# Cascades and Fluctuations in an Economy with an Endogenous Production Network*

Mathieu Taschereau-Dumouchel†

Cornell University

November 2024

**Abstract**

This paper studies the efficient allocation in an economy in which firms are connected through input-output linkages and must pay a fixed cost to produce. When economic conditions are poor, some firms might decide not to operate, thereby severing the links with their neighbors and changing the structure of the production network. Since producers benefit from having access to additional suppliers, nearby firms tend to operate, or not, together. As a result, the production network features clusters of operating firms, and the exit of a producer can create a cascade of firm shutdowns. While well-connected firms are better able to withstand shocks, they trigger larger cascades upon exit. The theory also predicts how the structure of the production network changes over the business cycle. As in the data, recessions are associated with more dispersed networks that feature fewer highly connected firms. In the calibrated economy, the endogenous reorganization of the network substantially dampens the impact of idiosyncratic shocks on aggregate fluctuations.

**JEL Classifications:** E23, D85, C67

*I thank Enghin Atalay, Julieta Caunedo, Jeremy Greenwood, Veronica Guerrieri, Christopher Huckfeldt, Philipp Kircher, Ben Lester, Kristoffer Nimark, Ezra Oberfield, Nicolas Petrosky-Nadeau, Carlos Ramírez, Stephen Redding, Gill Segal, Ali Shourideh, Gianluca Violante, three anonymous referees, and participants at various seminars and conferences for helpful comments and suggestions. I gratefully acknowledge support from the Cornell Center for Social Sciences.

†Email: mt763@cornell.edu; Address: 480 Uris Hall, Cornell University, Ithaca NY 14853

## 1 Introduction

Production in modern economies involves a complex network of specialized firms, each using inputs from suppliers and providing their own output to downstream producers. Recent research has shown that the structure of this production network is an important determinant of economic outcomes. At the micro level, it influences the size of firms and how they withstand shocks (Barrot and Sauvagnat, 2016), while at the macro level, it affects how idiosyncratic shocks contribute to aggregate fluctuations (Acemoglu et al., 2012). Yet, the production network is not a fixed object and also responds to changes in the environment. This paper proposes a theory to study how the production network is formed, how it responds to shocks, and how that response affects macroeconomic fluctuations.

One distinguishing feature of the theory is that it focuses on the firms' extensive margin of production as the key driver behind the formation of the network. Consider, for instance, a firm that goes out of business after facing a severe shock. Since it no longer supplies to customers or purchases from suppliers, the links with its previous neighbors are cut. Similarly, when a new firm begins production, new connections with customers and suppliers are created. This process plays an important role in shaping the production network in the data.[1] It is also responsible for creating cascades of firm shutdowns: a chain reaction through which a shock to a vulnerable firm can lead to the exit of many of its (perhaps removed) neighbors. Policymakers were worried about such cascades during the financial crisis and the model can shed light on their origin and propagation.

In the model, a finite number of firms produce differentiated goods using labor and inputs from other producers. Production requires the payment of a fixed cost so that firms operate or not as a function of economic conditions. When a firm does operate, it makes an additional input available to all of its customers, and it purchases intermediate goods from its suppliers, thereby creating new input-output relationships. Together, the operating decisions of the firms therefore determine the structure of the production network. The main goal of the paper is to study the efficient allocation in this environment. That allocation provides a natural benchmark as it captures some of the main forces at work in the environment.

In the model, firms combine intermediate inputs using a standard CES production technology. As a result, having access to an additional input lowers the marginal cost of production and makes the firm effectively more productive. Because of these gains from input variety, firms with multiple suppliers are more likely to operate. Similarly, firms with many customers provide a valuable input to multiple producers and are also more likely to operate.

These complementarities between the operating decisions of nearby firms have important implications for the structure of the production network. First, they lead to the creation of clusters of firms that are tightly connected with one another. By organizing production in this way, firms in-

[1] According to Factset Revere, a large dataset of firm-level input-output linkages in the U.S., about 40% of link destructions occur when either the supplier or the customer (or both) stops producing. See Section 6.1 for details.

crease their number of customers and suppliers to take full advantage of the gains from input variety. Second, cascades of firm failures can arise in the efficient allocation. If a firm faces a severe shock and stops production, its customers, having lost a valuable input, and its suppliers, now producing a less useful product, are also more likely to shut down. The same logic applies to the firm's second neighbors, which are more likely to shut down as well, and so on. As a result, the initial shock can trigger a cascade of firm shutdowns that propagates upstream and downstream through the production network.

The operating complementarities between neighbors can also lead to a large reorganization of the network after a small change in the environment. For instance, a small decline in the productivity of a centrally located firm can lead to the shutdown of its whole neighborhood, as economic activity moves to a more productive part of the network. Through that mechanism, large changes in the distribution of firm-level outcomes can occur in response to arbitrarily small shocks.

Two features of the environment make the problem of the social planner particularly challenging to solve. First, since the decision to operate a firm is binary, the planner's optimization problem has a non-convex feasible set. Second, the increasing returns to scale generated by the fixed costs break the concavity of the objective function. As a result, the objective function generally features multiple local maxima and standard algorithms are unable to identify the global maximum. This paper provides a novel solution approach that involves reshaping the original optimization problem such that 1) this reshaped problem can be solved easily, and 2) its solution coincides with that of the original problem. I establish sufficient conditions under which this approach is guaranteed to find the efficient allocation. But even when those conditions are not met, numerical simulations show that it provides a rapid and robust way of tackling a class of challenging network formation problems (Carvalho and Tahbaz-Salehi, 2018).

While this paper focuses on the efficient allocation, distortions such as market power might play a role in shaping the production network in reality. In particular, superstar firms such as Walmart and IKEA are believed to have a large amount of pricing power when dealing with suppliers (Bloom and Perry, 2001). I incorporate this feature in the model and show that it leads to inefficient entry decisions by firms. The forces at work in that equilibrium are similar to those of the efficient allocation, but there are differences in how the complementarities between firms operate and, as a result, how cascades propagate. The proposed solution method can also be used to solve for that inefficient equilibrium in a straightforward way.

I provide a basic calibration of the model using firm-level data for the United States economy. To better understand how economic forces shape the production network, I compare the efficient network, designed optimally by the planner, to a neutral benchmark whose structure is randomly determined. The efficient network features indegree (number of suppliers) and outdegree (number of customers) distributions with thicker right tails, as well as a higher amount of clustering between firms. These differences show that the planner takes advantage of the operating complementarities by creating tightly connected clusters of economic activity centered around well-connected firms.

I investigate how cascades of firm shutdowns arise and propagate through the calibrated network. As in the data, highly connected firms are more resilient to shocks but, upon shutting down, they create larger cascades that lead to the exit of several of their neighbors. Cascades also interact with macroeconomic aggregates, and those that originate from highly connected firms are often associated with substantial declines in GDP.

One contribution of this paper is to highlight novel business cycle correlations between aggregate output and the structure of the production network. In the data and in the model, recessions are periods in which, 1) the tails of the degree distributions are thinner, and 2) there is less clustering between firms. These correlations are naturally explained by the model. Expansions are periods in which it is easy to leverage the complementarities at work in the economy by creating productive clusters of firms. In contrast, recessions are periods in which creating these clusters would be too costly, perhaps because a few influential firms are facing bad shocks, and in which production therefore involves a more diffused, and less productive, network.

I also consider how the endogenous formation of the network interacts with firm-level shocks to influence aggregate fluctuations. To do so, I compare the benchmark economy, in which the production network reorganizes itself in response to shocks, to an alternative economy in which the network is kept fixed. Aggregate output is on average 11% lower and 20% more volatile under the fixed network. This last finding highlights the importance of considering how the production network adapts to shocks to better understand the microeconomic origin of aggregate fluctuations.

Finally, I compare the efficient allocation to the inefficient equilibrium in the calibrated economy to evaluate the quantitative impact of pricing distortions. While they are mostly similar, there are a few notable differences between the two allocations. For instance, the equilibrium production network is less correlated with GDP, suggesting that it is more rigid and less able to adapt to changing economic conditions. Cascades also tend to propagate more downstream than in the efficient allocation, as predicted by the theory.

