the valuations of the date. The Newton solve of all nodes is repeated until g_k satisfies the implicit-Euler form of (55),
$$g_k = \frac{\sum_n \mathbb{E}_{n,k+1}\big[(\Phi_n(\iota_k^*) - \delta_n)\kappa\big]\Delta_n}{\sum_n \mathbb{E}_{n,k+1}[\kappa]\Delta_n}, \tag{221}$$
at the policy ι_k^* just solved and the distribution $\mathfrak{m}_{k+1}$ of the previous forward sweep: a scalar root per date, found by the secant method in two or three solves, with the static block of the date not re-solved. Store the policy of each interval, the market consumption $\widehat{C}_k^M = q_k - \sum_n \mathbb{E}_{n,k}[\iota_k^*\kappa]\Delta_n$ against the distribution of the previous forward sweep, and the two aggregates S_k and c_k of item 5.
3. **Forward.** Solve (53) node by node from $\mathfrak{m}(0,\cdot)$ forwards by implicit Euler, with the forward operator $\mathsf{G}_k^{\mathsf{T}}$ of Step 4 rebuilt at each step from the stored policy of the interval. The growth rate is determined again *inside* this sweep, now with the distribution that the sweep itself produces: on the interval $[t_k, t_{k+1}]$, g_k is the value at which (221) holds for the distribution $\mathfrak{m}_{k+1}$ produced with that g_k (a scalar fixed point per step, a few passes), which keeps $\sum_n \mathbb{E}_n[\kappa_t]\Delta_n = 1$ at every step by construction. The two determinations of g_k differ only through the distribution, and coincide at the fixed point of the iteration.
4. **Aggregates and the numeraire.** Recompute $\mathcal{P}(t,\cdot)$, $\widehat{\mathcal{D}}(t,\cdot)$ and q_t from (54) pointwise in t, then rescale $(\widehat{\mathcal{D}}_t, q_t)$ jointly by one Newton step on the numeraire residual of date t, using the analytic elasticity of that residual to a joint rescaling along the scale family of Remark 14 at the fixed distribution $\mathfrak{m}_t$. Along a transition the numeraire is an equation, not a check: the distribution at each date is given by the forward sweep rather than re-solved, so the scale family of Remark 14 is no longer absorbed by adding-up, and at a fixed distribution the consistency map has, in the nominal-scale direction of $(\widehat{\mathcal{D}}, q)$, the eigenvalue $1 - e_\kappa > 1$ for $\eta < 1$ (with $e_\kappa < 0$ the capital elasticity of sales(131)) and exactly 1 under Cobb–Douglas, where q_t is pinned by the numeraire alone. The elasticity is the one place in the transition where the model enters, and it enters through the single array e_S of (210): the box below gives it in a form that covers both models, with the model without capacity constraints the case $e_S \equiv 1/\alpha_n$.

5. **Interest rate.** Integrate the Euler equation (145) backward from the terminal BGP, giving $\ln \widehat{C}^E_t$ from (r_t, g_t), and impose goods-market clearing, $R_k := \ln \widehat{C}^E_k - \ln \widehat{C}^M_k = 0$, as the equation for r_k, with $r_T = \rho + \gamma g_T$ at the terminal date. The Jacobian of R in r is upper-triangular: a change in r_k moves the valuations of date t_k and of every earlier date, and through them investment, growth and both consumption paths. Because g_k is solved with the valuations (item 2), the backward sweep returns the full response of $\mathfrak{q}$ to the rate path, and the Newton step is built on that response, taken for the aggregate. Let y_k be the change of the logarithm of the aggregate marginal value $\bar{\mathfrak{q}}_k$ caused by a change Δr of the path. Then, with $y_T = 0$,

$$y_k = \psi_k \left(\frac{\bar{\mathfrak{q}}_{k+1}}{\bar{\mathfrak{q}}_k}\, y_{k+1} - \Delta t_k\, \Delta r_k \right), \qquad \psi_k = \frac{1}{1 + (r_k - g_k)\Delta t_k},$$
$$S_k = \frac{\sum_n \mathbb{E}_{n,k}[\iota^{*\prime}(\mathfrak{q}_k)\, \mathfrak{q}_k\, \kappa]\Delta_n}{\widehat{C}^M_k}, \qquad c_k = \sum_n \mathbb{E}_{n,k}\big[\iota^{*\prime}(\mathfrak{q}_k)\, \kappa\big]\Delta_n, \qquad \bar{\mathfrak{q}}_k = \frac{S_k\, \widehat{C}^M_k}{c_k}, \tag{222}$$

where S_k is the response of aggregate investment to $\ln \bar{\mathfrak{q}}_k$ per unit of market consumption (so that $\ln \widehat{C}^M_k$ changes by $-S_k y_k$), c_k the response of the growth rate g_k to $\ln \bar{\mathfrak{q}}_k$, and $\bar{\mathfrak{q}}_k$ the mean of $\mathfrak{q}_k$ with the weights $\iota^{*\prime}\kappa$. The residual responds by $\Delta R_k = -V_k + W_k + S_k y_k$, with $V_k = \sum_{m \geq k} \Delta t_m \Delta r_m / \gamma$ the cumulated Euler term and $W_k = \sum_{m \geq k} \Delta t_m c_m y_m$ the cumulated growth response. Requiring $\Delta R_k = -(R_k - R_T)$ at every date, the terminal residual R_T being beyond the reach of any rate, gives the step by one backward recursion from $y_T = V_T = W_T = 0$:

$$\Delta r_k = \frac{R_k - R_T - V_{k+1} + W_{k+1} + \Lambda_k \dfrac{\bar{\mathfrak{q}}_{k+1}}{\bar{\mathfrak{q}}_k}\, y_{k+1}}{\Delta t_k\, (1/\gamma + \Lambda_k)}, \qquad \Lambda_k := (\Delta t_k\, c_k + S_k)\, \psi_k,$$
$$V_k = V_{k+1} + \Delta t_k\, \Delta r_k / \gamma, \qquad W_k = W_{k+1} + \Delta t_k\, c_k\, y_k, \tag{223}$$

with y_k updated by (222). For one representative industry (223) is the exact Newton step; with a cross-section it is an approximation, since the aggregate response stands in for the responses of the individual industries. It shapes the update only: the system is triangular with a non-zero diagonal, so $\Delta r = 0$ exactly when $R_k = R_T$ at every date. The gap $\ln \widehat{C}^E_t - \ln \widehat{C}^M_t$ at the converged path is the horizon check: it is a constant in t inherited from the terminal date and vanishes as $T \to \infty$.

6. Relax the paths $(\ln \mathcal{P}, \widehat{\mathcal{D}}, q)_t$ towards their recomputed values (damped, or Anderson-mixed on the stacked state), take the Newton step on r_t at full weight, and repeat 2–5 to convergence in the sup-norm over the paths. The tolerance is set relative to the sup-norm of the first pass, which measures the size of the disturbance, with an absolute cap; an absolute tolerance alone stops a small shock too early. Check insensitivity to T and to the time grid.

Configurations. The shock is to the sector level $\bar{Z}_n$, which enters Step 2 as a time-varying coefficient and leaves the z-grid untouched, since the initial ensemble sits on the BGP;

a shock to the initial firm-level productivity would instead make $z_n(t)$ a known decaying coefficient of Step 2 (*proposed*). Without capital shocks the forward operator is the transposed backward operator at each t.

The model without capacity constraints. Nothing in the scheme changes: the same forward–backward sweep, the same five closures, the same interest-rate step. Only two inputs depend on the model, the static block of Step 2 called at each date and the elasticity e_t of the numeraire scale step, and the box below gives the second in a form that reduces to that model's at $e_S \equiv 1/\alpha_n$. **[Derivation of the interest-rate step** (223)**.]** The step is derived for one representative industry and then read with aggregates.

- *The valuation step.* With one industry, $\kappa \equiv 1$, and with g_k solved inside the backward step the drift of κ vanishes, $g_k = \chi(\mathfrak{q}_k)$ with $\chi(\mathfrak{q}) := \Phi(\iota^*(\mathfrak{q})) - \delta$. The implicit step of item 2 is then, with π_k the flow of (52) at date t_k,

$$\pi_k - \iota^*(\mathfrak{q}_k) - [r_k - \chi(\mathfrak{q}_k)]\,\mathfrak{q}_k - \frac{\mathfrak{q}_k - \mathfrak{q}_{k+1}}{\Delta t_k} = 0.$$

- *Its linearisation.* By the first-order condition $\Phi'(\iota^*)\mathfrak{q} = 1$, $\chi'(\mathfrak{q})\,\mathfrak{q} = \iota^{*\prime}(\mathfrak{q})$, so the terms $-\iota^{*\prime}$ and $\chi'\mathfrak{q}$ cancel and the derivative with respect to $\mathfrak{q}_k$ is $-(r_k - g_k) - 1/\Delta t_k$. The derivatives with respect to r_k and $\mathfrak{q}_{k+1}$ are $-\mathfrak{q}_k$ and $1/\Delta t_k$. Hence $[1 + (r_k - g_k)\Delta t_k]\,\delta\mathfrak{q}_k = \delta\mathfrak{q}_{k+1} - \mathfrak{q}_k\,\Delta t_k\,\delta r_k$, and in logarithms, with $y_k = \delta \ln \mathfrak{q}_k$,

$$y_k = \psi_k\left(\frac{\mathfrak{q}_{k+1}}{\mathfrak{q}_k}\,y_{k+1} - \Delta t_k\,\delta r_k\right), \qquad y_T = 0.$$

  The ratio $\mathfrak{q}_{k+1}/\mathfrak{q}_k$ comes from the passage from levels to logarithms and equals one only on a stationary path.

- *Growth and the two consumption paths.* $\delta g_k = \chi'(\mathfrak{q}_k)\,\delta\mathfrak{q}_k = \iota^{*\prime}(\mathfrak{q}_k)\,y_k = c_k y_k$. Market consumption is $\widehat{C}^M_k = q_k - \iota^*(\mathfrak{q}_k)$, so $\delta \ln \widehat{C}^M_k = -\iota^{*\prime}(\mathfrak{q}_k)\,\mathfrak{q}_k\,y_k/\widehat{C}^M_k = -S_k y_k$. The Euler equation integrated backward, $\ln \widehat{C}^E_k = \ln \widehat{C}^E_{k+1} - \Delta t_k[(r_k - \rho)/\gamma - g_k]$, gives $\delta \ln \widehat{C}^E_k = -\sum_{m \geq k} \Delta t_m(\delta r_m/\gamma - c_m y_m) = -V_k + W_k$.

- *The residual and the recursion.* $\delta R_k = \delta \ln \widehat{C}^E_k - \delta \ln \widehat{C}^M_k = -V_k + W_k + S_k y_k$, which depends on δr_m for $m \geq k$ only. Substituting the recursions of y_k, V_k and W_k into $\delta R_k = -(R_k - R_T)$ leaves δr_k as the only unknown when the dates are taken backward, and solving for it gives (223). The diagonal coefficient is $-\Delta t_k(1/\gamma + \Lambda_k)$, which is negative, so the step exists and is unique.

- *With a cross-section.* The sums over nodes and the expectations of (222) replace the single industry's terms, and $\bar{\mathfrak{q}}_k = S_k \widehat{C}^M_k/c_k$ reduces to $\mathfrak{q}_k$ for one industry. The drift of κ then vanishes on average and not industry by industry, which is why the step is approximate there.

- *The step it replaces.* If g_k is instead read from the previous forward sweep, the valuations respond to the rate path only in part within a pass. Dropping the growth response ($c \equiv 0$), setting $\psi \equiv 1$ and holding S constant along each

column of the Jacobian gives $\Delta r_k = (R_k - R_{k+1})/[\Delta t_k(1/\gamma + S_k)]$. With g_k solved inside the backward step that step cannot be used: the growth response is the larger part of the Jacobian away from the diagonal.

[Transition closure: what fails, and the formulas, Step 6.] Each item was measured on the two-sector reference economy unless it says otherwise; the alternative named first was implemented first and diverged.

- *Semi-implicit backward step (policy frozen at $\mathfrak{q}^{k+1}$, one linear solve) is unstable.* Its matrix $[(1/\Delta t_k + R_{\text{eff}})\mathsf{I} - \mathsf{A}]$ is an M-matrix only where $1/\Delta t_k + R_{\text{eff}} > 0$, and at the low-$\kappa$ edge of a wide grid R_{eff} is negative and large at the BGP policy itself for any physical time step: at a zero shock the recursion left the BGP by an order of magnitude per backward step from rounding alone. The fully implicit step (one to three Newton steps from the warm start $\mathfrak{q}^{k+1}$, cached factorisation) is exact at a zero shock from the first pass.
- *Updating g_t from the previous iteration's distributions lags one time step per outer pass* (an $O(N_t)$ Gauss–Seidel wave, visible as the residual's argmax marching forward in t). The per-step fixed point above is also the Newton update for g_k, since $\frac{\mathrm{d}}{\mathrm{d}g_k}\sum_n \mathbb{E}_{n,k+1}[\kappa]\Delta_n = -\Delta t_k \sum_n \mathbb{E}_{n,k+1}[\kappa]\Delta_n$ to $O(\Delta t^2)$; two or three passes converge it to rounding.
- *Reading g_t in the backward sweep from the previous forward sweep does not contract under capacity.* This was the scheme first implemented, with the simpler interest-rate step named at the end of the derivation box above.
  - *Why.* g_k enters the valuation through the drift of κ. With g_k held at its old value, a change of the rate path moves $\mathfrak{q}$, investment and growth only in part within a pass (the firms act as if their capital moved relative to the aggregate), and the rest arrives over the following passes. The interest-rate step, built for the response within a pass, does not contract together with that lag.
  - *Where it shows.* Not at the two-sector reference economy with capital shares near one half, where the scheme converged to the tolerance in every run. It shows at small capital shares. On the two-node network with $\alpha_n = 0.083$, $\varepsilon = 10$ and a 0.1 % shock, the damped iteration fell to 2×10^{-6} in 36 passes and then wandered between 10^{-5} and 10^{-4} (highly persistent productivity), or diverged, reaching 8×10^{-3} after 300 passes (the reference economy's productivity process and dynamic parameters). What failed to converge was the pair (r_t, g_t) and the goods-market gap; the three blocks of Θ converged to about 10^{-6}. Anderson mixing jumped erratically, and after a jump the Newton solve of the backward step could fail to converge. The model without capacity constraints at the same parameters converged to 10^{-9}.
  - *What does not cure it.* Damping the interest-rate step (weights one half and one quarter) lowers the floor but is not a cure: at one quarter one of the two cases still stalled at 4×10^{-6}. The step (223) without the consistent growth

path fails at once.

  - *With g_k solved inside the backward step and the step* (223). The same transitions converge to 10^{-9} or better, monotonically under plain damping, for shocks of 0.1 % and 1 %, at within-industry elasticities of 8, 10 and 12, with $r - g$ between 0.012 and 0.034, and with and without capacity. The backward step takes 2.4 to 2.8 Newton solves per date instead of one.
  - *The fixed point is the same.* On the two-sector reference economy the four responses to a $+5$ % shock (either node, with and without capacity), resolved to 10^{-9}, differ from those of the first scheme at its tolerance of 10^{-5} by 1.5×10^{-4}, 2×10^{-8}, 2×10^{-6} and 6×10^{-7} relative. The first scheme, run again, reproduces its own numbers to eleven digits.
  - *Open.* Why the model without capacity constraints is not affected has not been established.

- *An absolute tolerance stops a small shock too early.* At a tolerance of 10^{-5} on the sup-norm of the update, a 0.1 % shock stopped at pass 25 with its peak response 1.3 % off. The damped map contracts at about 0.9 per pass, so the error left is about ten times the last update. With the tolerance at 10^{-4} of the first pass the peak response is within about 10^{-4} relative of its converged value.

- *The numeraire scale step, in the form that covers both models.* Summing the node demand relations $\widehat{m}_{S,n} = \beta_n \mathcal{N}_n q + \widehat{\mathcal{D}}_n$ against Δ and using $\sum_n \widehat{\mathcal{D}}_n \Delta_n = \sum_n \widehat{m}_{X,n} \Delta_n$ (row-stochasticity) gives $\sum_n [\widehat{m}_{S,n} - \widehat{m}_{X,n}] \Delta_n = q\,(P^f)^{1-\eta}$, so the raw consistency map returns $q^{\text{new}} = q\,(P^f)^{1-\eta}$ at a fixed distribution and the eigenvalue $1 - e_\kappa$ of item 4 follows; iterating without the rescaling diverges at a rate of about 1.3 per pass with r and g frozen. (Without capacity constraints $\widehat{m}_{X,n} = \xi_n \widehat{m}_{S,n}$ and the left-hand side is $\sum_n [1 - \xi_n] \widehat{m}_{S,n} \Delta_n$; the identity is the same.) The implemented step: with $\nu_t = \ln\left(P_t^f\right)^{1-\eta}$ (CES) or $\nu_t = \sum_n \beta_n \mathbb{E}_{n,t}[\ln P] \Delta_n$ (Cobb–Douglas), and the elasticity of ν_t to a joint rescaling $(\widehat{\mathcal{D}}, q) \mapsto u(\widehat{\mathcal{D}}, q)$ at fixed κ (a rescaling acts as $\kappa \mapsto \kappa / u$, so the elasticity is $-\partial \ln P / \partial \ln \kappa$ averaged with the weights of the numeraire),

$$\epsilon_t = \frac{1}{(P_t^f)^{1-\eta}} \sum_n \beta_n \, \mathbb{E}_{n,t} \left[\frac{(1-\eta) P^{1-\eta}}{e_S - (1-\eta)\vartheta} \right] \Delta_n \quad \text{(CES)},$$

$$\epsilon_t = \sum_n \beta_n \, \mathbb{E}_{n,t} \left[\frac{1}{e_S} \right] \Delta_n \quad \text{(Cobb–Douglas)},$$

  rescale the pair $(\widehat{\mathcal{D}}_t^{\text{new}}, q_t^{\text{new}})$ by the factor $\exp(-\nu_t / \epsilon_t)$ under Cobb–Douglas, and by $\exp[-\nu_t (1 + 1/\epsilon_t)]$ under CES, the extra factor undoing the $(P^f)^{1-\eta}$ that the raw map already applied to q. Both expectations are sums over the node's own e_S array (210) and its own $\vartheta = \beta_n P^{1-\eta} q / \widehat{s}$; at $e_S \equiv 1/\alpha_n$ they collapse to that model's $\epsilon_t = \sum_n \beta_n \mathbb{E}_{n,t}[(1-\eta)\alpha_n (1 - e_\kappa) P^{1-\eta}] \Delta_n / (P_t^f)^{1-\eta}$ and $\epsilon_t = \sum_n \beta_n \alpha_n \Delta_n$, since $\alpha_n (1 - e_\kappa) = 1 / [e_S - (1-\eta)\vartheta]$ there.

- *The Cobb–Douglas loading is $1/e_S$, not $1 - 1/e_S$.* The tempting error is to use the

loading of the node's log *price* on its *bundle* price, $1 - 1/e_S$ (Proposition 4(ii)), which is $1 - \alpha_n$ without capacity constraints. The numeraire scale step rescales $(\widehat{\mathcal{D}}, q)$, that is nominal *sales*, whose loading is $1/e_S$ — α_n without capacity constraints, which is the form the transition parity check confirms to 2×10^{-11}. The two differ by exactly $1 - 2\alpha_n$ of that model, so at α_n near one half the error is nearly invisible in a single pass and shows up only as a slowly diverging path.

- *The literal Euler update is backward-unstable.* Solving $r_t = \rho + \gamma(g_t + \frac{\mathrm{d}}{\mathrm{d}t} \ln \widehat{C}_t)$ with the market residual $\widehat{C}^M$ differentiated numerically along the path amplifies an error in $\widehat{C}_k$ by a factor of roughly $\gamma\, \iota^{*\prime} \mathfrak{q} / \widehat{C}$, close to ten, at $k - 1$ (r_{k-1} moves with $\widehat{C}_k$, feeds $\mathfrak{q}^{k-1}$, ι_{k-1} and $\widehat{C}_{k-1}$), whether r sits inside the sweep or is relaxed outside. The stable orientation integrates the Euler equation backward for $\widehat{C}^E$ given r, with no amplification, and uses clearing as the residual; the term S_k in the diagonal of the Newton step is what damps the $r \to \mathfrak{q} \to \iota \to \widehat{C}$ feedback, and with S and c both zero the step is exactly the Euler equation solved for r. The implementation clamps r_t to a wide band against transient overshoots.
- *Relaxation.* Plain damping of the Θ paths (weight one half on the old paths at the reference economy; heavier damping stalls) with the r-step at full weight is the robust default. Anderson mixing on the stacked (Θ, r) state is kept as an option, with the same positivity restart as in Step 5. Under the first scheme it converged faster on a large shock and stalled on a small one. With g_k solved inside the backward step it is reliable and about halves the count on the two-node network at small capital shares (39 to 56 iterations against 85 to 126 damped); on a network of sixty-five nodes it gained little (37 iterations).
- *Horizon.* The Euler gap at the converged path is a constant offset when the distribution has not returned to the BGP by T: lengthen the horizon, and confirm that the impact and peak responses are insensitive to T and to the number of time steps (the standing check of Step 7). The time grid is stretched towards $t = 0$.
- *What capacity costs and does to a transition.* The machinery is exact under capacity from the first pass: at a zero shock the forward–backward sweep reproduced the balanced growth path to 2×10^{-8} in every configuration, which is the check that the seven-array static cache and the generic scale step are consistent with the stationary solver. Capacity *dampens* the impact of an upstream productivity shock, and the sign is not an artefact of the shock or of the demand system. The two boxes below record the transition itself and the grid it runs on; the closing box of this appendix records what the dampening turns out to be.
- *How the outer iteration converges, and what it costs.* The numbers of this item, and every iteration count and Walras check in the two boxes that follow, were measured with the first scheme (g_t read from the previous forward sweep, an absolute tolerance of 10^{-5}). The responses stand, by the comparison above. The counts do not carry over: with g_k solved inside the backward step, the four $+5\,\%$ transitions of the two-sector reference economy take 35, 33, 30 and 36 iterations in place of 34, 113, 39 and 50, and the goods-market gap of the upstream capacity

run is 10^{-6} in place of 10^{-3}. The forward–backward sweep converges *linearly*, at a rate of about 0.78 per iteration: a run that shows nothing for minutes is usually converging, not stuck, and the way to tell is a small run with the residual printed, not a longer wait. Over the thirty-eight transitions measured here the outer count ran from 22 to 117, and the expensive ones are systematic rather than random: a shock on the *downstream* node took 113 iterations against 34 for the same shock on the upstream node of the same economy. Profiled at the converged paths on a two-node economy, one iteration costs 1.3 s and splits as forward Kolmogorov sweep 67 %, static threshold sweeps 23 %, backward Newton steps 0.5 %, and 10 % for the moments, the coupling, the numeraire and the interest-rate step: *the capacity block is not what makes a transition expensive, the forward sweep is*. Two cautions on that split. The backward step is nearly free only because the continuation value is warm-started from the converged one, so early iterations pay more; and the split has to be taken on an otherwise idle machine, because two concurrent jobs inflated the same transition elevenfold and the outer share, being measured as a residual, absorbed all of it.

- *The time grid governs whether a transition converges; the state grids govern only what it costs.* These are two different roles of "the grid" and conflating them wastes runs. Cutting the number of time steps from 100 to 60 sent every run to the iteration cap without reaching tolerance, at every state grid tried; cutting N_κ from 101 to 61, or N_z from 11 to 7, left the iteration count at exactly 34 and made the transition three to six times cheaper. *And the non-converged runs still reproduced the converged peak response to* 0.015 %. Agreement of a response across grids is therefore not evidence that a grid is usable, and a convergence test that reads only the response will certify a grid that never converged: read the iteration count and the residual first, the response second.

[Q23 | mynotes | ANSWERED] Record the two-dimensional capacity transition once it has been run: wall time and iteration count against those of the model without capacity constraints, the three Walras residuals along the path, and whether the dampening of the impact response seen in the capital-only test survives the productivity-only configuration — which is the configuration the experiments use. Until then the only capacity transition measured is one-dimensional, and the sign of the impact effect should be treated as provisional. **[The two-dimensional capacity transition, measured.]**

- *Cost and convergence.* On a two-node calibrated economy the whole transition is 44 s in 35 outer iterations, 1.3 s per iteration, against the 40 of its twin without capacity constraints; the split is the one in the box above. Neither model is systematically the slower: which of the two takes more iterations depends on the economy and on where the shock lands.
- *Residuals along the path.* On the grid the box below records, the three Walras checks are the numeraire at 1×10^{-5}, adding-up at 9×10^{-6} and the Euler gap at 5×10^{-4} (a horizon offset, constant along the converged path). On a wide coarse grid the same three sit between 10^{-3} and 10^{-1}, which is the clearest single

symptom that a grid is not resolving the distribution.

- *The dampening survives the productivity-only configuration, and is not an artefact.* On the paper's own configuration — CES, productivity only, a $+5\,\%$ shock to the most upstream node — the peak response of $\ln q$ is 0.0407 under capacity against 0.0531 without capacity constraints, a ratio of 0.77 at $N = 2$ and 0.71 at $N = 3$ on hand-calibrated economies, and 0.79 and 0.74 on calibrated ones. The Cobb–Douglas pair gives 0.76, so the sign is not a CES artefact. Aggregate utilization rises about 2.1 percentage points on impact from a level of 0.81, and the peak utilization move per unit of peak log GDP move is about one half. The mechanism, and the rest of the magnitudes, are in the closing box of this appendix: they are economics, not numerics.

[Q24 | mynotes | ANSWERED] Record the coarse time and state grid chosen for the two-dimensional capacity transitions and the convergence test that chose it: which responses were held to what tolerance as N_κ, N_z and the number of time steps were refined, and how much wall time the coarse grid saves per transition. The text's claim that a deliberately coarse grid is the right tool rests on that test. **[The grid the two-dimensional capacity transitions run on, and the test that chose it.]**

- *The grid.* $N_\kappa = 101$ on a half-width of 0.5 log units, $N_z = 11$, 100 time steps. It is coarse in points and *narrow* in window, and the window is the part that matters (Step 1).
- *What the test found, and what it did not.* Five grids were compared on the same shock. All five agreed on the peak responses to better than 0.01 %, an order of magnitude inside the tolerance the test was set, so *accuracy did not decide*: convergence and resolution did. A wide coarse grid (76 points on ± 2, 60 time steps) did not converge at all under either model, running to the iteration cap at a residual of 2×10^{-3} against a tolerance of 10^{-5}, in sixteen minutes. A wide fine grid (151 points on ± 2) converged and was accurate, but resolved the stationary distribution over 1.3 cells at the same wall time at which the chosen grid resolves it over 2.8.
- *What the coarse grid saves, and what a further cut costs.* Refining the chosen grid moves the peak response by 7×10^{-6} relative in N_κ and 7×10^{-8} in N_z. Cutting it further — $N_\kappa = 61$, $N_z = 7$, the time grid untouched — costs 0.002 % on the reported response and takes the transition from 50–80 s to 14 s, which is what makes a sweep of two dozen transitions affordable; the cut was validated against the uncut runs on the overlapping cases rather than assumed. Its one documented side effect is the fallback of Step 4, the distribution being narrower than a cell there, which is why a figure about the *shape* of the distribution is not drawn on the cut grid.

[Q25 | tag: hide | ANSWERED] subsection "Step 7: Validation"

D.8 Step 7: Validation

[Flag: set aside on 5 November 2026] This subsection is held, unchanged and with its heading, inside a maroon box until what to do with it is decided. Its black-text mentions, in the implementation item of the numerical appendix's conventions, in Step 1 (the checks by widening the grid and of the z-marginal) and in Step 2 (the closed-form elasticities), are in maroon boxes too, with the claims left black; its other mentions are in private boxes.

The checks are listed from the static layer outward; each re-derives the object it tests by a route independent of the solver. All are implemented, run from one call on the converged solution, except those marked *proposed*. The first four are the capacity block's own checks and are skipped under the mapping to the model without capacity constraints, where the objects they test are constants; the rest are common to the two models.

1. **The tail functions.** The closed forms (60) are checked against quadrature of their defining integrals in $\ln v$, and their log-derivatives — including the two the solver uses, ℓ_U and ℓ_Υ — against central differences of $\ln U_n$ and $\ln \Upsilon_n$; the two displayed forms of Γ_n (64) and of Υ_n (106) agree; and at a threshold far above the pole every tail object attains its constant of Lemma 4(i) for the model without capacity constraints. The last of these is the check that the capacity block and the block without capacity constraints are the same code path evaluated at different thresholds.

2. **Static identities at the solved node block.** The node equation's own residual sup $|H|$ of (208) relative to its leading term; the demand relation (141), which the solver never imposes directly because it solves the eliminated equation; the two forms of the flow (115), one through the share functions and one through $\Phi_n(\overline{v})$, which share no intermediate quantity; the rent identity $\xi_n + \zeta_n[1-\xi_n] = 1 - \frac{1}{\varepsilon}$ (77) and the four-way split of sales (78), both of which hold at every threshold and not only at the equilibrium one; and the eight elasticities of the solved block in $(\ln \kappa, z)$, (129)–(130), against central differences on the grid, which is the strongest check on the static block because it tests the root map rather than the algebra at a root.

3. **The capacity limit $x_n \to 0$.** Solving the same economy at x_n scaled down towards zero must reproduce the solution without capacity constraints (Lemma 4): g, q, the node distributions, the Domar weights and $\mathbb{E}_n[\kappa]$ converge to those of that model, the rationed share $\mathbb{E}_n[1 - G_n(\overline{v})]$ and the rationing rent share $1 - \mathbb{E}_n[\varrho_n]$ vanish, and node utilization $\mathbb{E}_n[U]$ falls to zero. This is the one check that exercises the entire capacity pipeline against an independently coded model, and it is the numerical form of the statement that the model without capacity constraints is a limit of the capacity model rather than a different model.

4. **Cobb–Douglas under capacity.** At $\eta = 1$ the threshold, sales, firm revenue, spending, the shares and the flow must be exactly free of z, and $\ln P + z$ and

q must be z-free, node by node and grid point by grid point (Lemma 17, Theorem 1); the sales moment must equal $\beta_n q + \widehat{\mathcal{D}}_n$ without integration; and Cobb–Douglas capacity solves must be invariant to the productivity variance ς_n^Z, which is the neutrality theorem in computable form.

5. **The static closed forms without capacity constraints** (under that mapping). The solved node block reproduces the elasticities (195) against central differences; the Cobb–Douglas initialiser, which both models use, satisfies $\sum_n [1 - \xi_n] h_n \Delta_n = 1$ and $\widehat{\mathcal{D}}_n = (h_n - \beta_n) q$ (Corollary 2).

6. **The z-discretisation.** The numerical z-marginal of every node distribution equals the analytic stationary distribution of the reflected Ornstein–Uhlenbeck process on $[-\bar{z}, \bar{z}]$, up to the order of the stencil; its variance matches the analytic truncated value.

7. **Operator structure.** A and G have zero row sums and nonnegative off-diagonals; the diagonal-dominance margin of the implicit matrix M equals $1/\Delta\tau_{\text{eff}} + R_{\text{eff}}$ row by row, and goes negative at a fixed step where $R_{\text{eff}} < -1/\Delta\tau_{\max}$ (the failure the cap (218) prevents).

8. **Reassembled residuals.** The residual of (157) on the grid with the policy and operator rebuilt from the stored q, and the exact policy evaluation $\mathsf{q} = (R_{\text{eff}}\mathsf{I} - \mathsf{A})^{-1}(\pi_k - \iota^*)$ at the converged policy; FPK mass and sup $\left|\mathsf{G}^{\mathsf{T}} p\right|$; the stationarity of $\mathbb{E}_n[f]$ for $f \in \{\kappa, z, \kappa^2, z^2, \kappa z\}$ under G; adding-up; the numeraire; the GDP identity; the diagnostic (220).

9. **Degenerate limit $\eta \to 1$.** The solution must converge, by continuation to η close to 1 rather than at $\eta = 1$ (Remark 8), to the rest-point system of the model being solved, with $(g, q, \mathbb{E}_n[\kappa], \mathbb{E}_n[\mathsf{q}])$ agreeing to $O(\eta - 1)$ and the z-loadings of $\widehat{s}$ vanishing linearly at the rates (131): the algebraic system (194) closed with (192) under the mapping to the model without capacity constraints, and under capacity the node-wise system (127) closed by the static fixed point of Proposition 4, which is solved directly and is the target the continuation is measured against.

10. **Transition consistency.** At a zero shock the forward–backward sweep of Step 6 must reproduce the BGP as an exact fixed point from the first pass; along a shocked path the three Walras checks (adding-up, the numeraire, the Euler gap) must be at tolerance, and the impact and peak responses insensitive to T and to the number of time steps.

11. **Grid and node refinement.** Double N_κ and N_z (preserving parity) and widen the κ-grid; aggregates should move within tolerance, and the edge mass of Step 1 should stay at rounding. Refinement in N at a fixed kernel (Step 0) is *proposed*.

12. **Finite-J back-test (ε-Nash in practice).** For small J solve the finite model exactly (with a finite-sector solver, which is not part of this paper) with block-constant primitives, and compare (i) the finite policy functions along simulated paths with the GMFG policies at the corresponding nodes, and (ii) welfare from using GMFG policies inside the finite economy. The gap should shrink in J at the

conjectured rate(Section 3.4); its aggregate-risk component can be isolated by shutting the shocks of all but one sector. *Proposed.*

13. **Multiplicity probe.** Restart the inner loop at the converged g from perturbed aggregates Θ and cold $\mathfrak{q}$ caches, and report the deviation of the re-converged fixed point(Section 3.4). *Proposed.*

[Validation knobs and cautions, Step 7.]

- Knobs: the degenerate limit is run by continuation at η just below one, never at $\eta = 1$; the grid check refines to $2N_\kappa - 1$ and $2N_z - 1$ and widens the κ-grid by one log unit; the finite-J back-test envisages $J = 2$–4.
- *Multiplicity.* No uniqueness result exists (Section 3.4); the multiplicity probe is not implemented. Report multiplicity if it is seen; do not assume it away.
- *The capacity static elasticities do not reach rounding, and that is correct.* Under capacity the sales elasticity in $\ln \kappa$ varies with the state, so the central-difference check of the static block has an $O(h^2)$ truncation gap: 3×10^{-4} for sales and 2×10^{-4} to 6×10^{-4} for the threshold, price and utilization elasticities on a 301-point κ-grid, falling by exactly the factor four the second order predicts when the grid is refined to $2N_\kappa - 1$. The 3×10^{-7} of the model without capacity constraints on the same check reflects an elasticity that is nearly constant in $\ln \kappa$, not a better solver. Do not tighten a tolerance to chase it.
- *The capacity limit as measured.* Scaling the calibrated capacity intensities by 1, 0.3 and 0.1 moved g from that of the model without capacity constraints by 8×10^{-4}, 2×10^{-6} and 3×10^{-11}, the last being the solver's own tolerance floor, and the rationed share $\mathbb{E}_n[1 - G_n]$ fell 0.15, 2×10^{-4}, 10^{-9}; the log-log slope of $|\Delta g|$ against the scale is 5 and then 10, the accelerating approach a Gaussian tail produces, not a fixed power. Mean utilization falls like a power over the same range (0.80, 0.47, 0.27, 0.15, 0.086), so utilization is the slow object and the rationed share the fast one: a run at a tenth of the calibrated intensity is already the model without capacity constraints for every aggregate while still showing 9% utilization. Below that the check reads the tolerance floor and nothing else.
- *Other capacity magnitudes at the two-node calibration*, for comparison on a later run: static identities and the demand relation at 7×10^{-13}; the two forms of the flow at 3×10^{-15}; the rent identity and the four-way split at 2×10^{-16}; z-freeness under Cobb–Douglas exact for the threshold, sales and spending and at 10^{-15} for $\ln P + z$ and $\mathfrak{q}$; the zero-shock transition fixed point at 2×10^{-8} from the first pass; grid refinement moving g by 7×10^{-5} at $2N_\kappa - 1$, 6×10^{-9} at $2N_z - 1$ and 4×10^{-5} for one more log unit of width, with edge mass never above 10^{-16}. A cold capacity solve of the two-node CES economy on 301×31 grids took about 1.5 times the wall time of the model without capacity constraints, and 0.05 s warm.
- *The capacity limit along a continuation, and its rate.* Read at four or five scales the limit is a pass or a fail; read along a twenty-six-point continuation in the capa-

city intensities it is a *rate*, and the rate is the informative object. The measured local slope of $\ln|\Delta g|$ against $\ln s$, s the scale on x_n, runs 9.7, 8.6, 7.5, 6.1, 4.4 as s falls from 0.50 to 0.25 — *steepening* as the scale falls, faster than any power law, which is the Gaussian tail of the log-normal shares. Below $s \approx 0.03$ the capacity economy reproduces its own limit without capacity constraints to 1×10^{-8} in g and 2×10^{-7} in $\ln q$, which is the solver's floor and not a property of the limit.

- *The same checks on the calibrated economies.* Rerun on the two- and three-node economies used for the transitions (on the narrow window): the scheme residual at 3×10^{-14}, the exact policy evaluation at 3×10^{-15}, the forward-equation identities at 6×10^{-17} with κ-edge mass 10^{-18} and the state drift pointing *inward* at both edges at every node, tail closed forms against quadrature 6×10^{-14}, (ℓ_U, ℓ_Y) against central differences 8×10^{-10}, the static elasticities 4×10^{-5} in $\ln\kappa$ and 7×10^{-5} in z, and the $\eta \to 1$ continuation $|\Delta g| = 1 \times 10^{-6}$ at $N = 2$ and 3×10^{-6} at $N = 3$, of order $\eta - 1$. Several of these are better than the figures above because the narrow window resolves the distribution far better than the default one, not because anything in the solver changed.
- *A warning is not a diagnosis.* Two things make a run's warnings misleading. They are written unbuffered while ordinary output is buffered, so in a redirected log they all appear at the end whatever produced them; and in a sweep most of them come from the intermediate evaluations of an outer loop rather than from the solutions actually reported. The remedy is to re-solve the reported configurations one at a time and *measure the object the warning names*. Done for a twenty-eight-transition sweep that had printed nineteen edge warnings: every reported baseline had the state drift pointing inward at both κ-edges at every node, by at least 0.12, with κ-edge mass below 5×10^{-18} — no reported result rested on a warned solve.
- The expected magnitudes of every check are recorded, and kept current, with the reference implementation; the numbers above are the ones this appendix needs to be read against, not a substitute for that record.

[Q26 | mynotes | ANSWERED] Record how the capacity solver behaves at $N = 3$: growth iterations and inner iterations against the two-node case, wall time, whether the initialiser without capacity constraints still converges without a fallback, and whether any node's threshold settles close enough to its pole for the cancelling margin of Step 2 to matter. Every capacity magnitude quoted in this appendix is from a two-node economy. **[$N = 3$ under capacity.]** The three-node capacity economy has since been solved, and the answer is that nothing about the solver changes with N at these sizes.

- *Iterations.* Four growth iterations, as at $N = 2$, in every one of the four capacity configurations of {CES, Cobb–Douglas} at $N = 2$ and at $N = 3$ — one for the degenerate branch, which is algebra — with the inner counts unchanged. *The initialiser without capacity constraints still converges without a fallback at $N = 3$*, which is the question the box above left open.
- *Wall time.* A cold three-node solve is 0.19 s on the narrow grid at $N_\kappa = 101$, and a

warm one 0.012 s; at these sizes the warm cost is dominated by fixed overheads rather than by N or by the point count, so a calibration of a hundred-odd solves is a matter of seconds.

- *Thresholds and the pole.* No node's threshold settled anywhere near its pole. The rationed shares at $N = 3$ are 0.145, 0.155, 0.157, the same range as at $N = 2$, where the thresholds sit about 0.3 log units above their poles; the shape checks of Step 2 warned at no node of any configuration; and the static residuals stay at 7×10^{-13} throughout, which is the operative evidence, since a threshold approaching its pole degrades exactly that residual first (the cancelling margin of Step 2).
- *What does change with N* is not the numerics but the economics: the dampening of a shock is stronger at $N = 3$ than at $N = 2$ and differs across the nodes, which is the subject of the closing box.