The model is motivated by an empirical literature documenting that losing a supplier is disruptive to a firm's operations. Carvalho et al. (2014) document that firms that stopped production because of the Great East Japan Earthquake of 2011 had a significant negative impact on their customers and suppliers. Hendricks and Singhal (2005) find that firms facing supply chain disturbances suffer from large and long-lasting negative abnormal stock returns. Wagner and Bode (2008) survey business executives in Germany who report that issues with supply chains, including the loss of a supplier, were responsible for significant disturbances to production.

This paper also relates to a literature that studies how shocks to interconnected sectors contribute to aggregate fluctuations in exogenous networks. In an influential paper, Acemoglu et al. (2012) find that sectoral shocks can lead to large aggregate fluctuations if there is enough asymmetry in the way sectors supply to each other.[2] Acemoglu et al. (2015) further show that inter-sectoral linkages can generate larger tail-risks in aggregate output. This literature emphasizes the importance of the (fixed) structure of the network in transmitting idiosyncratic shocks. In contrast, the current paper studies how endogenizing the network affects aggregate fluctuations.[3]

One of the first papers to study the macroeconomic impact of cascades is Baqaee (2018), which considers a model with an exogenous sectoral network in which the mass of firms in each sector can vary. This adjustment margin can lead to further amplification of sectoral shocks in the presence of external economies of scale. In contrast to that paper, the present work considers a discrete adjustment margin that leads to the creation and destruction of nodes and edges in the production network, something that we observe in the firm-level data. Another paper that emphasizes the role of discreteness is Elliott et al. (2020), which considers a supply network in which each link is at risk of failure. This discrete margin can make the network fragile, in the sense that aggregate output becomes very sensitive to small shocks. Acemoglu and Tahbaz-Salehi (2020) also look at the impact of supply chain disruptions in a model with bargaining and endogenous markups.

This paper contributes to a recent literature in which production networks are built endogenously by the decisions of economic agents. One of the first in that literature is Oberfield (2018) who builds a model in which producers optimally choose one input from a randomly evolving set of suppliers, thereby creating a production network. Lim (2018) studies sourcing decisions in a model with sticky relationships. Unlike the present work, these papers feature a continuum of firms so that aggregate fluctuations do not arise from individual idiosyncratic shocks, a margin whose importance has been emphasized by the granularity literature (Gabaix, 2011) but that has proven challenging to incorporate in network formation models (Carvalho and Tahbaz-Salehi, 2018).

Acemoglu and Azar (2018) consider a network of competitive industries in which firms select a production technique that involves different sets of suppliers. They show that the endogenous evolution of the network can generate long-run growth. Tintelnot et al. (2018) build a model of endogenous network formation and international trade. In contrast to the current paper, they only consider acyclic networks. Boehm and Oberfield (2018) estimate a model of network formation using Indian micro data to study misallocation in the inputs market. Kopytov et al. (2022) look at the impact of uncertainty on the structure of the production network.

This paper proposes a new solution technique for some nonconvex optimization problems with binary variables. Several heuristics have been developed to handle these problems (Li and Sun, 2006). Closest to the present work are smoothing algorithms that attempt to get rid of the local maxima that emerge in the relaxed problem (Murray and Ng, 2010). In practice, finding an appropriate smoother is usually done through trial and error and there is no guarantee that the algorithm converges to a

[2] In a related paper, Gabaix (2011) shows that when the tail of the firm size distribution is sufficiently thick, firm-level shocks can have large effects on aggregates.

[3] See Carvalho and Tahbaz-Salehi (2018) for an overview of the literature on production networks. Recent contributions include di Giovanni et al. (2014), Atalay (2017), Baqaee and Farhi (2017b), Bigio and La'O (2016), Caliendo et al. (2017b), Caliendo et al. (2017a), Ozdagli and Weber (2017), Liu (2019) and Chahrour et al. (2019).

global maximum. In contrast, the current work explicitly describes how to reshape the problem and proposes a rapid and robust solution method.

The next section introduces the model. Section 3 describes the solution method. Section 4 discusses equilibrium allocations. Section 5 explores the forces at work in the economy. Section 6 provides a calibration to U.S. data.

## 2 A model of production networks with entry

There is a set $\mathcal{N} = \{1, \ldots, n\}$ of firms, each producing a differentiated good that can be used as intermediate input by other firms or consumed by a representative household. The preferences of the household are given by the utility function

$$C = \left( \sum_{j=1}^{n} \beta_j^{\frac{1}{\sigma}} c_j^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \tag{1}$$

where $c_j$ is consumption of good $j$, $\sigma > 1$ is the elasticity of substitution between goods and $\beta_j \geq 0$ determines the household's taste for good $j$. Throughout, I refer to $C$ as aggregate consumption or GDP. The household also supplies one unit of labor inelastically.

To produce, a firm $j$ must employ $f_j \geq 0$ units of labor as a fixed cost, in which case $j$ is *operating*. This fixed cost captures overhead labor, such as managers and other non-production workers, that is necessary for production.[4] The vector $\theta \in \{0, 1\}^n$ keeps track of the operating decisions of the firms, such that $\theta_j = 1$ if $j$ operates and $\theta_j = 0$ otherwise.

When operating, firm $j$ can convert $l_j$ units of labor and a vector of intermediate inputs $x_j = (x_{1j}, \ldots, x_{nj})$ into $y_j$ units of good $j$ according to the production function

$$y_j = \frac{A}{\alpha_j^{\alpha_j} (1-\alpha_j)^{1-\alpha_j}} z_j \theta_j \left( \sum_{i=1}^{n} \Omega_{ij}^{\frac{1}{\varepsilon_j}} x_{ij}^{\frac{\varepsilon_j - 1}{\varepsilon_j}} \right)^{\frac{\varepsilon_j}{\varepsilon_j - 1}\alpha_j} l_j^{1-\alpha_j}, \tag{2}$$

where $\Omega_{ij} \geq 0$ denotes the factor intensity of input $i$, $\varepsilon_j > 1$ is the elasticity of substitution between inputs, $0 < 1 - \alpha_j < 1$ is the labor intensity, and $A > 0$ and $z_j > 0$ are aggregate and firm-specific total factor productivities.[5] Since $\varepsilon_j > 1$, intermediate inputs are substitutes in the production of good $j$.[6]

We see from (2) that firm $j$ can only use inputs from supplier $i$ if $\Omega_{ij} > 0$. As such, the matrix $\Omega$

[4] Empirical studies have found that fixed costs are important in explaining firm-level outcomes. For instance, Bresnahan and Ramey (1994) show that the extensive margin of operation they generate is responsible for about 80% of plant-level output fluctuations in the automobile industry.

[5] I assume that $\sum_i \Omega_{ij} > 0$ for all $j$, otherwise $j$ cannot produce and the economy can be redefined without it. The term $\alpha_j^{\alpha_j} (1-\alpha_j)^{1-\alpha_j}$ in (2) is a normalization to simplify some expressions.

[6] The restrictions $\sigma > 1$ and $\varepsilon_j > 1$ are necessary to avoid a complete shutdown of the economy, or of a customer, if a single producer does not operate.

describes a network of *potential connections* between firms. A potential connection $(i, j)$ is *active*—with goods being traded—if firms $i$ and $j$ both operate, otherwise it is *inactive*. The production network is therefore jointly determined by $\Omega$ and $\theta$, and economic conditions endogenously determine the structure of the network through their impact on operating decisions.

Panel (a) in Figure 1 provides an example of potential connections in an economy with six firms. Each arrow represents a connection $\Omega_{ij} > 0$, with the direction of the arrow indicating the potential movement of goods. The set of active connections, in blue in panel (b), is determined endogenously by the set of operating firms, also shown in blue.