[Flag: what the solved model says — economics, not numerics, and not yet paper text.] The items below are *results*, not settings, and they do not belong in a numerical appendix. They are recorded here because this is where the implementation that produced them is described, and they should be *promoted* into the paper's own voice — a quantitative section that does not yet exist, with the "Established" paragraph of Section 3.4 updated to match. Every number is a measured statistic of a calibrated two- or three-node step-kernel economy under CES demand with productivity shocks only, capacity intensities and sector productivities chosen so that $\mathbb{E}_n[U] = 0.80$ at every node with a balanced growth rate of 2 % and given Domar weights, the shock a temporary $+5$ % sector-TFP impulse unless stated. *None of it is a quantitative claim about data*: the input–output table is a stylised two- or three-block kernel, and the one parameter that scales the whole mechanism is not yet measured (last item).

1. *Capacity dampens the propagation of a positive sector-TFP shock, and the damping is a gradient in network position.* The peak response of $\ln q$ under capacity, relative to the same economy with the capacity block switched off, is 0.79 for a shock to the upstream node and 0.98 for the same shock downstream at $N = 2$; at $N = 3$ it is 0.74, 0.90, 0.98 from the most upstream node to the most downstream. *Capacity matters for aggregate fluctuations to the extent that it binds upstream.*
2. *The utilization margin absorbs about half of an upstream shock and almost none of a downstream one.* Peak $|\Delta \overline{U}|$ per unit of peak $|\Delta \ln q|$ is 0.52 upstream against 0.05 downstream at $N = 2$, and 0.56, 0.27, 0.03 at $N = 3$; aggregate utilization rises 2.2 percentage points on impact for the upstream shock and 0.2 for the downstream one. A downstream shock raises demand for its own output, not for the capacity of the sectors above it.
3. *Capacity also shortens a shock's reach.* Measured on real node activity, $\ln \mathbb{E}_n[\widehat{s}] - \mathbb{E}_n[\ln P]$ relative to the shocked node's, the downstream node picks up 0.39 of an upstream shock against 0.49 without capacity constraints at $N = 2$, and 0.31, 0.27 against 0.43, 0.38 at $N = 3$.

4. *The response is asymmetric in the sign of the shock; that of the model without capacity constraints is not.* The peak response per unit shock is exactly constant in a linear economy. Under capacity, upstream, it falls from 0.894 at $+1\,\%$ to 0.831 at $+10\,\%$ and rises to 0.966 at $-10\,\%$: an asymmetry of -0.150 at $\pm 10\,\%$ against the control's -0.011 without capacity constraints, and growing with the shock (-0.015, -0.076, -0.150 at ± 1, ± 5, $\pm 10\,\%$). Downstream, capacity adds essentially nothing (-0.017 against -0.015). The ratio of item 1 is therefore a function of the sign and the size of the shock as well as of where it lands.

5. *The mechanism is the node's own supply curve, and it is read off e_S.* The utilization share of a movement along a node's supply curve, $[e_S - 1]/e_S = \ell_U/[\ell_U + \ell_Y]$, is 0.350, 0.365, 0.379 at the calibrated nodes against $1 - \alpha_n = 0.50, 0.55, 0.60$ without capacity constraints — a ratio of 0.70, 0.66, 0.63, which is where the aggregate ratio of item 1 sits. The accounting identity $\Delta \ln(\text{real activity}) = \Delta \ln \mathbb{E}_n[U] + \Delta \ln \bar{Z}_n + \Delta \ln \mathbb{E}_n[\kappa]$ closes to 4×10^{-5} under capacity and leaves a residual of 5×10^{-2} without capacity constraints, where the utilization term is identically zero: without a ceiling the response has no capacity-side account at all.

6. *In levels, at identical primitives*, capacity lowers the balanced growth rate by 0.04 percentage points, GDP per unit of capital by 2.8 %, consumption by 3.6 % and welfare by 4.2 %, and raises the investment share of GDP by 2.9 % of itself (from 0.223 to 0.230), at $N = 2$; and it moves the network's mass away from the node that is expensive to fill, the upstream Domar weight falling from 0.587 to 0.543.

7. *The mechanism is a threshold in utilization, not a gradient.* Along a continuation in the capacity intensities, $|\ln(q/q^{\text{nc}})|$, with q^{nc} the GDP of the same economy without capacity constraints, first reaches $0.1\,\%$ at $\overline{U} = 0.66$, $1\,\%$ at 0.75 and $5\,\%$ at 0.83: almost nothing happens below $\overline{U} \approx 0.65$, and past the calibrated point the cost rises steeply — at $\overline{U} = 0.90$ the growth rate is $1.2\,\%$ against $2.0\,\%$ without capacity constraints and welfare is a third below it. A *little* capacity is mildly expansionary, g sitting 5×10^{-5} above the model without capacity constraints around $\overline{U} = 0.71$, which was checked on four grids and reproduces to four significant digits.

8. *At equal aggregate utilization an upstream bottleneck dampens most.* Recalibrating one node to $\mathbb{E}_n[U] = 0.88$ and loosening the others until $\overline{U}$ returns to 0.80, the ratio of item 1 is 0.66, 0.71, 0.74 as the bottleneck moves from the most upstream node to the most downstream at $N = 3$ — and the downstream-bottleneck economy has 22 % *more* firms rationed while damping less, so this is not a level effect. Spillover to the last node separates the two mechanisms, 0.34 with the bottleneck upstream against 0.13 with it downstream: a downstream bottleneck blocks transmission, an upstream one blunts the source.

9. *Capacity needs a network to matter, but not too much of one.* Sweeping the off-diagonal input–output intensity at fixed $\mathbb{E}_n[U]$ and fixed g, the ratio of item 1 is U-shaped in network density: 0.87 in the disconnected economy, a minimum of 0.76 below the calibrated density, 0.79 at it and 0.84 at 1.6 times it. In the dis-

connected economy the shock barely travels and the capacity margin is barely engaged.

10. *Business-cycle second moments, without ever solving the model with aggregate risk.* Reading the N sector-TFP impulse responses as the moving-average kernels of AR(1) sector productivities, capacity lowers the volatility of GDP by 12 %, of consumption by 17 % and of investment by 10 %, and *raises* price volatility by 8–11 %: it converts quantity fluctuations into price fluctuations. What does not depend on the assumed innovation size — every correlation, autocorrelation and ratio of standard deviations is invariant to a common scale — includes $\mathrm{sd}(\overline{U})/\mathrm{sd}(\ln q) = 0.31$ and $\mathrm{corr}(\ln q, \overline{U}) = 0.73$; and the one thing that does scale inverts cleanly, a time-series standard deviation of one percentage point in a sector's utilization implying sector-TFP innovations of 1.4 % a year. The persistence is not free: it is the decay at which the impulse responses were computed.

11. *Everything above scales with the within-industry dispersion of firm utilization, and that number is not yet a datum.* It is the second capacity moment of the calibration, a function of $(\sigma_{v,n}, \overline{v}_n)$ alone through $\Psi_{n,1}$ and $\Psi_{n,2}$. At a dispersion of 0.14 the upstream ratio of item 1 is 0.79, 13 % of firms are rationed and welfare is 4.2 % below that of the model without capacity constraints; at 0.08 they are 0.94, 1.2 % and 0.3 % — *even though aggregate utilization moves more* (+3.4 percentage points on impact against +2.2), because a utilization margin that moves is not the same thing as a capacity constraint that binds. Over a 2.4-fold range of that target, with $\overline{U}$ pinned at 0.80 throughout, the constrained share moves by a factor of a hundred and the rationing rent by a factor of several hundred. *The level of utilization says almost nothing about how much capacity binds; the dispersion says everything* — which is at once the sharpest prediction of the model and the reason the empirical side has to pin that moment before any magnitude above is quoted.

Table 7: Which layer breaks under which network pathology.

Network pathology	Bounded kernel limit exists?	Sectional convergence?	Caines–Huang-type theorems apply?	This model's static block?	This model's uniform bounds?
Sparse (bounded suppliers per buyer)	no (rung 1)	no	no	no along the sequence	no
Heavy-tailed weights with unbounded column masses	no: the limit is unbounded	open	no (a graphon is bounded by definition)	yes a.e., if entropy stays finite (rung 4)	no (rung 3)
Quasi-random rows (Example 1)	yes (uniform kernel, in the cut metric)	no (rung 2)	no ((H11) fails)	yes for the limit kernel, which however describes no node	yes
Dominant hub (Domar weight not vanishing)	no bounded limit (rung 3)	no (the hub's column is a diverging spike)	no	yes a.e.	no; a granular agent
Near-atomic sourcing rows	yes if the density stays bounded; otherwise entropy may be $-\infty$	open	open	no where $H(i) = -\infty$ (rung 4)	no if $\bar{a}$ fails
Self-loops with non-vanishing weight	only with the explicit self-loop extension (atom by design, rung 5)	open	not covered (no self-loop atom in their setting)	yes, with one extra scalar fixed point per node	yes, in the extended specification

Notes: Bases: rungs 1–5 of Appendix B.2.2, the Caines–Huang hypotheseslisted in Appendix B.9.2, and the zero-graphon limit of sparse sequences (Lovász, 2012). "Open" marks cells not established either way here.

Table 8: Borrowable versus to-be-developed, relative to Caines and Huang (2021).

Item	Status here	Why
GMFG equation-system format: one backward and one forward equation per node, coupled through kernel integrals of the ensemble	borrowed as architecture (Definition 6)	the format carries over verbatim; the coupling here runs through two moment functions and one scalar
Ensemble/consistency formulation; the sup-over-nodes fixed-point viewpoint	borrowed as architecture	the natural notion of "consistency at every network position at once"
ε-Nash programme (limit policies as approximate equilibria of the finite economy)	borrowed as a programme, not as a theorem	see the Theorem 4.4 row
Theorem 3.12 (existence and uniqueness)	does not apply	(H1)–(H2) fail (unbounded payoffs and control); finite horizon versus stationary BGP; drift coupling versus payoff/GE coupling; no contraction constants available
Theorem 4.4 (ε-Nash)	does not apply	in addition, the finite counterpart here carries aggregate risk (the J sector shocks do not wash out), so any honest statement must include an $O(J^{-1/2})$ aggregate-fluctuation term with no analogue in their setting; and the step kernel is not their (H9) regime
Exact law of large numbers on a Fubini extension	cited, not developed	Sun (2006), Sun and Zhang (2009); the graphon-interacting case is Bayraktar et al. (2023); the kernel-weighted form is argued in Lemma 13
Existence of the stationary step-kernel equilibrium	to develop (the existence conjecture, Appendix B.9)	a finite-dimensional fixed point over (Θ, g) plus N two-dimensional PDE pairs; a Schauder or degree-theoretic argument on the moment map is the natural route
ε-optimality of step-kernel policies in the finite J-sector economy	to develop (the ε-optimality conjecture, Appendix B.9)	both error sources must appear: within-block granularity and aggregate fluctuation
Uniqueness	open; probed numerically	no monotonicity structure in the sense of Lasry and Lions (2007): the mean field enters through general-equilibrium prices

Table 9: The three models side by side.

	Finite (J sectors)	**General-kernel GMFG**	**Step-kernel GMFG (N blocks)**
Generic agent	monopolist in sector i	monopolist at node $i \in [0,1]$	monopolist at an industry of block n
Individual state	own k	own k	own k
Environment state	full $(\boldsymbol{\kappa}, \mathbf{Z})$: sector i's shock moves q, c, the SDF, hence every sector's problem	own (κ, z) only: the node moves no moment, no aggregate	own (κ, z) only: the industry moves no moment, no aggregate
Taken as given	*stochastic* processes: SDF Λ_t, Q_t, $S_{i,t}$, sector policy $\bar{\iota}_i$	*deterministic* objects: r, g, q, $\mathcal{P}(i)$, $\widehat{\mathcal{D}}(i)$, node policy $\bar{\iota}$	*deterministic* objects: r, g, q, $\mathcal{P}_n$, $\widehat{\mathcal{D}}_n$, industry policy $\bar{\iota}$
Can influence	own k only	own k only	own k only
Marginal flow	$\zeta_i(\overline{v}_i)[1 - \xi_i(\overline{v}_i)]\, \widehat{s}^{\mathrm{f}}_i(\boldsymbol{\kappa}, \mathbf{Z})/\kappa_i$ (state-relevant through the whole realised state: GE feedback of own Z_i; without capacity constraints $\frac{\varepsilon-1}{\varepsilon}\alpha_i h_i q/\kappa_i$ with weight ω_i/α_i)	$\zeta_i(\overline{v})[1 - \xi_i(\overline{v})]\, \widehat{s}^{\mathrm{f}}(i, \kappa, z)/\kappa$ (state-relevant through κ; through z only if $\eta \neq 1$)	$\zeta_n(\overline{v})[1 - \xi_n(\overline{v})]\, \widehat{s}^{\mathrm{f}}_n(\kappa, z)/\kappa$ (state-relevant through κ; through z only if $\eta \neq 1$: no GE feedback survives)
Capacity and threshold utilization	$\overline{v}_i(\boldsymbol{\kappa}, \mathbf{Z})$: realised, on the shared state $U_i(\overline{v}_i)$, a random function of the shared state; aggregate (23) stochastic	$\overline{v}(i, \kappa, z)$: own state, the root of (139)–(141) $U(i, \kappa, z)$ with node distribution $\mathfrak{m}(i, \cdot)$; node mean $\mathbb{E}_i[U]$ and aggregate $\overline{U}$ (116) deterministic	$\overline{v}_n(\kappa, z)$: own state, the root of (156) $U_n(\kappa, z)$ with block distribution $\mathfrak{m}_n$; block mean $\mathbb{E}_n[U_n]$ deterministic
Discounting	stochastic SDF; risk premia priced	deterministic $r = \rho + \gamma g$; no premia	deterministic $r = \rho + \gamma g$; no premia
Backward objects	J marginal values $\widehat{\mathfrak{q}}_i$ on $2J-1$ dims	a continuum of marginal values $\mathfrak{q}(i, \cdot, \cdot)$ on 2 dims	N marginal values $\mathfrak{q}_n$ on 2 dims
Forward object	one joint ergodic density $\mathfrak{m}_\infty(\boldsymbol{\kappa}, \mathbf{Z})$	an ensemble of node distributions $\mathfrak{m}(i, \kappa, z)$	N marginal distributions $\mathfrak{m}_n(\kappa, z)$
Network object	matrix $\mathbf{A}$	kernel $a(i, j)$; operators $\mathbb{A}, \mathbb{A}^*$	block matrix A and its transpose
Coupling channel	shared stochastic state; $\mathbf{A}$ via the static block (thresholds, q; without capacity constraints h_i, $\boldsymbol{\omega}$), c, μ^K	moment fixed point; a via $\mathbb{A}, \mathbb{A}^*$ in $(\mathcal{P}, \widehat{\mathcal{D}}, q)$, from $(m_P, \widehat{m_X})$	moment fixed point; A in $(\mathcal{P}_n, \widehat{\mathcal{D}}_n, q)$, from $(m_{Pn}, \widehat{m}_{X,n})$
Consistency conditions	pointwise in the state: c and the drifts at each $(\boldsymbol{\kappa}, \mathbf{Z})$	two moment functions $(m_P, \widehat{m_X})$ on $[0,1]$ plus scalars	two N-vectors plus scalars
Exact for an observed IO table?	yes: it is the table	not applicable: a limit object	yes (Lemma 10); differs from the finite model by granularity only
Computable?	in principle, on one joint PDE grid, exponential in J; not computed in this paper	only through a step kernel	yes: linear in N (N decoupled 2D pairs given the moments, each with a scalar static root per grid point)
Theoretical status	standard	existence and ε-optimality conjectured(Section 3.4)	finite-dimensional fixed point; existence conjectured (Appendix B.9)

E Data and Calibration

This appendix explains the calibration of the model of Sections 2.1–3.3 to the industries of the United States.

E.1 Overview

The calibrated object is the stationary balanced growth path of the step-kernel systemof Section B.10, with CES final demand and with a period of one year. Its nodes are the $N = 63$ summary industries of the input–output accounts of the Bureau of Economic Analysis (BEA), less government and housing and with the three finance industries merged (Section E.2); each has mass $1/N$. A node contains a continuum of industries, which the productivity process spreads over states, and an industry contains a continuum of firms, which differ only in a transient demand shock. The three layers are matched to the three layers of the data: nodes to the summary industries of the accounts, the industries of a node to the detailed industries of the productivity panel, and firms to the plants that the capacity surveys sample, which enter only through published industry-level statistics. Only objects without units are matched: cost shares, expenditure shares, the ratio of operating surplus to operating surplus plus intermediate inputs, utilization rates, elasticities, the capital stock over sales and the growth rate.

[The firm layer and the plant surveys: what is used]

- No plant-level data enter the calibration. The plant surveys reach it through three published industry-level statistics: the Census survey's utilization rates by four- to six-digit industry (a check on the Federal Reserve's rates, and the validation); the share of an industry's capacity in plants above 95 %, which Boehm and Pandalai-Nayar (2022) computed from the Census plant data and published in their replication package, of which the mean over industries and quarters is the target of σ_v; and the Federal Reserve's industry rates, which are themselves anchored on the Census plant survey. The plant is the model's firm in interpretation only.
- Until 28 October 2026 the sentence above read "and firms to the plants of the capacity surveys", which could be read as if plants were observed.

The parameters fall into three groups. The *externally set* parameters (Section E.4) are read off the accounts or taken from the literature: the network a_{nm}, the depreciation rates δ_n, the growth rate g, the final-demand elasticity η, risk aversion γ, the productivity process within a node, and the accounts' ratio ω_n from which the capital shares follow. The *capacity block*, σ_v and ε (Section E.5.1), is estimated inside the model from two pieces of evidence on US manufacturing: the share of capacity in plants that operate close to capacity identifies the dispersion σ_v of firms' demand shocks, and how steeply industry supply curves rise with utilization identifies ε given σ_v. The remaining parameters are *calibrated internally*, in the order in which they depend on one

another: a node's utilization target fixes its threshold; the threshold and ω_n fix its capital share α_n; the network, the final shares and the intermediate shares fix the nodes' sales, and the capital stock over sales then fixes the interest rate r and $\rho = r - \gamma g$; the capital intensities x_n, the final-demand weights β_n and the common productivity level $\overline{Z}$ take the values at which the balanced growth path reproduces the utilization targets, the final-expenditure shares and the growth rate. With the productivity process off this chain is a closed form (the *first stage*); with it on, the chain is iterated on the solved model, and ε, σ_v and the installation curvature ϕ are found in outer loops.

Two calibrations are reported throughout, and they differ in the utilization targets and in where capacity is on, and in nothing else. The *baseline calibration* uses every utilization rate that exists, node by node: capacity is on in the 49 nodes with a rate, the 23 core nodes (manufacturing, mining, utilities) at the long-run means of the Federal Reserve's rates and 26 transfer nodes at rates transferred from European surveys, and off in the 14 nodes without a rate. The *secondary calibration* is a convention: capacity is on in every node, at the mean target of the manufacturing nodes. Every parameter of the other two groups is set by the same rule in both, and each calibration is computed separately, with its own ε, σ_v, capital shares and interest rate. Four conventions hold for both. Labour is outside the model, and the model's capital is physical capital, equipment and structures (Section E.3). A node's recorded revenue is the model's sales at the shadow price index, so the rationing rent is recorded inside the node's own output. One σ_v, one ε, one ϕ and one productivity process hold for all nodes. The accounts, the network, the final shares and the capital stock are means over 2007–2019; utilization targets are long-run means.

Table 10 lists every parameter. The data-side numbers of this appendix are final; the parameter values are those of the solved model, with the productivity process on: the baseline calibration's at $\phi = 2$ and $\gamma = 1$ (since 3 November 2026), the secondary calibration's at the interim value of ϕ and at $\gamma = 2$, to be replaced when it is re-run at the baseline's conventions; this clause is adjusted when that is done.

E.2 Data

Table 11 lists the sources. All US sources are public tables of the BEA, the Bureau of Labor Statistics (BLS), the Federal Reserve Board and the Census Bureau; the European source is the harmonised business-survey programme of the European Commission; three further sources are published estimates.

The node set. The BEA publishes its accounts for 71 summary industries, each a list of industries of the North American Industry Classification System (NAICS). Six of them are not nodes: the five government industries and housing, which is owner-occupied and tenant-occupied housing together. Their output is not produced by profit-maximising firms, or, for tenant-occupied housing, cannot be separated at this level; together they are 21.6 % of GDP. They are removed by one rule: their purchases from the other industries become those industries' final demand, and their sales to the other industries are taken out of the buyers' bundles, as imported inputs are. Two of

Table 10: The parameters of the two calibrations.

	meaning	how it is set; target or source	baseline calibration	secondary calibration
externally set (Section E.4)				
N	number of nodes	BEA summary industries, less government and housing; finance merged	63, mass 1/63	
a_{nm}	cost share of m in n's bundle	domestic direct requirements, each buyer's shares rescaled to sum to one	matrix	
δ_n	depreciation	depreciation over the net stock of equipment and structures	0.029 to 0.150	
g	growth rate	real value added of the nodes, 2007–2019	0.021	
η	elasticity between nodes' goods in final demand	literature	0.5	
γ	relative risk aversion	convention	1 (log utility)	2
Corr, Var	productivity process within a node	detailed manufacturing industries, with the model's weights	0.97, 0.057	
ω_n	accounts' ratio behind α_n	operating surplus of physical capital over the model's sales	0.042 to 0.604	
the capacity block (Section E.5.1)				
σ_v	dispersion of firms' demand shocks	share of manufacturing capacity in plants above 95 % of capacity, 0.162	0.185	0.199
ε	elasticity between firms' goods	inverse supply elasticity by bin of utilization, given σ_v	9.6	10.3
calibrated internally (Sections E.5.2–E.5.4)				
α_n	capital exponent	ω_n at the node's threshold, net of the markup rent; floor 0.02	0.065 in manufacturing	0.072
x_n	capital intensity	the node's mean utilization	49 targets, 0.72 to 0.93	every node at 0.786
β_n	weight in final demand	the node's share in final expenditure	63 targets	
$\overline{Z}$	common productivity level	the growth rate g	absorber	
r, ρ	interest rate, discount rate	capital stock over sales of all nodes, 1.356; $\rho = r - \gamma g$	0.039, 0.018	0.036, −0.006
ϕ	installation curvature	convention; the order of structural estimates of convex adjustment costs	2	11.5, to be set at 2

Notes: How each parameter is set, its target or source, and its value in the two calibrations. Values in colour are provisional: the secondary calibration's are the solved model's at the interim value of ϕ.

Table 11: Data sources, what each records, and where it is used.

source	what it records	coverage used	used for
BEA input–output accounts	use and make tables; domestic direct requirements	71 summary industries, 2007–2019	a_{nm}, β_n
BEA GDP by industry	gross output, intermediate inputs, compensation, taxes, operating surplus; quantity and price indexes	the same	ω_n, g, validation
BEA–BLS integrated production account	capital compensation by type of asset; the imputed pay of the self-employed	2007–2019	ω_n
BEA fixed-asset accounts	net stock and depreciation by class of asset, current cost	74 industries, 2007–2019	δ_n, r
NBER–CES manufacturing database	shipments, inputs, capital of detailed industries	six-digit manufacturing, 1958–2016	the productivity process
Federal Reserve G.17	production, capacity and utilization	manufacturing, mining, utilities, 1972–2019	utilization targets, ϕ, validation
European Commission business surveys	utilization; factors limiting production	NACE divisions, 2011–2019	utilization targets
Boehm and Pandalai-Nayar (2022)	inverse supply elasticity by bin of utilization; share of capacity in plants above 95 % (Census plant survey)	21 manufacturing industries, 1973–2011; 1976–2016	ε, σ_v
Census survey of plant capacity	plant utilization rates	four- to six-digit manufacturing, 2007–2019	check, validation
Amberg et al. (2026)	dispersion of firm utilization within industries	Swedish manufacturing firms, 2021–2024	check

the remaining industries, securities and funds and trusts, show almost no operating surplus of physical capital; they are merged with banking into one finance node. The 63 nodes produce 78.4 % of GDP. Table 14 at the end of this part lists them with the utilization target and source of each, the accounts' ratio ω_n and the depreciation rate δ_n. The nodes fall into four groups by the utilization data that exist for them: the 19 manufacturing nodes (10.4 % of GDP), the three mining nodes and utilities (3.4 %), together the 23 *core nodes* with a Federal Reserve rate; 26 *transfer nodes* (39.2 %), mostly transport, information, finance, business and personal services, whose rates are transferred from the European surveys; and 14 nodes without a rate (25.4 %).

The accounts. For every industry and year the GDP-by-industry accounts divide gross output twice,

$$\mathrm{GO}_n = \mathrm{II}_n + \mathrm{VA}_n, \qquad \mathrm{VA}_n = \mathrm{COMP}_n + \mathrm{TOPI}_n + \mathrm{GOS}_n : \tag{224}$$

gross output is intermediate inputs plus value added, and value added is the compensation of employees plus taxes on production less subsidies plus gross operating surplus. Operating surplus is a residual: the income of everything that is not an employee, which the accounts do not separate into the return to each asset, pure profit and the labour income of the self-employed. The integrated production account of the BEA and the BLS imputes a labour income SE_n to the self-employed and divides capital compensation among types of assets with one rate of return per industry, which gives the share s_n^{IP} of intellectual property (software, research and development, artistic originals) in it. The input–output accounts record, by commodity, what each industry uses and what goes to each final use, and which industries make each commodity; from them the BEA derives the domestic direct requirements of each industry on every other. The industry accounts, the fixed-asset accounts and the integrated account are published before redefinitions (secondary products left in the industry that makes them) and the requirement tables after. The fixed-asset accounts record the net stock of equipment, structures and intellectual property at current cost and its depreciation. Table 12 shows the identities in the data and, below each of the two items that the model uses, the part of it that the model has (Section E.3).

utilization. For manufacturing, mining and utilities the Federal Reserve Board publishes an index of production and an index of capacity, its estimate of sustainable maximum output, for each industry; the utilization rate is their ratio. Twenty-two summary industries have a rate with exactly the industry's NAICS codes, and utilities a rate for its electric power and gas part. The Census Bureau's survey of plant capacity asks manufacturing plants for their output and their full production capability and publishes rates at the four- to six-digit level; on the 20 nodes where both exist the Federal Reserve's rate exceeds the Census rate by 5.0 points on average over 2007–2019, as expected from the Census survey's higher benchmark. No US source measures utilization in services. The European Commission's business surveys ask manufacturing firms at what percentage of full capacity they operate, and services firms by how much they could increase their activity with their present resources, which the Commission converts into a rate; both are published by division of the European classification, for

Table 12: Gross output and its components, means of 2007–2019.

	all 63 nodes		manufacturing (19)	
billions of dollars a year, 2007–2019		%		%
gross output, GO	25,307	100	5,580	100
intermediate inputs, II	11,943	47	3,607	65
bought from the 63 nodes, II^{dom}	10,370	41	2,833	51
imported, or from government and housing	1,573	6	774	14
compensation of employees, COMP	7,366	29	967	17
taxes on production less subsidies, TOPI	969	4	80	1
gross operating surplus, GOS	5,029	20	926	17
income of intellectual property	1,075	4	325	6
imputed pay of the self-employed	502	2	13	0
of physical capital, GOS^K	3,452	14	588	11
the part that the model has: $GOS^K + II^{dom}$	13,823	55	3,421	61

Notes: The four unindented rows below gross output are the terms of (224). The indented rows divide intermediate inputs by origin and operating surplus by claimant, as in (225); the pay of the self-employed is its part attributed to physical capital. Value added is 13,364 and 1,972; the net stock of equipment and structures of the 63 nodes is 18,741 billion dollars and its depreciation 1,197 billion a year.

finance and for arts and recreation by country only, and the surveys also ask which factors limit production. US manufacturing utilization averages 0.784 over 1972–2019 and 0.750 over 2007–2019; over 2011–2019 the European services divisions report rates 1.10 times those of the manufacturing divisions.

Supply curves and the share of capacity above 95 per cent. Boehm and Pandalai-Nayar (2022) estimate how steep the supply curves of 21 three-digit manufacturing industries are over 1973–2011 (819 industry-years): the growth of an industry's price on the growth of its output, instrumented by shifts in demand (the growth of foreign demand and of the demand of large downstream buyers), with capacity growth and unit variable cost (materials, energy and production-worker pay) held fixed and with industry and year effects. The coefficient, the *inverse supply elasticity*, is estimated separately in four bins of the previous year's utilization, measured as the deviation from the industry's own mean and cut at the 15th, 50th and 85th percentiles (Table 13). From the plant records of the Census survey the same authors report the share of manufacturing capacity in plants operating above 95 % of their capacity: 0.162 on average over the fourth quarters of 1976–2016 (standard deviation 0.042), correlated 0.79 with the Federal Reserve's manufacturing rate, with a slope of 0.74 points per point of it.

Productivity, investment and the validation series. The NBER–CES database gives shipments, materials, energy and the real stock of equipment and structures for 355 six-digit manufacturing industries over 1958–2016, the only panel of detailed industries inside a node. The Federal Reserve's capacity indexes over 1958–2021 for the 21

Table 13: The four supply-curve estimates of Boehm and Pandalai-Nayar.

	bin 1	bin 2	bin 3	bin 4
deviation of lagged utilization from the industry's mean	below −0.065	−0.065 to 0.009	0.009 to 0.061	above 0.061
mean utilization in the bin	0.679	0.776	0.827	0.874
industry-years	122	287	287	123
inverse supply elasticity	0.116	0.282	0.309	0.573
(standard error)	(0.082)	(0.101)	(0.086)	(0.139)
model, baseline calibration	0.111	0.197	0.279	0.548

Notes: By bin of the previous year's utilization, from the authors' replication package after their own cleaning of the sample, and the model's counterparts at the calibrated values of the baseline calibration (Section E.5.1).

industries of the supply-curve sample give the moment that sets ϕ: a utilization rate one point above the industry's mean is followed by capacity growth that is 0.167 percentage points higher in the next year (standard error 0.025; industry and year effects; 1,154 industry-years). The GDP-by-industry and fixed-asset accounts give quantity and price indexes of every node's gross output, intermediate inputs and capital over 1997–2019, from which the validation's productivity series and demand shifts are built (Section E.6); over those years the standard deviation of utilization at the median core node is 0.058, with an autocorrelation of 0.65, and the within-node dispersion of the Census rates has a median of 0.086. Amberg et al. (2026) report, for Swedish manufacturing firms over 2021–2024, a standard deviation of firm utilization of 12.3 percentage points within industry-quarter cells and 15.2 within manufacturing as a whole.

[What is in the data part, and what is left out] The part keeps to what the appendix uses. Left out, as the brief allows: the summary statistics of the input–output tables (the median imported share of a bundle is 7.4 % and reaches 44.7 % in petroleum refining; seven nodes import more than a fifth of their bundle; the few negative cells of the requirement tables, at most 0.14 % of a row, are set to zero), the vintages, and the construction of the European transfer by country (the four by-country targets rest on countries that carry about a third of the European weight, Section E.5.3). All numbers of this part were recomputed from the data set in this session, except the four supply-curve estimates and the Census share, which are the data set's copies of the authors' published tables, and the Swedish numbers, which were read in the authors' paper.

Table 14: The 63 nodes.

node	code	group	target	source	ω_n	δ_n	α_n (b.)	α_n (s.)
Farms	111CA	none	off	—	0.174	0.076	0.078	0.071
Forestry, fishing, and related activities	113FF	none	off	—	0.328	0.094	0.250	0.251
Oil and gas extraction	211	core	0.930	G.17	0.570	0.074	0.485	0.536
Mining, except oil and gas	212	core	0.826	G.17	0.428	0.068	0.358	0.369
Support activities for mining	213	core	0.720	G.17	0.348	0.099	0.272	0.274
Utilities	22	core	0.853	G.17 (part)	0.470	0.035	0.404	0.418

node	code	group	target	source	ω_n	δ_n	α_n (b.)	α_n (s.)
Construction	23	none	off	—	0.186	0.126	0.091	0.084
Wood products	321	core	0.765	G.17	0.114	0.084	floor	floor
Nonmetallic mineral products	327	core	0.733	G.17	0.209	0.077	0.113	0.113
Primary metals	331	core	0.782	G.17	0.154	0.057	0.043	0.048
Fabricated metal products	332	core	0.786	G.17	0.167	0.069	0.057	0.063
Machinery	333	core	0.777	G.17	0.166	0.066	0.057	0.061
Computer and electronic products	334	core	0.774	G.17	0.216	0.082	0.116	0.120
Electrical equipment, appliances, and components	335	core	0.821	G.17	0.212	0.073	0.101	0.115
Motor vehicles, bodies and trailers, and parts	3361MV	core	0.755	G.17	0.093	0.103	floor	floor
Other transportation equipment	3364OT	core	0.746	G.17	0.210	0.070	0.112	0.112
Furniture and related products	337	core	0.774	G.17	0.097	0.082	floor	floor
Miscellaneous manufacturing	339	core	0.766	G.17	0.203	0.087	0.102	0.105
Food and beverage and tobacco products	311FT	core	0.804	G.17	0.144	0.067	0.026	0.036
Textile mills and textile product mills	313TT	core	0.787	G.17	0.089	0.065	floor	floor
Apparel and leather and allied products	315AL	core	0.760	G.17	0.044	0.059	floor	floor
Paper products	322	core	0.866	G.17	0.179	0.082	0.039	0.077
Printing and related support activities	323	core	0.796	G.17	0.206	0.089	0.101	0.109
Petroleum and coal products	324	core	0.855	G.17	0.259	0.067	0.144	0.171
Chemical products	325	core	0.770	G.17	0.195	0.076	0.092	0.096
Plastics and rubber products	326	core	0.820	G.17	0.150	0.094	0.027	0.043
Wholesale trade	42	none	off	—	0.347	0.086	0.271	0.273
Motor vehicle and parts dealers	441	none	off	—	0.341	0.046	0.264	0.265
Food and beverage stores	445	none	off	—	0.273	0.054	0.189	0.185
General merchandise stores	452	none	off	—	0.135	0.052	0.034	0.023
Other retail	4A0	none	off	—	0.306	0.061	0.225	0.225
Air transportation	481	transfer	0.837	EU	0.284	0.063	0.182	0.200
Rail transportation	482	transfer	0.885	EU	0.372	0.029	0.268	0.303
Water transportation	483	transfer	0.855	EU	0.202	0.065	0.073	0.103
Truck transportation	484	transfer	0.885	EU	0.200	0.150	0.050	0.102
Transit and ground passenger transportation	485	transfer	0.885	EU	0.345	0.077	0.233	0.271
Pipeline transportation	486	transfer	0.885	EU	0.604	0.043	0.563	0.576
Other transportation and support activities	487OS	transfer	0.862	EU	0.211	0.069	0.081	0.115
Warehousing and storage	493	transfer	0.863	EU	0.170	0.050	0.031	0.067
Publishing industries, except internet (includes software)	511	transfer	0.843	EU	0.164	0.096	0.034	0.060
Motion picture and sound recording industries	512	transfer	0.823	EU	0.124	0.059	floor	floor

node	code	group	target	source	ω_n	δ_n	α_n (b.)	α_n (s.)
Broadcasting and telecommunications	513	transfer	0.873	EU	0.382	0.056	0.287	0.314
Data processing, internet publishing, and other information services	514	transfer	0.863	EU	0.271	0.139	0.155	0.185
Banking, securities and funds	52BSF	transfer	0.846	EU, by country	0.271	0.091	0.163	0.185
Insurance carriers and related activities	524	transfer	0.856	EU, by country	0.255	0.063	0.139	0.166
Other real estate	ORE	transfer	0.889	EU	0.364	0.035	0.255	0.294
Rental and leasing services and lessors of intangible assets	532RL	transfer	0.852	EU	0.530	0.135	0.478	0.489
Legal services	5411	transfer	0.882	EU	0.431	0.091	0.344	0.372
Computer systems design and related services	5415	transfer	0.888	EU	0.115	0.142	floor	floor
Miscellaneous professional, scientific, and technical services	5412OP	transfer	0.868	EU	0.147	0.099	floor	0.040
Management of companies and enterprises	55	transfer	0.875	EU	0.128	0.045	floor	floor
Administrative and support services	561	transfer	0.864	EU	0.165	0.112	0.023	0.061
Waste management and remediation services	562	none	off	—	0.266	0.059	0.180	0.179
Educational services	61	none	off	—	0.239	0.033	0.151	0.145
Ambulatory health care services	621	none	off	—	0.214	0.083	0.123	0.116
Hospitals	622	none	off	—	0.167	0.053	0.071	0.061
Nursing and residential care facilities	623	none	off	—	0.084	0.048	floor	floor
Social assistance	624	none	off	—	0.042	0.051	floor	floor
Performing arts, spectator sports, museums, and related activities	711AS	transfer	0.848	EU, by country	0.124	0.048	floor	floor
Amusements, gambling, and recreation industries	713	transfer	0.846	EU, by country	0.200	0.054	0.075	0.099
Accommodation	721	transfer	0.841	EU	0.355	0.040	0.266	0.283
Food services and drinking places	722	transfer	0.838	EU	0.176	0.068	0.051	0.073
Other services, except government	81	transfer	0.850	EU	0.164	0.042	0.030	0.059

Notes: Group, utilization target of the baseline calibration and its source, the accounts' ratio ω_n and the depreciation rate δ_n. Group: core (own Federal Reserve rate), transfer (rate transferred from the European surveys; "by country" where the European rate is published by country only), none (no rate; capacity off in the baseline calibration). In the secondary calibration every node has the target 0.786. The last two columns are the capital shares α_n of the baseline calibration (b.), on the solved model at its final ϕ, and of the secondary calibration (s.), on the solved model at the interim ϕ (provisional, in colour); "floor" is $\alpha_n = 0.02$.

E.3 From the Data to the Model

What the model has. The model's firms produce with physical capital and a bundle of the other nodes' goods. Labour, taxes, imported inputs and intellectual property have no counterpart in it, and neither do their payments. Capital is the stock whose capacity binds: the utilization rates measure output against the equipment and structures in place, which software and research do not limit, so the model's capital is equipment and structures, and intellectual property is treated like labour, as an input the model does not have, left out of the capital stock, of depreciation and of capital income. The model's sales therefore correspond to the part of gross output that the model has, the operating surplus of physical capital plus the inputs bought from the other nodes,

$$\mathrm{GOS}_n^{\mathrm{K}} + \mathrm{II}_n^{\mathrm{dom}}, \qquad \mathrm{GOS}_n^{\mathrm{K}} := \left(1 - s_n^{\mathrm{IP}}\right)(\mathrm{GOS}_n - \mathrm{SE}_n), \qquad \omega_n := \frac{\mathrm{GOS}_n^{\mathrm{K}}}{\mathrm{GOS}_n^{\mathrm{K}} + \mathrm{II}_n^{\mathrm{dom}}}, \tag{225}$$

where the operating surplus of physical capital is gross operating surplus less the imputed pay of the self-employed, which is labour income, and less the share of intellectual property, and where $\mathrm{II}_n^{\mathrm{dom}}$ is the inputs that the node buys from the 63 nodes, from the domestic flows of the input–output accounts. Both are measured before redefinitions, the basis of the industry accounts, and ω_n is the mean of the annual ratios of 2007–2019. This is 55 % of the gross output of the 63 nodes (Table 12). The ratio ω_n is the accounts' counterpart of one minus the intermediate share of sales, from which the capital shares are read (Section E.5.2).

Two revenue concepts. In the model an industry's output reaches its buyers in two steps: the firms sell to competitive aggregators at the posted price, and the aggregators sell the industry's good at the shadow price index P, which exceeds the posted price when some firms are rationed. The same output therefore has two values, firm revenue $S^{\mathrm{f}} = \varrho_n(\overline{v})S$ and sales S(118), and the difference is the aggregators' rationing rent. The accounts record one value, and the calibration takes it to be the model's sales: the recorded value of a flow is its value at the shadow price index. The aggregators of a node are thus read as part of the node, and the rationing rent, earned inside the industry that is short of capacity, is recorded in that industry's output and operating surplus, as the model's own accounting has it. Three things follow. The recorded cost shares of a buyer's bundle and the recorded final-expenditure shares are those of the model, with no correction. A node's recorded operating surplus of physical capital contains the income of capital, the markup rent and the rationing rent, so that by the four-way split of sales (78) the ratio ω_n corresponds to one minus the intermediate share of sales, $1 - \xi_n\varrho_n$. And the price that statistics measure, revenue over real output, is the shadow price index, the price of the supply-curve regressions.

[The alternative reading, and the size of the rent] The accounts could instead record every unit at the posted price, so that nobody earns the rent. A recorded sale of node m would then be ϱ_m times its value in the model, and every cost share, final share and intermediate share would be corrected by the sellers' ϱ_m, which the utilization target fixes; the corrections are small because the rent is small. On the solved model the

rationing rent is 0.28 % of the capacity nodes' sales and 0.91 % of GDP in the baseline calibration at its final ϕ (0.52 % of GDP in the secondary calibration at the interim ϕ; computed once; the first stage had 0.30 and 1.00, cross-examined). It is a transfer to the aggregators, not a measure of lost output, and the accounts do not show it, so it is a prediction of the model. The reading at posted prices was tested in an earlier draft (every capital share moves by less than 0.005) and is not carried as a robustness row.

The network and final demand. In the model node n spends the share a_{nm} of its intermediate bill on the good of node m, so the industry-by-industry direct requirement of n on m, purchases per unit of the buyer's output, is $a_{nm}\chi_n$ with χ_n the node's intermediate share of sales. Rescaling each buyer's requirements over its suppliers to sum to one therefore gives a_{nm} and removes χ_n, whose level comes from ω_n. The model has no foreign producers, so the network is built from the domestic requirements, with imported inputs and purchases from government and housing removed from the bundle before the rescaling; exports are final demand. Adding the model's value clearing over the industries of a node gives the counterpart of the accounts' total requirements,

$$\big(\text{sales of node } m\big)_m = \Big(I - a^{\mathsf{T}}\,\mathrm{diag}(\chi)\Big)^{-1}\big(\text{final expenditure on node } m\big)_m, \qquad (226)$$

so that the final-use columns of the use table, net of imports and with the purchases of the six removed industries added, are the model's final expenditure node by node, the targets of the weights β_n: consumption, investment, government purchases and exports all buy the model's final good. Purchases among the nodes are 41 % of their gross output and final use the remaining 59 %; the targets are the nodes' shares in that final use, not its level: the largest are construction (0.086), wholesale trade (0.065) and ambulatory health care (0.058), and the smallest management of companies and warehousing (0.0003).

[What the 59 per cent is for]

- It is a description, not a target, and not a quantity that a right model would push to 100 per cent. Final use (net of imports, with the purchases of the six removed industries added) over the gross output of the 63 nodes is 0.589: of each dollar of the nodes' output, 59 cents go to final users and 41 cents to other nodes. At 100 per cent there would be no network; the 41 cents are the network's weight, which is what this paper is about. The sentence tells the reader how much of output passes through the network before any matching starts.
- What is matched is not this level but the nodes' shares in final expenditure (construction 0.086, wholesale 0.065, . . .), which sum to one by construction.
- The model's own version of the ratio is not comparable, which is why the model's levels are not compared with the accounts' (Section E.3). The model's final expenditure is its GDP, the operating surplus of physical capital, since capital is its only factor, and its sales are $\mathrm{GOS}^{\mathrm{K}} + \mathrm{II}^{\mathrm{dom}}$; so in the model final demand is $1/4.0 = 25$ per cent of sales, against 59 per cent of gross output in the accounts.

The gap is labour, taxes, imports and intellectual property, outside the model and inside the accounts' final demand.

- The sentence above was reworded on 28 October 2026; it read "Final demand is 59 % of the nodes' output in the aggregate". Evidence: $s_F = 0.589$ in the data set's node table and the documentation of the final shares; the 4.0 from the accounts' bridge (reproduced).

[Derivation of (226)**]** Write $\widehat{m}_{S,m} = \mathbb{E}_m[\widehat{s}]$ and $\widehat{m}_{X,m} = \mathbb{E}_m[\widehat{x}]$ for node m's mean sales and mean intermediate spending per industry, $\Delta_m = 1/N$ for its mass, $\chi_m := \widehat{m}_{X,m}/\widehat{m}_{S,m}$ for its intermediate share of sales and F_m for the final expenditure on its good.

- *The demand of the step-kernel system.* By the demand line of (156), summed over the industries of node m, node m's sales are final expenditure plus the spending that the other nodes direct to it, $\Delta_m \widehat{m}_{S,m} = F_m + \Delta_m \widehat{\mathcal{D}}_m$, with $\widehat{\mathcal{D}}_m = \frac{1}{\Delta_m}\sum_n a_{nm}\widehat{m}_{X,n}\Delta_n$ by the coupling of (159): node n spends the share a_{nm} of its bill $\Delta_n \widehat{m}_{X,n}$ on m.
- *In terms of sales.* Substituting $\Delta_n \widehat{m}_{X,n} = \chi_n \Delta_n \widehat{m}_{S,n}$ gives, with $s_m := \Delta_m \widehat{m}_{S,m}$,

$$s_m = F_m + \sum_n a_{nm}\chi_n s_n, \qquad \text{that is} \qquad s = F + a^{\mathsf{T}}\,\mathrm{diag}(\chi)\,s,$$

  and (226) follows if $I - a^{\mathsf{T}}\,\mathrm{diag}(\chi)$ is invertible. It is: $a^{\mathsf{T}}\,\mathrm{diag}(\chi)$ is non-negative, and its column sums are $\chi_n \sum_m a_{nm} = \chi_n < 1$ (the rows of a sum to one over the domestic nodes, and $\chi_n < 1$ because a unit of sales also pays the markup rent, $1 - \xi_n \varrho_n \geq \varrho_n/\varepsilon > 0$ by the four-way split; on the data side this is $\omega_n > 0$), so its spectral radius is below one and the Neumann series converges.
- *The accounts' counterpart.* The BEA's total-requirements matrix is $(I - D \cdot B)^{-1}$ with $D \cdot B$ the industry-by-industry direct requirements, purchases of the buyer from each seller per unit of the buyer's output; with rows the supplying and columns the using industry; in the model these are $a_{nm}\chi_n$, so $D \cdot B$ corresponds to $a^{\mathsf{T}}\,\mathrm{diag}(\chi)$ and the accounts' total requirements to the inverse in (226). The calibration takes only the cost shares a_{nm} and the final-expenditure shares from this correspondence. In the no-capacity benchmark χ_n is the constant $\xi(n)$ and, with Cobb–Douglas final demand, $F_m = \Delta_m \beta_m q$, so (226) is the benchmark's formula for the Domar weights; under capacity $\chi_n = \mathbb{E}_n[\xi_n \varrho_n \widehat{s}]/\mathbb{E}_n[\widehat{s}]$ moves with the node's utilization.

[Status] Proof-checked by a fresh reviewer: sound. The relation and the invertibility argument were re-derived and confirmed numerically on a random four-node example (residual 4×10^{-16}; the alternatives with the matrix transposed or the shares on the other side fail). Three corrections were folded in: the BEA's industry-by-industry requirements correspond to $a^{\mathsf{T}}\,\mathrm{diag}(\chi)$ without a transpose; the benchmark's final expenditure carries the node mass; and $\chi_n < 1$ is now justified.

What leaving labour out implies. The model's sales are 4.0 times the model's GDP, which is the operating surplus of physical capital; in the accounts gross output is 1.9 times value added. A productivity shock therefore moves the model's GDP by more

than it would move measured GDP, by a factor that the accounts cannot supply, since they say how much of value added is labour income but not how labour income moves with productivity. The model is for this reason compared with the data through shares, ratios and rates only, and its responses to shocks are reported relative to an economy without capacity constraints, where the scale cancels.

[The consequence of leaving labour out, in numbers] Of 100 dollars of the nodes' value added, the model has the 26 that are the operating surplus of physical capital. Rescaled in the same way, the data and the model agree on sales over GDP, so inside the model's world nothing is off. The difference matters in one place: a response of the model's GDP to a shock cannot be set beside an elasticity estimated on actual GDP. In the no-capacity benchmark at given capital, the response of GDP to node n's productivity is a cost-based Domar weight in which intermediate inputs enter with their output elasticity $1 - \alpha_n$, not with their smaller share of revenue; these weights satisfy $\sum_n \lambda_n \alpha_n = 1$, so they are of order $1/\alpha$ and sum to several times the Domar weights (reasoned and derived in the review of the calibration, and confirmed numerically there by three implementations at earlier values; not recomputed on the 63-node set). This is why the quantitative section reports impact responses as ratios against the same-parameters economy and never as elasticities of GDP.

[The correspondences in one table] Data element, the model object it corresponds to, and the choice the correspondence rests on.

data element	model object	rests on
summary industry, less government and housing, finance merged	node n, mass $1/N$	the removal rule
$\mathrm{GOS}_n^{\mathrm{K}} + \mathrm{II}_n^{\mathrm{dom}}$	sales $\Delta_n \mathbb{E}_n[\widehat{s}]$	labour outside the model; the rent recorded in the node's output
domestic direct requirements, rescaled over the buyer's suppliers	a_{nm}	the domestic closure
final use net of imports, plus purchases of government and housing, as shares of the total	final-expenditure shares $\beta_n \mathbb{E}_n[P^{1-\eta}]\Delta_n$	the removal rule
ω_n of (225)	$1 - \mathbb{E}_n[\widehat{x}]/\mathbb{E}_n[\widehat{s}]$	physical capital only; the two revenue concepts
depreciation over the net stock of equipment and structures	δ_n	physical capital only
net stock of equipment and structures over $\mathrm{GOS}^{\mathrm{K}} + \mathrm{II}^{\mathrm{dom}}$, all nodes	capital stock over sales, (232)	one interest rate
Federal Reserve rate; European services rate relative to manufacturing	$\mathbb{E}_n[U_n]$	the transfer
Census plant rate	$\mathbb{E}_n[\Psi_1]$, the mean of firms' output over capacity	—
supply-curve estimates by bin	weighted arc elasticities along the model's supply curve	the counterpart of Section E.5.1
share of capacity above 95 %	share of firms with output above 95 % of capacity, $1 - G(0.95\,\overline{v})$	equal capacity across firms
changes of detailed industries' productivity within a node	changes of z	the model's input weights
Swedish dispersion of firm utilization	$\sqrt{\Psi_2 - \Psi_1^2}$	—

E.4 Externally Set Parameters

The network, a_{nm}. For each year from 2007 to 2019 the BEA's industry-by-industry domestic direct requirements (the market-share matrix times the domestic direct-requirements matrix) are formed, the six removed industries are taken out as buyers and as suppliers, the three finance industries are summed on dollar flows, the few small negative entries are set to zero, and every buyer's row is divided by its sum; a_{nm} is the mean over the years. The Cobb–Douglas bundle holds the cost shares fixed at every state. In the data they are nearly fixed for most buyers and move with relative prices for the energy rows: over the window the share of oil and gas in petroleum refining's bundle has a range of 0.25 around its mean of 0.61, and in the utilities' bundle a range of 0.33 around 0.17.

[Where the negative cells come from]

- They are the BEA's own, not a product of the removal or the merger. The industry-by-industry domestic direct requirements are $D \cdot B$, the market-share

matrix times the commodity-by-industry domestic requirements. The use table carries a few negative cells by construction: the commodity "scrap, used and secondhand goods" is negative for an industry that is a net seller of it, and an import adjustment can exceed use in a cell; multiplied through D these spread as small negatives into a few industry-by-industry cells.

- Size: among the 65 industries over 2007–2019, 89 cells of about 55,000, total mass -0.0045 against rows that sum to one, the largest 0.14 per cent of a buyer's bundle in any year. Setting them to zero before the rescaling moves no cost share by more than that. (Over all 71 industries and 1997–2019: 221 cells, the most negative -0.009.)
- Evidence: the documentation of the data set's input–output table, written by the script that builds the network, which floors the flows before the finance merger and again in the capacity sub-economy (the counts reproduced from it; the mechanism reasoned from the BEA's method).

[Where the claim about moving cost shares comes from]

- Source: the data set's own yearly tables, which hold a_{nm} for each year 2007–2019 and summarise, per buyer and over the cells with a mean share above 0.01, the range of each share over the years (largest minus smallest), its mean and its standard deviation. The two numbers in the text above (utilities' share of oil and gas: range 0.33 around 0.17; petroleum refining's share of oil and gas: 0.26 around 0.61) are the two largest cells of that file (reproduced).
- What the file shows. The median buyer's largest range is 0.088, its mean range 0.024 and its mean standard deviation 0.008: for most buyers the shares are nearly fixed, as the Cobb–Douglas bundle assumes. The largest movements are the energy cells, and they are relative-price movements: the oil price fell by about two fifths in 2009 and by half in 2014–16, and at fixed quantities a cost share moves one for one with the price.
- But not only energy moves: water transport's share of finance (range 0.28 around 0.16), apparel (0.24), insurance's own share (0.23 around 0.70), computers (0.21) and trucking (0.20) follow. The sentence above ("nearly fixed for most buyers and move with relative prices for the energy rows") is right as far as it goes; a compressed version written on 24 October 2026 ("in the data they move only for the energy rows") overstated it and was withdrawn.
- What it means for the model. Shares that move with prices say that the elasticity of substitution among inputs is not one, as the Cobb–Douglas bundle has it; estimates in the literature put it well below one for intermediates. Changing it is a change of the model, not of the calibration. For the stationary calibration nothing is lost (the shares are means over the window), and to first order in a shock the bundle price responds through the cost shares whatever the elasticity, so the impact results would hardly move; what a low elasticity would change is how quantities reallocate after a large shock.

Depreciation, δ_n. The node's current-cost depreciation over its current-cost net stock of equipment and structures, summed over the fixed-asset industries of the node and averaged over the thirteen annual ratios of 2007–2019. Both are at replacement cost, so the ratio is the model's δ_n in the law of motion of capital (28). It runs from 0.029 (rail transportation) to 0.150 (truck transportation), with a median of 0.068; the differences reflect the mix of long-lived structures and short-lived equipment.

The growth rate, $g = 0.021$, the mean growth rate of the real value added of the 63 nodes over 2007–2019, a Törnqvist aggregate of the nodes' quantity indexes. It is the target of the common productivity level (Section E.5.4).

The final-demand elasticity, $\eta = 0.5$**.** η is the elasticity of substitution between the nodes' goods in final demand (24). It must differ from one: at $\eta = 1$ an industry's final sales do not depend on its own price, so its productivity does not affect its own payoff and the industry dynamics degenerate(Section B.7). Estimates of the elasticity across sectors in final expenditure lie below one (Herrendorf, Rogerson & Valentinyi, 2013; D. Comin, Lashkari & Mestieri, 2021), and production-network models use values in the same range (Baqaee & Farhi, 2019, 2020); 0.5 is in the middle of it, and 0.3 and 0.9 are the sensitivities. With $\eta < 1$ the nodes' goods are complements in final demand: a node whose price falls loses final spending. On the balanced growth path η moves only the prices and the absorbers x_n, β_n and $\overline{Z}$.

Risk aversion, $\gamma = 1$, log utility, a convention. On the balanced growth path γ and ρ enter only through the interest rate $r = \rho + \gamma g$ (145), and the interest rate is pinned by the capital stock (Section E.5.4), so γ moves only ρ: at the calibrated $r = 0.039$ and $g = 0.021$, $\rho = r - \gamma g$ is 0.018 at $\gamma = 1$ and -0.003 at $\gamma = 2$, and log utility is the conventional value at which the discount rate is positive. It matters in transitions, where the interest rate moves with the growth of consumption, and in the comparison with the economy without capacity, whose growth rate differs (Section E.6).

[The negative discount rate: the arithmetic, log utility, and the other allocations] [Flag, 3 November 2026: written when the baseline calibration was $\phi = 13$, $\gamma = 2$. Since 3 November 2026 it is $\phi = 2$ and $\gamma = 1$ with $r = 0.0392$, so $\rho > 0$ if and only if $\gamma < 1.87$ and $\rho = 0.0183$ at $\gamma = 1$; the arithmetic below is the earlier calibration's.]

- **The arithmetic.** On the balanced growth path $\rho = r - \gamma g$, with $r = 0.0334$ from the capital stock and $g = 0.0210$ from the accounts, so $\rho > 0$ if and only if $\gamma < r/g = 1.59$:

γ	1	1.5	2	3
ρ	0.0124	0.0019	-0.0086	-0.0296

  Yes: log utility gives $\rho = 1.24$ per cent, a conventional value, and it is the smallest change that does it.

- **What log utility leaves unchanged and what it moves.** Every balanced-growth-path number of the baseline calibration except ρ: the first and second stages use r and never γ or ρ separately, so the stored calibration, the thresholds, the capital

shares, x_n, β_n, $\overline{Z}$ and the baseline calibration's own impact responses (which are static) stay as they are; so does the recalibrated economy's balanced growth path, which is calibrated to the same g. What moves: (i) every transition, since households obey $\dot{C}/C = (r_t - \rho)/\gamma$ and the intertemporal elasticity doubles from 0.5 to 1; (ii) the economy without capacity constraints, whose growth rate differs from g and whose interest rate is $\rho + \gamma g_{\text{twin}}$, so its balanced growth path and every ratio to it move. The results already hold the growth effect at three values of γ (a stored output table of the static results, computed once): the economy without capacity constraints grows faster than the baseline calibration by 0.19 points at $\gamma = 2$, by 0.35 at $\gamma = 1$ and by 0.13 at $\gamma = 3$. So under log utility the growth part of the long-run cost of capacity nearly doubles; Section 5.1 would be recomputed, not only ρ restated.

- **The welfare measure does not care.** The balanced-growth value is finite if and only if $\nu = \rho + (\gamma - 1)g > 0$, and $\nu = r - g = 0.0124$ for every γ; so the consumption equivalent rests on the same thin margin $r - g$ under log utility as now (the caveat on the welfare number in Section 5.1 stands either way).
- **Other allocations.** The four objects obey one equation, so three are free. Today: g from the accounts, r from the capital stock, γ a convention, ρ the residual. The alternatives:

  1. *ρ and γ conventions, r implied.* At $\rho = 0.01$ and $\gamma = 2$, $r = 0.052$, and the model would hold far less capital than the accounts unless another parameter is freed; δ_n is measured and nothing else in the chain is free. This gives up the capital-stock target, the one level moment of the calibration, and with it the comparison of the model's return with outside returns in the fit section. Not recommended.
  2. *ρ a convention, r from the stock, γ implied.* $\gamma = (r - \rho)/g$: 1.11 at $\rho = 0.01$, 0.64 at $\rho = 0.02$. This is log utility by another name, with the parameter the audience knows (γ) hidden behind one they do not check (ρ). Say $\gamma = 1$ instead.
  3. *r a convention, the capital stock dropped.* At the usual $r = 0.04$ and $\gamma = 2$, $\rho = -0.002$: still not positive. The sign is set by $r < \gamma g$, and no return on capital that the accounts support reaches $2g = 0.042$.
  4. *A lower g.* It helps only a little: at $g = 0.017$ and $\gamma = 2$, $\rho = -0.0006$; and g is a data fact of the 63 nodes.

  So the allocation is right and the one lever is γ: the measured objects fix r and g, and $r - \gamma g$ is negative at $\gamma = 2$ because the measured return on capital, a pre-tax 3.3 per cent, is low against $2g$. Log utility is the clean remedy and a common one; its cost is the change in the transitions and in the economy without capacity constraints above, and the audience's other question, whether an intertemporal elasticity of one is high. When you ask for it, it is a separate calibration with an unchanged stationary economy and recomputed economy without capacity

constraints and transitions; the stored calibration is reused with $\gamma = 1$ and $\rho = 0.0124$.

[Q27 | mynotes | ANSWERED] Two questions:

1. Is the stored calibration with $\gamma = 1$ and $\rho = 0.0124$ complete and in its final usable version?

2. In general, and this is a clarification question, if I wanted to have a new calibration with a new gamma (and thus also a new rho) does that require the full 60-hour procedure or does it just need the "partial" (don't know what to call it) 2-hour procedure?

[The calibration at $\gamma = 1$: what exists, and what a new γ costs] [Superseded, 3 November 2026: the stored baseline calibration is at $\gamma = 1$ and $\rho = 0.018259$ (with $\phi = 2$), made exactly as the box describes, and every result of the paper is computed on it; the box is kept as the record.]

- **Nothing is stored at $\gamma = 1$.** The one stored calibration of the baseline calibration is at $\gamma = 2$ and $\rho = -0.0086$, complete and final since 11 October 2026 (a parameter file with every field of the model, the node table, the capacity block, the moments, a README). "The stored calibration is reused with $\gamma = 1$ and $\rho = 0.0124$" in the box above means: those files, with the two fields γ and ρ overwritten, *are* the complete calibration at $\gamma = 1$, because nothing in the stationary capacity economy depends on γ and ρ separately. Checked in the code on 27 October 2026: the first stage uses γ in one line, $\rho = r - \gamma g$; the static block's user cost is $[\rho + (\gamma - 1)g]\mathfrak{q}_n + \iota_n = (r - g)\mathfrak{q}_n + \iota_n$; the stationary solver forms $r = \rho + \gamma g$ and uses nothing else; the pair enters separately only in the welfare formulas ($\nu = \rho + (\gamma - 1)g = r - g$, and a log branch at $\gamma = 1$) and in the transitions' Euler equation. Confirmed numerically by the results script, which solved the economy without capacity constraints at $\gamma = 1$ and 3 with $\rho = r - \gamma g$ and recorded the capacity economy's growth rate unchanged to every digit (the stored growth-effect table: $g = 0.020972$ at all three values). The exact value to use is $\rho = r - g = 0.012382$.
- **What exists at $\gamma = 1$ is one number**: the economy without capacity constraints's growth rate, 0.02446 against the capacity economy's 0.02097 (the solve was done inside the growth-effect loop of the static results script and only g was kept). Not computed at $\gamma = 1$: the economy without capacity constraints's level objects (output and consumption per unit of capital, the decomposition of the gap, the allocation of capital), the welfare gap, the per-node results, every transition, the response ratios to the recalibrated economy. The capacity economy's own impact responses are static and identical.
- **How to run it.** Copy the stored calibration into a trial folder of its own, set $\gamma = 1$ and $\rho = 0.012382$ in its parameter file, and point the results scripts at it; no step of the calibration chain runs. At $\gamma = 1$ the welfare formulas switch to the log branch, $V_0 = \ln \widehat{C}/\rho + g/\rho^2$, which needs $\rho > 0$ and has it; the flat consumption equivalent of the capacity economy, invalid today because $\rho < 0$ (the text uses

the ratio of the two values, which is valid at any sign of ρ), becomes a usable number.

- **The cost of a new γ: neither procedure.** The chain without the ϕ loop (the first stage, seconds; the second stage, two to four rounds of five to eleven minutes; the ε–σ_v loop, 71 minutes) does not run at all, since none of its moments involves γ: the stationary capacity economy is the same economy. What is recomputed is what compares or moves: the economy without capacity constraints (a cold solve, about two minutes), the static results on both economies (seven minutes), the welfare numbers (in the same script), and the transitions (22 to 79 minutes each; the four of Section 6.3, two to five hours; the sweeps, hours, if wanted). None of this is calibration.
- **Where the 60 hours would come back.** The ϕ moment is a regression on the model's transitions, and the transitions obey the household's Euler equation $\dot{C}/C = (r_t - \rho)/\gamma$, so the kernels that ϕ is fitted to change with γ: strictly, a new γ moves the fitted ϕ, and re-fitting it is the three evaluations of nineteen hours each. Holding ϕ instead (the rule for trials, and your stated preference) makes the held value an approximation at the new γ, to be flagged in the section. Direction (reasoned, not computed): a smaller γ halves the rise of the interest rate along a transition for a given rise of consumption growth, so q and investment respond a little more to the same utilization gap and the coefficient at a held ϕ comes out a little above the one fitted at $\gamma = 2$; the demand shifts, which carry 77 per cent of the regressor's variance, hardly move aggregate consumption, so the change is small. If ϕ is fixed by convention anyway (the next question), a new γ costs only the recomputation above.

The productivity process within a node. The log productivity z of each industry of a node follows the Ornstein–Uhlenbeck process of Section 3.3, independent across industries, parameterised by its stationary variance $\mathrm{Var} = (\varsigma^Z)^2/(2\theta)$ and its annual autocorrelation $\mathrm{Corr} = e^{-\theta}$, which set how the industries of a node are spread over states. The data's counterpart is a panel of detailed industries inside a node, which exists for manufacturing only: the 355 six-digit industries of the NBER–CES database. Productivity growth is built from the raw series with the model's own weights,

$$\Delta z = \Delta \ln Y - \alpha_n \, \Delta \ln k - (1 - \alpha_n) \, \Delta \ln M,$$

with Y real shipments, k the real stock of equipment and structures, M real materials and energy, and α_n the node's slack-capacity capital share (229); the database's own productivity series treats labour as an input and gives pure profit to capital, and is not used. The node's mean in the year and the industry's own mean growth are removed. Only changes of z can be used, because the levels are index numbers. Let $V(h)$ be the mean squared change of an industry's z over h years; for the process sampled once a year,

$$V(h) = 2 \, \mathrm{Var} \left(1 - \mathrm{Corr}^{\,h}\right). \tag{227}$$

One-year changes measure the size of the innovations, and persistence shows only in changes over long horizons. One process is applied to all nodes: $\mathrm{Corr} = 0.97$, the

value at which the process reproduces the ratio $V(40)/V(1)$ over 1958–2016 (without computers and electronics, whose industries have trends), and $V(1)$ at the median manufacturing node over 1998–2016, a standard deviation of annual growth of 0.058, which gives Var $= V(1)/(2(1-\text{Corr})) = 0.057$. The persistence is matched and not estimated, because the data are close to a random walk; the sensitivities 0.95 and 0.99 move the variance to 0.034 and 0.171. The process is measured in manufacturing only and applied to every node, so that every difference across nodes in the results comes from the network, the capital shares, the final shares and the utilization targets, and not from the shock inputs. It depends weakly on α_n through the weights; rebuilt at the baseline calibration's final ε of 9.6 in place of the working value 10, the variance moved from 0.0568 to 0.0570.

[Derivation of (227)**]** The stationary Ornstein–Uhlenbeck process $\mathrm{d}z = -\theta z\,\mathrm{d}t + \varsigma^Z \mathrm{d}B$ has mean zero, variance Var $= (\varsigma^Z)^2/(2\theta)$ and autocovariance $\mathrm{Cov}(z_t, z_{t+h}) = \mathrm{Var}\, e^{-\theta h}$: from $z_{t+h} = e^{-\theta h} z_t + \int_t^{t+h} e^{-\theta(t+h-u)} \varsigma^Z \mathrm{d}B_u$ and the independence of the integral from z_t, $\mathrm{Cov}(z_t, z_{t+h}) = e^{-\theta h}\mathbb{V}\mathrm{ar}(z_t)$. Hence

$$
\begin{aligned}
V(h) &= \mathbb{E}\left[(z_{t+h} - z_t)^2\right] = \mathbb{V}\mathrm{ar}(z_{t+h}) + \mathbb{V}\mathrm{ar}(z_t) - 2\mathrm{Cov}(z_t, z_{t+h}) \\
&= 2\,\mathrm{Var}\left(1 - e^{-\theta h}\right) = 2\,\mathrm{Var}\left(1 - \mathrm{Corr}^h\right),
\end{aligned}
$$

with Corr $= e^{-\theta}$. At $h = 1$, $V(1) = 2\mathrm{Var}(1 - \mathrm{Corr}) = (\varsigma^Z)^2(1 - e^{-\theta})/\theta \approx (\varsigma^Z)^2$ for small θ, so one-year changes measure the innovation size, and $V(40)/V(1) = (1 - \mathrm{Corr}^{40})/(1 - \mathrm{Corr})$ depends on the persistence alone. Removing each industry's own mean removes permanent differences in levels, which the model does not have, and leaves every difference $z_{t+h} - z_t$ unchanged. Removing the node–year mean removes what is common to the node's industries in a year, which in the model is zero by the law of large numbers across the continuum of independently shocked industries; in a finite cell of I industries it scales $V(h)$ by $1 - 1/I$ at every h, which cancels in $V(40)/V(1)$ and understates the innovation size by the factor $\sqrt{1 - 1/I}$, a few per cent for the nodes' panels. $V(h)$ is estimated as the mean squared h-year difference of the demeaned series. Industry-specific trends would add to $V(h)$ at long horizons, which is why computers and electronics are left out of the persistence match. The formulas treat the annual observations as point samples of z; if the data are averages of z over the year, the same θ gives a smaller $V(1)$ and a larger $V(40)/V(1)$, so the persistence that a given ratio implies is lower than the point-sampling formula says.

[Status] Proof-checked by a fresh reviewer: sound, with gaps now stated in the box. The formulas were re-derived and confirmed by simulation of the exact annual process ($\theta = 0.03$, $\varsigma^Z = 0.0588$, 200,000 paths: empirical over formula 1.0002, 1.0006 and 0.9972 at $h = 1$, 10 and 40). The reviewer's two findings, the finite-cell effect of the node–year demeaning and the point-sampling assumption, are in the box and in the flag above.

[What is set rather than measured in the productivity process, and the robustness variant] Point sampling against annual averages: settled, point samples kept. The reviewer of the derivation above computed that if the annual productivity series

are averages over the year, inverting the point-sampling formulas on them returns a mean-reversion speed about a fifth of the true one and an innovation size about four fifths of it (at $\theta = 0.03$: the ratio $V(40)/V(1)$ would be about 35 instead of 24 for the same θ). The question was settled on the panel's pooled variogram (without computers and electronics; computed once, the first step reproducing the data script's number): matching $V(40)/V(1) = 22.5$ gives $\theta = 0.033$ (Corr $= 0.967$, Var $= 0.056$) under the point reading and $\theta = 0.063$ (0.939, 0.046) under the averaging reading; the point reading fits the short horizons better ($V(2)/V(1)$ implied 1.97 against 1.82 in the data, 2.43 under averaging; root-mean-square error of $\ln V$ over $h = 1, \ldots, 40$: 0.28 against 0.45). The data's $V(2)/V(1)$ below 2 is the signature of measurement error in the levels; with a white error added, both readings go to a random walk with the error at 60–80 % of $V(1)$ and cannot be told apart, so what the data leave uncertain is the innovation size, as said below. The point reading is kept: it fits better under the matching rule in force, and the model's own annual moments are point samples, so the comparison is internally consistent. Three things are set. The process holds outside manufacturing, where no panel exists (the candidates, productivity from the BEA's detailed accounts, a price-based measure and labour productivity, fail where they can be checked against the manufacturing panel). The persistence is a match to long-horizon dispersion and not an estimate of a stationary process: estimators that avoid the short-panel bias give values between 0.96 and 0.99 without computers and electronics, and exactly one with them, and the estimates rise with the length of the panel, which a random walk plus a short-lived component fits better than an autoregression. And the innovation size rests on the window, 1998–2016 (over 1958–2016 the variance would be about 30 % smaller), and about half of the one-year variance reverses within two to three years, which may be measurement error. The stakes for the targets are small: on a two-node economy a much lower persistence changed the refitted capital intensities by about 1 % and left the response of GDP to a productivity shock unchanged; what responds is the within-node spread of capital. The robustness variant is a node-specific innovation size for the manufacturing nodes by an empirical-Bayes rule, with the manufacturing mean for the other nodes. (The statements of this box are from the record of the calibration, at earlier values; the numbers of the black text are the data set's.)

[Is the productivity process as good as it can be?]

- **What is identified, and from what.** Two numbers, from the 355 six-digit manufacturing industries of the NBER–CES panel. Corr comes from the long-horizon ratio $V(40)/V(1) = 22.5$ in the pooled panel: for an Ornstein–Uhlenbeck process the ratio is $(1 - \text{Corr}^{40})/(1 - \text{Corr})$, which is 17.4 at 0.95, 23.5 at 0.97, 33.1 at 0.99 and 40 for a random walk. Var, the spread of log productivity across the industries of a node, then follows from the one-year growth variance, $\text{Var} = V(1)/(2(1 - \text{Corr}))$ with $\sqrt{V(1)} = 0.058$. The innovation size is measured; the persistence is a match to one long-horizon moment in a short panel, which is why the text above says "matched, not estimated"; the implied half-life of an industry's deviation is 23 years.

- **What is weak, in order.**
  1. Var *is the least well identified object.* With persistence near one, $V(1)/(2(1-\text{Corr}))$ is hypersensitive to Corr: 0.034, 0.057 and 0.171 at 0.95, 0.97 and 0.99. Choosing Corr is in effect choosing Var within a factor of two or more. Measurement error does not rescue it: in the data $V(2)/V(1) = 1.82 < 2$, the signature of white noise in the levels, and a fit of $V(h) = 2\text{Var}(1-\text{Corr}^h) + 2\sigma_e^2$ to $V(1)$, $V(2)$ and $V(40)$ gives Corr $= 0.978$, noise worth 16 % of $V(1)$ and Var $= 0.064$, up rather than down, because the corrected persistence is higher (computed on 25 October 2026 from the ratios above; a check, not a result); a fit of the full variogram drifts to a random walk, where Var is undefined. What the data do pin without a model is a floor, $V(40)/2 = 0.039$, which the stationary variance cannot be below. The two sensitivities, Corr $= 0.95$ and 0.99, are listed as robustness variants in Section E.6 and have not been run on 63 nodes.
  2. *One process for 63 nodes.* It is measured in manufacturing and applied to every node by assumption; the candidates for the other nodes failed where they could be checked. Within manufacturing $V(1)$ differs across the 19 nodes; the variant listed in Section E.6, a node-specific innovation size by an empirical-Bayes rule for the manufacturing nodes and the mean elsewhere, has not been run either.
  3. *What z is.* The counterpart is built with the model's weights on capital and materials only, so labour is inside "productivity". Consistent with a model without labour, but the measured innovation size contains variation in labour input.
  4. *Independence across industries* is imposed; the node–year means are removed, which takes the common component out. The common shock is used only in the validation and in the moment for ϕ.
- **How much it matters.** On two nodes, little: a much lower persistence moved the refitted capital intensities by about 1 per cent and left the response of GDP to a productivity shock unchanged; what responds is the within-node spread of capital and utilization (the box above, at earlier values). On 63 nodes this has not been shown.
- **Answer.** Under one process the two moments are the right ones and are matched correctly, so the point estimate is as good as the panel allows; what is missing is the evidence that Var, which the data leave loose within a factor of two, does not drive the results. The clear improvement is therefore not a better estimate but the two sensitivity runs at 0.95 and 0.99 (a second stage each, minutes, then the static block), and only if they move the aggregate results, a direct cross-sectional moment for Var (the within-node dispersion of utilization in the Census survey, median 0.086, is the candidate, with the caveat that it contains permanent differences across industries that the model does not have). The node-specific variances come after that. All of it reuses the existing panel.

E.5 Internally Calibrated Parameters

E.5.1 The dispersion of demand and the elasticity of substitution, σ_v and ε

Given a node's threshold, every object of the capacity block is a function of two parameters: the dispersion σ_v of firms' demand shocks (60) and the elasticity of substitution ε between firms' goods (2). The utilization target pins the threshold, so two further observables are needed, and both come from US manufacturing. Neither parameter can be taken from the literature or from the accounts. The model has no overhead cost, so ε is at once the elasticity of demand and the reciprocal of the markup rent's share of revenue (Lemma 5(iii)); estimated elasticities between firms' goods, typically 2 to 4, would make that rent larger than the whole operating surplus of a typical manufacturing industry. And the accounts give one relation between ε and the interest rate, since a node's operating surplus of physical capital is the user cost of its capital plus the markup rent; the capacity block pins ε, and the accounts then give the interest rate. One σ_v for all nodes is a data concession: node-specific dispersions could be read only off the European surveys, on which the transfer targets already rest.

The share of capacity above 95 per cent identifies σ_v. In the model all firms of an industry have the same capacity, so the share of an industry's capacity in firms that produce above 95 % of it is the share of such firms, $1 - G(0.95\,\overline{v})$, the firms whose demand shock exceeds $0.95\,\overline{v}$; the counterpart of the Census share is its mean over the 19 manufacturing nodes at their utilization targets (unweighted in the first stage, with capacity weights on the solved model). At a given utilization the share depends almost only on σ_v: the more firms differ, the more of them are close to capacity. The value at which the model reproduces 0.162 is between 0.184 and 0.185 for every ε between 8 and 15 in the baseline calibration.

The supply curves identify ε given σ_v. In the model, shift the demand for an industry's good while its capital is fixed. Output and the price move, and the ratio of the two responses is the model's inverse supply elasticity. It is flat while almost no firm is rationed and rises as firms reach capacity: with a small σ_v firms reach capacity almost together, so the curve stays flat and then turns steep; with a large σ_v some firms are rationed already at moderate utilization and the curve rises gradually. The four bins of Table 13 measure this rise. At a given σ_v the rise depends on ε as well, because a lower ε gives a larger markup response per rationed firm, and the supply curves alone are fitted almost equally well along a line of pairs on which σ_v rises with ε. The share selects the point on that line. The model's value for each bin mirrors the estimator in three steps: each industry-year of the authors' sample is put on the model's curve at its utilization at the start of the year and moved to its utilization at the end, and the model's elasticity is averaged along that path; the model's bin value weights the industry-years as the instrumental-variables estimator does; and the regression's control for unit variable cost is applied to the model with the estimated pass-through and with production-worker pay proportional to output. The estimates are in deviations from each industry's own mean, and the model's curve depends on

the level; every industry is placed at one common level plus its deviation, the mean target of the 19 manufacturing nodes, 0.786, so that the economy that is fitted is the economy that is calibrated.

[The model's inverse supply elasticities and the counterpart of the four estimates]

- *Objects.* Fix an industry's capital, productivity and bundle price, so that $c := \mathcal{P}/(\overline{Z}e^{z}x^{\alpha})$ and capacity $\overline{Y}$ are fixed, and write $\ell = \ln \overline{v}$. The static block (139)–(140) gives the price $P = c/\Upsilon(\overline{v})$, output $Y = \overline{Y}U(\overline{v})$, sales $S = PY$ and intermediate spending $X = \chi(\overline{v})S$ with $\chi := \xi\varrho$. Recorded revenue is sales (Section E.3), so the measured price, revenue over real output, is P, and the intermediate part of unit variable cost is $\mathrm{UVC}^{x} := X/Y = \chi P$.
- *Log-derivatives.* Let $\ell_U := -\,\mathrm{d}\ln U/\mathrm{d}\ell$, $\ell_\Upsilon := \mathrm{d}\ln\Upsilon/\mathrm{d}\ell$ and $\chi_\ell := \mathrm{d}\ln\chi/\mathrm{d}\ell$. A movement of the threshold at fixed $(c, \overline{Y})$ gives $\mathrm{d}\ln Y = -\ell_U \mathrm{d}\ell$, $\mathrm{d}\ln P = -\ell_\Upsilon \mathrm{d}\ell$ and $\mathrm{d}\ln \mathrm{UVC}^{x} = (\chi_\ell - \ell_\Upsilon)\mathrm{d}\ell$.
- *Two inverse elasticities.* The slope of the price at fixed input prices and the slope at fixed intermediate unit cost are

$$m^{a} := \frac{\mathrm{d}\ln P}{\mathrm{d}\ln Y} = \frac{\ell_\Upsilon}{\ell_U}, \qquad m^{x} := \left.\frac{\mathrm{d}\ln P}{\mathrm{d}\ln Y}\right|_{\mathrm{UVC}^{x}} = \frac{\chi_\ell}{\ell_U}.$$

  $m^{a} = 1/(e_S - 1)$ with $e_S = 1 + \ell_U/\ell_\Upsilon$ the supply elasticity of sales (111); it tends to $\alpha/(1-\alpha)$ with slack capacity. m^{x} is the response of $1/\chi$, the part of the price that is not intermediate unit cost, and tends to zero: holding unit variable cost fixed removes the marginal-cost channel.
- *The coefficient that is estimated.* The regression controls for the growth of measured unit variable cost, which includes production-worker pay, with a freely estimated coefficient $\widehat{\tau}$ (0.89). Along a demand shift the measured cost moves with output at the rate e_{UVC}, so the output coefficient identifies $m^{a} - \widehat{\tau}e_{\mathrm{UVC}}$. If pay is proportional to output and s_W is pay's share of measured unit variable cost (0.146), then $\mathrm{d}\ln\mathrm{UVC} = (1 - s_W)\mathrm{d}\ln\mathrm{UVC}^{x}$ and $e_{\mathrm{UVC}} = (1 - s_W)(m^{a} - m^{x})$, so the model's counterpart of the coefficient is

$$m^{w} = m^{a} - \widehat{\tau}\, e_{\mathrm{UVC}} = m^{x} + [1 - \widehat{\tau}(1 - s_W)](m^{a} - m^{x}).$$

  If instead pay were fixed within the year, $e_{\mathrm{UVC}} = (1 - s_W)(m^{a} - m^{x}) - s_W$; the data reject this reading (pay per unit of output does not respond to demand-driven output in any bin).
- *From elasticities to bins.* Each observation is a change over a year, so the counterpart for industry i in year t is the arc elasticity, the average of m^{w} along the path of $\ln U_n$ from $U_{i,t-1}$ to $U_{i,t}$, with every industry placed at 0.786 plus its deviation. The bin estimates are two-stage least squares coefficients: if every observation obeyed $\Delta\ln P_i - \widehat{\tau}\Delta\ln\mathrm{UVC}_i = \overline{m}_i^{w}\Delta\ln Y_i$, the coefficient of bin b would be $\sum_i A_{bi}\,\overline{m}_i^{w}$, a weighted sum over the whole sample of the arc elasticities, with weights A_{bi} that are determined by the instruments, the controls and the output

changes, sum to one over the industry-years of bin b and to zero over each other bin, and can be negative. This is the model's counterpart of bin b. The counterpart holds the industry's capital and bundle price fixed within the year; the control for unit cost removes most of what a movement of input prices would add to the coefficient, and the check on simulated data (below) tests the rest. The distance to the four estimates is the quadratic form in the differences with the inverse of the estimates' covariance matrix.