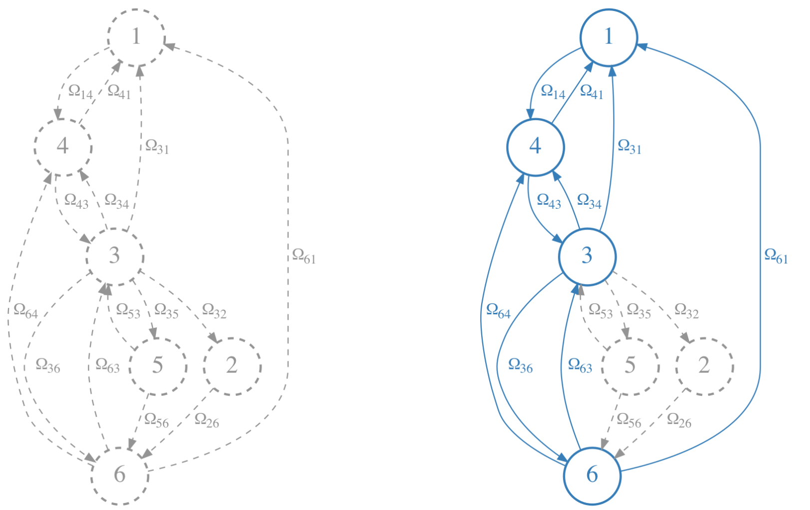


(a) Potential connections (b) Active firms and connections

Figure 1: The firms' operating decisions determine the production network

While this paper focuses on the role of the firms' extensive margin of operation in the formation of the network, the model can also accommodate the formation of individual links. Specifically, a link between any two firms $i$ and $k$ can be interpreted as a *pseudo* firm $j$ that 1) only has potential connections in $\Omega$ with $i$ as a supplier and $k$ as a customer, and 2) produces a good that is not included in the production of the final good ($\beta_j = 0$). $\theta_j \in \{0, 1\}$ then indicates whether the link between $i$ and $k$ is active or not.

It is useful to describe the set of firms that *can* produce under a given operating vector $\theta$. For a firm to produce, it must receive some intermediate input from at least one supplier, and this supplier must also receive some input from a supplier, and so on. Since $n$ is finite, this sequence of suppliers must contain a cycle for production to take place. As a result, a firm without access to such an operating cycle simply cannot produce.[7] We will use this fact later on to characterize an allocation under a given vector $\theta$.

[7]Formally, an *operating cycle* is a sequence of operating firms $\{s_1, \ldots, s_k\}$, for some $k \geq 1$, such that 1) $\Omega_{s_i, s_{i+1}} > 0$ for all $i \in \{1, \ldots, k-1\}$, and 2) $\Omega_{s_k, s_1} > 0$. A firm $s_j$ *has access to an operating cycle* if there exists a sequence of operating firms $\{s_1, \ldots, s_j\}$ such that 1) $s_1$ is part of an operating cycle, and 2) $\Omega_{s_i, s_{i+1}} > 0$ for all $i \in \{1, \ldots, j-1\}$.

## 3 The efficient allocation and how to find it

Consider the problem $\mathcal{P}$ of a social planner that maximizes the utility of the household

$$\mathcal{P}: \max_{\substack{c\geq 0,\, x\geq 0,\, l\geq 0\\ \theta\in\{0,1\}^n}} \left(\sum_{j=1}^{n} \beta_j^{\frac{1}{\sigma}} c_j^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}, \tag{3}$$

subject to a resource constraint for each intermediate good $j$,

$$c_j + \sum_{k=1}^{n} x_{jk} \leq y_j, \tag{4}$$

where $y_j$ is given by (2), and a resource constraint for labor,

$$\sum_{j=1}^{n} l_j + \sum_{j=1}^{n} \theta_j f_j \leq 1. \tag{5}$$

An allocation is *efficient* if it solves $\mathcal{P}$. We will tackle the planner's problem in two steps: first assuming that the production network, and therefore $\theta$, is fixed, and then looking at the full problem with an endogenous network.

### 3.1 Planner's problem with exogenous operating decisions $\theta$

For a fixed $\theta$, $\mathcal{P}$ is a standard convex maximization problem and the usual Karush-Kuhn-Tucker conditions are necessary and sufficient to characterize its solution. Denote by $\lambda_j$ the Lagrange multiplier on good $j$'s resource constraint (4) and by $w$ the multiplier on the labor resource constraint (5). The first-order conditions imply that $(1-\alpha_j)\, y_j \lambda_j = w l_j$ so that, as in Oberfield (2018), we can define $q_j = w/\lambda_j$ as a measure of firm $j$'s labor productivity. From the planner's first-order conditions, we can then characterize the vector $q = (q_1, \ldots, q_n)$ as a function of $\theta$.

**Lemma 1.** *In the efficient allocation, the labor productivity vector $q$ satisfies*

$$q_j = z_j \theta_j A \left(\sum_{i=1}^{n} \Omega_{ij} q_i^{\varepsilon_j - 1}\right)^{\frac{\alpha_j}{\varepsilon_j - 1}}, \tag{6}$$

*for all $j$. Furthermore, there is a unique $q$ that solves (6) such that $q_j > 0$ if $j$ operates and has access to an operating cycle, and $q_j = 0$ otherwise.*

The proof of Lemma 1 shows that $q$ can be found by iterating on (6). Several features of (6) are worth emphasizing. First, its recursive structure implies that a change in the labor productivity $q_j$ of a firm $j$ propagates downstream through supply chains. For instance, if $j$ faces a negative TFP shock, the amount of labor needed to produce one unit of good $j$ increases, which leads to a lower labor productivity for $j$'s customers, and for its customers' customers and so on. Second, (6) implies that a firm that has access to a greater set of active suppliers (more positive terms in the summation) is more productive (higher $q_j$). Intuitively, with a more diverse set of inputs, a firm might be able to use better production techniques that would otherwise be unavailable. The elasticity $\varepsilon_j$ governs how substitutable these inputs are and is the key parameter determining the strength of this mechanism. When $\varepsilon_j$ is small, intermediate inputs are poor substitutes, and the benefit of having an additional supplier is large. In contrast, when $\varepsilon_j$ is large, firm $j$'s labor productivity is almost entirely driven by its most productive supplier. As we will see, these mechanisms have important implications for the structure of the network and the propagation of shocks.

With $q$ in hand, it is straightforward to derive all other quantities in the efficient allocation. In particular, the next lemma shows that GDP $C$ can be computed as the product of aggregate productivity

$$Q = \left( \sum_{j=1}^{n} \beta_j q_j^{\sigma-1} \right)^{\frac{1}{\sigma-1}}. \tag{7}$$

and the amount of labor available after fixed costs have been paid.

**Lemma 2.** *In the efficient allocation, GDP is given by*

$$C = Q \left( 1 - \sum_{j=1}^{n} \theta_j f_j \right). \tag{8}$$

We see from (7) that aggregate productivity $Q$ is a CES aggregator of the underlying firm-level labor productivities $q_j$, with an elasticity of substitution controlled by $\sigma$. When $\sigma$ is small, the differentiated goods are poor substitutes, and each additional good is more highly valued by the planner. Conversely, when $\sigma$ is large, the household derives utility primarily from the most productive firm. As a result, $\sigma$ affects the planner's incentives to operate more firms and will also play an important role in shaping the structure of the network.

### 3.2 The planner's problem with $\theta$ as a choice variable

Lemma 2 completes the solution of the planner's problem with a fixed $\theta$, and we can now take a step back to consider the full problem $\mathcal{P}$, in which the network itself is a choice variable. By combining Lemmas 1 and 2, $\mathcal{P}$ can be written as the problem of finding the vector of operating decisions $\theta^*$ that maximizes consumption (8), with $q$ determined by (6). These equations highlight the key trade-off faced by the planner when deciding whether to operate firm $j$. Because of the recursive structure of (6), operating $j$ improves the labor productivity $q$, not only of $j$ itself, but also of all its downstream customers, which benefits aggregate productivity $Q$. On the other hand, operating $j$ takes $f_j$ units of labor away from other productive uses.