[Status] Proof-checked by a fresh reviewer: sound, no algebraic error. The three log-derivative relations, $m^a = 1/(e_S - 1)$, the slack limits ($m^a \to \alpha/(1-\alpha)$, $m^x \to 0$) and the algebra of m^w under both readings of pay were re-derived and confirmed numerically to eight decimals; the two-stage-least-squares identity behind the bin counterpart was confirmed by simulation. Two findings were folded in: the weights A_{bi} depend on the output changes as well and sum to zero, not vanish, outside the bin; and the counterpart holds capital and the bundle price fixed within the year, which is stated and which the check on simulated data tests. The reviewer also noted that $\ell_Y > 0$ is used without a lemma; it holds numerically over the whole admissible range.

The rule and the values. At every ε, σ_v is the value at which the model's share above 95 % equals the Census share; ε is the value at which the distance between the four model counterparts and the four estimates is then smallest. Each calibration has its own two values by this rule, because the spread of the manufacturing targets enters the share: in the baseline calibration $\varepsilon = 9.6$ with $\sigma_v = 0.185$, where the model's four bins are 0.11, 0.20, 0.28 and 0.55 against 0.12, 0.28, 0.31 and 0.57 in the data ($\chi^2 = 1.1$ with three degrees of freedom); in the secondary calibration $\varepsilon = 10.3$ with $\sigma_v = 0.199$. The baseline calibration's values are the solved model's at its final ϕ; the secondary calibration's are at the interim ϕ, since its simulation of ϕ is deferred (Section E.5.4). At $\varepsilon = 10$ the markup rent is a tenth of firm revenue and the markup with slack capacity is 1.11.

ε is not sharply estimated. The sampling error of the four estimates puts the upper end of the 95 % interval at 15.5 in the baseline calibration and 16.6 in the secondary calibration, and the lower end at 7.7 in the secondary calibration; in the baseline calibration the distance has not reached its 95 % level when, at ε between 7.3 and 7.5, no interest rate above the growth rate reproduces the capital stock of the accounts, so the interval's lower end is this bound. There is more room above than below; results are reported at both ends, with σ_v at the value that the share requires there. Three choices of the mapping move the estimate by more than its interval does. Leaving out one industry, wood products, raises it to 12.4; leaving out any other industry moves it much less. Re-estimating the supply curves with the materials price held fixed instead of unit variable cost gives 7.0; this rests on the tightest bin, is noisier, and counts the rise of unit materials cost in that bin as a markup, so it is a check and not a better estimate (Section E.6). And the share is measured against the plants' full production capability in the Census survey, while the supply curves and the targets are on the Federal Reserve's scale, about five points higher; with utilization five points lower everywhere the estimate is 7.2, with $\sigma_v = 0.27$. The counterpart itself was checked by running the

authors' regression on data simulated from the model after shifts of each manufacturing node's weight in final demand, at the shock date, in the 63-node economy, in two forms of the experiment. The general-equilibrium response of input prices, which the formula holds fixed, moves the four coefficients by about one standard error of the estimates, downward in one form and upward in the other, so the difference is not clear and the recipe stands. The dispersion of firm utilization that the values imply is compared with Swedish firm data as a check, not a target (Section E.6).

[Flag closed: the check on simulated data, two forms, and the decision] The recipe as agreed, every manufacturing node's weight in final demand shifted so that its utilization moves from the start-of-year level (0.786 plus the industry's deviation) to the end-of-year level of each industry-year, with every other node responding in general equilibrium and capital fixed, is infeasible for the nodes whose final demand is a small share of their demand (primary metals, paper, textiles, wood): no admissible weight moves them ten points. The bounded form limits the shift of a node's final demand to a factor of $e^{1.5}$ either way, drops the node-years that cannot be reached (572 of the 819 industry-years are kept; the bins hold 62, 189, 224 and 97 of them) and runs the authors' regression on the kept rows twice: on the general-equilibrium changes and on the partial-equilibrium changes of the formula's recipe for the same rows. Result (baseline calibration, solved model at the interim ϕ, computed once): general equilibrium -0.02, 0.03, 0.10, 0.26 against partial equilibrium 0.07, 0.18, 0.27, 0.47, a difference of 0.09 to 0.22 in every bin, 1.1 to 1.9 of the data's standard errors; the unit-cost pass-through is 1.77 in general equilibrium against 0.52 in partial equilibrium and 0.89 in the data. The channel is the one the formula holds fixed: a node's demand shift raises its input prices, and the control for unit variable cost absorbs that part of the price rise. The check has two weaknesses of its own, the large simultaneous level placements and the thinned outer bins, so a second form was run (each node moved by its within-year change from its own calibrated level; 682 industry-years kept, the bins holding 91, 244, 254 and 93): there general equilibrium gives 0.33, 0.24, 0.17, 0.12 against partial equilibrium 0.22, 0.19, 0.13, 0.10 on the same rows, above rather than below, by 1.4 standard errors in the first bin and within half a standard error in the others, with pass-throughs 1.39 and 1.27 (the bins' order is inverted in this form because every industry sits at its own level, so the bins differ only by the direction and size of the within-year change). So the general-equilibrium channel exists, the pass-through exceeds one in both forms, but its effect on the four coefficients changes sign with the state the simulation is run at and is not robust. A decisive form would move one node at a time to its placed level (no cross-contamination of simultaneous placements), at about 1,600 fixed-law solves. Decision (11 October 2026): the difference is not clear in the protocol's sense, the recipe stands, and the decisive form is not run, in favour of results from the baseline calibration; both forms ran at the interim ϕ and are computed once. The black text above states the outcome.

[What each choice is worth in ε, both calibrations; the identification; status] Regenerated in this session on the 63-node data set at the decided capital income, with the audit's estimator in the decided order (every row: σ_v from the share at each ε, then

the minimum distance):

change, one at a time	baseline calibration	secondary calibration
none (the estimate; σ_v 0.193 / 0.200; χ^2 1.0 / 1.0)	9.9	10.3
without wood products	12.4	12.9
the supply curves re-estimated with the materials price held fixed	7.0	7.2
utilization on the Census scale, five points lower (σ_v 0.27 / 0.28)	7.2	7.4
pay fixed within the year (rejected by the data)	10.9	11.3
firm revenue per unit of output as the measured price	9.6	9.9
95 % interval; one-standard-error interval	7.5–15.7; 8.5–12.1	7.7–16.2; 8.8–12.5

- **Identification.** The share alone requires σ_v between 0.1941 ($\varepsilon = 7$) and 0.1928 ($\varepsilon = 13$) in the baseline calibration (0.2009 to 0.1995 in the secondary calibration); every change of the mapping moves ε and leaves σ_v there. So the share pins σ_v and the supply curves pin ε given σ_v, which is how the black text states it. In the other order (fit σ_v to the curves at each ε, then match the share) the estimates were a little lower; both orders match the share exactly, and the decided order is the one that the description says.
- **Levels: tightness or convention.** The fit places every manufacturing industry at the common level 0.786 plus its deviation and reads an industry's mean level as a convention of how its capacity is measured; the baseline calibration's targets read the core nodes' levels as tightness. Within manufacturing the supply curves speak for the convention: evaluating each industry at its own level fits the four estimates much worse, and industries with a higher mean utilization do not have steeper supply curves at the same deviation from their own mean. The estimator puts much of its weight on two industries with low mean utilization, wood products and nonmetallic minerals. Aggregates hardly depend on the reading; differences between manufacturing nodes do, which is why they are shown under both calibrations (Section E.5.3). For mining and utilities there is no evidence either way.
- **The intervals** hold σ_v at the value that the share requires and leave out the sampling error of the share itself, which is smaller. They are computed on the first stage and are computed again on the solved model.
- **Status.** The table is first stage. Its estimate row and its intervals are superseded by the solved model (computed once; the baseline calibration at $\phi = 2$, the secondary calibration at the interim ϕ): $\varepsilon = 9.61$ / 10.34, $\sigma_v = 0.1851$ / 0.1991, χ^2 1.08 / 0.97, one-standard-error interval 8.2–11.8 / 8.9–12.7, 95 % upper end 15.4 / 16.6, lower end the admissibility bound (between 7.3 and 7.5) / 7.7. The variant rows are the first stage's (the audit's estimator re-run once on the decided data set) until they are re-run on the solved model. The first-stage estimates themselves reproduce the record exactly, and the σ_v half is cross-examined (the inde-

pendent implementation reproduces the Census share 0.1623 at $\sigma_v = 0.1933$ to four decimals). The derivation of the counterpart (the box above) was re-derived in the review of the calibration and checked there against numerical derivatives of the static block; it is proof-checked again for this appendix (status below the box).

Why the two are identified separately (Figure 11). The Census share responds to σ_v and hardly to ε: it is 0.162–0.164 at every ε between 7.5 and 21 and runs from 0.05 to 0.21 along σ_v, through the data's 0.162 at the calibrated value. The supply curves respond to ε given σ_v: the χ^2 distance is 1.1 at the estimate, 2.4 at $\varepsilon = 8$ and 8.5 at 21. The χ^2 also moves along σ_v (8.3 at 0.105, 1.9 at 0.22) and is smallest at the estimate in both directions. So σ_v is pinned by the Census share alone, and ε by the supply curves given σ_v. The figure is drawn from a sweep of the two parameters, the calibration re-done at every point.

[Flag, 3 November 2026] The sweep behind this paragraph and Figure 11 was run on the previous baseline calibration ($\phi = 13$, $\gamma = 2$; $\varepsilon = 9.60$, $\sigma_v = 0.184$) and is not re-run on the baseline of 3 November 2026 ($\phi = 2$, $\gamma = 1$; $\varepsilon = 9.61$, $\sigma_v = 0.185$). The two moments are static objects of the threshold, σ_v and ε, so the figure stands to the eye; the numbers of the paragraph are the earlier sweep's.

E.5.2 The capital shares, α_n

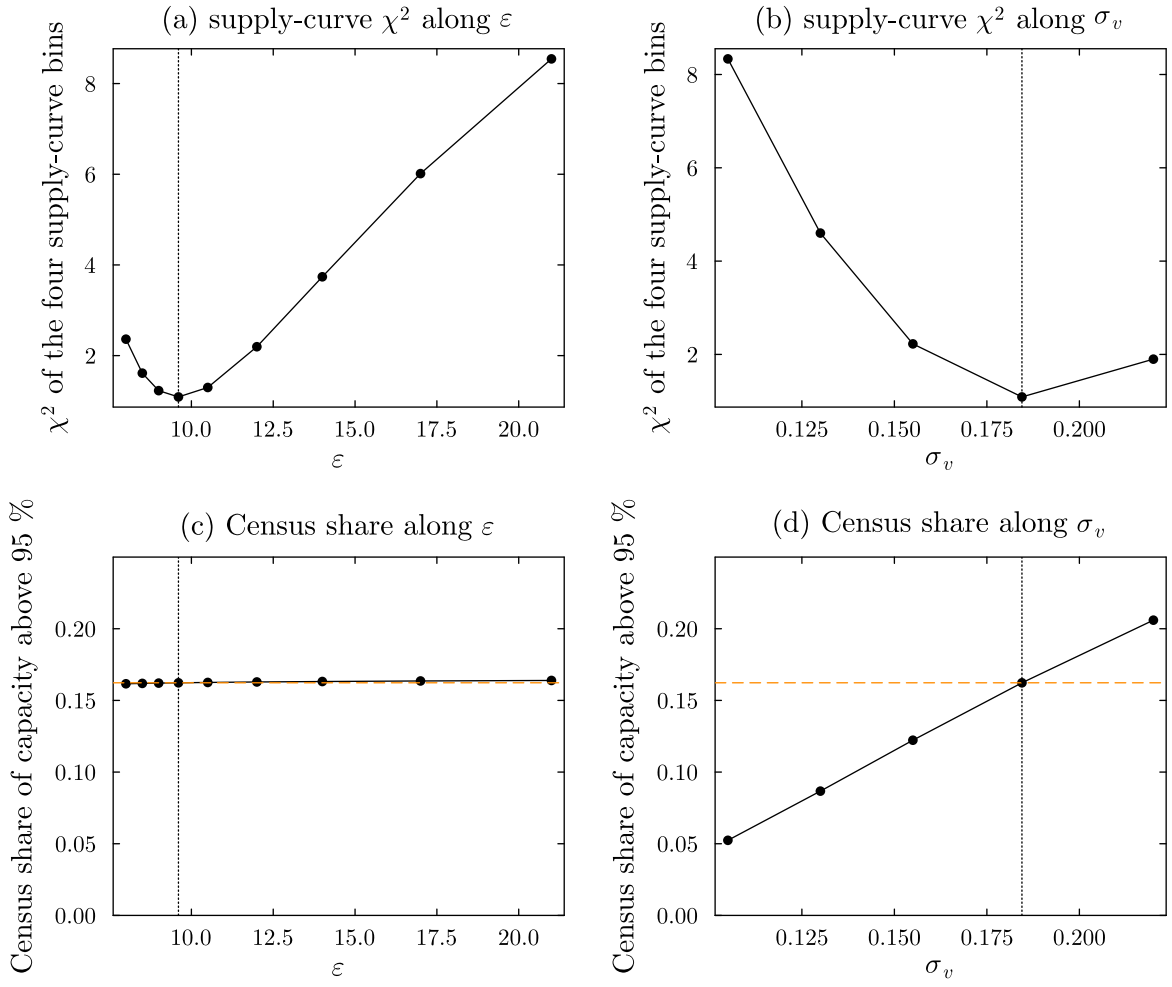
The capital share α_n, the exponent of capital in (8), is inferred from the division of a node's sales between intermediate inputs and the income of capital, which the accounts record and the model determines given α_n. On the accounts' side the object is the ratio ω_n of (225): the operating surplus of physical capital, net of the imputed pay of the self-employed and of the income of intellectual property, over the model's sales, the mean of the annual ratios of 2007–2019. The imputed pay of the self-employed is large in a few service industries, and the authors of the account note that the imputation may not hold in social assistance and performing arts; one rule is applied to all nodes nonetheless. The ratio runs from 0.042 (social assistance) to 0.604 (pipelines); it averages 0.164 over the 19 manufacturing nodes and 0.235 over all 63.

In the model a unit of a node's sales divides into four parts (Lemma 5(iv)): intermediate inputs, $\xi_n \varrho_n$; the income of capital; the markup rent, ϱ_n/ε; and the rationing rent, $1 - \varrho_n$. Here $\varrho_n(\overline{v}_n)$of (118) is the share of sales that accrues to the node's firms and $\xi_n(\overline{v}_n; \alpha_n)$ of (75) the share of their revenue that they spend on inputs; both are functions of the node's threshold $\overline{v}_n$, and ξ_n of the capital share as well. Recorded revenue is the model's sales at the shadow price index, and the operating surplus of physical capital contains the income of capital and both rents (Section E.3). The accounts' ratio therefore corresponds to one minus the intermediate share of sales, and α_n is the value that solves

$$\xi_n(\overline{v}_n; \alpha_n)\, \varrho_n(\overline{v}_n) = 1 - \omega_n. \tag{228}$$

The threshold is set differently in the two kinds of nodes.

Figure 11: Identification of the two capacity parameters.



Notes: The two calibration moments along the two parameters, the calibration re-done at every point. Top row, the χ^2 distance between the model's and the estimated inverse supply elasticities by utilization bin, (a) along ε and (b) along σ_v. Bottom row, the model's share of manufacturing capacity in plants above 95 per cent of capacity, (c) along ε and (d) along σ_v; the dashed line is the data, 0.162. The dotted line marks the calibrated value.

[The two functions of (228) under the log-normal specification] Firms' demand shocks have mean one and $\ln v \sim \mathcal{N}(-\sigma_v^2/2, \sigma_v^2)$, as in (60); a firm is rationed when $v > \overline{v}$. Write $\ell := \ln \overline{v}$, $\Phi_{\mathcal{N}}$ for the standard normal distribution function and $\mu_\theta := \mathbb{E}[v^\theta] = e^{\sigma_v^2 \theta(\theta-1)/2}$ for the moments of the shock.

1. *The tail moments of* (57)–(58). The share of firms that are not rationed is $G(\overline{v}) = \Phi_{\mathcal{N}}(\ell/\sigma_v + \sigma_v/2)$. For $\theta \geq 0$, the change of variable $\ln v = -\sigma_v^2/2 + \sigma_v w$ gives $\int_0^{\overline{v}} v^\theta \mathrm{d}G(v) = \mu_\theta \, \Phi_{\mathcal{N}}(\ell/\sigma_v + \sigma_v/2 - \theta\sigma_v)$, so that

$$\mathcal{G}_\theta^-(\overline{v}) = \mu_\theta \, \overline{v}^{-\theta} \, \Phi_{\mathcal{N}}\left(\frac{\ell}{\sigma_v} - \left(\theta - \tfrac{1}{2}\right)\sigma_v\right), \qquad \mathcal{G}_\theta^+(\overline{v}) = \mu_\theta \, \overline{v}^{-\theta} - \mathcal{G}_\theta^-(\overline{v}),$$

   and $\Psi_\theta = \mathcal{G}_\theta^- + 1 - G$, $\eta_\theta = \mathcal{G}_\theta^- / \Psi_\theta$. (The node index is dropped: one σ_v serves every node, so the functions differ across nodes only through the threshold and α_n.)

2. *The firms' share of sales and utilization* depend on (σ_v, ε) and the threshold only. From (118) and (22), with $\mu_{1/\varepsilon}^{-\varepsilon/(\varepsilon-1)} = e^{\sigma_v^2/(2\varepsilon)}$,

$$\varrho_n(\overline{v}) = \frac{\Psi_1(\overline{v})}{\mathcal{G}_1^-(\overline{v}) + \mathcal{G}_{1/\varepsilon}^+(\overline{v})}, \qquad U_n(\overline{v}) = e^{\sigma_v^2/(2\varepsilon)} \left[\overline{v}^{1/\varepsilon}\left(\mathcal{G}_1^-(\overline{v}) + \mathcal{G}_{1/\varepsilon}^+(\overline{v})\right)\right]^{\frac{\varepsilon}{\varepsilon-1}}.$$

   U_n is a strictly decreasing bijection of $(0,\infty)$ onto $(0,1)$ (Lemma 3(ii)), so every target in $(0,1)$ fixes one threshold.

3. *The intermediate share of firm revenue* is (75), with $\theta_n := 1/(1-\alpha_n)$:

$$\xi_n(\overline{v}; \alpha_n) = \frac{\eta_1(\overline{v}) - \frac{1}{\varepsilon}}{\theta_n \, \eta_{\theta_n}(\overline{v})}.$$

4. *The slack limit.* As $\overline{v} \to \infty$ every $\eta_\theta \to 1$ and $\mathcal{G}_{1/\varepsilon}^+ \to 0$, so $\varrho_n \to 1$ and $\xi_n \to (1 - \frac{1}{\varepsilon})(1 - \alpha_n)$ (Lemma 5(v)). The left side of (228) then equals $(1 - \frac{1}{\varepsilon})(1 - \alpha_n)$, which gives (229); its root lies in $(0,1)$ if and only if $1/\varepsilon < \omega_n < 1$.

[Status] Proof-checked by a fresh reviewer: sound. Every closed form was re-derived and matched against quadrature of the log-normal density over 36 combinations of $(\sigma_v, \theta, \overline{v})$ to 10^{-15}; U_n strictly decreasing was proved from the closed form; the slack reduction was confirmed to 10^{-14} at large thresholds. One condition was added: the slack formula gives a share in $(0,1)$ only for $1/\varepsilon < \omega_n < 1$.

A node in which capacity is off sits at the slack-capacity limit: no firm is rationed, $\varrho_n = 1$, and $\xi_n = (1 - \frac{1}{\varepsilon})(1 - \alpha_n)$ (Lemma 5(v)). Equation (228) has the closed-form solution

$$\alpha_n = \frac{\omega_n - 1/\varepsilon}{1 - 1/\varepsilon}. \tag{229}$$

The markup rent, the share $1/\varepsilon$ of sales, is taken out of the operating surplus, and the income of capital is expressed as a share of sales net of that rent. The data do not separate the rent from the income of capital; the value of ε does, so every value of ε implies its own capital shares. (Setting $\alpha_n = \omega_n$ would count the rent twice, inside ω_n

and in the model.) Formula (229) is used for the 14 nodes without a utilization rate in the baseline calibration; in the secondary calibration capacity is on in every node.

In a node in which capacity is on, (228) is solved at the node's utilization target $\overline{U}_n$ of Section E.5.3. utilization $U_n(\overline{v})$ of (22) is a decreasing function of the threshold that depends on (σ_v, ε) only, so the target fixes the threshold, $U_n(\overline{v}_n) = \overline{U}_n$; the threshold fixes ϱ_n; and α_n is the root of (228), found numerically. Under the log-normal law the left side of (228) is strictly decreasing in α_n, so the root is unique when it exists, and a root in $(0,1)$ exists if and only if

$$\omega_n > 1 - \varrho_n(\overline{v}_n) + \frac{\varrho_n(\overline{v}_n)}{\varepsilon\, \eta_{n,1}(\overline{v}_n)}, \tag{230}$$

the share of sales that the rationing rent and the margin of the posted price would take if the technology used no capital. Here $\varepsilon\eta_{n,1}(\overline{v}_n)$ is the effective demand elasticity of Lemma 1(iii), which is below ε when capacity binds. With slack capacity the condition is $\omega_n > 1/\varepsilon$.

At a given capital share, binding capacity lowers the intermediate share of sales below its slack value, for two reasons. Firms that expect to be rationed post a higher margin: the share $1/(\varepsilon\eta_{n,1})$ of revenue instead of $1/\varepsilon$, so less of each unit of revenue is spent on inputs; and a small part of sales is rationing rent. The accounts fix the intermediate share at $1-\omega_n$, so a node with binding capacity needs a technology that uses more inputs, that is, a smaller α_n, and the more so the tighter the node. Over the capacity nodes that are not floored, the pinned share is below the slack-capacity formula (229) by 0.023 at the median and by 0.057 at most in the baseline calibration, and by 0.012 and 0.018 in the secondary calibration, whose targets are lower; in the nodes with the largest capital shares the difference is small and of either sign. In the solved model, with the productivity process on, (228) is imposed on the node's mean intermediate spending over its mean sales (Section E.5.4).

[Uniqueness and existence of the root of (228)] Fix the threshold $\overline{v} = \overline{v}_n$ and write $\theta = 1/(1-\alpha_n) \in (1,\infty)$, so that $\alpha_n \mapsto \theta$ is increasing. By item 3 of the previous box the left side of (228) is $\varrho_n(\overline{v})\,(\eta_1 - \frac{1}{\varepsilon})/(\theta\eta_\theta)$, and $\eta_1 > \frac{1}{\varepsilon}$ at every threshold of the symmetric equilibrium (Lemma 5(i)); at the calibrated values this holds for every utilization target below 0.99, so for all targets used. The left side is therefore decreasing in α_n if and only if $\theta\eta_\theta(\overline{v})$ is increasing in θ.

1. *The derivative.* Substituting $u := \ln(\overline{v}/v)$ in (57) gives $\mathcal{G}_\theta^-(\overline{v}) = \int_0^\infty e^{-\theta u} f(u)\mathrm{d}u$ with $f(u) := \overline{v}e^{-u}g(\overline{v}e^{-u})$, where g is the density of v; f is the density of u over the firms that are not rationed, not normalised. Hence $\mathrm{d}\ln\mathcal{G}_\theta^-/\mathrm{d}\theta = -\mathbb{E}_w[u]$, the mean of u under the weights $w(u) \propto e^{-\theta u}f(u)$, and since $\Psi_\theta = \mathcal{G}_\theta^- + 1 - G$ with $1-G$ free of θ,

$$\frac{\mathrm{d}\ln(\theta\eta_\theta)}{\mathrm{d}\theta} = \frac{1}{\theta} + (1-\eta_\theta)\frac{\mathrm{d}\ln\mathcal{G}_\theta^-}{\mathrm{d}\theta} = \frac{1}{\theta} - (1-\eta_\theta)\,\mathbb{E}_w[u].$$

Under the log-normal specification, with $z := \ell/\sigma_v - (\theta - \frac{1}{2})\sigma_v$ and $\lambda(z) := \phi(z)/\Phi_{\mathcal{N}}(z)$ the inverse Mills ratio, differentiating the closed form of item 1 gives $\mathbb{E}_w[u] = \sigma_v[z + \lambda(z)]$.

2. *The two ends.* As $\alpha_n \to 0$, $\theta \to 1$ and the left side of (228) tends to $\varrho_n\,(1 - \frac{1}{\varepsilon\eta_1})$: the firms' share of sales less the margin of the posted price, which is the share $1/(\varepsilon\eta_1)$ of revenue by (62). As $\alpha_n \to 1$, $\theta \to \infty$ and $\theta\mathcal{G}_\theta^- \to f(0^+) = \overline{v}g(\overline{v})$ (Watson's lemma: $\theta\int_0^\infty e^{-\theta u}f(u)\mathrm{d}u \to f(0^+)$ for f bounded with a limit at 0), so $\theta\eta_\theta \to \overline{v}g(\overline{v})/[1 - G(\overline{v})]$ and the left side tends to $\varrho_n\,(\eta_1 - \frac{1}{\varepsilon})\,[1 - G(\overline{v})]/[\overline{v}g(\overline{v})] > 0$: as the technology uses no inputs when demand is below capacity, what remains is the bundle use of the rationed firms, which produce at capacity.

3. *Existence.* Given monotonicity, a root in $(0,1)$ exists if and only if $1 - \omega_n$ lies strictly between the two limits of item 2. The lower bound on ω_n is (230); the upper bound is far above every node's ratio. A node whose ω_n violates (230) has no positive capital share and is floored. With slack capacity $\varrho_n = \eta_1 = 1$ and (230) reads $\omega_n > 1/\varepsilon$.

4. *Monotonicity, under the log-normal law.* Write $A(\theta) := \mathcal{G}_\theta^-(\overline{v}) = \int_0^\infty e^{-\theta u}f(u)\mathrm{d}u$, $B := -A'(\theta) = \int_0^\infty ue^{-\theta u}f(u)\mathrm{d}u$ and $c := 1 - G(\overline{v})$. Then $\mathrm{d}(\theta\eta_\theta)/\mathrm{d}\theta = [A^2 + c(A - \theta B)]/(A + c)^2$, so the sign is that of $N := A^2 + c(A - \theta B)$. Under the log-normal law $f(u) = \phi((u - \sigma_v d)/\sigma_v)/\sigma_v$ with $d := \ell/\sigma_v + \sigma_v/2$, so $f'(u) = -((u - \sigma_v d)/\sigma_v^2)f(u)$, and integrating by parts, $A - \theta B = \int_0^\infty \mathrm{d}(ue^{-\theta u})/\mathrm{d}u\ f\mathrm{d}u = -\int_0^\infty ue^{-\theta u}f'(u)\mathrm{d}u = (M_2 - \sigma_v d M_1)/\sigma_v^2$ with $M_k := \int_0^\infty u^k e^{-\theta u}f\mathrm{d}u$. Completing the square, $e^{-\theta u}f(u)$ is K times the density of $\mathcal{N}(-\sigma_v a, \sigma_v^2)$ with $a := \theta\sigma_v - d$ and $K = \phi(d)/\phi(a)$, so with $w := (u + \sigma_v a)/\sigma_v$ a standard normal truncated to $w > a$: $A = K\Phi_{\mathcal{N}}(-a)$, $M_1 = K\sigma_v[\phi(a) - a\Phi_{\mathcal{N}}(-a)]$, $M_2 = K\sigma_v^2[(1 + a^2)\Phi_{\mathcal{N}}(-a) - a\phi(a)]$. Dividing N by $K\phi(a)\Phi_{\mathcal{N}}(-d) > 0$ and writing $R(x) := \Phi_{\mathcal{N}}(-x)/\phi(x)$ for the Mills ratio,

$$\operatorname{sign} N = \operatorname{sign} S, \qquad S := \frac{R(a)^2}{R(d)} + \Big(1 + a^2 + ad\Big)R(a) - (a + d).$$

Three facts about the Mills ratio hold for every real x: (i) $1 - xR(x) > 0$; (ii) $R(x)^2 + xR(x) - 1 > 0$, the variance of a standard normal truncated above x; (iii) $(1 + x^2)R(x) - x > 0$, which follows from (ii). If $d \geq 0$, (i) at d gives $1/R(d) \geq d$, hence $S \geq [(1 + a^2)R(a) - a] + d[R(a)^2 + aR(a) - 1] > 0$ by (iii) and (ii); if $d < 0$, $S = R(a)^2/R(d) + [(1 + a^2)R(a) - a] + (-d)[1 - aR(a)] > 0$ by (iii) and (i). So $\theta\eta_\theta(\overline{v})$ is strictly increasing in θ for every $\overline{v} > 0$ and $\sigma_v > 0$, and the left side of (228) is strictly decreasing in α_n on $(0,1)$. The property is one of the log-normal law and not of the definitions: for a two-point law it fails.

[Status of the box above] Proof-checked by two fresh reviewers: one audited the box line by line, one derived the result forward from the definitions without seeing it. Both confirmed items 1–3 (the derivative identity against finite differences, the closed form of $\mathbb{E}_w[u]$ against quadrature, the limits against direct evaluation at $\alpha_n = 10^{-9}$ and 0.999; the auditor also derived the $\alpha_n \to 1$ limit from the static block as the rationed firms' bundle spending over firm revenue, which agrees to nine digits). Item 4, the proof of monotonicity, is the forward reviewer's: this session had established the sign numerically only, and the auditor's grid of 43 million points and the reviewer's grid to sixty digits found no violation; the proof was checked step by step by this session

(the differentiation of $\theta A/(A+c)$, the integration by parts, the tilted truncated-normal moments, the reduction to S, the three inequalities and the two cases). Two cautions from the auditor: the closed form $\sigma_v[z+\lambda(z)]$ cancels catastrophically in double precision for $|z|$ above a few hundred (the stable form is a continued fraction), which the numerical checks of this session avoided by working with $\ln(\theta\eta_\theta)$ directly; and the upper bound on ω_n is about 0.91 to 0.94 at the calibrated σ_v, above the largest ratio in the data, 0.604.

A node whose root is below 0.02, or whose ratio violates (230), is given $\alpha_n = 0.02$. Such a node spends less of its sales on intermediates in the model than in the accounts. In the baseline calibration twelve nodes are floored: ten capacity nodes (wood products; motor vehicles; furniture; textiles; apparel and leather; motion pictures and sound recording; computer systems design; other professional, scientific and technical services; management of companies; and performing arts and spectator sports) and two nodes without a rate (nursing and residential care; social assistance). In the secondary calibration eleven nodes are floored: the same, except that other professional services is not floored at the lower common target, and the two nodes without a rate are now capacity nodes. The floored capacity nodes miss their intermediate share by 3.7 percentage points at the median and by 8.7 (9.0) at most; floored nodes are flagged in every table of long-run results.

A node is floored because its accounts leave no room for a markup rent of about a tenth of sales beside the income of its capital: its operating surplus is small next to its purchases (textiles, apparel, furniture, motor vehicles, wood products), or much of it is the income of intellectual property or the pay of the self-employed, which are outside the model (computer systems design, motion pictures, performing arts, social assistance), or its utilization target is high, which lowers the pinned share (four of the ten capacity nodes of the baseline calibration have a slack-capacity share above the floor and are floored by their target alone). The root of this is that the model has no overhead cost: the markup rent is pure profit, and it must fit inside the operating surplus of physical capital. This is the main limitation of the calibration. With an overhead cost that absorbs the rent, so that no pure profit remains, the whole of $\mathrm{GOS}_n^{\mathrm{K}}$ would be the income of capital, no capacity node would be floored, and the capital share of manufacturing would be about 0.17 instead of 0.07. The two versions are two corners: all of the markup rent is pure profit, about a tenth of value added, or none of it is; Barkai (2020) finds a profit share of 7.9 % of value added for 2014.

Table 15 gives the capital shares by group of nodes under the two calibrations. The variants reported in the robustness part are capital income measured as operating surplus and as the capital compensation of the integrated production account, ω_n measured after redefinitions, and intellectual property counted as capital consistently, in income, in the stock and in depreciation.

The capital shares are small next to the capital share of a one-sector model, which is operating surplus over value added, 0.38 in the 63 industries, for three reasons that the accounts make explicit. The base is sales, not value added, and sales contain the intermediate inputs: operating surplus is 0.20 of gross output. Only part of operating

Table 15: Capital shares by group of nodes.

			baseline calibration			secondary calibration		
group of nodes	n	mean ω_n	target	α_n	floored	target	α_n	floored
manufacturing	19	0.164	0.786	0.065	5	0.786	0.072	5
mining and utilities	4	0.454	0.832	0.380	0	0.786	0.399	0
transfer nodes	26	0.260	0.862	0.149	5	0.786	0.173	4
nodes without a rate	14	0.221	off	0.140	2	0.786	0.137	2
all nodes	63	0.235		0.136	12		0.149	11

Notes: Unweighted means over the nodes of each group, floored nodes counted at 0.02. "Target" is the group's mean utilization target. The groups are those of Section E.5.3: the 23 core nodes (manufacturing, mining, utilities) with their own rates, the 26 transfer nodes with rates transferred from European surveys, and the 14 nodes without a rate. The slack-capacity formula (229) gives means of 0.066, 0.390, 0.174, 0.131 and 0.145 at the baseline calibration's ε.

surplus is the income of physical capital: of 100 dollars of value added, 55 are the compensation of employees, 7 taxes on production and 38 operating surplus, and of these 38, 8 are the income of intellectual property, 4 the imputed pay of the self-employed and 26 the operating surplus of physical capital. And the markup rent is taken out of the 26: at the baseline calibration's ε it takes 11 of the 100 dollars and leaves 15 for physical capital, a capital share of 0.163 as a ratio of sums over all nodes with slack capacity (0.076 in manufacturing); pinning lowers it further in the capacity nodes. With an overhead cost and no pure profit the same sums give 0.279 (0.192).

[The bridge from the textbook capital share to α_n, step by step] Ratios of sums over industries, billions of dollars a year, means of 2007–2019, on the 63-node data set; the markup rent at each calibration's own ε (solved model; the baseline calibration's ε is the same at its final ϕ as at the interim one to every digit printed). The identities are those of the accounts: gross output GO is intermediate inputs II plus value added VA, and value added is compensation COMP plus taxes on production TOPI plus operating surplus GOS. The two terms of ω_n differ from these: II^{dom} leaves out imported inputs and purchases from government and housing, and GOS^{K} is defined in (225).

	all 63 nodes	manufacturing (19)
gross output = inputs + compensation + taxes + operating surplus	$25{,}307 = 11{,}943 + 7{,}366$ $+969 + 5{,}029$	$5{,}580 = 3{,}607 + 967$ $+80 + 926$
(a) operating surplus over value added: the textbook share	0.376	0.469
(b) value added over gross output; operating surplus over gross output	0.528; 0.199	0.353; 0.166
(c) operating surplus = intellectual property + pay of the self-employed + GOS^{K}	$5{,}029 = 1{,}075 + 502 + 3{,}452$	$926 = 325 + 13 + 588$
of 100 dollars of value added: compensation, taxes, IP, self-employed, GOS^{K}	55, 7, 8, 4, 26	49, 4, 16, 1, 30
(d) the model's sales $\text{GOS}^{\text{K}} + \text{II}^{\text{dom}}$; ω	13,823; 0.250	3,421; 0.172
markup rent = sales/ε (baseline calibration); capital income	1,439; 2,013	356; 232
capital share α with slack capacity, (229)	**0.163**	**0.076**
the same at the secondary calibration's ε	0.169	0.083
(f) with an overhead cost and no pure profit: $\alpha = \omega/(1 - \frac{1}{\varepsilon})$	0.279	0.192

The pay of the self-employed in row (c) is its part attributed to physical capital, $(1 - s_n^{\mathrm{IP}})\mathrm{SE}_n$, so that the three parts add up to operating surplus; the imputed pay itself is 634 billion a year over the nodes, 12.6 % of their operating surplus and 1.8 % in manufacturing. The ratios of sums of row (d) differ from the means over nodes of Table 15, which weight every node equally and include the pinned and floored nodes. If labour were an input of the bundle, sales would contain compensation and the rent would be a tenth of larger sales; the capital share would fall further (about 0.06 over all nodes at the accounts' totals), which is why that extension would not remove the floors.

[Weak spot: the overhead cost] The black text names it as the main limitation. The evidence behind the two numbers: with an overhead cost of the usual form (every firm buys a fixed amount of the bundle, of a size that leaves no pure profit on the balanced growth path) the matching equation becomes $\varrho_n(1 - \frac{1}{\varepsilon} - \bar{\xi}_n) = \omega_n$; solving it at the round working values of the earlier node set gave about 0.17 at the manufacturing mean instead of 0.07, no floored capacity node in either calibration, an interest rate of about 10 % (all of the operating surplus is then the return on capital) and much smaller effects of capacity, because the comparison economy is less elastic when capital shares are larger. Computed once, at the round working values and the operating-surplus definition, not at the decided rule nor on the 63-node set; the sentence of the black text is written with "about" for that reason. The decision was not to add an overhead cost now. The row (f) of the bridge table above is the same object as a ratio of sums on the current data (0.279 over all nodes, 0.192 in manufacturing, at the solved model's ε).

[Weak spot: the imputed pay of the self-employed] It is an imputation: the hours of the self-employed valued at the pay of comparable employees. Largest shares of operating surplus on the 63-node set: performing arts 0.70, social assistance 0.70, apparel 0.64, farms 0.59, other services 0.48, construction 0.44, ambulatory health care 0.38, computer systems design 0.33, other professional services 0.32, forestry 0.32; thirteen nodes are above 0.30; in manufacturing only apparel (0.64) and furniture (0.26) exceed a quarter. The authors of the integrated account note that the imputation may not hold in social assistance and performing arts, where it takes two thirds of the operating surplus; both nodes are floored under the decided rule and neither would be with operating surplus as capital income; together they are about 1 % of the nodes' operating surplus. One rule is kept for all nodes; operating surplus as capital income is the robustness row. (Shares recomputed from the data set in this session; the authors' caveat is from the account's documentation.)

[Worked example and the pinning table] Chemicals in the baseline calibration, on the solved model at its estimates and its final ϕ (computed once; every number unchanged from the interim ϕ at this precision): $\omega = 0.195$ and the target 0.770 give the threshold $\overline{v} = 1.29$ at the node's mean, so a firm is rationed when its demand is 29 % above that of the average firm, which happens to 7.0 % of the firms; $\varrho = 0.9992$, so the rationing rent is 0.08 % of sales; the pinned share is 0.092 against 0.101 from (229). At the pinned share the node's mean sales divide into materials 0.805 (which is $1 - \omega$ exactly), markup rent 0.104, rationing rent 0.001 and capital income 0.090,

against $(1-\frac{1}{\varepsilon})\alpha = 0.083$ with slack capacity (the first stage, at a point law, had 8.0 % rationed, the pinned share 0.095 and the split 0.805, 0.100, 0.001, 0.094): when capacity binds, capital is the scarce factor and earns more than its slack-capacity share. The next table solves (228) at the manufacturing mean of the ratio, $\omega = 0.164$, for a range of targets (baseline calibration estimates):

utilization target	slack	0.70	0.75	0.786	0.83	0.864	0.90
firms rationed, %	0	2.6	6.0	9.9	17.1	25.1	36.4
rationing rent, % of sales	0	0.03	0.07	0.12	0.22	0.34	0.54
capital share that solves (228)	0.070	0.067	0.062	0.055	0.042	0.023	none

At a target of 0.90 no positive capital share reproduces this ω: condition (230) fails. The table is first stage, at the first-stage estimates (the independent implementation, cross-examined against the first-stage calculator); the chemicals example above is the solved model's.