Problem $\mathcal{P}$ is challenging for two reasons. First, $\theta$ is limited to the *corners* $\{0,1\}^n$ of the $n$-dimensional unit hypercube—a non-convex set. But even if $\theta$ could move freely within $[0,1]^n$, the fixed costs of operation create firm-level increasing returns to scale that break the concavity of the objective function. As a result, there are usually multiple local maxima, and the standard Karush-Kuhn-Tucker conditions are not sufficient to find the global maximum.[8]

There is, however, a brute-force method of solving $\mathcal{P}$. Since there are only a finite number of vectors $\theta$ in the feasible set $\{0,1\}^n$, we can try them all. For each $\theta$, we can iterate on (6) to find $q$, and the objective function can then be computed using (8). While this *exhaustive search* strategy is guaranteed to find the correct solution, it is in practice limited to economies with only a few firms. Since there are $2^n$ possible $\theta$'s in $\{0,1\}^n$, the number of potential vectors explodes as $n$ grows.

**Reshaping the planner's problem**

To handle economies with large $n$, this paper proposes a novel solution method that is less computationally intensive. The key idea is to consider an alternative optimization problem that is easy to solve and whose solution coincides with that of $\mathcal{P}$. This alternative problem, denoted by $\mathcal{R}$, is obtained by *relaxing* and *reshaping* $\mathcal{P}$ and is defined in the right column of Figure 2.

$\mathcal{P}$: Original planner's problem

$$\max_{\theta \in \{0,1\}^n} Q\left(1 - \sum_{j=1}^{n} \theta_j f_j\right) \tag{8}$$

where $q$ solves, for each $j \in \mathcal{N}$,

$$q_j = z_j \theta_j A \left(\sum_{i=1}^{n} \Omega_{ij} q_i^{\varepsilon_j - 1}\right)^{\frac{\alpha_j}{\varepsilon_j - 1}} \tag{6}$$

$\mathcal{R}$: Relaxed and reshaped problem

$$\max_{\theta \in \mathbf{[0,1]^n}} Q\left(1 - \sum_{j=1}^{n} \theta_j f_j\right) \tag{8}$$

where $q$ solves, for each $j \in \mathcal{N}$,

$$q_j = z_j \theta_j^{\boldsymbol{a_j}} A \left(\sum_{i=1}^{n} \Omega_{ij} \left(\boldsymbol{\theta_i^{b_{ij}}} q_i\right)^{\varepsilon_j - 1}\right)^{\frac{\alpha_j}{\varepsilon_j - 1}} \tag{9}$$

Figure 2: Differences between the original and reshaped problems

$\mathcal{R}$ differs from $\mathcal{P}$ in two important ways—emphasized in blue in Figure 2. First, the binary constraint $\theta \in \{0,1\}^n$ is relaxed, and $\theta$ can now take values *inside* the unit hypercube $[0,1]^n$. While this relaxation has the advantage of convexifying $\mathcal{P}$'s feasible set, it also adds points that have no real economic meaning to the planner's problem. For instance, $\theta_j = 0.5$ does not correspond to any physical reality in the economic environment. One of the key insights of the paper is that, since they have no interest on their own, we can change the value of the objective function over these new points to help us solve $\mathcal{P}$.

[8] Problem $\mathcal{P}$ belongs to the class of Mixed Integer Nonlinear Problems (MINLP). Their combinatorial nature makes these problems notoriously challenging to solve and they are, in general, NP-Hard (Garey and Johnson, 1990).

This is done in (9), which is a transformed version of (6) that includes the *shape parameters* $a_j > 0$ and $b_{ij}$.[9] These parameters modify the shape of the optimization problem everywhere *except* over $\mathcal{P}$'s original feasible set $\{0,1\}^n$. Indeed, for $\theta \in \{0,1\}^n$, $\theta_j^{a_j} = \theta_j$ for all $j$. Similarly, for $b_{ij}$, if $\theta_i = 0$ then $q_i = 0$ anyway, and if $\theta_i = 1$ then $\theta_i^{b_{ij}} = 1$. In both cases, the term in the summation is unchanged. This reshaping procedure therefore preserves the ranking, in terms of household utility, of the corners $\{0,1\}^n$—the only points with actual economic meaning—while changing the shape of the optimization problem elsewhere. As a consequence, if we solve the reshaped problem $\mathcal{R}$ and that its solution belongs to $\{0,1\}^n$, then that solution must necessarily also solve $\mathcal{P}$. The following proposition formalizes that idea.

**Proposition 1.** *If $\theta^* \in \{0,1\}^n$ solves $\mathcal{R}$, then $\theta^*$ also solves $\mathcal{P}$.*

This result provides a clear way of solving $\mathcal{P}$, but it requires that 1) $\mathcal{R}$ can be solved easily, and 2) its solution belongs to $\{0,1\}^n$. This is not the case in general, but we can pick the shape parameters to help us achieve both objectives. In what follows, I first provide some intuition about how $a_j$ and $b_{ij}$ affect the shape of the objective function. I then establish conditions on these parameters such that solutions to $\mathcal{R}$ belong to $\{0,1\}^n$. Finally, I provide sufficient conditions on the environment for $\mathcal{R}$ to be a convex optimization problem. In this case, standard algorithms can solve $\mathcal{R}$ rapidly even in economies with many firms.

### The role of the shape parameters $a_j$ and $b_{ij}$

Equation (9) shows how $a_j$ and $b_{ij}$ affect the recursive mapping that determines the labor productivity vector $q$, but we can equivalently think of the reshaping procedure in terms of $\mathcal{P}$'s original formulation given by (3) to (5). In that case, that procedure simply involves transforming the production function (2) into

$$y_j = \frac{A}{\alpha_j^{\alpha_j}\left(1-\alpha_j\right)^{1-\alpha_j}} z_j \theta_j^{\boldsymbol{a_j}} \left( \sum_{i=1}^{n} \Omega_{ij}^{\frac{1}{\varepsilon_j}} \left( \boldsymbol{\theta_i^{b_{ij}}} x_{ij} \right)^{\frac{\varepsilon_j - 1}{\varepsilon_j}} \right)^{\frac{\varepsilon_j}{\varepsilon_j - 1}\alpha_j} l_j^{1-\alpha_j}, \tag{10}$$

where the affected terms are highlighted in blue. We see that $\theta_j$ plays the role of a TFP shifter in the production of good $j$, with $a_j$ controlling its influence. Similarly, with $b_{ij} > 0$ a larger $\theta_i$ provides an input-specific productivity increase to firm $j$. From this reshaped production function, it is straightforward to interpret the first-order conditions of problem $\mathcal{R}$.The first-order condition of problem $\mathcal{R}$ with respect to $\theta_j$ can be written as

$$a_j \lambda_j \theta_j^{-1} y_j + \sum_k b_{jk} \lambda_k \theta_j^{-1} x_{jk} - w f_j = \Delta\mu_j, \tag{11}$$

[9] Note that $\theta_i^{b_{ij}} q_i \propto \theta_i^{b_{ij}} \theta_i^{a_i}$. For (9) to be well-defined as $\theta_i \to 0$, I therefore impose that $a_i + b_{ij} \geq 0$ for all $i, j$.

where $\Delta\mu_j = \overline{\mu}_j - \underline{\mu}_j$ is the difference between the Lagrange multipliers on the constraints $\theta_j \leq 1$ and $\theta_j \geq 0$, respectively.

As we can see, the shape parameters change the marginal benefit of increasing $\theta_j$ (first two terms in (11)). From the perspective of the reshaped production function (10), a marginal increase in $\theta_j$ leads to the production of an additional $a_j \theta_j^{-1} y_j$ units of good $j$, each of which with a social value $\lambda_j$. This effect is captured by the first term in (11). Similarly, an increase in $\theta_j$ provides firm $k$ an input-specific increase in productivity (for $b_{jk} > 0$) that is proportional to $k$'s use of input $j$. Accounting for good $k$'s social value $\lambda_k$ and summing across all firms that use good $j$ yields the second term in (11). The first-order condition (11) implies that the planner trades off these marginal benefits of increasing $\theta_j$ with the marginal labor cost given by $wf_j$.[10]

**When do solutions to $\mathcal{R}$ belong to $\{0,1\}^n$?**

Proposition 1 implies that a solution to $\mathcal{R}$ also solves $\mathcal{P}$ only if it belongs to $\{0,1\}^n$. Here, we look at conditions on $a_j$ and $b_{ij}$ for this to happen. These conditions are derived by looking at how the first-order condition (11) varies with $\theta_j$. Since several terms in that equation depend implicitly on $\theta_j$, it is convenient to rewrite these terms to make this dependence explicit. The following lemma provides a version of (11) that depends exclusively on $\theta_j$ and aggregate quantities.