[The 19 manufacturing nodes] The solved model at each calibration's own estimates, the baseline calibration at its final ϕ (its shares equal those at the interim ϕ to three decimals) and the secondary calibration at the interim ϕ (provisional; computed once); "floor" means $\alpha_n = 0.02$. The slack-capacity share is (229).

		baseline calibration			secondary calibration	
node	ω_n	target	slack share	α_n	slack share	α_n
wood products	0.114	0.765	0.011	floor	0.020	floor
nonmetallic minerals	0.209	0.733	0.118	0.113	0.125	0.113
primary metals	0.154	0.782	0.056	0.043	0.064	0.048
fabricated metals	0.167	0.786	0.070	0.057	0.078	0.063
machinery	0.166	0.777	0.069	0.057	0.076	0.061
computers and electronics	0.216	0.774	0.124	0.115	0.131	0.120
electrical equipment	0.212	0.821	0.120	0.101	0.127	0.115
motor vehicles	0.093	0.755	−0.012	floor	−0.004	floor
other transportation equipment	0.210	0.746	0.117	0.112	0.125	0.112
furniture	0.097	0.774	−0.008	floor	0.000	floor
miscellaneous manufacturing	0.203	0.766	0.110	0.102	0.117	0.105
food, beverages and tobacco	0.144	0.804	0.045	0.025	0.053	0.036
textiles	0.089	0.787	−0.017	floor	−0.008	floor
apparel and leather	0.044	0.760	−0.067	floor	−0.058	floor
paper	0.179	0.866	0.083	0.039	0.090	0.077
printing	0.206	0.796	0.114	0.101	0.121	0.109
petroleum and coal products	0.259	0.855	0.173	0.144	0.180	0.171
chemicals	0.195	0.770	0.101	0.092	0.109	0.096
plastics and rubber	0.150	0.820	0.051	0.027	0.058	0.043
mean	0.164	0.786	0.066	0.065	0.074	0.072

In the secondary calibration every node has the target 0.786. The merged finance node has $\omega = 0.271$, the target 0.846 (0.786) and the capital share 0.163 (0.185). The table is printed in black inside this private box only because every entry of it is provisional; it is replaced again when ϕ is final.

[The unit-materials-cost moment, not targeted] In the model without capacity, at given capital and input prices, the cost of intermediates per unit of output rises with output with the elasticity $\alpha_n/(1-\alpha_n)$: from (11), $M \propto y^{1/(1-\alpha_n)}$ at fixed k, so $\mathcal{P}M/y \propto y^{\alpha_n/(1-\alpha_n)}$. At the manufacturing mean of the baseline calibration this is 0.069. The supply-curve sample allows a regression of unit materials cost on demand-driven output at given materials prices, which speaks to the size of the capital share and was run once in the review of the calibration; the same regression finds that pay per unit of output does not fall when demand raises output, as the unit cost of a fixed factor would, which speaks against counting labour as part of capital. **[Flag: the estimate of that regression is not among the numbers that may be printed; it is regenerated before any number is written here.]** The moment is reported for reference only and is not targeted.

[How far the numbers of this sub-part are checked]

- ω_n, its range and its means by group: final data-side numbers, recomputed from the data set in this session (reproduced).
- The capital shares, the floored nodes, the misses and the gaps between slack and pinned shares, at each calibration's first-stage estimates: computed by the first-stage calculator and reproduced in this session by an independent implementation written from the paper's formulas (57), (75), (118) and (22); the two agree to four decimals for every manufacturing node, every group mean, the finance node, the sets of floored nodes and the misses (cross-examined). The split of the floored nodes into those floored by their accounts alone and those floored by their target: this session's implementation, computed once.
- The first-stage estimates of ε and σ_v themselves: computed once (Section E.5.1).
- The bridge and the self-employed shares: recomputed from the data set in this session; the totals equal those of the accounts table of Section E.2.
- The overhead-cost numbers: computed once at earlier values, see the box above.
- Since 3 November 2026 the baseline calibration is at $\phi = 2$ and $\gamma = 1$ (its outer loop run on 31 October 2026: $\varepsilon = 9.614$, $\sigma_v = 0.1851$); against the solved model at $\phi = 13$ the capital shares moved by 0.003 at most (ten cells of the node table by one unit in the third decimal, one by three), the floored nodes, the group means and the bridge not at all to the printed digits, r from 0.033 to 0.039. Before that, the baseline calibration's numbers of this sub-part were the solved model at its final $\phi = 13.0$ (11 October 2026; computed once; the second stage meets every target to 10^{-6}; the group means, the floored nodes, the misses and the gaps read off the stored calibration by a script that reproduces the interim numbers exactly). From the interim ϕ the capital shares moved by 0.0003 at most (two cells of the node table by one unit in the third decimal), ε and the bridge not at all to the printed digits, r from 0.034 to 0.033. The secondary calibration's numbers stay those of the solved model at the interim ϕ (computed once) and stay in colour: its simulation of ϕ is deferred. Still in colour in the baseline calibration:

the overhead-cost numbers (computed once at earlier values, see the box above) and the first-stage variants of the capacity block (Section E.5.1).

E.5.3 The utilization targets and the two calibrations

A node's capital intensity x_n of (10) is chosen so that the node's mean utilization on the balanced growth path, $\mathbb{E}_n[U_n(\overline{v})]$ with U_n of (22), equals its target. The target fixes the node's threshold, and with it how much of the node is rationed: at the manufacturing targets the constrained share is 10 % of the firms, and the markup of the posted price is 1.14 against 1.12 with slack capacity. A higher target also means a larger rationing rent, a steeper supply curve at the node's normal operating point and, through pinning, a lower capital share. The level of the targets is therefore the most direct lever on how much capacity matters, and the two calibrations differ in it.

The baseline calibration uses every utilization rate that exists, node by node. For the 23 core nodes the target is the mean of the Federal Reserve's annual rate over 1972–2019: its concept, output over a capacity index, is the model's utilization, and the supply-curve estimates are regressions on this rate. The long-run mean is used, and not the mean over the window of the accounts, because utilization is cyclical and 2007–2019 contains the deepest post-war slump (the manufacturing nodes average 0.742 over 2007–2019 and 0.786 over 1972–2019), and because the supply-curve estimates describe industries around their long-run means. The core targets run from 0.720 (mining support) to 0.930 (oil and gas extraction), the 19 manufacturing nodes from 0.733 to 0.866. For the 26 transfer nodes the European ratio of the services rate of the node's industry to the manufacturing rate is applied to the US manufacturing rate,

$$\overline{U}_n = 0.784 \times \frac{\text{European services rate of the node's industry, 2011–2019}}{\text{European manufacturing rate, 2011–2019}}, \qquad (231)$$

where 0.784 is the Federal Reserve's rate for US manufacturing over 1972–2019. European services rates average 1.10 times the manufacturing rate, so the 22 transfer nodes with a European aggregate have targets between 0.823 and 0.889, mean 0.864, about eight points above manufacturing; the four nodes whose European rate is published by country only (the finance node, insurance, performing arts, recreation) are treated likewise. Capacity is off in the 14 nodes without a rate: x_n is set so that capacity never binds, and the node behaves as in the model without capacity. The transfer rests on an assumption that the data cannot test. The European services question asks firms by how much they could raise their activity with their present resources, which for a services firm include staff, while the model's capacity is that of physical capital: at the transfer targets the model has about a quarter of a service node's firms at capacity, whereas in the same surveys about 2 % of services firms say that space or equipment limits them, against about 8 % of manufacturers. The level is kept nonetheless: the model has no labour, and the calibration abstracts from that here as it does elsewhere; the secondary calibration is the alternative. Within manufacturing, reading an industry's mean level as tightness is also not what the supply curves suggest

(Section E.5.1); the aggregate results hardly depend on this reading, but differences between manufacturing nodes do, so they are shown under both calibrations.

The secondary calibration replaces the three choices above by one convention: capacity is on in every node, and every node, mining and utilities included, has the target 0.786, the mean target of the 19 manufacturing nodes. It is a convention and not a measurement: it puts the same share of firms at capacity in every node, also in construction, where the European building survey finds that equipment limits few firms. Each calibration reads a utilization level in one way, the baseline calibration as tightness and the secondary calibration as a convention. Two variants of the baseline calibration are reported in the robustness part: capacity on in the 23 core nodes only, and capacity on in all 63 nodes with the 14 others at the mean transfer target.

[Evidence behind the two weak spots, and the thin bridges]

- **The level of the transfer nodes.** The limiting-factor questions of the same European surveys are the test of the transfer's scale. Across services industries the services rate rises with the share of firms limited by labour, falls with the share limited by demand and is unrelated to the share limited by equipment; the Commission itself lists human against technical resources as an open issue of its question. Relative to manufacturing, the model at the transfer targets has several times too many firms limited by physical capital. If only space and equipment count as capacity, the survey shares would place services at or below the manufacturing level; a reviewer who ranked the options without seeing results ranked the level in force last. Kian keeps it and defends it as in the black text. The evidence was computed and cross-examined in the review of the calibration; its numbers beyond the two shares in the black text are not reprinted here.
- **Where capacity is on.** The rule is simple and rests on data: capacity is on wherever a rate exists; legal services is on and construction is off because the European services survey has a utilization question and the building survey has none. Keeping capacity off in the 14 nodes is the conservative choice: switching it on makes capacity matter more.
- **Thin bridges.** Air transport rests on one European division with few quarters; other services on the repair division alone; management of companies on a division that includes consultancy; pipelines on land transport; four transport nodes share one division. The four by-country targets rest on countries that carry about three tenths of the European weight (Germany and France do not report finance). The finance node's target is the unweighted mean of its two divisions, banking and auxiliary finance, as every node with several divisions is treated.
- **What the secondary calibration is worth.** Almost all of the difference between the two calibrations is the level of the service nodes; with only the transfer nodes of the baseline moved to 0.786, the baseline's results are close to the secondary calibration's (first stage, computed once).

E.5.4 The remaining parameters and the algorithm

The installation curvature, ϕ. Investment ι per unit of capital adds $\Phi(\iota) = (1 + \phi\iota)^{1/\phi} - 1$ units of capital, and the firm invests until $\Phi'(\iota)\mathfrak{q} = 1$, so that $\mathrm{d}\ln\iota/\mathrm{d}\ln\mathfrak{q} = (1+\phi\iota)/((\phi-1)\iota)$: at an investment rate of 0.05 the elasticity is 22 at $\phi = 2$ and 3 at $\phi = 11.5$. ϕ is set at 2 by convention, the order of the convex adjustment costs that structural estimates find (Hall, 2004; Cooper & Haltiwanger, 2006): at that value installation costs absorb about 5 per cent of investment spending and Tobin's q lies between 1.05 and 1.17 across nodes. The same industry data that discipline the capacity block contain a moment that would identify ϕ inside the model: a utilization rate one point above the industry's mean is followed by capacity growth 0.167 percentage points higher in the next year (Section E.2), and in the model higher utilization raises the return on installed capital, its value and investment, the more so the smaller ϕ is, so the same regression run on data simulated from the model gives a coefficient that falls in ϕ. Targeting it gives $\phi = 13.0$ (three evaluations, at 11.5, 12.35 and 13.02, gave 0.181, 0.173 and 0.167; capacity in the simulated regression is the node's physical capacity, its capital, as the Federal Reserve's index is anchored on plant surveys of physical capability), a value at which 45 per cent of investment spending is installation cost and the marginal unit of capital costs 2.8 units of goods at the median node, which is why the convention is preferred. At $\phi = 2$ the model's coefficient is not computed; the slope near 13 and two-node runs suggest about 0.3, so the model's capital would respond to utilization about twice as fast as the Federal Reserve's capacity index, part of whose coefficient may be mechanical because the index is built in part from capital stocks. In the secondary calibration the interim $\phi = 11.5$ stands until it is re-run at 2. No response at the date of a shock depends on ϕ; it moves the interest rate, the absorbers and everything that involves time.

[To do: the ϕ moment at the baseline calibration] One evaluation of the simulated regression at $\phi = 2$, $\gamma = 1$ (about nineteen hours of transitions) would give the model's coefficient to report beside the data's 0.167 as an untargeted moment. Not run as of 3 November 2026; Kian decides when.

[Flag closed: the simulated regression for ϕ, its three evaluations, and what is deferred] The simulation feeds the three measured shock series of the validation (the common productivity component, the node productivity residuals and the demand shifts, the latter as permanent changes of the weights, since the measured levels are a random walk), reads capacity growth and utilization off the model's impulse responses as yearly averages, and computes the regression's coefficient over a 64-year window exactly from those responses and the shocks' covariance; the loop on ϕ runs on the solved economy by a secant on the coefficient, to a tolerance of 0.002. The baseline calibration (computed once, 8–11 October 2026; about nineteen hours of transitions per evaluation): at $\phi = 11.5$ the coefficient is 0.1805 (the demand shifts contribute 0.221 and carry 77 % of the regressor's variance, the node productivity shocks 0.033 and 16 %, the common shock 0.063 and 6 %; a 64-year panel of the model would scatter with a standard deviation of 0.014 around it), at 12.35 it is 0.1730, and at 13.02 it

is 0.1675, within the tolerance of the data's 0.167; the shares by type are the same at every ϕ. With capacity including productivity (capital times the node's productivity level) the coefficients are 0.192, 0.185 and 0.179: that concept lets a productivity shock raise measured capacity on impact, so its node-productivity channel is four times larger (0.12 against 0.03) and in a two-node economy of an earlier run it was negative; the physical concept is the decision of 9 October 2026. A check on the measured 1998–2019 history itself (the static block on each year's realised shock vector, at $\phi = 11.5$): 0.190 with physical capacity and 0.354 with the inclusive concept, against 0.167 (standard error 0.025) in the data. Two closures are left open and recorded here: the loop ran at the interim ε and σ_v, and their re-estimation at the final ϕ moved them by 0.0004 and 0.00006, so the coefficient at the final pair is the one reported to the digits given, but it has not been re-evaluated there (one more nineteen-hour run, postponed on 11 October 2026); and the secondary calibration's loop is deferred, so its ϕ is the interim value. The interim value was near the data in the one comparison that existed before: in runs of a two-node economy, a node's capital grew in the first year after a productivity shock by about the data's coefficient per point of utilization at $\phi = 11.5$ and by twice as much at $\phi = 2$ (computed at earlier values, once). An approximate mapping of the coefficient through the investment condition put ϕ in the upper single digits to about eleven; the mapping was cross-examined and corrected once. Which value of that range is used matters little for the first-stage numbers.

[$\phi = 2$: not run (as of 27 October 2026)] [Superseded, 3 November 2026: $\phi = 2$ was run on 31 October 2026 as a trial, with ε and σ_v re-estimated by the outer loop on the wide capital grid (201 points, log-width 1.0), and together with $\gamma = 1$ it is the baseline calibration since 3 November 2026; the results chain ran on it on 2–3 November 2026.]

- Nothing at $\phi = 2$ exists on 63 nodes: no stored calibration (the store holds the baseline calibration at the interim and the final ϕ, the secondary calibration at the interim ϕ, and old two- and three-node calibrations) and no results. The only evidence at $\phi = 2$ is from two-node runs of the numerical programme: a node's capital grew in the first year after a productivity shock by about twice the data's coefficient per point of utilization (computed once, at earlier values). Section E.6 lists $\phi = 2$ among the robustness variants; of those only the ε–σ_v sweep (the comparative-statics box at the end of Section 5) exists on the final calibration.
- What it takes. ϕ enters the balanced growth path through the installation condition, so the second stage is rerun at $\phi = 2$ (minutes; ε and σ_v can be held, since the move from $\phi = 11.5$ to 13.0 changed them by 0.0004 and 0.00006), then the results scripts: the static block (7 minutes), the transitions (hours), the sweeps if wanted. At $\phi = 2$ the interest rate, the absorbers and everything that involves time move; no response at the date of a shock does.
- It is a sensitivity of the baseline calibration, not a new calibration, run on a folder of its own; the rule that trial calibrations hold ϕ at 13.0 is about new calibrations and does not bar it.

[Q28 | mynotes | ANSWERED] I will state here that for possible new calibrations I am only interested in those that do *not* involve the time-consuming 60-hour calibration of phi and I prefer to focus on improving/changing other parts of the calibration (e.g. gamma) and instead fixing phi at some "reasonable" value such as 2 because 2 is very common in the literature and nobody will object to it. But before I ask you to actually perform the calibration with phi equal to 2, can you answer the question below to the extend that you can predict things (if you can't just say so):

How would $\phi = 2$ change (1) the calibration of other parameters and (2) the fit? Would these changes be desirable, underisable or neutral?

[$\phi = 2$: what it would change in the calibration and in the fit, predicted from the first stage] [Flag, 3 November 2026: run. The solved model at $\phi = 2$ (31 October 2026) confirmed every prediction below to the second digit: $r = 0.0392$, $\rho = -0.0027$ at $\gamma = 2$ and 0.0183 at $\gamma = 1$, Tobin's q 1.05–1.17, $\overline{q} = 0.193$, $\overline{Z} = 0.322$; ε and σ_v moved by 0.009 and 0.0006 (to 9.614 and 0.1851); the growth gap came out 0.30 points at $\gamma = 2$ and 0.60 at $\gamma = 1$, as reasoned; the ϕ moment at $\phi = 2$ is not computed.]

- **Basis.** The first stage of the baseline calibration (the closed form of the protocol: productivity process off, one state per node) at the stored ε and σ_v, run at $\phi = 13.02$ and at $\phi = 2$ on 27 October 2026 (a scratch computation, seconds; at $\phi = 13.02$ it reproduces the stored interest rate, 0.03334 against 0.03335). It predicts everything that enters through the rest point of capital; what it cannot predict, the second stage's cross-section within nodes and everything dynamic, is listed at the end. The solved model at $\phi = 2$ has not been computed. Evidence status of every number below: first stage, computed once.
- **Where ϕ enters.** Only through the rest point of a node's capital: $\Phi(\iota_n) = g + \delta_n$ gives $\iota_n = [(1 + g + \delta_n)^{\phi} - 1]/\phi$ and $\mathfrak{q}_n = (1 + g + \delta_n)^{\phi - 1}$, hence the user cost $u_n = (r - g)\mathfrak{q}_n + \iota_n$ and the accounts' stock per unit of installed capital, $\iota_n/(g + \delta_n)$. Everything that precedes r in the chain does not see ϕ: at $\phi = 2$ the thresholds, the capital shares, the floors (12 nodes), ϱ_n, ξ_n, the sales and the intermediate shares are reproduced to the last digit. Everything after r is an absorber (x_n, β_n, $\overline{Z}$) and moves so that the targets are hit again.
- **(1) What moves in the calibration.**

	$\phi = 13.02$	$\phi = 2$
interest rate r	0.0333	0.0392
ρ at $\gamma = 2$	-0.0086	-0.0028
ρ at $\gamma = 1$	0.0124	0.0182
Tobin's q_n, median (range over nodes)	2.79 (1.79–6.69)	1.09 (1.05–1.17)
investment per unit of capital ι_n, median	0.157	0.093
installation cost as a share of investment spending, $1 - \Phi(\iota)/\iota$	0.45	0.05
user cost u_n, median (largest over smallest node)	0.19 (6.8)	0.11 (3.0)
elasticity of ι to q at the rest point, median	1.6	12.7
physical capital per unit of sales	0.79	1.30
GDP per unit of physical capital, $\overline{q}$	0.315	0.193
$\overline{Z}$	0.348	0.322
consumption over GDP	0.530	0.527
the fit's return: r plus the markup rent per unit of the accounts' stock	0.110	0.116
capital stock over sales by group (manufacturing; mining and utilities; transfer; none)	0.67; 4.20; 1.38; 1.31	0.67; 4.16; 1.38; 1.33

The mechanism: at $\phi = 2$ installation is nearly linear, so a dollar of investment buys 0.95 units of capital instead of 0.55, and the accounts' stock (cumulated spending) per unit of physical capital falls from 1.71 to 1.04. To reproduce the stock over sales of 1.356 the economy must then hold 1.6 times more physical capital per unit of sales, which it does because the user cost of a unit falls from 0.19 to 0.11; at the old r the moment would overshoot (1.444), so r rises by 0.6 points. The intensities x_n absorb the rest (median $\times 1.7$, between 1.2 and 2.8 across nodes, most where δ_n is high); relative prices move by 1 per cent at the median node and by up to a third at the extreme (the high-depreciation nodes, whose user cost falls most); the weights β_n move by a few per cent, one node's by 22. ε and σ_v are not predicted to move materially: the Census share and the supply-curve counterpart are static objects of the threshold, σ_v and ε, and the outer loop sees ϕ only through the second stage's cross-section (the move from 11.5 to 13.0 moved them by 0.0004 and 0.00006; $\phi = 2$ moves the cross-section far more, so expect more, still small; reasoned).

- **Desirable.** r at 3.9 per cent, and ρ at -0.3 per cent ($\gamma = 2$) or $+1.8$ ($\gamma = 1$): with both changes the pair is conventional. Tobin's q at 1.09 instead of 2.8 at the median node and 6.7 in trucking: at $\phi = 13$ the marginal unit of installed capital costs 2.8 units of goods, which is hard to defend, and 45 per cent of all investment spending is installation cost, against 5 per cent at $\phi = 2$, the order of the structural estimates cited above for $\phi = 2$ (Hall, 2004; Cooper & Haltiwanger, 2006). The user costs across nodes become those of a standard model (largest over smallest 3.0, against about 2.8 for $(r + \delta_{\max})/(r + \delta_{\min})$ without installation

costs) instead of 6.8: at $\phi = 13$ the high-depreciation nodes pay user costs that the installation function imposes, not the data, and the intensities x_n carry that distortion; at $\phi = 2$ they do not. The fit's return moves from 11.0 to 11.6 per cent, against 12.0 in the accounts and 10.2 in Gomme, Ravikumar and Rupert (2011): neutral to slightly better.

- **Neutral.** Every targeted moment is hit as now (the floored nodes miss their intermediate share as now); the capital stock by group, the dispersion of firm utilization, the constrained share, the slope of the Census share and the input-price check are static objects and do not move; the rents, the shares and consumption over GDP do not move. The impact responses of the capacity economy are static and do not move at the first stage; through the second stage's cross-section they move a little. The level part of the long-run gap (the wedges at the capacity economy's allocation, about 8 of the 10 points) is static at the pinned thresholds; the reallocation part (2 points) and the comparison with the economy without capacity constraints can move, because the economy without capacity constraints's allocation of capital follows the user costs, whose pattern across nodes changes. Expect the headline to move by a point or two, not more (reasoned).
- **To flag.** GDP per unit of physical capital falls from 0.315 to 0.193: a change of units (the same GDP on 1.6 times more capital), but every table that prints $\bar{q}$ changes. The growth part of the long-run cost rises: the economy without capacity constraints grows faster because its return on capital is higher, and at $\phi = 2$ growth responds about 2.7 times more to a given change in the return ($\mathrm{d}u/\mathrm{d}g$ at the rest point is 6.0 at $\phi = 13$ and 2.2 at $\phi = 2$, from $u = [\rho + (\gamma - 1)g]\mathfrak{q} + \iota$ with $\mathfrak{q}$ and ι functions of $g + \delta_n$); since the user cost itself is 0.6 times as large, the same proportional gap in the return gives a growth gap about 1.6 times as large: roughly 0.3 points instead of 0.19 at $\gamma = 2$, and roughly 0.6 instead of 0.35 at $\gamma = 1$ (first order, reasoned; the solved economy without capacity constraints decides). The welfare number with the growth gap (26.5 per cent today as the table states it, 36 from the capacity side) grows with it and stays fragile; the level numbers stand.
- **(2) The fit: one moment missed.** The targeted moments stand and the untargeted static moments of the fit stand (above). What $\phi = 2$ gives up is the response of capacity to utilization, 0.167 (s.e. 0.025), the one dynamic moment of the calibration and the only one that disciplines how fast capital follows. Its value at $\phi = 2$ on 63 nodes is not computed. Two pieces of evidence: in the two-node runs a node's capital grew in the first year after a productivity shock by about twice the data's coefficient per point of utilization at $\phi = 2$ and by about the data's at 11.5 (the box on the simulated regression above; computed once, at earlier values); near 13 the 63-node coefficient falls by 0.009 per unit of ϕ, which extrapolated linearly (which it is not) would give 0.27 at $\phi = 2$. So expect about 0.3, four to seven standard errors above the data: the model's capital would respond to utilization about twice as fast as the Federal Reserve's capacity index does. This is undesirable, and it is the price of the convention. Three things to say

beside it: the text above notes that part of the data's coefficient may be mechanical (the Federal Reserve builds its capacity index partly from capital stocks); the data's slow response may reflect time to build and lumpy investment, which a convex installation function does not have; and reporting the model's coefficient at $\phi = 2$ as an untargeted moment costs one evaluation (about nineteen hours of transitions), not the loop. The elasticity of investment to q of 13 at $\phi = 2$ (the text above: 22 at an investment rate of 0.05) is high against regression estimates of investment on q and in line with the structural estimates of small convex costs; a referee from the first camp will ask, and the citation is the answer.

- **What changes in the time-series results** (reasoned, not computed). With investment eight times as elastic, capital closes a given gap several times faster: in the transitions "utilization absorbs the shock first and capital follows" keeps its sign, and the horizon over which utilization carries the adjustment shrinks from years to a year or two. The model's yearly utilization becomes less persistent (capital catches up), which bears on the volatility and persistence comparison of the validation; where the model stands at $\phi = 13$ can be read off the stored kernels, at $\phi = 2$ only from new ones (nineteen hours).
- **What the first stage cannot predict.** The second stage's cross-section: at $\phi = 2$ capital tracks productivity within a node, so the distribution of capital within a node widens and that of utilization narrows; the capital grid of log-width 0.5 may need widening (check the edge mass; the sweep's rescue grid of width 1.0 exists), and the second stage's corrections and the small move of ε and σ_v follow from it. The solved responses relative to the economy without capacity constraints. The speed of the transitions and the ϕ-moment coefficient.
- **In one line.** On every count the first stage can see, $\phi = 2$ is the cleaner calibration (r, ρ, q, installation costs, relative user costs) at no cost to the targets or to the static fit; its price is one dynamic moment, which the model would then miss by about a factor of two and should report as an untargeted moment. If it is run: the second stage at $\phi = 2$ with ε and σ_v held (minutes to an hour), the static block (seven minutes), the transitions (hours), the coefficient if reported (nineteen hours); a sensitivity of the baseline calibration, as the box above says, or a new baseline if later calibrations are to build on it, which is your call.

The interest rate and the discount rate. The model has no aggregate risk, so r is at once the riskless rate and the return that firms require on capital, with $r = \rho + \gamma g$ on the balanced growth path. A node's capital income is fixed by its sales and its capital share, so the higher the interest rate, the less capital the node holds per unit of sales. The target is the capital stock of the 63 nodes over their sales, the current-cost net stock of equipment and structures over $\text{GOS}^{\text{K}} + \text{II}^{\text{dom}}$, both means of 2007–2019: 1.356; all nodes enter, because one interest rate prices the capital of every node. The accounts build a stock by cumulating investment spending less depreciation, which on the balanced growth path equals investment spending over $\delta_n + g$, so the model's

counterpart is

$$\frac{\sum_n \Delta_n \mathbb{E}_n[\iota^* \kappa] / (\delta_n + g)}{\sum_n \Delta_n \mathbb{E}_n[\widehat{s}]}, \tag{232}$$

which exceeds installed capital, $\sum_n \Delta_n \mathbb{E}_n[\kappa] = 1$, because installation is concave. At a point node law the moment falls in r over the admissible range $r > g$, so the target gives one interest rate: $r = 0.033$ in the baseline calibration and 0.036 in the secondary calibration, a pre-tax return, since operating surplus is measured before the corporate income tax. With $\gamma = 2$ the discount rate is $\rho = r - \gamma g = -0.009$ (-0.006). A negative discount rate is admissible: what the model needs is $r > g$, and the margin is $r - g = 0.012$ (0.015). The interest rate is the residual of everything else: it moves with ε, with ϕ and with the definition of capital income, and it moves no response at the date of a shock. The capital stock is matched in total and not group by group, and the return it implies is compared with outside returns (Section E.6).

The final-demand weights and the productivity level. Node n's share in final expenditure is $\beta_n \mathbb{E}_n[P^{1-\eta}] \Delta_n$, which depends on the model's own prices, so the weights are the solution of this equation at the equilibrium prices, with $\sum_n \beta_n \Delta_n = 1$. Relative productivity levels across nodes are a choice of the units in which goods are measured: scaling the unit of good n by λ_n maps the economy with parameters $(\overline{Z}_n, x_n, \beta_n)$ into the same economy with

$$\ln \tilde{\overline{Z}} = \ln \overline{Z} + (I - \operatorname{diag}(1-\alpha)\, a) \ln \lambda, \qquad \ln \tilde{x} = \ln x - a \ln \lambda, \qquad \tilde{\beta}_n = \beta_n \lambda_n^{1-\eta}, \tag{233}$$

and every other parameter unchanged. Every node therefore has the common level $\overline{Z}$, and x_n and β_n absorb what relative levels would have done; $\overline{Z}$ is chosen so that the growth rate of the balanced growth path is $g = 0.021$.

[Derivation of (233)**]** Scale the unit of good n by λ_n: quantities of good n are multiplied by λ_n, its price P_n divided by λ_n, nominal values unchanged.

- *A buyer m of good n* ($n = m$ included if $a_{nn} > 0$). Its bundle $M_m = \prod_j m_{mj}^{a_{mj}}$ of (9) in the new units is $\tilde{M}_m = \lambda_n^{a_{mn}} M_m$, and its bundle price falls accordingly, $\ln \tilde{\mathcal{P}}_m = \ln \mathcal{P}_m - a_{mn} \ln \lambda_n$ (the bundle price of (95) is the cost-share-weighted geometric mean of the input prices). Its production function $y = Z_m k^{\alpha_m} M_m^{1-\alpha_m}$ reads $y = \tilde{Z}_m k^{\alpha_m} \tilde{M}_m^{1-\alpha_m}$ with $\tilde{Z}_m = Z_m \lambda_n^{-(1-\alpha_m) a_{mn}}$. Its capacity, at most k / x_m units of the bundle, is at most $\lambda_n^{a_{mn}} k / x_m$ new units, so $\tilde{x}_m = x_m \lambda_n^{-a_{mn}}$; capacity in units of m's own output, $Z_m k x_m^{\alpha_m - 1}$ by (10), is unchanged.
- *The producer n.* Its output in new units is λ_n times the old, so $\tilde{Z}_n$ carries the additional factor λ_n.
- *The final good.* The CES aggregator (24) is unchanged with $\tilde{Q}_n = \lambda_n Q_n$ if $\tilde{\beta}_n = \beta_n \lambda_n^{1-\eta}$, so the final good, and with it capital and investment, keep their units.
- *The static block.* In the pricing line (139) of a buyer $m \neq n$ the right-hand side $\mathcal{P}_m / (Z_m x_m^{\alpha_m})$ is unchanged, since its three factors change by $\lambda_n^{-a_{mn}}$, $\lambda_n^{-(1-\alpha_m) a_{mn}}$ and $\lambda_n^{-\alpha_m a_{mn}}$; for n itself it is divided by λ_n, as its price is. In the supply line

(140) the ratio $\mathcal{P}_m/x_m$ is unchanged, and in the demand line $\tilde{\beta}_n\tilde{P}_n^{1-\eta} = \beta_n P_n^{1-\eta}$. So every threshold, every nominal quantity and the flow (115) that prices capital are unchanged, and so is the whole equilibrium.

- Collecting the effects over all goods gives (233), for $\eta \neq 1$ (at $\eta = 1$ the weights are unchanged and the final good keeps its unit only under $\sum_n \beta_n \Delta_n \ln \lambda_n = 0$). The matrix $\mathrm{diag}(1-\alpha)\,a$ has row sums $1-\alpha_m < 1$ because every $\alpha_m > 0$, which the floor guarantees, so its spectral radius is below one and $I - \mathrm{diag}(1-\alpha)\,a$ is invertible: any vector of relative levels can be produced or removed by a choice of units. The N changes of units are independent, since the final good's unit is the numeraire and is unchanged. One of them is used for the normalisation $\sum_n \beta_n \Delta_n = 1$ and $N-1$ to equalise the levels, which leaves the common $\overline{Z}$; so of the $3N$ parameters $(\overline{Z}_n, x_n, \beta_n)$, $2N$ remain, the N intensities x_n, the $N-1$ free weights and $\overline{Z}$, and they are pinned by the N utilization targets, the $N-1$ independent final shares and the growth rate. (Capacity-off nodes have no utilization target, but their x_n is then irrelevant.)

[Status] Proof-checked by a fresh reviewer: sound. Every step was re-derived; on a three-node economy with the model's own tail functions the old equilibrium re-expressed in the new units satisfies the transformed equations to 3×10^{-16}, and re-solving from scratch gives the same thresholds, sales and GDP. Two conditions were added: $\alpha_m > 0$ for the invertibility, and $\eta \neq 1$ for the transformation of the weights; the counting was reworded.

The algorithm. In the *first stage* the productivity process is off: all industries of a node sit at one state, and the parameters follow from the targets without iteration, in the order of Section E.1; the condition that a node's capital grows at the rate g fixes its investment rate and, with r, its user cost, and equating the user cost to capital income per unit of capital gives x_n in proportion to the node's bundle price; given $\overline{Z}$ the pricing equations are linear in the log prices, and the final shares give the weights, whose normalisation determines $\overline{Z}$. In the *second stage* the process is on, the node laws spread out and the node means of the nonlinear objects move; the first stage is repeated with each target shifted by the miss, and a few rounds suffice. The capital shares are then those at which (228) holds for the node's mean intermediate spending over its mean sales. Two *outer loops* close the calibration: ε and σ_v by the rule of Section E.5.1 on the solved model, with the share above 95 % computed node by node with capacity weights; and ϕ by the simulated regression. The secondary calibration is computed by the same steps with its own targets.

[The first stage in closed form] At a point law node n's industries all have $z = 0$ and the same capital κ_n, so $\mathbb{E}_n[f] = f(\kappa_n, 0)$. Write $\Delta_n = 1/N$, f_n for the final-expenditure shares, and evaluate every capacity function at the node's threshold.

- *Thresholds.* $U_n(\overline{v}_n)$ equals the node's target; for a capacity-off node the threshold is the one at which utilization is 0.25, below which the node behaves to numerical precision like its no-capacity limit at the calibrated σ_v (at a much larger σ_v a lower utilization would be needed). The target must lie below the utilization at

which the markup becomes infinite, $\varepsilon\eta_{n,1}(\overline{v}) = 1$ ($\eta_{n,1}$ increases in the threshold under the log-normal law), above which $\mathrm{Y}_n \leq 0$ and the pricing equation has no solution.

- *Capital shares.* α_n solves (228), or takes the floor (Section E.5.2); for a capacity-off node the root coincides with (229) to numerical precision. In every later step ξ_n and ϱ_n are evaluated at the α_n that is used, so the steps are exact for floored nodes as well, which miss only their intermediate share.

- *The interest rate.* By (226) the nodes' sales per unit of final expenditure are $s = (I - a^{\mathsf{T}}\operatorname{diag}(\chi))^{-1} f$ with $\chi_n = \xi_n \varrho_n$; they do not depend on r. Node n's capital income is $\zeta_n[1-\xi_n]\varrho_n s_n$ with $\zeta_n[1-\xi_n] = 1 - \frac{1}{\varepsilon} - \xi_n$ by (77). At the rest point of capital, $\Phi_n(\iota_n) = \delta_n + g$ and $\mathfrak{q}_n = 1/\Phi_n'(\iota_n)$, neither of which depends on r, and by (52) at $\varsigma^K = 0$ capital income equals $[(r-g)\mathfrak{q}_n + \iota_n]$ times installed capital. The stock of the accounts is $\iota_n/(\delta_n + g)$ times installed capital. Hence (232) is

$$\frac{1}{\sum_n s_n} \sum_n \frac{\iota_n}{\delta_n + g} \cdot \frac{\zeta_n[1-\xi_n]\,\varrho_n\, s_n}{(r-g)\,\mathfrak{q}_n + \iota_n},$$

  which falls monotonically in r on the range where every user cost is positive, because every node's capital income is positive (Lemma 6(iii)), and tends to zero as $r \to \infty$. On the admissible range $r > g$ (detrended values are finite only if $\rho + (\gamma - 1)g > 0$) it falls from its finite value at $r = g$, $\sum_n s_n \zeta_n[1-\xi_n]\varrho_n/(\delta_n + g)\big/\sum_n s_n$, so the target gives exactly one admissible r if and only if it lies below that value, and then $\rho = r - \gamma g$.

- *Capital intensities.* At the rest point the user cost equals capital income per unit of capital, $[(r-g)\mathfrak{q}_n + \iota_n] = \zeta_n[1-\xi_n]\widehat{s}^{\mathrm{f}}_n/\kappa_n$. By the supply line (140) and $\widehat{s}^{\mathrm{f}} = \varrho_n \widehat{s}$,

$$x_n = h_n\,\mathcal{P}_n, \qquad h_n := \frac{\zeta_n[1-\xi_n]\,\varrho_n\, e^{-\sigma_v^2/(2\varepsilon)}\, U_n/\mathrm{Y}_n}{(r-g)\,\mathfrak{q}_n + \iota_n},$$

  with $e^{-\sigma_v^2/(2\varepsilon)} = \mu_{1/\varepsilon}^{\varepsilon/(\varepsilon-1)}$ the log-normal constant of the supply line.

- *Prices.* With $x_n = h_n \mathcal{P}_n$ the pricing line (139) is $\ln P_n = (1-\alpha_n)\ln \mathcal{P}_n - \alpha_n \ln h_n - \ln \mathrm{Y}_n - \ln \overline{Z}$, and $\ln \mathcal{P}_n = \sum_m a_{nm} \ln P_m - c_n$ with c_n the network constant of (95). In vector form

$$(I - \operatorname{diag}(1-\alpha)\,a)\ln P = -\operatorname{diag}(1-\alpha)\,c - \alpha \circ \ln h - \ln \mathrm{Y} - (\ln \overline{Z})\,\mathbf{1},$$

  a linear system whose matrix is invertible because every $\alpha_n \in (0,1]$.

- *Weights and the productivity level.* $\beta_n = f_n P_n^{\eta-1}/\Delta_n$, and the normalisation $\sum_n \beta_n \Delta_n = 1$ requires $\sum_n f_n P_n^{\eta-1} = 1$. Every $\ln P_n$ falls linearly in $\ln \overline{Z}$ with the node-specific slope $-\lambda_n$, $\lambda = (I - \operatorname{diag}(1-\alpha)a)^{-1}\mathbf{1} \geq \mathbf{1}$, so for $\eta \neq 1$ the left-hand side is strictly monotone in $\overline{Z}$ and runs over $(0,\infty)$ (increasing for $\eta < 1$, decreasing for $\eta > 1$): exactly one root. (At $\eta = 1$ it equals $\sum_n f_n = 1$ for every $\overline{Z}$, and the numeraire of the Cobb–Douglas case, Section B.7, pins $\overline{Z}$.)

- *Quantities.* Capital adding up gives GDP, $q = 1/\sum_n \zeta_n[1-\xi_n]\varrho_n s_n/[(r-g)\mathfrak{q}_n + \iota_n]$; sales are $\Delta_n \widehat{s}_n = q s_n$; capital is $\kappa_n = \zeta_n[1-\xi_n]\varrho_n \widehat{s}_n/[(r-g)\mathfrak{q}_n + \iota_n]$; and $x_n = h_n \mathcal{P}_n$ at the solved prices. The growth closure, the numeraire and the definition of GDP then hold identically: the growth target enters through the investment rates ι_n, and $\overline{Z}$ is determined by the normalisation of the weights.

[Status of the box above, and of the first-stage numbers] The closed form was proof-checked in an earlier draft by two fresh reviewers and again here, after its labels and wording were adapted, by a third: sound with gaps, now closed in the box. Every equation re-derives from the model's definitions, and a three-node test economy built from the closed form satisfies all seventeen equations and targets with a largest residual of 1.5×10^{-13}. Four corrections were folded in: the capital-stock moment is finite at $r = g$ (it tends to infinity only below g), so the admissible range $r > g$ and the bound on the target are stated; the direction of the root in $\overline{Z}$ depends on $\eta \lessgtr 1$; the capacity-off threshold at utilization 0.25 is exact to numerical precision at the calibrated σ_v only; and that $\eta_{n,1}$ increases in the threshold is stated. The interest-rate step was reproduced in this session by an independent implementation: $r = 0.03534$ against 0.03535 in the first-stage record (baseline calibration) and 0.03584 against 0.03585 (secondary calibration), so the first-stage interest rate is cross-examined. The second stage and the outer loop for ε and σ_v have run on the decided data set (every target met to 10^{-6} in two to four rounds; the solved model's estimates $\varepsilon = 9.60$ / 10.34 at the interim ϕ, computed once); the loop for ϕ is running. Every number through `prov` is the solved model at the interim ϕ. The model code gained the two shocks, to several nodes at once and to a node's weight in final demand, and a permanent shock.

E.6 Fit

Targeted moments. Table 16 lists them with the solved model's values, the baseline calibration at $\phi = 2$ and the secondary calibration at the interim ϕ. The utilization targets, the final-expenditure shares, the growth rate and the capital stock over sales are reproduced exactly by construction, each by its own parameter. The Census share is reproduced exactly by σ_v. The four supply-curve estimates are reproduced up to the distance that ε minimises. The intermediate shares of sales are reproduced exactly except in the floored nodes, which miss theirs by 3.7 percentage points at the median.

Untargeted moments. The following are reported against their data counterparts without being targeted; the first two concern the capacity block, the next four the rest of the calibration.

- *The slope of the share above 95 per cent on utilization: a miss.* In the data the share rises by 0.74 points per point of manufacturing utilization (fourth quarters, 1976–2016); in the model, along the threshold at the manufacturing target, it rises by 1.7. With one σ_v the model cannot match both the level and the slope of the share, and it is calibrated to the level.
- *The supply curves with input prices held fixed: a pass.* The published regression

Table 16: The targeted moments and the solved model's values.

moment	data	baseline calibration	secondary calibration
share of manufacturing capacity above 95 %	0.162	0.162	0.162
inverse supply elasticity, bins 1 to 4	0.12, 0.28, 0.31, 0.57	0.11, 0.20, 0.28, 0.55	0.12, 0.20, 0.29, 0.55
mean utilization, by node	49 targets	exact	exact (63 at 0.786)
intermediate share of sales, by node	$1 - \omega_n$	exact in 51 nodes; 12 floored	exact in 52; 11 floored
share in final expenditure, by node	63 shares	exact	exact
growth rate	0.021	exact	exact
capital stock over sales, all nodes	1.356	exact	exact

Notes: The baseline calibration at $\phi = 2$ and the secondary calibration at the interim ϕ. "Exact" means to numerical precision on the solved model; the floored nodes miss their intermediate share (Section E.5.2).

holds unit variable cost fixed; the model's own experiment of Section E.5.1 holds input prices fixed. Re-estimated with the materials price in the place of unit cost, the four coefficients are 0.14, 0.33, 0.30 and 0.84 (standard errors 0.11 to 0.17); the model's total slope at fixed input prices, at the baseline calibration's estimates on the solved model, is 0.15, 0.24, 0.33 and 0.60, a distance of $\chi^2 = 2.4$ with four degrees of freedom. The re-estimated coefficients do not reject the calibrated values, although on their own they would give a lower ε.

- *The dispersion of firm utilization.* At the fitted values the standard deviation of output over capacity across the firms of a node at the manufacturing targets is 0.129 (0.137 in the secondary calibration), between the 12.3 points that Amberg et al. (2026) report within industry-quarter cells and the 15.2 within Swedish manufacturing as a whole, as a node should be. It is a check and not a target; one run at $\sigma_v = 0.13$, the value that would match the within-industry dispersion, is in the robustness part.
- *Firms limited by equipment.* In the European surveys about 8 % of manufacturers name material or equipment as the factor limiting production (2011–2019); the model's constrained share at the manufacturing targets is 9.8 % (10.7 %).
- *The capital stock by group: a miss.* The stock is matched in total. By group, the model's capital stock over sales against the accounts' is 0.67 against 0.71 in manufacturing, 4.2 against 5.6 in mining and utilities, 1.38 against 1.28 in the transfer nodes and 1.31 against 1.25 in the nodes without a rate (baseline calibration); across nodes the correlation of the two is 0.77. With one interest rate and one markup rent the model cannot reproduce the spread of recorded returns across industries.
- *The return on capital, data against data.* The interest rate that the capital stock implies is a pre-tax return without pure profit. The accounts' own net return of

physical capital, operating surplus of physical capital less depreciation over the stock, is 12.0 % over 2007–2019; Gomme et al. (2011) find a pre-tax return on business capital of 10.2 % over the same years (11.9 % without capital gains). Both contain pure profit. The model's counterpart, the interest rate plus the markup rent per unit of capital, is 11.6 % (10.8 %).

- *utilization per unit of a node's own output.* How much a manufacturing node's utilization moves per per cent of its own real output, in the data over 1997–2019 and in the model under a common shock, is reported as a conditional moment of the time series.

[Flag: two untargeted moments are not yet computed on the decided protocol.] The movement of utilization per unit of a node's own output in the model (under a common shock, with the static block) and its data counterpart were computed in the review of the calibration at earlier values and found close; neither number is reprinted until it is regenerated. The break-adjusted slope of the Census share (the survey became quarterly in 2007) was also computed there and is not reprinted. Everything else in the list above is now the solved model's (computed once; the baseline calibration's at its final $\phi = 13$, 11 October 2026, and not recomputed at $\phi = 2$, where the first stage says they stand, being static objects of the thresholds, except the return, which is the first stage's value at $\phi = 2$, 11.6 against the solved model's 11.0 at $\phi = 13$, and the capital stock by group, which the first stage puts at 0.67, 4.16, 1.38 and 1.33 at $\phi = 2$; the secondary calibration's values of the dispersion and the constrained share from its stored ledger at the interim ϕ; the secondary calibration's return, 10.8 %, is still the first stage's): the slope along the threshold at the solved σ_v and the manufacturing mean capital share, the constrained share and dispersion as means over the 19 nodes at their targets, the capital stock by group and the model's return from the solved economy, the input-price check from the nodes' own curves. The accounts' side is the data set's.

[Evidence status of the rows of Table 16 and of the untargeted list]

- **Targeted.** The Census share and the four bins: the solved model at the final ϕ (the stored capacity block; computed once; the σ_v half cross-examined by an independent implementation). utilization, the intermediate shares, the final shares, the growth rate and the capital stock: the second stage's misses, all below 10^{-6} (the stored moments). No longer targeted (since 3 November 2026 $\phi = 2$ by convention; at $\phi = 2$ the moment is not computed, one evaluation of about nineteen hours, to be decided): the ϕ moment, 0.1675 at $\phi = 13.02$ against the data's 0.167 when it was targeted, the loop's third evaluation within its tolerance of 0.002 (computed once; the loop ran at the interim ε, σ_v, whose re-estimation at the final ϕ moved them by 0.0004 and 0.00006; not re-evaluated there); and the two moments of the productivity process ($V(40)/V(1)$, 22.5 in the panel against 23.5 at Corr $= 0.97$, matched not estimated, 17.4 at 0.95 and 33.1 at 0.99; the median manufacturing growth volatility 0.058, which Var reproduces by construction; data side, reproduced).
- **Untargeted.** The first six items of the list: the solved model at $\phi = 13$ (computed

once, 11 October 2026; the data counterparts reproduced from the data set and the accounts), not recomputed at $\phi = 2$ (the box on $\phi = 2$ in Section E.5.4: the five static items do not move with ϕ; the return is the first stage's 11.6 at $\phi = 2$). Two untargeted moments exist beyond the list. The capacity-growth regression run on the measured 1998–2019 history gives 0.190 against the data's 0.167, but only at $\phi = 11.5$ (computed once; its capital paths depend on ϕ, so the number is recomputed at 13.0, one kernel set of about nineteen hours, or dropped). The standard deviation and autocorrelation of utilization at the median core node (0.058 and 0.65 in the data, 1997–2019) are not computed in the model, although the stored kernels at the final ϕ give them without a new transition. The last item of the list (utilization per unit of own output) was computed in the review of the calibration at earlier values and found close; not reprinted until regenerated. Left out on purpose: the rationing rent (0.91 per cent of GDP), a prediction with no data counterpart, and the average markup (1.155), which needs a chosen literature value beside it.

The time series. The model has no aggregate risk, but its responses to one-time shocks describe its behaviour under a stream of shocks. Two series of shocks are measured for every node and year from 1998 to 2019 and fed in at their measured size, without scaling. The first is productivity with the model's own weights, the growth of a node's real gross output less the growth of its real stock of equipment and structures and of its real intermediate inputs, weighted by α_n and $1 - \alpha_n$, split into a component common to all nodes (standard deviation 0.7 % a year) and node residuals (standard deviation 4.0 % at the median node). The second is shifts of final demand: a node's share in final expenditure is $\beta_n P_n^{1-\eta}$, so the change in its weight is

$$\Delta \ln \beta_{n,t} = \Delta \ln \left(\text{share in final expenditure}\right)_{n,t} - (1 - \eta)\, \Delta \ln P_{n,t},$$

with the weights rescaled in every year so that a movement common to all prices drops out (standard deviation 4.7 % a year at the median node); it depends on η and measures shifts of final demand between the nodes only. At each date the static block is solved on the realised vector of the nodes' shocks, with capital given, rather than by adding responses up, because responses to common shocks are not linear; the part of the response that runs through capital comes from the model's paths after one-time shocks. The standard deviations, autocorrelations and cross-correlations of utilization, output and prices by node are then compared with the data over 1997–2019 for the 23 core nodes; none of them is targeted. utilization in the data is driven by demand: with productivity shocks alone the model reproduces about half of the volatility of utilization and not its path, and the measured demand shifts move with utilization.

[What is known about the time-series validation, and what is not]

- **Productivity shocks alone** (the measured shocks of every year fed into the static block with capital given, first stage, computed once): the yearly change in utilization of the median core node has about half the standard deviation it has in the data, and its path does not follow the data's; in 2009 the model's manufacturing utilization rises while the data's falls by eleven points. The common com-

ponent is small in the data, about what independent series give by chance, and unrelated to utilization outside 2008–2010; its 2009 value is the largest (+2.1 %). A node's own productivity shock hardly moves its own utilization, because it raises capacity and demand together; utilization moves when other nodes are shocked, and most of all under a common shock (model side, cross-examined by a reviewer who rebuilt the static block). The old plan's series, the integrated account's productivity, rises with demand and is not the model's productivity (data side, cross-examined).

- **The demand shifts** (checked once, this data set, the 23 core nodes, yearly changes 1998–2019): the measured shift is correlated +0.51 over time with the change in the Federal Reserve's rate at the median node and positive in 20 of 23; without 2008–2010 +0.34, positive in 22; the productivity shock with the model's weights +0.04 and −0.02. In 2009 the manufacturing nodes' mean shift is −10.6 % and their utilization falls by 11.2 points; in 2010 +7.2 % and +5.6 points. Limits: the measure depends on η; it measures the composition of final demand only; a node's final sales are part of its own output, so part of the relation is built in. **[Flag: what the model does with these shifts is not computed; the static block has no final-demand shock yet, and how the static block on realised vectors is combined with the capital paths is worked out in the implementation.]**
- **Not settled**: how the size of the node shocks should be measured (with the model's weights the node residual has a standard deviation of 4.0 % against 2.6 % in the account's own series, and most of the difference is the real intermediate-input series, which is noisy outside manufacturing); the value of the floor under shifts of demand (it was decided on the evidence of productivity shocks, under which it does not move the responses at the shock date); whether the G.17 data must be refreshed with the revision of November 2025 before any comparison after 2019.

[Q29 | text | ANSWERED] Please clarify what you mean by "The measured productivity shocks are Solow residuals with the utilization bias, so the model's path against the data's is not a validation yet." and what the potential solution(s) to this could be.

[What "Solow residuals with the utilization bias" means, and what would make the time-series comparison a validation]

- **The series.** The validation's productivity shock is a Solow residual with the model's weights, $\Delta z_n = \Delta \ln Y_n - \alpha_n \Delta \ln K_n - (1-\alpha_n)\Delta \ln M_n$, with Y real gross output, K the real stock of equipment and structures, M real intermediate inputs, and $1-\alpha_n$ about 0.86 on average and 0.93 in manufacturing. Inside the model the formula is exact for a firm with spare capacity: output is $z\,k^{\alpha}m^{1-\alpha}$ with k the stock, so the residual recovers z and utilization never enters it. In the data it is not exact: output is also produced with labour, hours and the services of the installed stock, none of which the residual sees, and purchased inputs are about 0.6 of gross output rather than 0.9. When demand falls, firms cut purchases faster than output and keep their workers and their plant, and a residual that puts a

weight of 0.9 on purchases reads the gap as a rise in productivity. Relative to a residual with the accounts' own shares (intermediates about 0.6 and labour about 0.2 of gross output in manufacturing), the model's exceeds it by about $-0.3\,\Delta \ln M + 0.2\,\Delta \ln L$, positive in a downturn in which purchases fall faster than hours. The sign is the opposite of the familiar bias of measured productivity, which falls with demand because the hoarded factors' utilization falls, but the cause is the same: the factors the residual holds fixed are being used less. That is the "utilization bias" of the warning.

- **What it does to the comparison.** The common component is +2.1 per cent in 2009, the largest value of the series, and in the model a common rise in productivity raises utilization: it raises every node's demand by much more than its capacity, since a common shock moves GDP about five to one at fixed capital (Section 6.1). Fed the measured 1998–2019 history (the three series together, the static block on each year's realised vector), the model's mean manufacturing utilization rises from 0.723 in 2008 to 0.778 in 2009 and 0.759 in 2010; in the data it falls by eleven points in 2009 and recovers six in 2010. The disagreement is in the shock, not in the propagation, so the path comparison would reject the measurement before it tests the model. A second gap stands behind the first: the demand shifts measure the composition of final demand only (a movement common to all nodes drops out by construction), and the model has no aggregate demand shock, so the economy-wide fall of utilization in 2009 has no shock in the model's menu that could produce it, whatever the productivity series says.
- **What would make it a validation**, in order of cost.
  1. *Compare what the bias cannot reach.* The node-level moments under the demand shifts and the node residuals: the volatility and persistence of utilization by node, which nodes move and by how much, and the pass-through of arriving demand shifts into utilization, model against data (the regression of Appendix G does the last), with the common component dropped, or kept with 2008–2010 set aside. The cheapest, and honest about what the model does not explain, the aggregate movements.
  2. *Replace the common component.* By a utilization-adjusted aggregate productivity series (Fernald's quarterly series, which corrects the aggregate residual for factor utilization with hours per worker by the method of Basu, Fernald and Kimball, annualised), so that 2009 enters as the technology movement it was and not as the measurement artefact; the node residuals stay (their 2009 values have both signs).
  3. *Treat 2009 as the aggregate demand shock the model lacks.* Add a shock to the household's discount rate, or to the level of final demand, back out its path from aggregate utilization, and test the model on the cross-section of the nodes' responses: which nodes' utilization falls most. This turns the path comparison into a test of the propagation given the aggregate, which is what this paper claims, and is the form a year-by-year comparison should take.

4. *Measure node productivity consistently.* A residual with the accounts' full set of inputs at their cost shares (capital, labour, intermediates), corrected for utilization with the Federal Reserve's rates and hours for the core nodes, in the manner of Basu, Fernald and Kimball; the model's z then absorbs labour's long-run movement and not its cyclical one. The standard remedy and the most work, a data-side project of its own.

[Evidence behind the answer above]

- The series and its 2009 value: the data set's table of validation productivity shocks and its documentation (the common component $+0.0210$ in 2009, sd 0.0069 over the years; the node residuals in 2009 from -0.25 to $+0.28$; in 42 of the 63 nodes real intermediate inputs fell by more than real output, most where they fell far more, oil and gas -50 against $+8$ per cent, computers -36 against -14). Reproduced from the file on 27 October 2026.
- The simulated path: the check on the measured history at $\phi = 11.5$ (the static block on each year's realised vector at the shifted distributions, all 22 solves converged; computed once): mean manufacturing utilization 2008 0.723, 2009 0.778, 2010 0.759. ϕ does not enter the static block at the date, so the numbers stand at $\phi = 13.0$. The data's eleven-point fall and six-point recovery, and the manufacturing nodes' mean demand shift of -10.6 per cent in 2009 and $+7.2$ in 2010, are the check of the demand-shift series in the box above (reproduced from the data set).
- "About five to one": the common-shock response of GDP with capacity, 4.9 per unit (the static results; computed once). "About half the volatility" of utilization under productivity shocks alone, quoted in the warning's neighbourhood: the first-stage check recorded in the box above (computed once).
- The sign of the bias is reasoned from the weights (the accounts' intermediate share of manufacturing gross output 0.62 and compensation about 0.2 of it, from the accounts; the model's $1 - \alpha_n$ from the stored calibration), not measured; a node-by-node comparison of the model-weight residual with the integrated account's productivity (in the data set's node parameters; it rises with demand) would measure it. The 2009 node residuals also carry deflator noise outside manufacturing (the real intermediate-input series is noisy there).
- A cheap variant of remedy 1, not listed above: project the node residuals on the measured demand shifts (own and arriving) and use the orthogonal part as the productivity series; it removes demand-driven variation at the node level but cannot touch 2009, because the shifts have no common component by construction. Remedy 2's series is public (Federal Reserve Bank of San Francisco); its 2009 value is to be read, not assumed. Remedy 3 needs a new shock type in the transition solver (a time path of ρ or of aggregate final demand), the one code change among the four. Remedy 4 would also change the size of the node shocks, listed as unsettled in the box above (4.0 per cent with the model's weights against 2.6 in the account's own series).

Robustness. The following variants are reported beside the baseline calibration: capacity on in the 23 core nodes only; capacity on in all 63 nodes, with the 14 others at the mean transfer target; ε at the two ends of its 95 % interval, with σ_v from the share; $\sigma_v = 0.13$; intellectual property counted as capital consistently, in income, in the stock and in depreciation; capital income measured as operating surplus, and as the capital compensation of the integrated production account; ω_n after redefinitions; $\phi = 2$; $\eta = 0.3$ and 0.9; the autocorrelation of the productivity process at 0.95 and 0.99. Beside the secondary calibration: its own ends of the interval and the same definitional variants. The one change that would remove the floors, an overhead cost, is a change of the model and is taken up after the first calibration (Section E.5.2).

[For the quantitative section: the first-stage results of the two calibrations, and the two reporting decisions] Not for this appendix. All numbers are first stage (all industries of a node at one state), at each calibration's own first-stage estimates, computed once with the first-stage calculator on the 63-node data set; the three ratios against the same-parameters economy are the capacity economy's own parameters with capacity switched off.

	baseline calibration	secondary calibration
GDP / gross-output / other-nodes ratio, same-parameters economy	0.77 / 0.53 / 0.45	0.88 / 0.70 / 0.62
the same against the recalibrated economy	0.85 / 0.59 / 0.49	0.92 / 0.73 / 0.65
asymmetry of a common shock of 1, 2, 5 %	1.04, 1.08, 1.21	1.04, 1.09, 1.21
rationing rent, % of GDP	1.00	0.52
long-run loss of GDP per unit of capital; capital in other nodes	9.7 %; 6.3 %	1.8 %; 3.3 %
growth effect at $\gamma = 1, 2, 3$, points	0.39, 0.21, 0.14	0.06, 0.03, 0.02

- **The statistics.** GDP ratio: the peak response of GDP (per unit of aggregate capital) to a $+1$ % productivity shock at one node with capacity over without, the mean over the nodes shocked in turn. Gross-output ratio: the same for real gross output. Other-nodes ratio: the same for the real output of the nodes that are not shocked. Asymmetry: all nodes shocked together, the response of GDP to $-x$ % over the response to $+x$ %. Long-run loss: how much lower GDP per unit of aggregate capital is on the balanced growth path with capacity than in the same-parameters economy. Reallocation: the share of aggregate capital that is in different nodes. Growth effect: the difference in the growth rate.
- **Three facts behind the numbers** (cross-examined in the review of the calibration). The response at the shock date is a property of the static block at the state before the shock, so ϕ, δ_n, r, γ and the persistence of the shock cannot move it when the model is recalibrated to the same targets. A node's own shock hardly moves its own utilization; capacity acts on the other nodes and on gross output much more than on GDP, and when all nodes are shocked together the response is asymmetric. The answer depends on which no-capacity economy is used: the same-parameters economy carries the pinned capital shares and misses the accounts, and is more elastic than a no-capacity economy calibrated to the same

targets; the recalibrated economy fits the accounts with a different technology.

- **Reporting decisions.** The main comparison is the same-parameters economy, for impact responses (read as the effect of capacity) and for the long run; the recalibrated economy is a second column of every table of impact responses (read as the error of a network model without capacity fitted to the same accounts) and has no long-run column, because it hits the same long-run targets by construction. The lead statistics are the other-nodes ratio, the gross-output ratio and the asymmetry of common shocks; the GDP ratio is given as support, with its numerator and its denominator. No magnitude is taken from the two-node economy. One dynamic statistic is added: the cumulative response of GDP over ten years, with capacity over without (not computed yet). In the long run the lead statistics are the loss, the reallocation of capital and the rationing rent; the growth effect is reported for $\gamma = 1, 2$ and 3 with the statement that its size follows from a convention, because the economy without capacity keeps ρ and its interest rate moves with its own growth rate. GDP is the model's GDP, which has no labour income. Floored nodes are flagged.
- **What the difference between the two calibrations rests on**: almost entirely the level of the service nodes. A lower ε gives larger effects, mainly through the comparison economy, whose capital shares fall; counting intellectual property as capital, or leaving the pay of the self-employed inside capital income, gives smaller ones. The overhead-cost version would move the results to about the upper end of the reported range.

[Evidence status of the calibration, and the cost of a change]

- The stored baseline calibration (a parameter file, the node table, the capacity block, the moments, a README): $\varepsilon = 9.6044$, $\sigma_v = 0.18448$, $\phi = 13.0202$, $r = 0.03335$, $\rho = -0.00859$, $g = 0.020972$, $\overline{Z} = 0.3482$, Var $= 0.05703$, Corr $= 0.97$; second stage two rounds, worst miss 4.3×10^{-7}; grid $N_\kappa = 101$, log-width 0.5, $N_z = 11$. The data side is the calibration data set (every table with its documentation).
- Data-side numbers (accounts, ω_n, δ_n, targets, the shock series): reproduced from the data set. The first-stage interest rate and the capital shares: cross-examined by an independent implementation. The solved model's parameters: computed once (the second stage meets every target to 10^{-6}). The seven derivations of this appendix (the Leontief relation, the variogram, the tail functions, the uniqueness of the pinned share, the units transformation, the first stage in closed form, the bin counterpart): each proof-checked by a fresh reviewer.
- **Cost of a change.** The ϕ moment is about 60 hours of transitions per calibration (nineteen hours per evaluation, three evaluations); the rest of the chain is under two hours; a cold stationary solve with its twin about two minutes; a static solve at the fixed distribution four to eight seconds; one transition 22 to 79 minutes. Trial calibrations run at $\phi = 13.0$ unless they are promoted (rule agreed on 24

October 2026); any change to ε, σ_v or the targets strictly changes the kernels ϕ is fitted to, so the held ϕ is an approximation to be flagged in that section.

- **Open**: the two closures (the ϕ loop ran at the interim ε, σ_v, whose re-estimation at the final ϕ moved them by 0.0004 and 0.00006; not re-evaluated), and the robustness variants promised in Section E.6 (none run).

F The Results by Node

Table 17 lists, for every node of the baseline calibration, its position in the network, its utilization, rationed share and markup on the balanced growth path, its share of aggregate capital with and without capacity, its long-run output gap, and the ratio of the responses of GDP and of the other nodes' output to a +1 per cent shock at the node, with capacity over without. Nodes are ordered from the most upstream to the most downstream.

Table 17: Statistics by node, baseline calibration.

code	group	upstr.	U	rationed %	markup	K share cap	K share no cap	output gap log pts	GDP ratio	others ratio
113FF	off	2.90	–	–	–	0.0036	0.0042	17.07	0.73	0.45
331	mfg	2.70	0.782	8.5	1.132	0.0056	0.0050	18.96	0.71	0.40
55[f]	transfer	2.62	0.875	26.9	1.177	0.0192	0.0064	25.03	0.68	0.36
487OS	transfer	2.56	0.862	23.3	1.166	0.0100	0.0083	23.39	0.66	0.34
561	transfer	2.52	0.864	23.6	1.167	0.0147	0.0064	26.56	0.65	0.33
493	transfer	2.48	0.863	23.4	1.167	0.0036	0.0018	20.67	0.72	0.40
ORE	transfer	2.43	0.889	31.1	1.191	0.1765	0.1814	20.72	0.66	0.33
321[f]	mfg	2.43	0.765	6.6	1.128	0.0008	0.0006	19.19	0.71	0.36
211	min/util	2.42	0.930	47.7	1.272	0.0456	0.0510	19.13	0.69	0.38
212	min/util	2.40	0.826	15.2	1.146	0.0109	0.0125	16.95	0.71	0.41
322	mfg	2.39	0.866	24.3	1.169	0.0044	0.0025	24.11	0.66	0.33
332	mfg	2.35	0.786	8.9	1.133	0.0074	0.0069	19.01	0.72	0.38
486	transfer	2.33	0.885	29.8	1.186	0.0094	0.0111	17.77	0.67	0.36
111CA	off	2.30	–	–	–	0.0057	0.0064	15.70	0.78	0.44
327	mfg	2.24	0.733	4.0	1.124	0.0035	0.0039	16.96	0.74	0.39
482	transfer	2.21	0.885	29.8	1.186	0.0100	0.0102	17.84	0.71	0.40
326	mfg	2.20	0.820	14.0	1.144	0.0025	0.0016	19.41	0.73	0.38
562	off	2.18	–	–	–	0.0059	0.0068	20.48	0.68	0.34
323	mfg	2.16	0.796	10.3	1.136	0.0033	0.0035	22.11	0.67	0.34
524	transfer	2.08	0.856	21.7	1.162	0.0517	0.0515	22.75	0.64	0.30
532RL	transfer	2.05	0.852	20.8	1.160	0.0319	0.0386	19.93	0.67	0.34
325	mfg	2.04	0.770	7.1	1.129	0.0181	0.0190	18.13	0.72	0.37
5411	transfer	2.01	0.882	28.8	1.183	0.0325	0.0370	21.37	0.66	0.32
22	min/util	1.98	0.853	20.8	1.160	0.0796	0.0898	16.54	0.69	0.35
514	transfer	1.96	0.863	23.4	1.167	0.0093	0.0095	23.04	0.66	0.33
5412OP[f]	transfer	1.91	0.868	24.7	1.170	0.0235	0.0087	24.86	0.68	0.34
484	transfer	1.90	0.885	29.8	1.186	0.0042	0.0024	20.39	0.74	0.41
313TT[f]	mfg	1.89	0.787	9.0	1.133	0.0004	0.0003	17.95	0.75	0.38
711AS[f]	transfer	1.87	0.848	19.7	1.157	0.0031	0.0014	25.04	0.66	0.31
52BSF	transfer	1.83	0.846	19.4	1.156	0.0608	0.0633	20.96	0.68	0.33
42	off	1.71	–	–	–	0.0841	0.0954	17.04	0.75	0.40
512[f]	transfer	1.71	0.823	14.6	1.145	0.0017	0.0009	22.87	0.70	0.32
513	transfer	1.67	0.873	26.2	1.175	0.0629	0.0678	19.34	0.69	0.33
324	mfg	1.65	0.855	21.5	1.162	0.0239	0.0230	16.57	0.73	0.37
485	transfer	1.61	0.885	29.8	1.186	0.0038	0.0039	20.37	0.68	0.33
335	mfg	1.60	0.821	14.1	1.144	0.0028	0.0027	16.86	0.77	0.38
5415[f]	transfer	1.53	0.888	31.0	1.190	0.0044	0.0013	24.42	0.72	0.35
3361MV[f]	mfg	1.51	0.755	5.7	1.127	0.0019	0.0015	16.07	0.80	0.35

code	group	upstr.	U	rationed	markup	K share	K share	output gap	GDP	others
721	transfer	1.51	0.841	18.3	1.153	0.0161	0.0176	18.94	0.69	0.32
311FT	mfg	1.50	0.804	11.4	1.138	0.0061	0.0039	15.24	0.83	0.38
481	transfer	1.48	0.837	17.3	1.151	0.0089	0.0093	18.52	0.71	0.34
334	mfg	1.44	0.774	7.6	1.130	0.0070	0.0073	16.36	0.78	0.38
511	transfer	1.42	0.843	18.7	1.154	0.0032	0.0019	20.87	0.74	0.34
81	transfer	1.41	0.850	20.3	1.158	0.0123	0.0064	18.97	0.75	0.37
333	mfg	1.41	0.777	7.9	1.131	0.0041	0.0038	15.44	0.81	0.40
339	mfg	1.38	0.766	6.7	1.129	0.0024	0.0025	13.93	0.85	0.43
722	transfer	1.35	0.838	17.7	1.152	0.0111	0.0081	18.77	0.75	0.34
337[f]	mfg	1.34	0.774	7.5	1.130	0.0003	0.0002	14.00	0.86	0.44
315AL[f]	mfg	1.31	0.760	6.2	1.128	0.0001	0.0001	15.64	0.81	0.40
483	transfer	1.31	0.855	21.5	1.162	0.0008	0.0006	16.23	0.80	0.41
213	min/util	1.27	0.720	3.2	1.122	0.0040	0.0046	15.85	0.77	0.30
3364OT	mfg	1.20	0.746	5.0	1.126	0.0041	0.0042	12.24	0.91	0.38
23	off	1.17	–	–	–	0.0105	0.0118	14.29	0.82	0.32
4A0	off	1.16	–	–	–	0.0257	0.0276	12.36	0.89	0.46
441	off	1.15	–	–	–	0.0085	0.0091	13.05	0.86	0.41
713	transfer	1.06	0.846	19.4	1.157	0.0021	0.0016	13.26	0.88	0.40
61	off	1.06	–	–	–	0.0066	0.0069	11.32	0.92	0.46
452	off	1.03	–	–	–	0.0009	0.0010	11.90	0.97	0.69
621	off	1.03	–	–	–	0.0102	0.0108	10.68	1.01	1.53
445	off	1.01	–	–	–	0.0045	0.0048	10.50	0.97	0.71
622	off	1.00	–	–	–	0.0061	0.0063	10.65	1.01	1.53
623[f]	off	1.00	–	–	–	0.0005	0.0006	10.74	1.00	1.31
624[f]	off	1.00	–	–	–	0.0005	0.0005	10.72	1.00	1.47

Notes: Utilisation, the rationed share and the markup are means over the node's industries; the capital share is the node's share of aggregate capital with capacity and in the economy without capacity constraints; the output gap is ln(real output without capacity / with); the last two columns are the response of GDP and of the other nodes' real output to a $+1\,\%$ productivity shock at the node, with capacity over without. Floored nodes are marked [f].

G A Regression the Model Implies: State-Dependent Pass-Through of Network Shocks

[Flag: set aside on 30 October 2026] This appendix is held, unchanged and with its heading, inside a maroon box, until what to do with the regression exercise is decided. Its one mention in the main text, the parenthesis closing the sentence on the node-level price response in Section 6.1, is in a maroon box too; its two other mentions were already inside private boxes. With the maroon boxes switched off the appendix and the mention vanish together, so no reference dangles, and no other appendix is renumbered because this is the last one.

[Status and how to read this appendix]

- This appendix is the complete record of one exercise: a regression on the sector data that comes from the model, so that a motivating fact and the model speak about the same object. It was written on 24 October 2026 so that the whole exercise can be read in one place and judged; nothing in it is yet referred to from the main text, and where its result should go (Section 6.1, the introduction, or nowhere) is decided in Appendix G.7.
- Three layers, as in Appendix E. Black text is written as a candidate for the paper. Maroon boxes hold derivations and proposed sentences for the main text. Dark-teal boxes, like this one, hold what is private: the evidence status of every number, the warnings, the alternatives not taken and the decisions that are open. With the dark-teal boxes switched off, what remains is the appendix as it would stand in the paper.
- Evidence status, in the register of Appendix E: the log-linearised supply relation of Appendix G.2 was checked line by line against the lemmas it cites, in session, not by a fresh reviewer; the model's responses of Appendix G.3 were computed once, except the capacity economy's two stored impact matrices of Section 6.1, which the re-solve of that appendix reproduced exactly; the panel and the regressions of Appendices G.4–G.6 were computed once, with the checks stated there. Every number in this appendix is read from the stored output of the computation it belongs to.
- The figures are the exploratory ones of the computation, not drawn in the paper's style; the tables are generated from the stored output. If the appendix stays, the figures are redrawn and the tables regenerated by the paper's own figure and table scripts.

G.1 Overview

Section 6.1 shows that a shock arriving at a node through the network moves the node's quantities by less when the node is tight, and its price by about as much as it would without capacity. That is a statement about the solved model. This appendix

asks whether the same pattern is in the data, and it does so with a regression that the model itself implies, rather than with a correlation chosen by eye. The exercise has five parts. Appendix G.2 derives from the model's supply and demand relations which regression carries a prediction with a sign, and which does not. Appendix G.3 reads the model's own prediction for that regression off the solved economy, so that the data are compared with a number and not with a sentence. Appendix G.4 describes the panel and the two series of shocks, one of which, the shift of final demand between nodes, is used as a regressor for the first time here and is looked at before it is used. Appendix G.5 reports the estimates, and Appendix G.6 runs the same regression on the model's responses, in the capacity economy and in the economy without capacity constraints, to see how much of a cross-sectional interaction is capacity and how much is node heterogeneity.

What comes out, in one paragraph. The regression the model implies is estimable on the 24 nodes that have a utilization series and, in a cross-sectional form, on the 49 capacity-on nodes. On the quantity side it comes out with the model's sign: the pass-through of a demand shift arriving from a node's customers into the node's real output growth and into its utilization is smaller, and in the data about zero, when the node is tight, in every form and variant, at one to three standard errors. On the price side it comes out as the model says once the oil node is set aside: a pass-through of 0.4 to 0.8 with no interaction with tightness. Two things limit what the result can carry. The cross-sectional form of the interaction is two thirds node heterogeneity by the model's own economy without capacity constraints, so the within-node form, with 24 nodes, is the one that speaks for capacity; and the identification sits in a handful of upstream nodes, one of which, oil and gas extraction, is an oil-price story on the price side.

G.2 The Regression the Model Implies

Throughout, i is a node of the calibrated economy, which is a sector of the finite-sector model of Section 2; U_i is its utilization, P_i its shadow price index, $p_i = P_i \mathcal{R}_i(\overline{v}_i)$ its posted price (Lemma 3(i)), Z_i its productivity, $\overline{v}_i$ its threshold, x_i its capital intensity, and a_{ij} the cost share of input j in its intermediate bundle, with $\sum_j a_{ij} = 1$ (9). The log-derivatives of the threshold map and of utilization along the threshold are those of the proof of Lemma 15(iv),

$$\ell_{Y,i}(\overline{v}) := \frac{\mathrm{d}\ln Y_i}{\mathrm{d}\ln \overline{v}} \geq 0, \qquad \ell_{U,i}(\overline{v}) := -\frac{\mathrm{d}\ln U_i}{\mathrm{d}\ln \overline{v}} > 0,$$

and the node's supply elasticity is $e_{S,i} = 1 + \ell_{U,i}/\ell_{Y,i}$ (111): the elasticity of the node's sales to its own price at given capital and productivity.

The supply relation. At the node's symmetric equilibrium the price and the threshold satisfy the pricing relation (107), $P_i\, Y_i(\overline{v}_i) = \mathcal{P}_i(\mathbf{P})/(Z_i x_i^{\alpha_i})$, with the unit cost of the bundle $\ln \mathcal{P}_i = \sum_j a_{ij} \ln P_j - \sum_j a_{ij} \ln a_{ij}$ (13), and utilization is a strictly decreasing function of the threshold alone, $U_i = U_i(\overline{v}_i)$ (Definition 1). Neither relation contains the

node's capital: capital decides where the threshold lands, not the link between the threshold, the price and utilization. Log-differentiating both for any change in the state and eliminating the threshold gives

$$\Delta \ln U_i = (e_{S,i} - 1)\left[\Delta \ln P_i - \sum_j a_{ij}\,\Delta \ln P_j + \Delta \ln Z_i\right]. \tag{234}$$

utilization rises with the node's own price relative to its input-cost index and with its productivity, and the slope $e_{S,i} - 1 = \ell_{U,i}/\ell_{Y,i}$ is the elasticity of utilization, equivalently of output at given capital and productivity, to the price. The slope is a function of the threshold alone, so of the node's own utilization. Written with the price on the left, (234) is the linearised supply curve that Boehm and Pandalai-Nayar (2022) estimate on manufacturing industries: the change in the price equals the inverse slope times the change in utilization plus the change in marginal cost, and here the marginal-cost term is $\sum_j a_{ij}\Delta \ln P_j - \Delta \ln Z_i$, so the network sits inside it. Their estimating form interacts the change in output with the deviation of the previous year's utilization from the industry's mean, which is the within-node form of tightness used below.

[Derivation of (234), and its posted-price form]

- Take logs of the pricing relation: $\ln P_i + \ln \mathrm{Y}_i(\overline{v}_i) = \ln \mathcal{P}_i - \ln Z_i - \alpha_i \ln x_i$. For a change in the state at fixed x_i, $\Delta \ln P_i + \ell_{Y,i}\,\Delta \ln \overline{v}_i = \sum_j a_{ij}\Delta \ln P_j - \Delta \ln Z_i$, using $\Delta \ln \mathcal{P}_i = \sum_j a_{ij}\Delta \ln P_j$ at fixed cost shares.
- Take logs of $U_i = U_i(\overline{v}_i)$: $\Delta \ln U_i = -\ell_{U,i}\,\Delta \ln \overline{v}_i$, so $\Delta \ln \overline{v}_i = -\Delta \ln U_i/\ell_{U,i}$.
- Substitute: $\Delta \ln P_i - (\ell_{Y,i}/\ell_{U,i})\,\Delta \ln U_i = \sum_j a_{ij}\Delta \ln P_j - \Delta \ln Z_i$, and $\ell_{U,i}/\ell_{Y,i} = e_{S,i} - 1$ by (111). This is (234). The relation holds at every level of the node's capital, because neither the pricing relation nor the utilization map contains K_i; the statement "at given capital" in the definition of $e_{S,i}$ is about the elasticity, not about the relation.
- *Posted prices.* With $\ln p_i = \ln P_i + \ln \mathcal{R}_i(\overline{v}_i)$ and $\ell_{\mathcal{R},i} := \mathrm{d} \ln \mathcal{R}_i/\mathrm{d} \ln \overline{v} > 0$ (Lemma 3(i): $\mathcal{R}_i$ is increasing), the same steps give

$$\Delta \ln U_i = \frac{\ell_{U,i}}{\ell_{Y,i} - \ell_{\mathcal{R},i}}\left[\Delta \ln p_i - \sum_j a_{ij}\,\Delta \ln P_j + \Delta \ln Z_i\right],$$

  the input prices staying shadow indices because the aggregators sell at P_j. When the node tightens, the posted price moves less than the shadow index, the difference being the rationing rent. Which of the two the gross-output price index of the accounts measures is the question settled in Appendix E in favour of the shadow index; Appendix G.3 shows that at the sizes of shock that matter here the two respond alike, so the price column of the regression does not depend on the answer.
- *The relation to the supply-curve estimates of Appendix E.* The slope of (234) by bin of lagged utilization is what the four estimates of Table 13 measure, and those

estimates are the target of σ_v. Re-estimating (234) on the node data would therefore motivate "capacity" with worse data than the calibration already uses, and would say nothing about the network.

Where the neighbours enter. Stacking the supply relations over the nodes, with $\mathbf{D} = \operatorname{diag}(e_{S,i} - 1)$ and $A = (a_{ij})$,

$$(I - A)\,\Delta \ln \mathbf{P} = \mathbf{D}^{-1} \Delta \ln \mathbf{U} - \Delta \ln \mathbf{Z}. \tag{235}$$

The matrix $I - A$ is singular, because every row of A sums to one: the supply relations are homogeneous of degree zero in prices and fix the price vector only up to a common component, which the numeraire fixes. Off that component the system is invertible and its solution carries every power of A, so a node's price depends on the tightness of the whole chain of its suppliers, each scaled by its slope, and not on the direct suppliers alone. Downstream tightness enters through demand. The value-clearing identity (29), $S_i = \beta_i P_i^{1-\eta} Q + \sum_k a_{ki} X_k$, log-differentiated, gives

$$\Delta \ln S_i = s_i^F \left[\Delta \ln \beta_i + (1-\eta)\,\Delta \ln P_i + \Delta \ln Q\right] + \sum_{k \neq i} s_{ki}^D \,\Delta \ln X_k, \tag{236}$$

where $s_i^F = \beta_i P_i^{1-\eta} Q / S_i$ is the share of node i's sales that goes to final demand, $s_{ki}^D = a_{ki} X_k / S_i$ the share that goes to node k, and X_k the customer's intermediate spending, a threshold-dependent share of its sales, $X_k = \xi_k(\overline{v}_k)\, \varrho_k(\overline{v}_k)\, S_k$ (Lemma 5(i) and (118)). So the model's weights are cost shares upstream and sales shares downstream, and what the model puts on the right-hand side of a node's quantity equation is prices and spending, not utilization rates.