**Lemma 3.** *The first-order condition of problem $\mathcal{R}$ with respect to $\theta_j$ can be written as*

$$a_j \frac{\lambda_j}{\theta_j} \underbrace{\beta_j \lambda_j^{-\sigma} C}_{c_j} + \sum_k (a_j + b_{jk}) \frac{\lambda_j}{\theta_j} \underbrace{\theta_j^{b_{jk}(\varepsilon_k - 1)} \Omega_{jk} X_k \left(\frac{\lambda_j}{\Lambda_k}\right)^{-\varepsilon_k}}_{x_{jk}} - wf_j = \Delta\mu_j, \tag{12}$$

*where*

$$X_j = \left(\sum_k \theta_k^{b_{kj}\frac{\varepsilon_j - 1}{\varepsilon_j}} \Omega_{kj}^{\frac{1}{\varepsilon_j}} x_{kj}^{\frac{\varepsilon_j - 1}{\varepsilon_j}}\right)^{\frac{\varepsilon_j}{\varepsilon_j - 1}} \text{ and } \Lambda_j = \left(\sum_k \theta_k^{b_{kj}(\varepsilon_j - 1)} \Omega_{kj} \lambda_k^{1-\varepsilon_j}\right)^{\frac{1}{1-\varepsilon_j}}, \tag{13}$$

*are, respectively, the reshaped intermediate input bundle of firm $j$ and the social value of that bundle, and where*

$$\lambda_j = \frac{1}{z_j \theta_j^{a_j} A} \Lambda_j^{\alpha_j} w^{1-\alpha_j}, \tag{14}$$

*is the social value of a unit of good $j$.*

Equation (12) is derived by combining (11) with the first-order conditions for $c_j$ and $x_{jk}$. It provides the accounting of resources that go into operating firm $j$. Operating $j$ requires $f_j$ units of labor to pay the fixed cost and produces $c_j$ units for consumption and $\sum_k x_{jk}$ units for other firms.[11]

[10] As we will see in Figure 3, some terms in (11) can go to infinity as $\theta_j \to 0$ depending on $a_j$ and $b_{ij}$. This does not happen under the specific shape parameter values that we adopt below to solve $\mathcal{P}$.

[11] See Lemma 7 in Appendix A for a version of that equation that explicitly accounts for the intermediate inputs and the non-fixed-cost labor used by $j$.

By replacing (14) in (12), we see that we can write this first-order condition with respect to $\theta_j$ as a function of only $\theta_j$ itself and the *aggregate* quantities $w$, $C$, $(X_1, \ldots, X_n)$ and $(\Lambda_1, \ldots, \Lambda_n)$. These aggregates involve summations over many firms and become more and more independent of $\theta_j$ as the number of firms $n$ increases.

We will later consider situations in which changes in $\theta_j$ can impact these aggregates, but when they are independent we can characterize how the shape parameters matter for the first-order condition in a simple way. By substituting (14) into (12), we see that the first term's dependence on $\theta_j$ operates through $\theta_j^{a_j(\sigma-1)-1}$. Since the exponent involves the product of $a_j$ and the elasticity of substitution $\sigma$, it follows that a higher $a_j$ makes the first-order condition behave as if goods are more substitutable in the consumption bundle, taking all aggregates as constant. Similarly, the second term in (12) depends on $\theta_j$ through $\theta_j^{(a_j+b_{jk})(\varepsilon_k-1)-1}$, so that $b_{jk}$ plays a similar role for the elasticity of substitution in good $k$'s intermediate bundle.

Since $a_j$ and $b_{jk}$ influence the marginal benefit of increasing $\theta_j$ we can pick their values to help us solve $\mathcal{R}$. For instance, specific values for $a_j$ and $b_{jk}$ might remove local maxima in which optimization algorithms might get stuck. To explain the intuition, it is useful to consider a simple economy in which there are no intermediate inputs ($\alpha_j = 0$ for all $j$, so that $b_{ij}$ does not matter).[12] Figure 3 shows the marginal benefit of increasing $\theta_j$ (first two terms in (12)) as a function of $\theta_j$ itself, and for different values of $a_j$. The solid blue line shows that without any reshaping ($a_j = 1$) this quantity is *increasing* in $\theta_j$. As a result, there are two local maxima, and the first-order conditions are not sufficient to guarantee optimality. Indeed, suppose that we use a simple gradient ascent algorithm starting from $\theta_j = 0$. At this point, the marginal benefit of increasing $\theta_j$ is low compared to the marginal cost $wf_j$ (black line in the figure). It follows that the algorithm does not move and that $\theta_j = 0$ is a local maximum. If, instead, we use the same algorithm but starting from $\theta_j = 1$, the marginal benefit of increasing $\theta_j$ is larger than $wf_j$, such that $\theta_j = 1$ also satisfies the first-order condition, and we have a second local maximum.

To understand intuitively why multiple local maxima exist, recall that goods are substitutes in the consumption and intermediate input bundles. It follows that the household and the firms tend to rely disproportionately on high-productivity producers. When $\theta_j \approx 0$, we see from (10) that firm $j$ has effectively a low productivity, so that $\theta_j$'s *marginal* impact on consumption is small. If the substitution forces are strong enough, that marginal impact is lower than $wf_j$ so that the local incentives to operate $j$ around $\theta_j = 0$ push to keep it inoperative. But $j$ might be extremely productive when $\theta_j = 1$, in which case the planner is happy to keep $j$ operating at $\theta_j = 1$, and we have a second local maximum at that point.

Figure 3 also shows the marginal benefit curve under different values of the shape parameter $a_j$ (dashed green lines). We see that this curve remains increasing for high values of $a_j$, such that the local maxima problem remains. As explained earlier, a higher $a_j$ effectively increases the substitution

[12] This is only to simplify the exposition. The intuition for $b_{ij}$ is analogous to that for $a_j$.

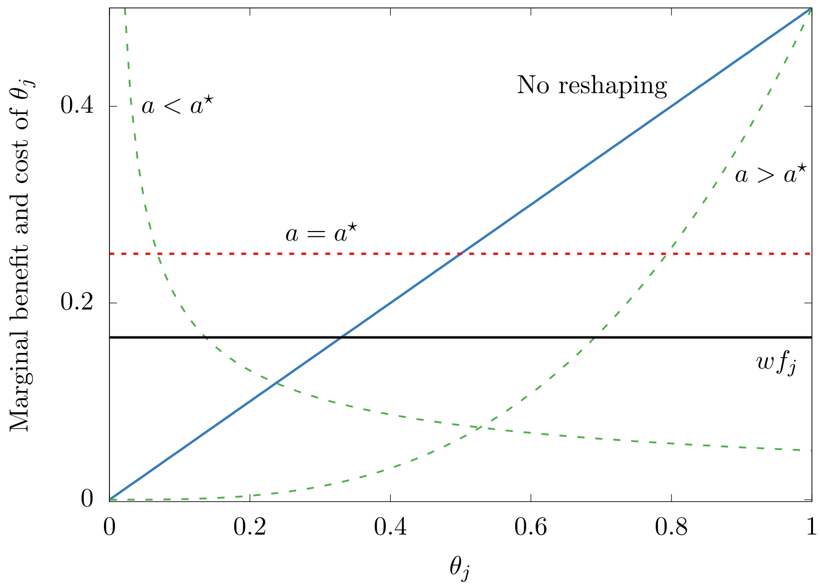


Notes: Marginal benefit (first two terms in (12), all lines except the black one) and marginal cost (third term $wf_j$ in (12), black line) of increasing $\theta_j$ under different values of $a_j$. Parameters: $\alpha_j = 0$, $\sigma_j = 3$, $\beta_j = 1/2$, $f_j = 0.15$, $C = z = w = 1$.