[The singular system (235): the numeraire, the solvability condition and a check]

- $(I - A)\mathbf{1} = 0$ by (9), so $I - A$ has a one-dimensional null space spanned by the common component of the price vector. The numeraire (26), $\sum_j \beta_j P_j^{1-\eta} = 1$, differentiated at fixed β, pins it: $\sum_j \beta_j P_j^{1-\eta}\, \Delta \ln P_j = 0$.
- A solution exists only if the right-hand side is orthogonal to the left null vector of $I - A$: with $\boldsymbol{\pi}$ the left Perron vector of A ($\boldsymbol{\pi}^\top A = \boldsymbol{\pi}^\top$), the condition is $\boldsymbol{\pi}^\top [\mathbf{D}^{-1} \Delta \ln \mathbf{U} - \Delta \ln \mathbf{Z}] = 0$: the π-weighted sum of the utilization changes, each scaled by its inverse slope, equals the π-weighted sum of the productivity changes. The usual Leontief inverse exists when a primary factor with a given price sits in the cost shares; here capital is quasi-fixed and outside the bundle, so the leak out of the input-output loop is zero and the inverse does not exist.
- *Check.* Under a common productivity shock δ in the no-capacity limit of a node, $e_S - 1 = (1-\alpha)/\alpha$ (Lemma 15(iv)), and with a common α the relative prices do not move, so (234) gives $\Delta \ln U = \delta(1-\alpha)/\alpha$. Capacity is proportional to productivity (Definition 1), so output moves by $\Delta \ln U + \delta = \delta/\alpha$ at fixed capital, the Cobb–Douglas multiplier of productivity at given capital with an optimally chosen bundle. The solvability condition holds with equality on both sides.

- *The customer's spending.* From $X_k = \xi_k \varrho_k S_k$, $\Delta \ln X_k = \Delta \ln S_k + \ell_{\xi\varrho,k} \Delta \ln \overline{v}_k$ with $\ell_{\xi\varrho,k}$ the log-derivative of the product of the two shares along the threshold. That a tight customer pulls in less when its demand rises rests on two things: its sales respond less (the computed fact of Section 6.1), and the materials share of its sales does not rise with its tightness, which is reasoned and not computed; the sign of $\ell_{\xi\varrho,k}$ is computable from the tail functions.

What the model does not predict. A regression of a node's utilization on a weighted average of its neighbours' utilization, in levels or in changes, has no sign in the model. In levels the calibration fits x_i to each node's utilization target, so the cross-section of levels is used up by the calibration and not predicted by it. In changes the sign depends on the mix of shocks: in years driven by demand all nodes tighten together, which gives a positive coefficient; in years driven by costs a tight supplier means dearer inputs, so the node raises its price and sells less, which gives a negative one. A positive coefficient would therefore not say that capacity matters in the network. This is the sharp form of the worry that a price sits on both sides of such a regression.

[The earlier regressions of utilization on the neighbours' utilization] Two such regressions were run before the model was calibrated, and this is what they found, read from their stored output (not re-run). (i) One cross-section, the year 2002, on 279 detailed manufacturing industries of the benchmark input-output table, utilization from the Census plant survey: OLS in levels without controls gave a coefficient of 0.53 (s.e. 0.22, 100 five-digit industries) and 0.64 (s.e. 0.11, 162 six-digit industries) on the cost-share-weighted utilization of the suppliers, and about zero on the customers'. (ii) A panel of the 19 Federal Reserve manufacturing groups, 1997–2024, with industry and year effects: nothing on utilization was significant (suppliers -0.04, s.e. 0.14; customers -0.03, s.e. 0.13). The level result is what the model says a level regression cannot carry, and the null result with fixed effects is consistent with a coefficient that has no model sign. Neither is worth re-running; the regression below replaces them.

What the model predicts with a sign: state-dependent pass-through. Section 6.1 says that a shock arriving at a node through the network moves the node's quantities by less when the node is tight, and that the dampening sits at the tight nodes a shock passes through. The regression that carries this is the one of Acemoglu et al. (2016), a node's growth on the shocks of its customers and of its suppliers built with the input-output table, with the interaction of Boehm and Pandalai-Nayar (2022) between each shock and the node's lagged utilization:

$$\begin{aligned} y_{it} = \alpha_i + \lambda_t + b_1 \,\mathrm{Down}_{it} + b_2 \,\mathrm{Down}_{it}\, \tau_{it} + b_3 \,\mathrm{Down}_{it}\, \tau_{it}^{D} \\ + c_1 \,\mathrm{Up}_{it} + c_2 \,\mathrm{Up}_{it}\, \tau_{it} + d\, \tau_{it} + e\, \Delta \ln \beta_{it} + \varepsilon_{it}, \end{aligned} \tag{237}$$

where y_{it} is the node's real output growth, its price growth or the change in its utilization; Down_{it} is the demand shift arriving from the node's customers, the customers' shifts of final demand weighted by the share of the node's sales that goes to each;

Up_{it} is the productivity shock arriving from its suppliers, the suppliers' productivity growth weighted by cost shares; τ_{it} is the node's tightness, its lagged utilization less a mean; τ_{it}^{D} is the tightness of its customers, weighted the same way as the shift; and $\Delta \ln \beta_{it}$, the node's own shift of final demand, is the control for the node's own shock. The predictions are set out in Appendix G.3, where they are read off the solved economy; in words they are these. In quantities, $b_1 > 0$ and $b_2 < 0$: a demand shift arriving through the network raises the node's output, and by less when the node is tight. This is the sentence of Section 6.1 in regression form, and it is the fully covered column, because every node in the panel has its own utilization series. $b_3 < 0$ is the same mechanism one step down the chain, a shift arriving through tight customers is damped because they do not pull in more inputs, and it is poorly covered, because most customers of manufacturing have no utilization series. In prices the two columns differ: for arriving demand shifts the receiving node's price responds as it would without capacity and the response does not vary with the node's tightness, so the Down column predicts no price interaction; for arriving productivity shocks the pass-through of a supplier's productivity into the receiving node's price is damped when the node is tight. Four cautions, stated once. Only 24 of the 63 nodes have a utilization series, so the within-node form of the regression runs on 24 nodes over 22 years. The measured productivity shocks are Solow residuals, contaminated by utilization, which is why the Down column, built from demand shifters, is the main one and the Up column is secondary. A calibration moment is not a motivating fact: the supply-curve bins and the response of capacity to lagged utilization are inside the calibration, the network interactions are outside it. And the regression is a motivating one, without instruments, with lagged utilization as the modifier: the claim it can make is that the pattern the model predicts is in the data, not that the model is tested.

G.3 The Model's Own Prediction for the Interaction

Before the regression is run on the data, this appendix asks what the solved economy itself says about its interaction terms, in the same units. The object is the response of a *receiving* node to a shock at one of the other 62 nodes, the arriving shock, read in the capacity economy and in the economy without capacity constraints. For each receiving node the ratio of its response with capacity to its response without is fitted as a slope through the origin across its 62 arriving shocks, so that one number per node summarises how much capacity dampens what reaches it; the R^2 of that fit says how well one number does. "Real activity" is the node's sales deflated by its shadow price index, "price" the shadow price index, "posted price" $P_i \mathcal{R}_i(\overline{v}_i)$, all per unit of the shock, and the node's utilization is its mean on the balanced growth path of the capacity economy. Two kinds of arriving shock are used. The first is the $+1$ per cent productivity shock at one node of Section 6.1, whose stored responses are re-read here; the capacity economy was also re-solved from its stored state and the 63 shocks repeated, and the stored response matrices were reproduced to every digit. The second is new: a $+1$ per cent shift of one node's weight in final demand, the weights of the other nodes res-

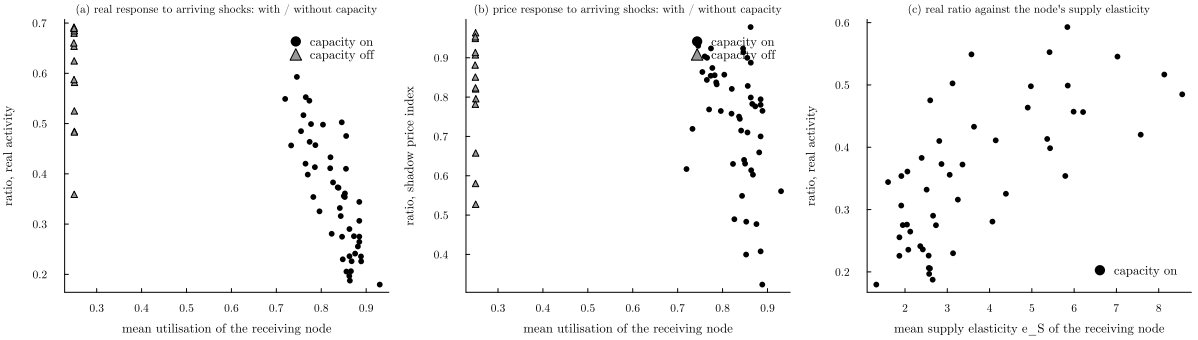
caled so that they still sum to one, solved with the static block at the balanced-growth distribution of each economy; all 126 solves converged. This second shock is the channel of the regression's main column, since a demand shift at a customer reaches its suppliers as demand. Table 18 collects the results and Figures 12 and 14 draw them.

Table 18: The model's prediction for the interaction, by receiving node.

	mean	range	median R^2	slope on the node's utilisation (s.e.)	correlation
arriving productivity shocks, 49 capacity-on receiving nodes					
real activity, with capacity over without	0.36	0.18–0.59	0.96	−1.85 (0.20)	−0.80
shadow price index, with over without	0.74	0.32–0.98	0.94	−1.33 (0.42)	−0.42
posted price, with over without	0.74	0.36–0.97	—	−1.34 (0.42)	−0.42
shadow price index, demeaned across receivers (year-effect analogue)	0.59	0.10–0.88	—	−2.71 (0.51)	−0.61
utilisation response, points per unit of the no-capacity real response	0.29	0.16–0.44	—	−1.14 (0.17)	−0.70
the 14 capacity-off receiving nodes: real 0.60, shadow price 0.81					
arriving demand shifts, 49 capacity-on receiving nodes					
real activity, with capacity over without	0.66	0.24–0.99	0.93	−2.89 (0.38)	−0.74
shadow price index, with over without	0.98	0.90–1.10	0.97	0.07 (0.10)	0.09
posted price, with over without	0.98	0.90–1.08	0.98	0.05 (0.10)	0.08
shadow price index, demeaned across receivers (year-effect analogue)	1.43	0.23–4.31	—	−0.58 (1.79)	−0.05
utilisation response, points per unit of the no-capacity real response	0.54	0.20–0.73	0.93	−1.69 (0.32)	−0.60
the 14 capacity-off receiving nodes: real 0.96, shadow price 0.99					

Notes: For each of the 49 capacity-on receiving nodes, the ratio of its response with capacity to its response without, fitted as a slope through the origin across the 62 arriving shocks; the mean and range over the nodes, the median R^2 of the per-node fits, the OLS slope of the ratio on the node's mean utilization across the nodes with its standard error, and the correlation. Productivity shocks are the +1 per cent shocks of Section 6.1; demand shifts are +1 per cent shifts of one node's weight in final demand. The demeaned rows remove, for each shocked node, the mean price response across the receivers before the ratio is fitted, as the year effects of a regression would.

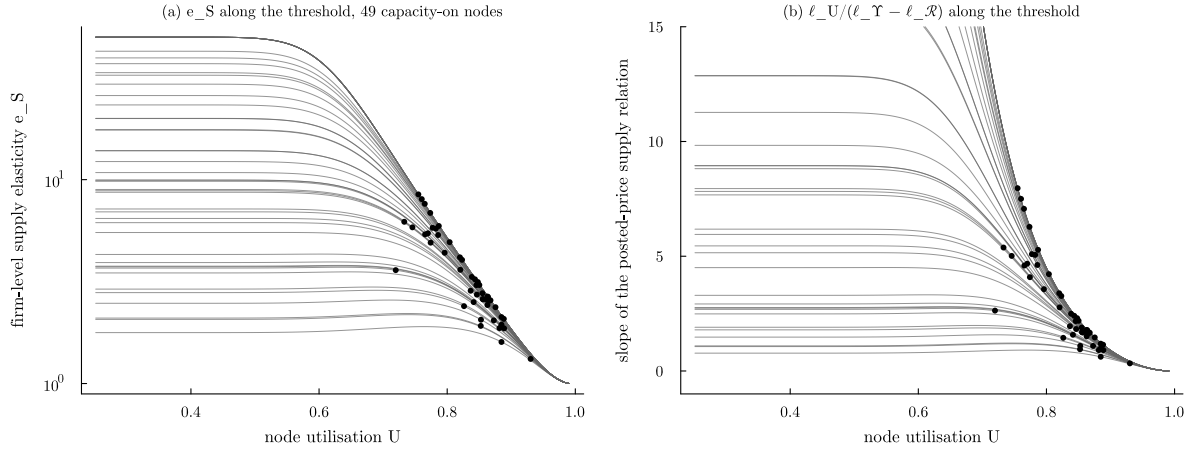
Figure 12: Arriving productivity shocks, by receiving node.



Notes: (a) The ratio of the node's real response with capacity to its response without, against the node's mean utilization; (b) the same for the shadow price index; (c) the real ratio against the node's mean supply elasticity. Circles are the 49 capacity-on nodes, triangles the 14 capacity-off nodes.

Arriving productivity shocks. One ratio describes a node's 62 responses well: the median R^2 is 0.96 for real activity and 0.94 for the price. The real ratio is 0.36 on average over the 49 capacity-on nodes and runs from 0.18 to 0.59; it falls with the node's utilization, with a slope of -1.85 (s.e. 0.20) per unit of utilization and a correlation of -0.80, falls with the node's rationed share (correlation -0.77), and rises with its supply elasticity (correlation 0.69). Pooled over the 3038 pairs of a shocked node and a capacity-on receiver, the regression of the response with capacity on the response without and on that response times the receiver's utilization less its mean gives 0.355 on the first and -1.97 (s.e. 0.04) on the second: the model's own version of the regression's quantity interaction is negative and large. The price ratio is 0.74 on average, from 0.32 to 0.98, and it also falls with utilization (slope -1.33, s.e. 0.42, correlation -0.42). So at the level of a node the price response to an arriving shock is damped too, and more at tight nodes. This does not contradict Section 6.1, whose statement that "the price response is the same" is about the sales-weighted sum of the other nodes' price responses: that sum is 0.999 of its no-capacity value on average over the 63 shocked nodes, because it nets the shocked node's customers, whose prices fall as their inputs get cheaper, against everyone else, whose prices rise in units of the final good. Split by the sign of the no-capacity price response, 18 per cent of the pairs have a negative one, the receiver being a customer of the shocked node; on those pairs, over the 41 receivers with at least three of them, the price ratio is 0.60 on average (0.74 at the median) and falls with utilization (correlation -0.56), and the real ratio is 0.45 (correlation -0.70); on the pairs with a positive price response the price ratio is 0.96 and unrelated to utilization (correlation 0.03), while the real ratio is 0.32 (correlation -0.75). The pass-through of a cost shock into a receiving node's price is damped when the node is tight; the price response to the rest is not. A shock reaching one of the 14 capacity-off nodes is still damped by what it passed through: their real ratio is 0.60 and their price ratio 0.81.

Figure 13: The slope of the supply relation along the threshold.

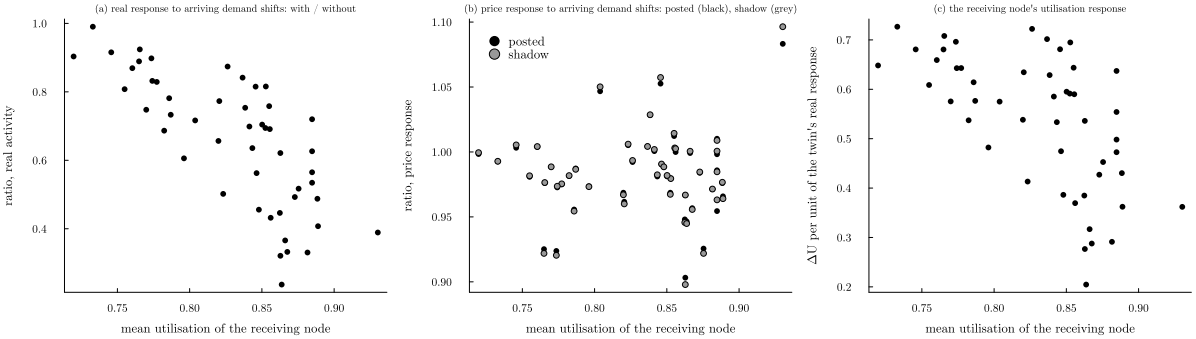


Notes: (a) The firm-level supply elasticity e_S of (111) as a function of the node's utilization, one curve per capacity-on node at the calibrated parameters; (b) the slope of the posted-price form, $\ell_U/(\ell_\Upsilon - \ell_\mathcal{R})$. Dots mark each node's mean utilization on the balanced growth path.

The slope of the supply relation along the threshold. The interaction of (237) presumes that the slope $e_{S,i} - 1$ of (234) falls as the node tightens. Section 6.1 reports this as a regularity across nodes; Lemma 15(iv) gives only the limit $e_S \to 1/\alpha_i$ as the node becomes slack and says nothing about monotonicity along the threshold. Figure 13 computes the firm-level elasticity from the tail functions at the calibrated parameters for each of the 49 capacity-on nodes along the whole range of utilization. The elasticity equals one at the maximum utilization, 0.991 at every node, because the pole depends on ε and σ_v only; it rises as utilization falls and flattens to the no-capacity value $1/\alpha_i$ in the slack region, overshooting it slightly at the five nodes with the highest capital shares, whose maximum sits at a utilization of about 0.65 to 0.74. At each node's mean utilization the firm-level elasticity matches the node mean reported in Section 6.1 to within three per cent (median 2.85, from 1.32 to 8.47), and its derivative in utilization is negative at all 49 nodes. The elasticity is monotone within ± 0.15 of the node's mean at 44 of the 49 nodes; the five exceptions are the high-capital-share nodes whose overshoot lies inside that window. So within a node, where it lives, the slope falls with tightness, and the proportional fall is steep: the log of $e_S - 1$ falls by 13 per unit of utilization at the median node (quartiles 10 and 15), which is 58 per cent per within-node standard deviation of utilization (0.067, Appendix G.4). The posted-price slope $\ell_U/(\ell_\Upsilon - \ell_\mathcal{R})$ is positive at every node near its mean and larger than $e_S - 1$: per unit of utilization the posted price moves less than the shadow index.

Arriving demand shifts. The real ratio is 0.66 on average over the capacity-on nodes, from 0.24 to 0.99 (median R^2 0.93), and 0.96 at the capacity-off nodes; it falls with the receiving node's utilization with a slope of -2.89 (s.e. 0.38), correlation -0.74, and with its rationed share (correlation -0.72). The node's own utilization rises by 0.54 points per unit of its no-capacity real response on average, and by less at the tight nodes (correlation -0.60). So for arriving demand shifts the quantity interaction is negative,

Figure 14: Arriving demand shifts, by receiving node.



Notes: The 49 capacity-on nodes. (a) The ratio of the node's real response with capacity to its response without, against the node's mean utilization; (b) the same for the posted price (black) and the shadow price index (grey); (c) the response of the node's utilization, in points per unit of its no-capacity real response.

as for productivity shocks, and larger. The price is different: the ratio of the shadow price index is 0.98, from 0.90 to 1.10, with a median R^2 of 0.97, and it is unrelated to the node's utilization (slope 0.07, s.e. 0.10, correlation 0.09). The posted price gives the same to two decimals, because the gap between the posted and the shadow price barely moves: the own response of the node's mean log price ratio $\mathbb{E}[\ln \mathcal{R}_i]$ to a one per cent shift of its own weight is -0.004 per unit against 0.10 for the shadow index, and across the other pairs the change in the log price ratio averages 0.0001 in absolute value against 0.009 for the price. The price ratio moves with the receiving node's utilization across pairs (correlation -0.79: more rationing when utilization rises), but by an order of magnitude too little to matter. So the question whether the accounts' price index is the posted or the shadow price is numerically moot at these shock sizes, and the price column of the regression does not depend on it.

The year-effect analogue. A regression with year effects removes a price movement common to all nodes. The analogue on the model's responses demeans, for each shocked node, the price responses across the receivers before the ratio is fitted. For demand shifts the demeaned price ratio averages 1.43 (1.35 at the median) with a wide spread, from 0.23 to 4.31, and no relation to utilization (slope -0.58, s.e. 1.79, correlation -0.05): relative prices respond somewhat more with capacity, and not in a way that tracks the receiver's tightness. For productivity shocks the demeaned ratio is 0.59, from 0.10 to 0.88, and falls with utilization (slope -2.71, s.e. 0.51, correlation -0.61).

What this fixes for the regression. (i) *Quantities.* The response of a receiving node to a network-built shock is damped when the node is tight, for both kinds of shock, in the model's cross-section (correlations -0.74 and -0.80) and in its pooled interaction, and the within-node slope of the supply relation falls with utilization at every node where it lives; so $b_2 < 0$ for the Down column and $c_2 < 0$ for the Up column in quantities. (ii) *Prices, the* Down *column.* No interaction with the receiver's tightness, for the

shadow and the posted price alike, raw or demeaned. (iii) *Prices, the* Up *column.* The pass-through of a supplier's productivity into the receiver's price is damped when the receiver is tight: a negative pass-through whose magnitude is smaller at tight nodes, so a positive interaction coefficient. (iv) The sentence "the price response is the same" of Section 6.1 holds for the sales-weighted aggregate and for the Down channel at the node level; it does not hold for cost pass-through at the node level, which is what a regression sees.

[Two sentences for Section 6.1, proposed here and applied on 28 October 2026]

- *The supply elasticity paragraph.* Section 6.1 calls the node's supply elasticity "the response of its output to its price at given capital" and gives $1/\alpha_n$ without capacity. Lemma 15(iv) defines e_S as the elasticity of *sales* to the price; the elasticity of output, equivalently of utilization, is $e_S - 1$, which is $1/\alpha_n - 1$ without capacity. Numerically close at $\alpha_n \approx 0.08$; the word is wrong. Proposed: "the node's supply elasticity: the response of its sales to its price at given capital, which the model gives in closed form as a function of the threshold (Lemma 15); the response of its output is one less."
- *The price sentence.* After "The price part of that response, the response of their shadow price indices, is the same in the two economies", add: "This is a statement about the sales-weighted sum of the other nodes' price responses, which nets the shocked node's customers, whose prices fall, against everyone else, whose prices rise in units of the final good. At the level of a single receiving node the pass-through of cheaper inputs into its price is damped when the node is tight, by about a quarter on average (Appendix G); the response of a node's price to a shift in demand is not."

[Evidence status of this subsection, and a note on the warnings in the solves] The productivity-shock ratios are a re-reading of the stored response matrices of Section 6.1, which the re-solve reproduced with a maximum absolute difference of zero over the 63×63 entries, so those two matrices are reproduced rather than computed once; the economy without capacity constraints's matrices were not re-run. The supply-slope curves use the tail functions of the model's code, which its verification suite checks against central differences. The demand-shift matrices, the posted-price and utilization responses and the demeaned ratios are computed once. The stationary solves of both economies emitted the solver's usual fallback warnings at one node (the pinned Kolmogorov fast path falling back to the general solver); the same warnings appear in every stored results run of the paper and do not affect the solution. The sign split of the productivity pairs was recomputed from the stored matrices when this appendix was written; the positive-pair numbers differ slightly from a first in-session reading (price ratio 0.96 and correlation 0.03 here, against 1.03 and 0.22 then), the negative-pair numbers are identical.

G.4 The Data

The panel. Nodes $i = 1, \ldots, 63$ and years $t = 1998, \ldots, 2019$, the years of the measured shocks of Appendix E.6. The outcomes are, from the accounts, the growth of the node's real gross output $\Delta \ln Y_{it}$ and of its gross-output price index $\Delta \ln P_{it}$, and, for the 24 nodes that have a Federal Reserve utilization series, the change in the node's utilization rate $\Delta U_{it} = U_{it} - U_{i,t-1}$ in points, utilization as a fraction. The 24 are the 23 core nodes of Appendix E.2 and publishing, whose Federal Reserve series exists although its calibration target is the European transfer.

The shocks. The own demand shift $\Delta \ln \beta_{it}$ is the series of Appendix E.6: the change in the node's weight in final demand,

$$\Delta \ln \beta_{it} = [\Delta \ln b_{it} - (1-\eta)\Delta \ln P_{it}] - \sum_j w_{jt} \left[\Delta \ln b_{jt} - (1-\eta)\Delta \ln P_{jt}\right], \tag{238}$$

with b_{it} the node's share in nominal final expenditure, $\eta = 0.5$ (the values 0.3 and 0.9 are variants) and w_{jt} the nodes' mean final shares of the two years, so that the share-weighted mean of the shifts is zero in every year: the series measures shifts of final demand *between* nodes, and a fall of all final demand together is removed. The arriving demand shift, the Down column, is the sum of the customers' shifts weighted by the share of the node's sales that goes to each,

$$\text{Down}_{it} = \sum_{k \neq i} \bar{w}^D_{ki} \, \Delta \ln \beta_{kt}, \qquad \bar{w}^D_{ki} = \frac{1}{13} \sum_{s=2007}^{2019} \frac{\text{sales of } i \text{ to } k \text{ in } s}{\text{output of } i \text{ in } s}, \tag{239}$$

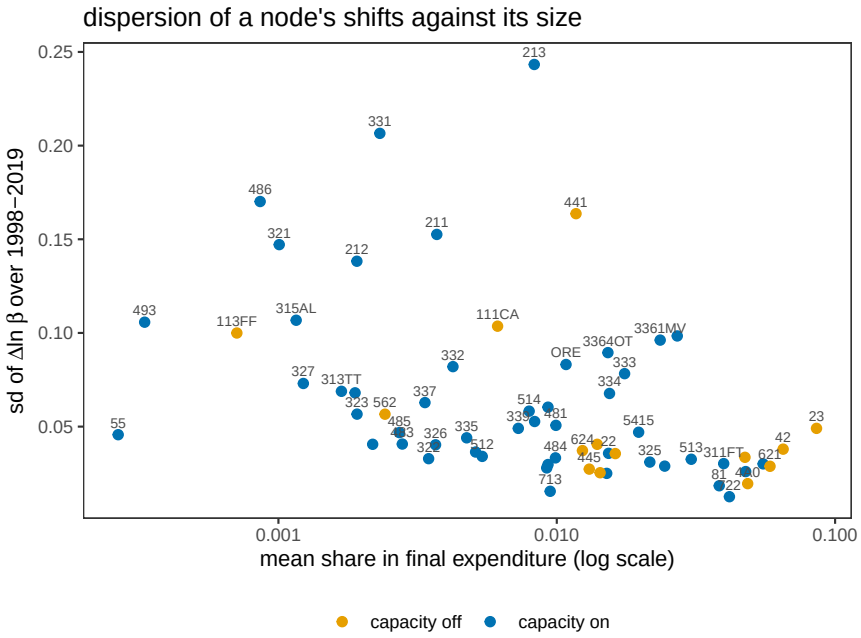
from the domestic input-output flows, the government and housing industries removed and the three finance industries merged as in Appendix E.2, own purchases excluded; a variant uses each year's own shares. The arriving productivity shock, the Up column, is the suppliers' productivity growth of Appendix E.6 weighted by cost shares, $\text{Up}_{it} = \sum_{j \neq i} a_{ij} \, \Delta z_{jt}$, with a_{ij} the cost shares of the calibration; a variant uses the node residual, the common component and the node means removed. The weights of the two columns are the model's (236): sales shares downstream, cost shares upstream.

The modifiers. Tightness takes two forms. The *within* form, $\tau_{it} = U_{i,t-1} - \bar{U}_i$, is the node's lagged Federal Reserve utilization less its mean over the sample years, the form of Boehm and Pandalai-Nayar (2022); it exists for the 24 nodes. The *cross-sectional* form, $\tau_i = U_i - \bar{U}$, is the node's mean utilization of the calibration (its target of Appendix E) less the mean over the sample nodes; it exists for the 49 capacity-on nodes, and it is the form that Section 6.1 and Appendix G.3 speak about, since the model's nodes differ in their mean utilization and not over time. The customers' tightness τ^D_{it} is the sales-weighted lagged utilization of the customers that have a series, less its node mean, with the covered share of the node's sales recorded beside it. One standard deviation of the within modifier is 0.067 and of the cross-sectional one 0.049; utilization being a fraction, an interaction coefficient is per unit of utilization.

The demand-shift series, looked at before it is used. The series was built for the validation of the time series and had not been used as a regressor. Its dispersion is what the construction expected: a standard deviation over the years of 0.047 at the median node, 0.069 at the median manufacturing node and 0.073 at the median core node, up to 0.24 at support activities for mining; the cross-node dispersion by year runs between 0.06 and 0.13 and peaks in 2001, 2009–2010 and 2015–2016. The manufacturing nodes together lose 6.2 per cent of their weight in 2009 and regain 3.9 per cent in 2010; the core nodes lose 6.9 and 6.8 per cent in 2015 and 2016, the collapse of the oil price. The price term and η do not matter: the series is 0.985 correlated at the median node with the change of the final share alone, and the versions at $\eta = 0.3$ and 0.9 are 0.992 and 0.974 correlated with it. Four features matter for a regression. First, the outliers are the small nodes (Figure 15): 53 of the 1386 node-years exceed 0.20 in absolute value, in 17 nodes whose mean final share is 0.6 per cent against 1.6 per cent over all node-years; for an upstream node the final share is a thin residual of exports and inventories, and its yearly change is noisy. Second, 21 nodes carry a trend of more than 0.02 a year (apparel -0.11, data processing and internet publishing $+0.08$, printing -0.05, oil and gas $+0.05$): the secular reallocation of final demand, not shocks; a node effect absorbs the mean, and the remaining persistence is moderate (autocorrelation 0.05 at the median node, above 0.3 at 16 nodes). Third, the first principal component of the 63 series carries 28 per cent of their variance and the second 12; the mean pairwise correlation is zero by construction and 0.10 within manufacturing, with the durable-goods nodes moving together (machinery, electrical equipment, primary metals and fabricated metals are correlated 0.74 to 0.86 with the manufacturing aggregate); year effects take out what is common, the durables' cycle stays in the cross-section. Fourth, the arriving shift is small and concentrated (Figure 16): its standard deviation is 0.0046 at the median node and 0.0071 at the median core node, about an eighth of the own shift's, because the weights sum to 0.36 of sales at the median core node and the customers' shifts partly cancel. It is sizeable where the customer base is concentrated, oil and gas extraction (0.058, mostly through refining), primary metals (0.035), mining (0.026), nonmetallic minerals and wood, and negligible for the nodes that sell to final demand (other transportation equipment 0.001, motor vehicles 0.002). Fixed and year-by-year weights give the same series (median correlation 0.96). The coverage of the customers' utilization is thin: the part of a core node's sales that goes to core customers is 0.13 at the median, from 0.008 to 0.71 (oil and gas, whose customer is refining), so the customers'-tightness interaction b_3 can be built for a handful of upstream nodes only. Finally, the shifts move utilization: over the 24 nodes, with node and year effects, a shift of one log point in a node's own weight raises its utilization by 0.23 points (s.e. 0.015) in the same year and by nothing the year after, and the arriving shift by 0.13 points (s.e. 0.09). The own shift is close to the node's own final-sales growth, so part of that link is the same output movement seen twice; it is the own-shock control, not a clean shock.

[The demand-shift series as a regressor: what the look settles and what it leaves to decide] The series was built under a working rule that asked for it to be looked at

Figure 15: The dispersion of a node's demand shifts against its size.



Notes: The standard deviation over 1998–2019 of $\Delta \ln \beta_{it}$ against the node's mean share in final expenditure, log scale, 63 nodes; capacity-on nodes in blue, capacity-off in orange.

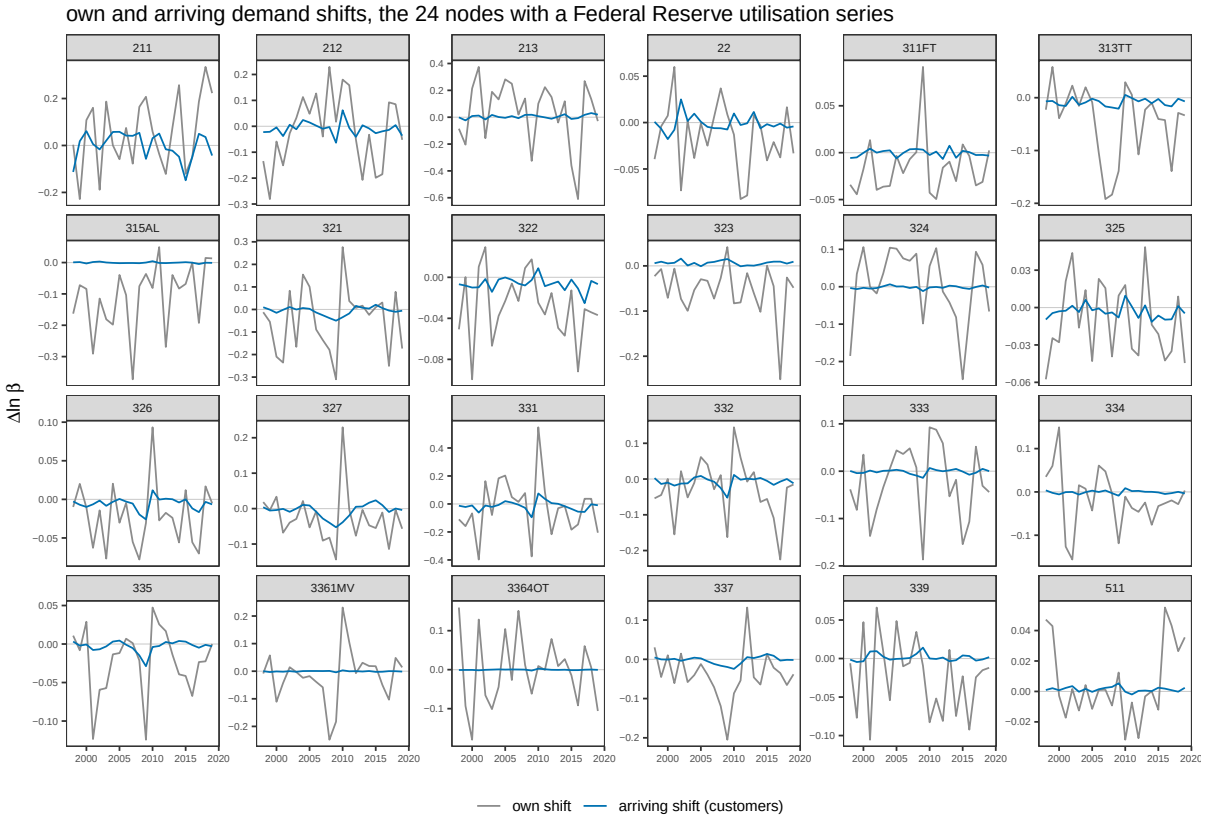
before use. The look settles that it behaves as a shock series should (dispersion, the 2009 and 2015–2016 episodes, independence from η and from the price term) and that two features must be handled: the outliers of the small nodes (winsorising at ± 0.20, or weighting by the final share) and the trends (node effects, with a check without the trending nodes). Both handlings are used in Appendix G.5 and the results survive them. **Recommendation:** confirm the series as the demand shock of the validation, with those two handlings written into its description. **Alternatives:** (a) rebuild the shifts from final *uses* by category (consumption, investment, exports, government) rather than from the share in total final expenditure, which would separate exports and inventories from domestic demand and tame the small-node outliers at the cost of a larger build; (b) use only the manufacturing and core nodes' shifts, since the small final shares that produce the outliers belong mostly to upstream non-manufacturing nodes. Neither has been tried.

G.5 Estimates

The specification. For each outcome $y \in \{\Delta \ln Y, \Delta \ln P, \Delta U\}$ the regression is (237) without the customers' term,

$$y_{it} = \alpha_i + \lambda_t + b_1 \,\mathrm{Down}_{it} + b_2 \,\mathrm{Down}_{it}\,\tau_{it} + c_1 \,\mathrm{Up}_{it} + c_2 \,\mathrm{Up}_{it}\,\tau_{it} + d\,\tau_{it} + e\,\Delta \ln \beta_{it} + \varepsilon_{it}, \tag{240}$$

Figure 16: Own and arriving demand shifts, 1998–2019.



Notes: The 24 nodes with a Federal Reserve utilization series. Grey: the node's own shift $\Delta \ln \beta_{it}$; blue: the shift arriving from its customers, Down_{it}. The vertical scales differ by panel.

with node and year effects and standard errors clustered by node, by OLS; in the cross-sectional form the modifier is τ_i and its level term is absorbed by the node effect. The within form runs on the 24 nodes with a series (528 node-years), the cross-sectional form on the 49 capacity-on nodes (1078 node-years) and, as a check, on the 24. The customers' term b_3 is added in a separate regression. The shares used as weights were checked to sum to one with the own purchases (on average 0.94 of a node's sales go to final demand and to the other nodes, the rest being own purchases), the lag of utilization to be aligned, and the sample sizes to be the same across the outcomes of a form. A split-sample version drops the interactions and fits b_1, c_1 and e separately on the node-years with $\tau_{it} \leq 0$ (slack) and $\tau_{it} > 0$ (tight). Table 19 reports the main regressions, Table 20 the halves and the pass-through implied by the interaction, Table 21 and Figure 18 seventeen one-at-a-time variants, and Table 22 where the Down column's variation sits.

[Boehm and Pandalai-Nayar's estimating equation in this appendix's notation, beside (240)**]** Equation numbers are those of their 2020 working paper; their notation is translated as follows: their industry utilization $u_{i,t}$ is U_{it}, their industry mean $\bar{u}_i$ is $\bar{U}_i$, so their modifier $u_{i,t-1} - \bar{u}_i$ is exactly the within-form tightness τ_{it} of this appendix; their price index $P^Y_{i,t}$ is P_{it}; their output $Y_{i,t}$ is Y_{it} and their capacity $Q_{i,t}$ is $\overline{Y}_{it}$, so $\Delta \ln Y_{it} - \Delta \ln \overline{Y}_{it} = \Delta \ln U_{it}$; their log average markup $M(\ln u)$ is, up to a constant, $-\ln \Upsilon_i(\overline{v}_i(U_i))$, and their marginal cost $mc_{i,t}$ is the cost term $\mathcal{P}_i/(Z_i x_i^{\alpha_i})$.

- *Their supply curve (11) and its linearisation (12).* In levels, $\ln P_{it} = M(\ln U_{it}) + \ln mc_{it}$, which is the pricing relation (107) with the threshold inverted into utilization. Linearised around the previous year,

$$\Delta \ln P_{it} = M'(\ln U_{i,t-1})\, \Delta \ln U_{it} + \Delta \ln mc_{it},$$
$$M'(\ln U) = \frac{\ell_{Y,i}}{\ell_{U,i}} = \frac{1}{e_{S,i} - 1}, \qquad \Delta \ln mc_{it} = \sum_j a_{ij}\, \Delta \ln P_{jt} - \Delta \ln Z_{it},$$

  which is (234) solved for the price. Their M' is the inverse of this appendix's slope, evaluated at the previous year's utilization; the network sits inside their marginal-cost term.

- *Their expansion of the slope.* $M'(\ln U_{i,t-1}) \approx M'(\ln \bar{U}_i) + \frac{M''(\ln \bar{U}_i)}{\bar{U}_i}\, \tau_{it}$: the slope at the industry's mean plus its derivative times the within-industry deviation of lagged utilization. In the model the two coefficients are $1/(e_{S,i} - 1)$ at $\bar{U}_i$ and the derivative of $1/(e_{S,i} - 1)$ in U at $\bar{U}_i$, which is positive because e_S falls with utilization (Figure 13).