Figure 3: Impact of $a_j$ on the marginal productivity of $\theta_j$

in the consumption bundle, which exacerbates the local maxima problem. For low values of $a_j$, the opposite happens. Goods become effectively more complementary, and the marginal product curve is decreasing. This implies, in this example, an *interior* solution $\theta_j^*$ where the marginal benefit curve intersects the $wf_j$ line. As $\theta^* \notin \{0,1\}^n$, the solution to $\mathcal{R}$ in this case clearly does not coincide with that of $\mathcal{P}$.

Figure 3 also shows an intermediate value of $a_j$, denoted by $a^\star$, for which the marginal product curve is completely flat (red dashed line). In this case, the marginal benefit of increasing $\theta_j$ is always larger than the marginal cost, and the gradient ascent method will converge, regardless of its starting point, to $\theta_j = 1$, which is indeed the solution. In general, we can identify the shape parameters $a_j^\star$ and $b_{ij}^\star$ that make the marginal benefit of operating $\theta_j$ constant by looking at the values $a_j$ and $b_{ij}$ for which $\theta_j$ drops out of (12). Simple algebra shows that

$$a_j^\star = \frac{1}{\sigma - 1} \quad \text{and} \quad b_{ij}^\star = \frac{1}{\varepsilon_j - 1} - \frac{1}{\sigma - 1}. \tag{$\star$}$$

Under these values, the first-order condition for $\theta_j$ does not directly depend on $\theta_j$ itself and, as a result, any solution to that first-order condition must be such that $\theta_j \in \{0,1\}$, as long as the aggregates also present in (12) are independent of $\theta_j$. In practice, this logic will also work as long as any dependence is weak. Indeed, if the impact of $\theta_j$ on the marginal benefit curve is small, it is likely to never cross the marginal cost curve, in which case the solution $\theta_j$ will be at a corner. I will show in numerical simulations below that in practice the solutions to $\mathcal{R}$ frequently belong to $\{0,1\}^n$ under ($\star$) even in economies with only a few firms.

Intuitively, we can understand ($\star$) by noticing that by construction $a_j$ and $b_{ij}$ affect the *marginal* benefit of increasing $\theta_j$ but not the *average* benefit of moving $\theta_j$ discretely from 0 to 1, which us the relevant quantity for the planner. Indeed, if we focus on the first term in (12), we can compute its integral as[13]

$$\int_0^1 a_j \frac{\lambda_j(\theta_j)}{\theta_j} c_j(\theta_j) \mathrm{d}\theta_j = \frac{1}{\sigma - 1} \lambda_j(1) c_j(1),$$

which is independent of $a_j$ and $b_{ij}$. But notice that this average product is also equal to the first term in (12) under ($\star$). A similar reasoning applies to the second term in (12). It follows that the shape parameters ($\star$) are the ones that make the marginal and average benefits of operating $j$ equal to each other. They therefore align the local incentives to move $\theta_j$ at the margin with the global incentives to operate firm $j$ or not, and lead to the global maximum.

**Sufficiency of the first-order conditions**

Proposition 1 shows that if $\mathcal{R}$'s solution belongs to $\{0,1\}^n$ then it must also solve $\mathcal{P}$. In the context of Lemma 3, we have described conditions under which any point that satisfies $\mathcal{R}$'s first-order conditions belongs to $\{0,1\}^n$. The global maximum of $\mathcal{R}$ does satisfy the first-order conditions, but other points might as well. In this section, we derive sufficient conditions on primitives such that first-order conditions are satisfied *only* at the global maximum. In this case, standard algorithms can solve $\mathcal{R}$ readily.

The next result shows that the first-order conditions are sufficient to characterize $\mathcal{R}$'s solution when the heterogeneity across firms is limited and the matrix $\Omega$ has rank one.

**Proposition 2.** *Let $\varepsilon_j = \varepsilon$ and $\alpha_j = \alpha$ for all $j$. If $\Omega_{ij} = d_i e_j$ for some vectors $d$ and $e$ then the Karush-Kuhn-Tucker conditions are necessary and sufficient to characterize a solution to $\mathcal{R}$.*

Under the conditions of the proposition, $\mathcal{R}$'s objective function can be written in closed form and we can verify that it is strictly concave under ($\star$), which leads to the result.

A similar result holds for a different set of $\Omega$ matrices. Define $\bar{\Omega} = \omega(O_n - I_n)$ where $O_n$ is the $n \times n$ matrix full of ones, $I_n$ is the $n \times n$ identity matrix and $\omega > 0$. The matrix $\bar{\Omega}$ describes a network of potential connections in which firms are connected to each other, but not to themselves, with the same intensity $\omega$. The following result shows that $\mathcal{R}$ is easy to solve when $\Omega$ is close to $\bar{\Omega}$.

**Proposition 3.** *Let $\sigma = \varepsilon_j$ for all $j$. Suppose that the $\{\beta_j\}_{j\in\mathcal{N}}$ are not too far from each other and that the matrix $\Omega$ is close enough to $\bar{\Omega}$. Then there exists a threshold $\bar{f} > 0$ such that if $f_j < \bar{f}$ for all $j$ the Karush-Kuhn-Tucker conditions are necessary and sufficient to characterize a solution to $\mathcal{R}$.*[14]

[13] In computing this integral, all aggregate quantities are taken as independent of $\theta_j$,

[14] To be precise, let $\bar{\beta}$ be a $n \times 1$ vector with identical elements. Then there exists a ball $\mathcal{B} = \{(\Omega,\beta) : \|(\Omega,\beta) - (\bar{\Omega},\bar{\beta})\| < \nu\}$ for some $\nu > 0$ such that the statement holds for $(\Omega,\beta) \in \mathcal{B}$.

Propositions 2 and 3 establish sufficient conditions under which a feasible point $\theta^*$ that satisfies the first-order conditions and the complementary slackness condition solves $\mathcal{R}$. As a result, standard algorithms, such as gradient ascent, can rapidly solve $\mathcal{R}$ even for economies with thousands of firms. If that solution belongs to $\{0,1\}^n$, we know that it also solves $\mathcal{P}$ by Proposition 1. We will discuss in the next section that the key to Propositions 2 and 3 is that the restrictions they impose on $\Omega$ rule out isolated groups of firms that can create local maxima in the objective function.

### 3.3 Example with two firms

To better understand how the reshaping procedure works, consider a simple economy with two firms $j \in \{1,2\}$ and a complete set of potential connections between them ($\Omega = O_2$). The objective function $V(\theta)$ of the relaxed planner's problem without any reshaping ($a_j = 1, b_{ij} = 0$) is shown in Figure 4a, where warmer colors represent higher utility levels for the household. The horizontal and vertical axes refer to the operating decisions $\theta_1$ and $\theta_2$. We see that $V$ is shaped like a saddle with local maxima at $(\theta_1, \theta_2) = (1,0)$ and $(0,1)$, and local minima at $(0,0)$ and $(1,1)$. The global maximum is at $(1,0)$.

Since $V$ is not concave, first-order conditions are not sufficient to characterize the global maximum–they are indeed satisfied at both $(0,1)$ and $(1,0)$. As a result, this problem cannot be solved reliably with standard algorithms. Starting from an initial point, these algorithms move locally by following the steepest slope, so they can easily converge to the local maximum at $(0,1)$.

Figure 4b shows the objective function $V_R(\theta)$ of the same optimization problem but, this time, reshaped according to ($\star$). Three things are worth noticing. First, $V$ and $V_R$ coincide, by construction, at the corners $\{0,1\}^2$, such that the ranking of these corners, in terms of utility, is the same in both problems. Second, the reshaping procedure stretches the objective function so that $V_R$ is concave. The first-order conditions are therefore sufficient to characterize the global maximum. Third, the procedure did not create another maximum somewhere inside $[0,1]^2$. As a result, starting from any initial $\theta_0$ in $[0,1]^2$, a simple gradient ascent algorithm will converge to the global maximum at $(1,0)$. This point also solves $\mathcal{P}$ by Proposition 1.

### 3.4 Numerical tests

The theoretical results of the last section give us *sufficient* conditions under which reshaping the planner's problem provides the solution to $\mathcal{P}$, but these conditions are not necessary. In this section, I show through numerical simulations that the solution approach also works well when these conditions are not satisfied. I first consider economies with only a few firms and then present results for economies with a large number of firms.