- *Their estimating equation (13).* Substituting, with a constant, industry and year effects, and the error,

$$\begin{aligned}\Delta \ln P_{it} = {} & \alpha + \beta_Y\, \Delta \ln Y_{it} + \beta_{Yu}\, \Delta \ln Y_{it}\, \tau_{it} + \beta_u\, \tau_{it} \\ & + \beta_Q\, \Delta \ln \overline{Y}_{it} + \beta_{Qu}\, \Delta \ln \overline{Y}_{it}\, \tau_{it} + \beta_{mc}\, \Delta \ln \mathrm{UVC}_{it} + \varepsilon_{it},\end{aligned}$$

  with the model's values $\beta_Y = 1/(e_{S,i} - 1)$ at $\bar{U}_i$, $\beta_{Yu} = \mathrm{d}[1/(e_{S,i} - 1)]/\mathrm{d}U$ at $\bar{U}_i$, $\beta_Q = -\beta_Y$, $\beta_{Qu} = -\beta_{Yu}$ and $\beta_{mc} = 1$; marginal cost is proxied by unit

variable cost (materials, energy and production-worker pay), whose coefficient is estimated freely ($\widehat{\tau} = 0.89$, Appendix E), and the level term $\beta_u \tau_{it}$ is included so that the interaction has its usual reading. The object of interest is $\beta_{Yu} > 0$: a convex supply curve. The four bins of Table 13, the target of σ_v, are the same equation with β_Y estimated separately in quantile bins of τ_{it} instead of the linear interaction.

- *Their instruments (15)–(17).* $\Delta \ln Y_{it}$ and its interaction are instrumented by observed shifts of the industry's demand and their interactions with τ_{it}: the growth of the materials purchases of each downstream industry, of consumption, investment, government purchases and exports by destination, weighted by the previous year's share of i's sales to each (the downstream part is the instrument of Shea (2002)), the matching buyers' price indices, and exchange rates. In this appendix's notation the downstream part is $\sum_k s^D_{ki} \Delta \ln X_{kt}$ and the final-demand part is the $s^F_i[\cdots]$ term, the two pieces of the demand identity (236).
- *How* (240) *differs.* (i) *What is on the left.* Theirs is the structural inverse supply curve, price on output, with output instrumented; (240) is the reduced form, the outcome (output growth, price growth or the change in utilization) on the demand shifter itself. The quantity column is the analogue of their first stage with the interaction, the price column of their reduced form; the ratio of the two would be their β_Y. (ii) *The modifier.* The within form uses exactly their τ_{it}; the cross-sectional form (the calibration's mean utilization) has no counterpart in their design, where industry effects absorb the level. (iii) *The shifter.* Theirs is the sum of every demand component, with the customers' total materials purchases as the downstream piece; Down_{it} is one component, the customers' shifts of final-demand weight $\Delta \ln \beta_{kt}$, weighted by the same sales shares, and the own shift $\Delta \ln \beta_{it}$ is a control rather than part of the shifter. A customer's materials purchases respond to its own tightness (the $\ell_{\xi\varrho,k}$ term of Appendix G.2), which is why this appendix uses the customers' final-demand shifts rather than their spending. (iv) *Capacity.* They hold capacity growth fixed, so their regressor is in effect $\Delta \ln U_{it}$; (240) has no capacity control in the output column and uses the change in utilization as a separate outcome. (v) *Costs.* They control for unit variable cost; (240) carries the cost side through the Up column (the suppliers' productivity with cost-share weights, the $-\Delta \ln Z$ part of marginal cost at the suppliers) and the year effects, and has no own-cost control, since own productivity is a Solow residual contaminated by utilization. (vi) *Sample and estimator.* Theirs 21 three-digit manufacturing industries, 1973–2011, instrumented; this appendix 24 nodes, 1998–2019, OLS with the shifter as its own regressor. Both have industry and year effects. (vii) *The customers' tightness* τ^D_{it} has no counterpart in their equation.

Quantities: no average pass-through, a negative interaction. In the within form the arriving demand shift has no average pass-through into real output growth ($b_1 = -0.15$, s.e. 0.22) and a negative interaction with the node's tightness ($b_2 = -5.07$, s.e.

Table 19: The regression (240) on the node panel, 1998–2019.

	within form, 24 nodes			cross-sectional form, 49 nodes	
	real output growth	price growth	change in utilisation	real output growth	price growth
arriving demand shift (customers), Down_{it}	−0.15 (0.22)	2.51 (0.72)	0.07 (0.12)	0.41 (0.28)	1.00 (0.39)
$\times$ the node's tightness	−5.07 (2.97)	9.05 (2.84)	−2.80 (1.23)	−5.10 (2.85)	19.15 (6.34)
arriving productivity shock (suppliers), Up_{it}	0.35 (0.21)	−1.36 (0.50)	0.11 (0.08)	0.24 (0.21)	−0.98 (0.47)
$\times$ the node's tightness	−2.13 (1.36)	4.89 (2.39)	−1.13 (1.18)	1.96 (3.44)	−1.97 (4.48)
own demand shift, $\Delta \ln \beta_{it}$	0.56 (0.12)	0.05 (0.06)	0.21 (0.05)	0.46 (0.11)	0.12 (0.04)
the node's tightness (lagged utilisation less its mean)	−0.04 (0.03)	0.08 (0.04)	−0.22 (0.03)	—	—
with the customers' tightness added (within form, a separate regression)					
$\text{Down}_{it} \times$ the customers' tightness	3.44 (2.86)	−0.97 (2.08)	—	—	—
$\text{Down}_{it} \times$ the node's tightness, in that regression	−6.44 (2.92)	9.44 (3.04)	—	—	—
node-years	528	528	528	1078	1078
nodes	24	24	24	49	49
within R^2	0.59	0.42	0.46	0.48	0.39

Notes: The within form uses the node's lagged Federal Reserve utilization less its mean as the modifier (24 nodes); the cross-sectional form uses the node's mean utilization of the calibration less the mean over the nodes (49 capacity-on nodes). Node and year effects; standard errors clustered by node in parentheses; utilization as a fraction, so an interaction is per unit of utilization. The model's prediction: in quantities and utilization a positive pass-through of the arriving demand shift and a negative interaction with tightness; in prices a positive pass-through of the demand shift with no interaction, and a negative pass-through of the suppliers' productivity whose magnitude is smaller at tight nodes.

2.97); the own shift passes through at 0.56 (s.e. 0.12). In the cross-sectional form on the 49 nodes $b_1 = 0.41$ (s.e. 0.28) and $b_2 = -5.10$ (s.e. 2.85); with the customers' shifts winsorised at ± 0.20, which removes the small-node outliers, $b_1 = 0.63$ (s.e. 0.18) and $b_2 = -8.41$ (s.e. 2.13), and in the within form $b_2 = -9.19$ (s.e. 2.69). The interaction keeps its sign in every variant but one, the sample without the trending receivers, where nine nodes remain and nothing is estimated; it is tighter without oil and gas (-7.00, s.e. 3.05) and on the manufacturing receivers alone (-4.89, s.e. 2.22). The split halves say the same without a functional form: the pass-through is 0.04 (s.e. 0.32) in slack node-years and -0.25 (s.e. 0.43) in tight ones. In the cross-section the implied pass-through runs from 0.66 one standard deviation slacker than the mean node to 0.16 one standard deviation tighter, a fall of 61 per cent per standard deviation; in the within form from $+0.18$ to -0.49, around a mean of zero. Figure 17 shows the two halves as binned scatters of the residualised outcomes.

utilization. The arriving shift raises the node's utilization when the node is slack and not when it is tight: $b_1 = 0.07$ (s.e. 0.12) with $b_2 = -2.80$ (s.e. 1.23), and the halves give 0.21 points per log point (s.e. 0.18) in slack node-years against 0.03 (s.e. 0.19) in tight ones; the own shift gives 0.24 and 0.18. Lagged utilization itself mean-reverts ($d = -0.22$, s.e. 0.03).

Prices, the Down **column: a strong pass-through, and an interaction that is one node.** The headline estimates are a pass-through of 2.51 (s.e. 0.72) in the within form and 1.00 (s.e. 0.39) across, with interactions of 9.05 (s.e. 2.84) and 19.15 (s.e. 6.34), and both are oil and gas extraction. Table 22 shows why: it has the largest arriving shift of the 24 nodes (standard deviation 0.058; 0.71 of its sales go to refining), it is the tightest node (mean utilization 0.94), and its arriving shift is 0.84 correlated with its own price growth, because refining's shift of final demand and the price of crude both follow the oil price. Without that node the pass-through is 0.68 (s.e. 0.23) within and 0.44 (s.e. 0.11) across, and the interaction is 1.38 (s.e. 1.32) and -1.06 (s.e. 2.50): zero. The manufacturing receivers alone give 0.80 (s.e. 0.32) and 1.85 (s.e. 1.72). So on the nodes that are not oil the price pass-through of an arriving demand shift is 0.4 to 0.8 and does not vary with tightness, which is the model's prediction (ii); the headline interaction is one node's oil-price story.

The Up **column (secondary: the productivity shocks are Solow residuals).** A supplier's productivity growth lowers the receiver's price, $c_1 = -1.36$ (s.e. 0.50) within and -0.98 (s.e. 0.47) across (-1.38, s.e. 0.15, on the manufacturing receivers); the interaction is $+4.89$ (s.e. 2.39) within, the damping of prediction (iii), but -1.97 (s.e. 4.48) across and -10.54 (s.e. 3.96) in the cross-section on the 24 nodes: not robust across forms. In quantities the pass-through is 0.35 (s.e. 0.21) within and 0.24 (s.e. 0.21) across, with interactions of -2.13 (s.e. 1.36) and $+1.96$ (s.e. 3.44): nothing. Using the node residual of productivity in place of the series with its common component changes nothing.

The customers' tightness. The thinly covered term gives $b_3 = 3.44$ (s.e. 2.86) in quantities and -0.97 (s.e. 2.08) in prices: nothing, as the coverage of Appendix G.4 said to expect.

Table 20: The two halves of the within sample and the pass-through implied by the interaction.

the two halves of the within sample, interactions dropped (254 slack and 272 tight node-years)	slack half	tight half
real output growth		
arriving demand shift	0.04 (0.32)	−0.25 (0.43)
arriving productivity shock	0.43 (0.27)	0.22 (0.26)
own demand shift	0.61 (0.13)	0.56 (0.13)
price growth		
arriving demand shift	2.23 (1.09)	2.09 (0.58)
arriving productivity shock	−1.69 (0.44)	−1.07 (0.48)
own demand shift	0.02 (0.08)	0.12 (0.06)
change in utilisation		
arriving demand shift	0.21 (0.18)	0.03 (0.19)
arriving productivity shock	0.13 (0.11)	0.09 (0.09)
own demand shift	0.24 (0.04)	0.18 (0.06)

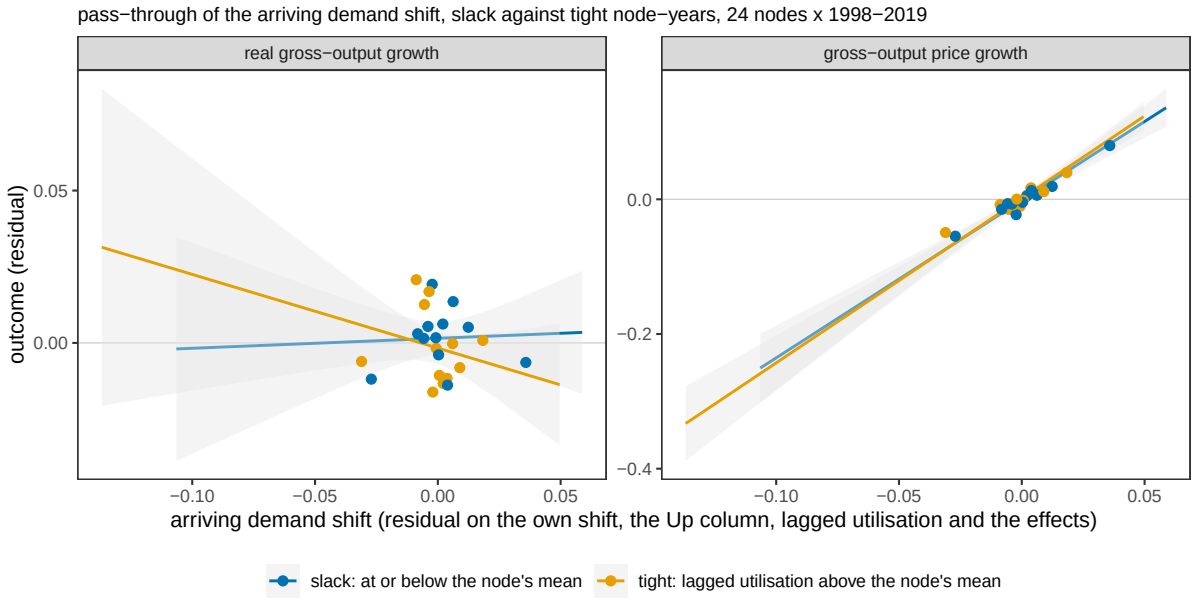
the implied pass-through of the arriving demand shift, $b_1 \pm b_2$ sd *(sd 0.067 within, 0.049 across)*	one sd slacker	at the mean	one sd tighter
within, 24 nodes, real output growth	0.18	−0.15	−0.49
within, 24 nodes, price growth	1.90	2.51	3.11
cross-section, 49 nodes, real output growth	0.66	0.41	0.16
cross-section, 49 nodes, price growth	0.07	1.00	1.94

Notes: Top: the regression (240) without its interactions, fitted separately on the slack and the tight node-years; standard errors clustered by node in parentheses. Bottom: the pass-through of the arriving demand shift one standard deviation of the modifier below the mean, at the mean and one above, from the main estimates of Table 19.

Robustness. The quantity interaction stays between -3.8 and -9.2 in the within form and between -1.9 and -9.0 in the cross-sectional form under the two other values of η, the previous year's sales shares, the shifts demeaned by node (which takes the trends out), dropping 2009, Domar weights, the Down column alone and the own shift interacted too. Dropping the own shift doubles the cross-sectional interaction (-15.3, s.e. 4.1) and raises the pass-through to 1.27: the own shift is 0.30 correlated with the arriving one and must stay. Two-way clustering by node and year returns a within standard error of 0.58 for the quantity interaction, which 22 year clusters cannot support; the node-clustered 2.97 is the one to read.

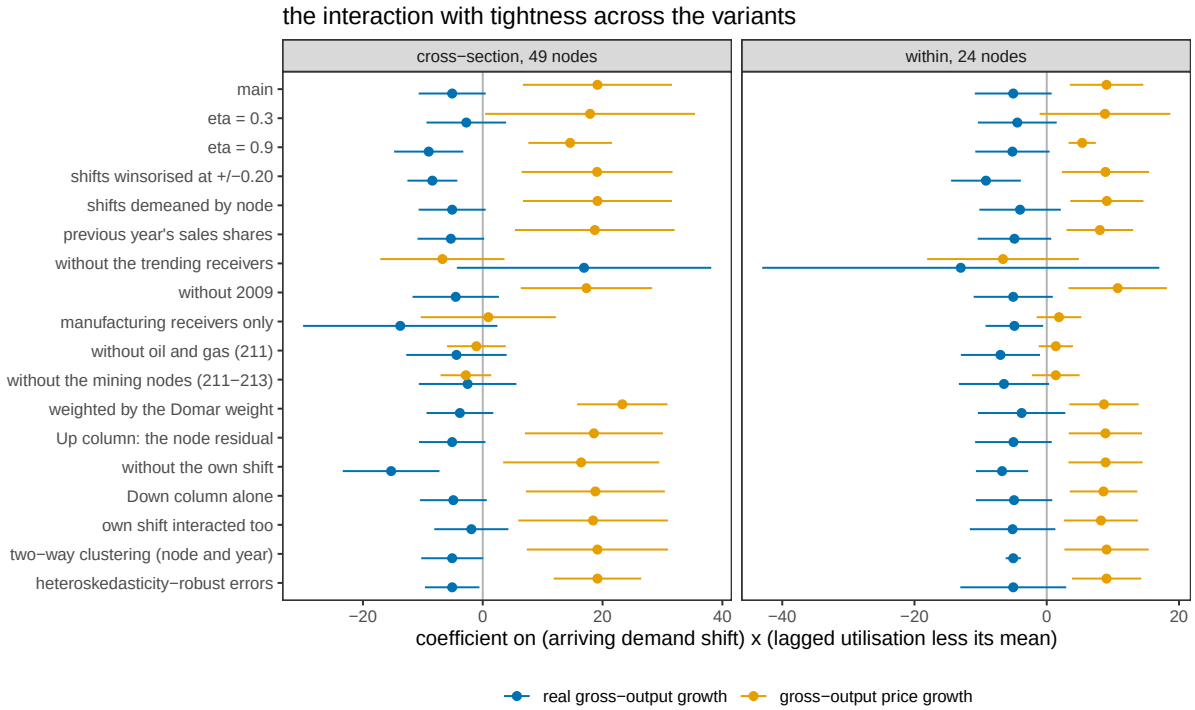
Where the identification sits. Five upstream nodes, oil and gas extraction, primary metals, mining, nonmetallic minerals and wood products, carry 87 per cent of the variance of the arriving shift over the 24 nodes. Fifteen of the 24 own shifts trend; the demeaned variant, which removes the trends, gives the main numbers to within a

Figure 17: The pass-through of the arriving demand shift, slack against tight node-years.



Notes: 24 nodes, 1998–2019. The outcome and the arriving shift are both residualised on the own shift, the Up column, lagged utilization and the two effects, then shown in twelve quantile bins per half with the OLS line of each half and its 95 per cent band.

Figure 18: The interaction across the variants of the regression.



Notes: The interaction of the arriving demand shift with the node's tightness across the variants of Table 21, with 95 per cent bands: real output growth (blue) and price growth (orange), cross-sectional form on the 49 nodes (left) and within form on the 24 (right).

Table 21: Seventeen one-at-a-time variants of the regression (240).

	real output growth				price growth	
	within, 24 nodes		cross-section, 49 nodes		within	cross-section
	Down	× tightness	Down	× tightness	× tightness	× tightness
main specification	−0.15 (0.22)	−5.07 (2.97)	0.41 (0.28)	−5.10 (2.85)	9.05 (2.84)	19.15 (6.34)
$\eta = 0.3$	0.03 (0.39)	−4.45 (3.04)	0.37 (0.28)	−2.74 (3.39)	8.81 (5.04)	17.92 (8.94)
$\eta = 0.9$	−0.44 (0.14)	−5.19 (2.87)	0.53 (0.27)	−9.02 (2.95)	5.36 (1.05)	14.59 (3.56)
shifts winsorised at ± 0.20	−0.03 (0.14)	−9.19 (2.69)	0.63 (0.18)	−8.41 (2.13)	8.88 (3.36)	19.09 (6.43)
shifts demeaned by node	−0.14 (0.21)	−4.03 (3.14)	0.41 (0.28)	−5.10 (2.85)	9.09 (2.82)	19.15 (6.34)
previous year's sales shares	−0.12 (0.23)	−4.88 (2.84)	0.38 (0.27)	−5.32 (2.84)	8.04 (2.56)	18.71 (6.80)
without the trending receivers (198, 638 node-years)	−1.89 (0.88)	−13.00 (15.32)	−0.67 (0.37)	16.91 (10.83)	−6.62 (5.85)	−6.73 (5.29)
without 2009 (504, 1029 node-years)	−0.11 (0.26)	−5.07 (3.05)	0.41 (0.32)	−4.51 (3.68)	10.73 (3.80)	17.31 (5.59)
manufacturing receivers only (418, 418 node-years)	0.12 (0.43)	−4.89 (2.22)	−0.35 (0.74)	−13.77 (8.28)	1.85 (1.72)	0.94 (5.75)
without oil and gas (211) (506, 1056 node-years)	−0.52 (0.36)	−7.00 (3.05)	0.39 (0.32)	−4.39 (4.27)	1.38 (1.32)	−1.06 (2.50)
without the mining nodes (211–213) (462, 1012 node-years)	−0.16 (0.49)	−6.47 (3.48)	0.57 (0.25)	−2.53 (4.15)	1.37 (1.84)	−2.82 (2.15)
weighted by the Domar weight	−0.21 (0.26)	−3.80 (3.37)	0.14 (0.29)	−3.81 (2.84)	8.66 (2.67)	23.31 (3.85)
Up column: the node residual	−0.16 (0.22)	−5.03 (2.96)	0.41 (0.28)	−5.11 (2.83)	8.87 (2.83)	18.56 (5.87)
without the own shift	0.61 (0.35)	−6.76 (2.02)	1.27 (0.27)	−15.29 (4.12)	8.90 (2.86)	16.43 (6.64)
Down column alone	−0.10 (0.23)	−4.94 (2.95)	0.43 (0.28)	−4.91 (2.85)	8.59 (2.60)	18.81 (5.92)
own shift interacted too	−0.16 (0.22)	−5.17 (3.30)	0.28 (0.26)	−1.90 (3.16)	8.20 (2.86)	18.42 (6.37)
two-way clustering (node and year)	−0.15 (0.22)	−5.07 (0.58)	0.41 (0.27)	−5.10 (2.63)	9.05 (3.25)	19.15 (6.02)
heteroskedasticity-robust errors	−0.15 (0.20)	−5.07 (4.08)	0.41 (0.20)	−5.10 (2.32)	9.05 (2.67)	19.15 (3.73)

Notes: The pass-through of the arriving demand shift and its interaction with tightness in real output growth, both forms, and the interaction in price growth. Standard errors clustered by node in parentheses, except in the two rows that change the standard errors. "Without the trending receivers" keeps the nine nodes whose own shift does not trend. The sample sizes are the main ones (528 and 1078 node-years) unless stated.

Table 22: Where the Down column's variation sits.

		mean utilis.	cover-age	sd of the arriving shift	sd of the own shift	corr. with output growth	corr. with price growth	own shift trends
211	Oil and gas extraction	0.94	0.71	0.058	0.15	0.19	0.84	yes
331	Primary metals	0.76	0.64	0.035	0.21	0.84	0.62	yes
212	Mining, except oil and gas	0.81	0.42	0.026	0.14	0.47	0.52	
327	Nonmetallic mineral products	0.68	0.19	0.019	0.07	0.70	0.32	yes
321	Wood products	0.73	0.17	0.018	0.15	0.77	0.27	yes
213	Support activities for mining	0.72	0.10	0.014	0.24	0.40	0.06	yes
332	Fabricated metal products	0.80	0.40	0.014	0.08	0.88	0.22	yes
337	Furniture and related products	0.75	0.03	0.009	0.06	0.70	−0.16	yes
22	Utilities	0.84	0.14	0.009	0.04	−0.55	−0.19	
326	Plastics and rubber products	0.79	0.34	0.008	0.04	0.80	0.14	yes
335	Electrical equipment and appliances	0.80	0.13	0.007	0.04	0.81	−0.01	yes
313TT	Textiles	0.73	0.27	0.007	0.07	0.72	0.19	yes
322	Paper products	0.84	0.32	0.006	0.03	0.14	0.03	yes
325	Chemical products	0.74	0.19	0.005	0.03	0.49	0.46	
333	Machinery	0.75	0.10	0.005	0.08	0.60	−0.16	yes
339	Miscellaneous manufacturing	0.76	0.03	0.005	0.05	−0.38	−0.03	yes
323	Printing	0.73	0.11	0.005	0.06	−0.34	−0.38	yes
311FT	Food, beverage and tobacco	0.78	0.01	0.004	0.03	0.01	0.16	
324	Petroleum and coal products	0.87	0.09	0.004	0.10	0.15	0.55	
334	Computer and electronic products	0.74	0.11	0.004	0.07	0.17	0.04	
315AL	Apparel and leather	0.71	0.03	0.002	0.11	0.02	−0.34	yes
3361MV	Motor vehicles and parts	0.74	0.04	0.002	0.10	0.33	0.15	
511	Publishing (except internet)	0.74	0.01	0.002	0.03	−0.38	0.20	
3364OT	Other transportation equipment	0.75	0.01	0.001	0.09	0.18	0.26	

Notes: The 24 nodes with a Federal Reserve utilization series, ordered by the standard deviation of their arriving demand shift. Mean utilization is the Federal Reserve mean over 1997–2019; coverage is the share of the node's sales that goes to customers with a utilization series; the correlations are over the years with the node's own real output growth and price growth; the last column marks the 15 nodes whose own shift trends by more than 0.02 a year.

standard error. A node's price and its customers' shifts can share a common cause, and the oil price is the case found. Instruments are not used, the interactions use lagged utilization, and the claim is that the pattern the model predicts is in the data, not that the model is tested.

[Caveats on the estimates, and what each would take to remove]

- *Clusters.* The within form has 24 clusters, and the standard errors are the usual small-sample-corrected cluster ones; a wild cluster bootstrap would be the check. The cross-sectional form on the 49 nodes has 49.
- *The price of the accounts.* The data's price is the gross-output price index, read as the shadow index in Appendix E; Appendix G.3 shows that the posted and the shadow price respond alike at these shock sizes, so the reading does not matter here.
- *The cross-sectional modifier* is the node's mean utilization, which the calibration fits with x_i; Appendix G.6 shows that in the model a cross-sectional interaction of the same sign appears without capacity, so the within form is the one that speaks for capacity.
- *The oil node.* Two handlings are in the tables: dropping it, and the manufacturing receivers. A third, a control for the oil price (its growth, or its growth interacted with the node's energy cost share), was not run because no oil price series is in the data set; it would need one shipped from the data library. A fourth is to instrument the arriving shift (next box).
- *The finer panel.* The Federal Reserve publishes utilization for three- and four-digit manufacturing from 1972, which with the manufacturing productivity database and the detailed input-output table would give a panel several times larger; it is a build from scratch of about a week and was not started.

G.6 The Same Regression on the Model's Responses

The predictions of Appendix G.3 were read off ratios, the capacity economy's response over the economy without capacity constraints's. The regression is not the ratio: it compares the responses of tight and slack nodes in absolute terms, and a tight node may differ from a slack one in its capital intensity and its position as well as in its tightness. So the question is what the data regression would return on model-generated responses, in the capacity economy and in the economy without capacity constraints; the economy without capacity constraints's coefficients say what node heterogeneity alone gives, and the difference is capacity's part. For every pair of a shocked node k and a receiving node $i \neq k$, the model's response per unit shock, to a demand shift at k and to a productivity shock at k, in real activity, in the shadow price and, in the capacity economy, in utilization, is regressed on the arriving shift-share built with the data's weights and on its interaction with the node's mean utilization less the sample mean,

$$y_{ik} = \alpha_i + \lambda_k + b_1\, x_{ik} + b_2\, x_{ik}\, \tau_i + \varepsilon_{ik}, \qquad x_{ik} = \text{Down}_{ik} \text{ or } \text{Up}_{ik}, \tag{241}$$

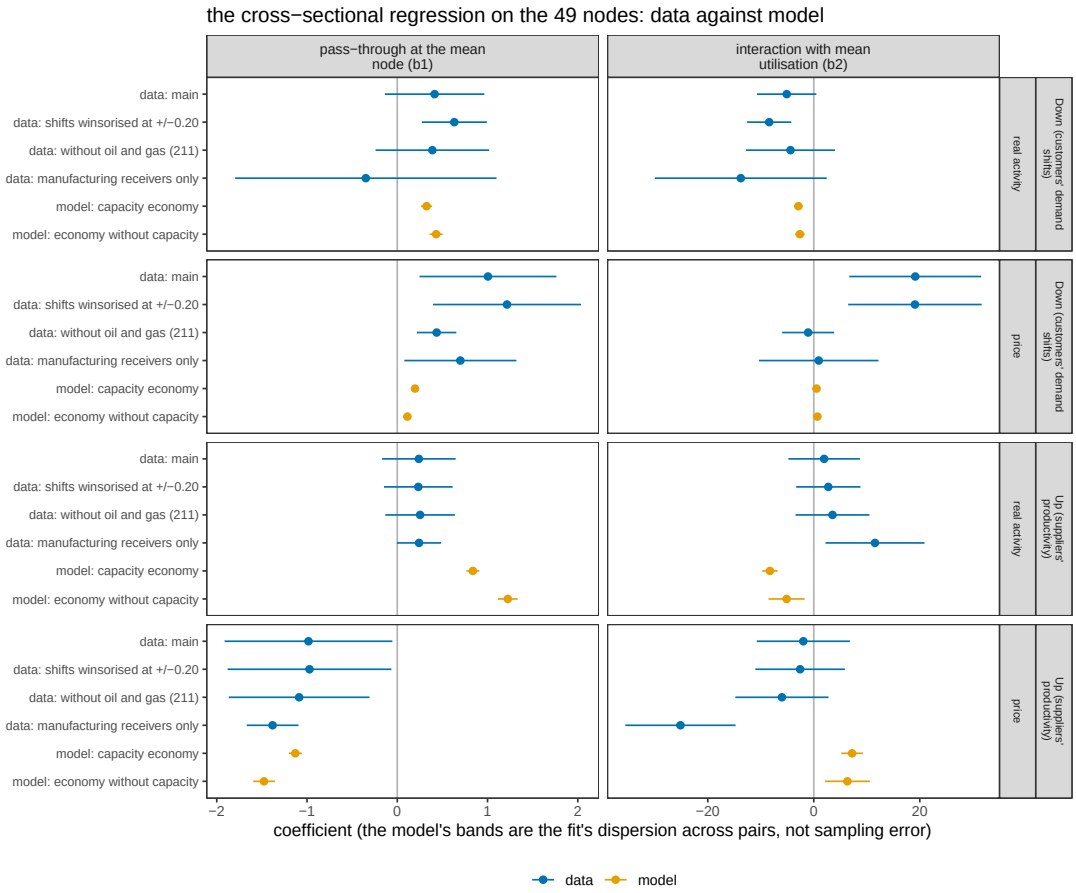
with receiver and shocked-node effects in the place of the node and year effects, on the 49 capacity-on receivers. For a demand shift every weight moves, so $\text{Down}_{ik} = \bar{w}^D_{ki} - \tilde{\beta}_k \sum_{m \neq i} \bar{w}^D_{mi}$ per unit of the shift, with $\tilde{\beta}_k$ the shocked node's weight in final demand (the second term is the rescaling and is small); for a productivity shock $\text{Up}_{ik} = a_{ik}$. The node's mean utilization in the capacity economy equals the calibration's level to every printed digit, so the modifier is the data regression's. Standard errors are clustered by receiver and measure how well one linear interaction describes a node's 62 responses; the model has no sampling error. Table 23 and Figure 19 set the data's cross-sectional estimates beside the model's.

Table 23: The cross-sectional regression on the 49 capacity-on nodes: data against model.

	pass-through at the mean node, b_1 (s.e.)	interaction with mean utilisation, b_2 (s.e.)	change per sd of the modifier
arriving demand shifts into real output growth			
data, main	0.41 (0.28)	−5.10 (2.85)	−0.61
data, shifts winsorised at ±0.20	0.63 (0.18)	−8.41 (2.13)	−0.65
data, without oil and gas (211)	0.39 (0.32)	−4.39 (4.27)	−0.55
data, manufacturing receivers	−0.35 (0.74)	−13.77 (8.28)	1.94
model, capacity economy	0.32 (0.03)	−2.90 (0.41)	−0.44
model, no-capacity economy	0.43 (0.04)	−2.62 (0.46)	−0.30
arriving demand shifts into price growth			
data, main	1.00 (0.39)	19.15 (6.34)	0.94
data, shifts winsorised at ±0.20	1.22 (0.42)	19.09 (6.43)	0.77
data, without oil and gas (211)	0.44 (0.11)	−1.06 (2.50)	−0.12
data, manufacturing receivers	0.70 (0.32)	0.94 (5.75)	0.07
model, capacity economy	0.20 (0.01)	0.53 (0.18)	0.13
model, no-capacity economy	0.11 (0.02)	0.69 (0.27)	0.30
arriving productivity shocks into real output growth			
data, main	0.24 (0.21)	1.96 (3.44)	0.40
data, shifts winsorised at ±0.20	0.23 (0.19)	2.75 (3.10)	0.58
data, without oil and gas (211)	0.25 (0.20)	3.53 (3.56)	0.69
data, manufacturing receivers	0.24 (0.12)	11.55 (4.76)	2.35
model, capacity economy	0.84 (0.04)	−8.27 (0.74)	−0.48
model, no-capacity economy	1.23 (0.06)	−5.13 (1.74)	−0.21
arriving productivity shocks into price growth			
data, main	−0.98 (0.47)	−1.97 (4.48)	0.10
data, shifts winsorised at ±0.20	−0.97 (0.46)	−2.58 (4.31)	0.13
data, without oil and gas (211)	−1.09 (0.40)	−6.01 (4.48)	0.27
data, manufacturing receivers	−1.38 (0.15)	−25.16 (5.31)	0.89
model, capacity economy	−1.13 (0.04)	7.23 (1.04)	−0.31
model, no-capacity economy	−1.48 (0.06)	6.35 (2.16)	−0.21

Notes: Data rows: the regression (240) in its cross-sectional form, main and three variants, standard errors clustered by node. Model rows: the regression (241) on the model's pair responses, in the capacity economy and in the no-capacity economy, standard errors clustered by receiver. The last column is the interaction times the standard deviation of the modifier (0.049) over the pass-through at the mean: the proportional change in the pass-through per standard deviation of mean utilization.

Figure 19: The cross-sectional regression, data against model.



Notes: The 49 capacity-on nodes; the pass-through at the mean node (left) and the interaction with mean utilization (right), for the four combinations of shock and outcome. Blue: the data with 95 per cent bands; orange: the model, whose bands are the dispersion of the fit across pairs, not sampling error.

Demand shifts into real activity. In the capacity economy $b_1 = 0.32$ and $b_2 = -2.90$, a fall of 0.44 per standard deviation of the modifier; in the economy without capacity constraints $b_1 = 0.43$ and $b_2 = -2.62$, a fall of 0.30; capacity's part is -0.11 in the pass-through and -0.28 (s.e. 0.24) in the interaction. The data's cross-section gives $b_1 = 0.41$ (s.e. 0.28) and $b_2 = -5.10$ (s.e. 2.85), a fall of 0.61 per standard deviation. So the data's pass-through at the mean node matches the model's, and the data's interaction is larger than the model's; but two thirds of the model's cross-sectional interaction is there without capacity, because the nodes with high mean utilization are the ones whose quantities respond less to demand for other reasons, their capital intensity and their position. The cross-sectional form does not isolate capacity, in the model or in the data. On the 24 nodes with a series the same holds (-2.77 with capacity against -2.41 without).

Demand shifts into prices. Capacity economy $b_1 = 0.20$, $b_2 = +0.53$; economy without capacity constraints 0.11, $+0.69$; capacity's part in the interaction -0.16 (s.e. 0.17): none, which is prediction (ii) in regression form. The data without the oil node give 0.44 (s.e. 0.11) and -1.06 (s.e. 2.50). The data's price pass-through of a demand shift, 0.4 to 1.0, is two to five times the model's shadow-price pass-through of 0.20.

Demand shifts into utilization (capacity economy only): $b_1 = 0.26$, $b_2 = -2.09$, a fall of 0.40 per standard deviation; the data's within form gives 0.07 (s.e. 0.12) and -2.80 (s.e. 1.23). The gradient agrees; the data's level is lower.

Productivity shocks. In prices the capacity economy gives $c_1 = -1.13$ and $c_2 = +7.23$, the economy without capacity constraints -1.48 and $+6.35$, and the data -0.98 (s.e. 0.47) and -1.97 (s.e. 4.48) across, $+4.89$ (s.e. 2.39) within: the pass-through matches, and the model's cross-sectional damping is mostly node heterogeneity again (capacity's part $+0.88$, s.e. 1.40). In real activity the capacity economy gives 0.84 and -8.27, the economy without capacity constraints 1.23 and -5.13, against the data's 0.24 (s.e. 0.21) and $+1.96$ (s.e. 3.44): the data's quantity pass-through of the suppliers' productivity is a third of the model's and shows no interaction.

The within form has no counterpart on these responses, which hold one state per node. Its partial-equilibrium analogue is the slope of the supply relation along the node's threshold (Figure 13): $e_S - 1$ falls by 58 per cent per within-node standard deviation of utilization at the median node. That is the own-price elasticity of utilization, not the general-equilibrium pass-through of an arriving shift, so it bounds the within interaction from above in magnitude rather than predicting it. A within-node counterpart would need the model solved at several states of the same node, for instance along the path of a common shock, with the arriving shocks repeated at each state.

[Caveats on the analogue] The model's shocks are one at a time and the regression pools 62 of them per receiver, whereas the data's shocks hit every node every year; the

weights are the data's, not the model's stationary sales shares (close by construction of the calibration); the economy without capacity constraints is the paper's own comparison economy, with the same capital intensities and no capacity, so "capacity's part" is the paper's definition of it. The within R^2 of the model regressions runs from 0.17 to 0.88: one linear interaction describes the real responses well and the price responses less well. Computed once, on the response matrices of Appendix G.3.

G.7 What the Result Can Carry, and the Open Decisions

The regression the model implies is estimable, and its quantity side comes out with the model's sign: the pass-through of an arriving demand shift into a node's output growth and into its utilization is smaller, and in the data about zero, when the node is tight, in every form and variant, at one to three standard errors. Its price side comes out as the model says once the oil node is set aside: a pass-through of 0.4 to 0.8 with no interaction. What the result cannot carry is also clear. The cross-sectional interaction is two thirds node heterogeneity by the model's own economy without capacity constraints, so the within-node form, 24 nodes with an interaction of -5 to -9 and a standard error of 2 to 3, and with no average pass-through to compare it to, is the result; and the identification sits in five upstream nodes, one of which is an oil-price story on the price side. The facts that match Section 6.1 are therefore the within-form interactions in quantities and in utilization, and the absence of a price interaction once oil is out; the cross-sectional form should be shown with the economy without capacity constraints's number beside it, or not at all.

[Decision 1: whether the result goes into the paper, and where]

- **Recommended:** into Section 6.1, as the data counterpart of the impact result, in one paragraph with a small table holding the within-form quantity and utilization interactions, the split halves, and the price pass-through without the oil node; and one sentence in the introduction that points to it. The reason: it is a within-node interaction identified off a handful of nodes, with no average pass-through to compare it to; that reads as corroboration of a stated mechanism, not as a headline fact that can open the paper. This appendix stays as its record, with the dark-teal boxes removed.
- **Alternative A:** the introduction's motivating fact, in the place of or beside the scatter of Figure 1. It would say: a demand shift that reaches a sector through its customers raises the sector's output and utilization when the sector is slack and not when it is tight. This is the paper's point in one sentence and it is what the model predicts; against it, the fact rests on 24 nodes and five of them, and a referee will ask for the instrument first.
- **Alternative B:** the results only, as a robustness appendix pointed to from Section 6.1 in one sentence, with nothing in the text. Cheapest; it wastes the agreement of the quantity side.

- **Alternative C:** nowhere, until the finer panel (three- and four-digit manufacturing from 1972) is built and the result holds there.

[Decision 2: the estimator, and the oil node]

- **Recommended:** OLS with lagged utilization as the modifier, as in Acemoglu et al. (2016), for the use recommended above (a corroborating fact); the oil node handled by reporting the estimates without it and on the manufacturing receivers, as Table 21 does.
- **Alternative A:** add a control for the oil price (its growth, and its growth interacted with the node's energy cost share). Cheap once a crude-oil price series is in the data set, which it is not; it keeps the oil node in the sample.
- **Alternative B:** instrument the arriving demand shift with the downstream instrument of Shea (2002), as Boehm and Pandalai-Nayar (2022) do, built from the same input-output table: the customers' demand predicted by their own customers' shifts, excluding the node's own. It is the natural instrument if the result is to carry evidential weight (Alternative A of Decision 1), and the exclusion restriction is the usual one of shift-share designs; a day's work on the existing panel.
- **Alternative C:** the finer panel, which would also allow the interaction with the customers' tightness to be estimated, since the detailed table has many more manufacturing-to-manufacturing links.

[Decision 3: the demand-shift series] Recommended: confirm it as the demand shock of the validation, with the two handlings (winsorising at ± 0.20 and node effects with a check without the trending nodes) written into its description in Appendix E.6; the alternatives are in the box of Appendix G.4.

[Decision 4: the two sentences for Section 6.1] The proposed text is in the maroon box of Appendix G.3: the word "output" for the sales elasticity, and the sentence on the aggregate against the node-level price response. Both were applied on 28 October 2026, in a shorter form; neither changes a number.

[If the appendix stays: what is redone] The eight figures are redrawn in the paper's style by the results scripts (one panel per file for the talk); the six tables are regenerated by the paper's table script from the same stored output; the demand-shift response matrices join the stored results of Section 6.1; the regression's code moves from the item's folder into the empirical pipeline; and the evidence boxes of this appendix are reduced to one status line each, as in Appendix E.