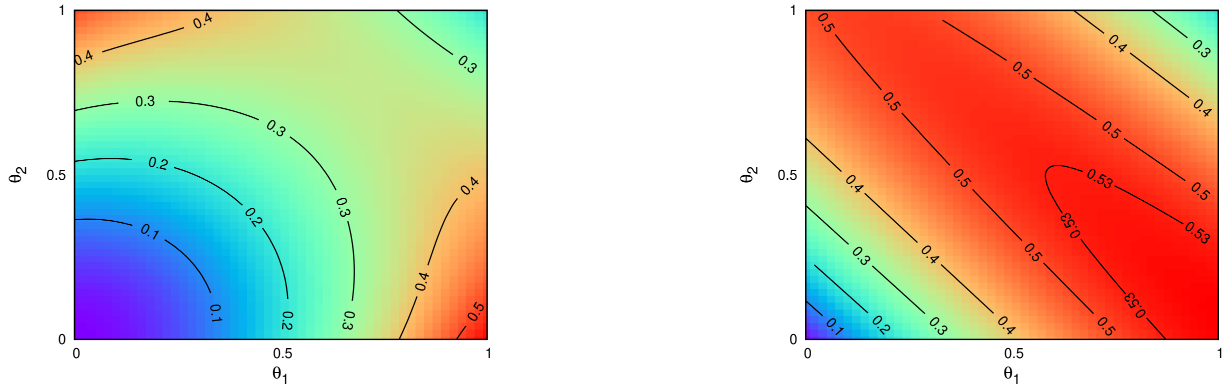


(a) The objective function $V(\theta)$ of the relaxed (but not reshaped) problem is not concave

(b) The objective function $V_R(\theta)$ of the relaxed and reshaped problem is concave

Notes: The parameters are $n = 2$, $\alpha_j = 0.5$, $L = 1$, $f = 0.45$, $\sigma = \varepsilon_j = 5$, $z_1 = 1$, $z_2 = 0.95$, $\Omega_{ij} = 1$ for all $i, j$.

Figure 4: Reshaping the planner's problem in a simple economy

### Economies with few firms

With a small number of firms, we can find the true solution to the planner's problem by comparing the utility provided by the $2^n$ possible vectors $\theta \in \{0,1\}^n$ using the exhaustive search algorithm described above. We can then compare this allocation to the solution of the relaxed problem with and without reshaping. Appendix B provides the details of the simulations. They involve a broad range of economies with firms that differ along all the dimensions of heterogeneity allowed by the model. They also cover matrices $\Omega$ with various shapes and degrees of sparsity.

The results for economies with up to $n = 14$ firms are presented in Table 1. We see that reshaping the planner's problem (first two columns) attributes the correct status $\theta$ to more than 99.9% of the firms across simulations. It also finds output levels that are within 0.001% of their correct values. In contrast, without reshaping the problem (last two columns), over 15% of the firms can be assigned the wrong status $\theta$, and the average error in output can reach above 0.9%, a large number when studying aggregate fluctuations. The table also shows that the performance of the reshaping algorithm stays relatively constant as $n$ increases, in contrast to the non-reshaped solution which performs worse as the number of firms increases.[15]

### Economies with many firms

When $n$ is large, finding the true solution to $\mathcal{P}$ through an exhaustive search would take an infeasibly long time. We can, however, verify whether there exist welfare-improving deviations from the solutions to the relaxed problems. To do so, I change the operational status $\theta_j$ of each firm to see if it improves the utility of the planner. I keep repeating this procedure as long as there are

[15] With reshaping 99.7% of tested economies have the correct $\theta_j$ *for all* $j$. Without reshaping that number is 19.9%.

Table 1: Testing the solution approach for small $n$

| | With reshaping | | Without reshaping | |
|---|---|---|---|---|
| $n$ | Correct $\theta$ | Error in $C$ | Correct $\theta$ | Error in $C$ |
| 4 | 99.9% | 0.000% | 91.5% | 0.502% |
| 6 | 99.9% | 0.000% | 88.1% | 0.692% |
| 8 | 99.9% | 0.000% | 86.5% | 0.791% |
| 10 | 99.9% | 0.001% | 85.2% | 0.855% |
| 12 | 99.9% | 0.001% | 84.5% | 0.903% |
| 14 | 99.9% | 0.001% | 84.0% | 0.928% |

Notes: See Appendix B for the details of the simulations.

deviations to be found. I then compare this deviation-free solution to the original one. The precise algorithm is described in Appendix E.3.

Since this procedure is computationally costly, I only consider economies that follow the calibration of Section 6. The results are presented in Table 2. Again, the reshaping approach performs well. After all the possible deviations are accounted for, more than 99.9% of the firms have kept the same operating status $\theta_j$ and aggregate output has changed by a negligible amount.[16] In contrast, without reshaping more than 30% of the firms are assigned the wrong operating status, and the error in aggregate output amounts to 0.56%. While this test does not guarantee that the reshaping strategy finds the correct efficient allocation, it provides a good indication that there are no obvious mistakes in its solution.

Table 2: Testing the reshaping approach for $n$ large

| | With reshaping | | Without reshaping | |
|---|---|---|---|---|
| $n$ | Correct $\theta$ | Error in $C$ | Correct $\theta$ | Error in $C$ |
| 1000 | 99.9% | $< 0.001\%$ | 66.5% | 0.56% |

Notes: Parameters as in the calibrated economy of Section 6.2. I simulate 100 different matrices $\Omega$ and, for each $\Omega$, draw 100 productivity vectors $z$. I run the procedure described in Appendix E.3 on each of them and report average results. $x < 0.001\%$ indicates that $x > 0$ but proper rounding would yield 0.

Appendix B provides additional exercises to test the robustness of the solution method. It considers economies 1) with very sparse matrices $\Omega$, 2) in which the production network is created through individual link formation, and 3) for which the solution to $\mathcal{R}$ is not in $\{0,1\}^n$. These additional tests show that the solution method performs well in a broad set of economic environments.

[16] When the reshaping approach fails it is often because it gets the wrong operating status for a firm that is fairly isolated from the rest of the network. Since these firms are in general small, they only have little influence on aggregate production, which explains why the error in output is very small in Table 2.

# 4 Reshaping and equilibria

The previous section described how we can reshape the planner's problem to find the efficient allocation. In the current section, we will instead consider equilibrium allocations and focus on two questions. First, can we decentralize the efficient allocation as an equilibrium and, second, can we use the reshaping strategy to find equilibrium allocations that are inefficient? In reality, distortions like market power or coordination failure might lead to inefficiencies, and it would be useful if the solution method could also find these allocations.

In a production network setting, the properties of an equilibrium depend on how extensively firms are allowed to interact with each other. At one extreme, we can allow for rich interactions, even between firms that are far apart in the network. This is the idea behind the notion of a *stable equilibrium*. In a nutshell, firms are facing contractual obligations to purchase and deliver goods, and a stable equilibrium is an allocation in which no groups of firms want to deviate from the terms of their contracts. Rich interactions between producers are allowed, implying that firms can internalize any externalities that they impose on each other and, as a result, stable equilibria are efficient.[17] This equilibrium concept therefore provides a decentralization of the efficient allocation as the outcome of individual decisions and market forces. The details of that equilibrium definition and the formal result about efficiency can be found in Appendix C.

In contrast, we can also think of a different equilibrium concept, closer to the standard monopolistic competition benchmark, in which firms simply maximize their individual profit without internalizing their impact on other producers. To explore both efficient and inefficient equilibria in that setting, I consider two versions of that equilibrium that differ in how prices are set in firm-to-firm transactions. In the first version, firms have some amount of market power, and offsetting subsidies are assumed to be in place. I show that the equilibrium can be efficient in this case, and so I refer to this version as the *undistorted equilibrium*. In the second version, which I refer to as the *distorted equilibrium*, prices in firm-to-firm transactions are set by a take-it-or-leave-it offer by the purchasing firm. The idea is to capture what might be an important source of distortions in reality: the presence of superstar firms such as Walmart and IKEA that have a lot of pricing power with their suppliers (Bloom and Perry, 2001). In this case, the entry/exit decisions are inefficient, but we will see that the reshaping methodology can also be used as a tool to find such equilibria.

## 4.1 Two equilibrium definitions

I begin by describing some elements that are common to both equilibrium definitions. The representative household owns the firms, supplies labor and purchases consumption goods from individual firms. It takes all prices as given and maximizes the utility function (1) subject to the

[17] Oberfield (2018) uses a version of this equilibrium concept in an economy with a continuum of firms. He shows that an equilibrium that is stable with respect to deviations by countable coalitions, which have measure zero, is efficient.

budget constraint

$$\sum_{j=1}^{n} p_j^c c_j \leq w^e + \Pi + T, \tag{15}$$

where $p_j^c$ is the price of good $j$, $w^e$ is the wage, $\Pi$ is profit and $T$ is a lump-sum transfer from the government. The household's maximization problem yields the demand curve for good $j$,

$$c_j = \beta_j C \left(\frac{p_j^c}{P^c}\right)^{-\sigma}, \tag{16}$$

where $P^c = \left(\sum_j \beta_j \left(p_j^c\right)^{1-\sigma}\right)^{\frac{1}{1-\sigma}}$ is a price index that is taken as given by all agents.

The decisions made by the firms take place over two sub-periods. In the first sub-period, each firm $j$ decides whether to pay a fixed cost $f_j$ to operate. In the second sub-period, firms that have paid that cost can produce. They hire labor, purchase intermediate inputs, and sell their good to the household and to other firms.

### Pricing decisions

To describe how prices are set, it is convenient to first define an operating firm's marginal cost of production in the second sub-period. If a firm $j$ purchases input $i$ at a price $p_{ij}^x$, its marginal cost $\delta_j$ is the outcome of the cost-minimization problem[18]

$$\delta_j := \min_{x,l} \sum_{i=1}^{n} p_{ij}^x x_{ij} + w^e l_j, \tag{17}$$

subject to $y_j \geq 1$, where $y_j$ is given by the production function (2).

Under both versions of the equilibrium, firms take into account the demand curve (16) when selling to the household, where $C$ and $P^c$ are taken as given. They therefore have some amount of monopoly power and earn positive profit from these transactions. In contrast, the way prices in firm-to-firm transactions are set varies between the two equilibrium definitions. In an undistorted equilibrium, firms also have some amount of market power when setting these prices. Specifically, firm $j$ selling goods to firm $k$ faces the demand curve

$$x_{jk} = \Omega_{jk} X_k \left(\frac{p_{jk}^x}{P_k^x}\right)^{-\varepsilon_k}. \tag{18}$$

When making decisions firm $j$ takes $k$'s input bundle $X_k = \left(\sum_j \Omega_{jk}^{\frac{1}{\varepsilon_k}} x_{jk}^{\frac{\varepsilon_k - 1}{\varepsilon_k}}\right)^{\frac{\varepsilon_k}{\varepsilon_k - 1}}$ and the price of that bundle $P_k^x = \left(\sum_j \Omega_{jk} \left(p_{jk}^x\right)^{1-\varepsilon_k}\right)^{\frac{1}{1-\varepsilon_k}}$ as given.

[18] Prices $p_{ij}^x$ are allowed to depend on the buyer $j$ for now but this will not be the case in equilibrium.

We can then write $j$'s problem in the second sub-period as maximizing profits

$$\pi_j^{\text{undist}} = p_j^c c_j + \sum_{k=1}^{n} \left(1 + s_{jk}^x\right) p_{jk}^x x_{jk} - \sum_{i=1}^{n} p_{ij}^x x_{ij} - w^e l_j - w^e f_j \theta_j, \tag{19}$$

subject to a resource constraint

$$c_j + \sum_{k=1}^{n} x_{jk} \leq y_j, \tag{20}$$

and to the demand curves (16) and (18). Since the goal here is to find an equilibrium that coincides with the efficient allocation, I also allow for a subsidy $s^x$ that increases the revenue from firm-to-firm sales, and set it to its efficient level $s_{jk}^x = \frac{1}{\varepsilon_k - 1}$ to compensate for $j$'s market power when selling to $k$. As a result, in an undistorted equilibrium firm $j$ sells to the household at a price $p_j^c = \frac{\sigma}{\sigma-1}\delta_j$ and to an intermediate producer $k$ at a price $p_{jk}^x = \frac{\varepsilon_k}{\varepsilon_k - 1}\frac{1}{1+s_{jk}^x}\delta_j = \delta_j$.[19]

In contrast, in the distorted equilibrium, the price of goods sold to other firms is determined by a take-it-or-leave-it offer made by the customer to the supplier. That offer specifies a price at which any amount of goods can then be purchased. Offers are made taking all other equilibrium prices as given. Profit maximization implies that the supplier is not willing to accept an offer with a price below its marginal cost of production, otherwise each unit sold would reduce profits. As a result, the customer offers a price exactly equal to that marginal cost so as to maximize its own profit. The profit maximization problem of a firm $j$ in the distorted equilibrium is therefore

$$\pi_j^{\text{dist}} = p_j^c c_j + \sum_{k=1}^{n} p_{jk}^x x_{jk} - \sum_{i=1}^{n} p_{ij}^x x_{ij} - w^e l_j - w^e f_j \theta_j, \tag{21}$$

subject to the resource constraint (20), to the household's demand curve (16) and to the fact that $p_{kl}^x = \delta_k$ for all $k, l$. Notice that there are no subsidies in (21). The goal here is to show that the reshaping technique can help to find an inefficient equilibrium, and as a result, I do not introduce subsidies or taxes that could undo any inefficiency.

**Entry decisions**

The entry problem of the firm in the first sub-period is similar under both versions of the equilibrium. A firm $j$ pays the fixed cost $w^e f_j$ if and only if $\pi_j \geq 0$. To compute $\pi_j$, firms use the equilibrium prices to calculate their own marginal cost of production $\delta_j$ conditional on entry, which is given by (17). With $\delta_j$, they can then compute their own prices and the demand for their goods from the household and other firms. As a result, they can also compute $\pi_j$ and make their optimal

[19]These assumptions imply that the household and the firms pay different prices for a given good. While this is a bit unappealing, this pricing structure has the advantage that trades between firms are undistorted by markups. To justify that buyers cannot resell goods to the household (in which case the price difference would disappear), we can assume that each good must be customized by the seller for each specific buyer. Alternatively, goods might be perishable so that they can only be transported once.

entry decision. As with pricing decisions, firms do not internalize the impact of their entry decisions on the demand shifters: $C$ and $P^c$ in (16), and $X_k$ and $P_k^x$ for all $k$ in (18).

We are now ready to define an undistorted equilibrium.

**Definition 1.** An *undistorted equilibrium* is a set of prices $(p^c, p^x, w^e)$ and an allocation $(c, l, x, \theta)$ such that: 1) given input prices and the demand curves (16) and (18), firms pick $(c, l, x, \theta)$ to maximize profit (19) subject to the resource constraint (20); 2) given prices the household maximizes utility (1) subject to (15), where $\Pi = \sum_j \pi_j$ and $T$ pays for the subsidies; and 3) all markets clear.

The definition of a distorted equilibrium is the same, except that firms are not subject to the demand curve (18). Instead, the prices $p_{jk}^x$ are equal to the marginal cost $\delta_j$, given by (17).

### 4.2 Characterizing undistorted and distorted equilibria

I first characterize the equilibrium decisions under an exogenously given vector of entry decisions $\theta$, and show that in this case the equilibrium allocations $(c, l, x)$ under both the undistorted and distorted definitions are efficient. I then consider the equilibrium entry decisions of the firms and compare them to the efficient allocation.

#### Equilibrium decisions under a fixed $\theta$

To better understand the links between an equilibrium and the efficient allocation, it is useful to first characterize the vector of equilibrium unit costs $\delta_j$. As discussed above, the pricing mechanisms imply that $p_{jk}^x = \delta_j$ under both versions of the equilibrium. Together with (17) this implies that[20]

$$\delta_j = \frac{1}{z_j A} \left( \sum_{i \in \mathcal{N}} \Omega_{ij} \theta_i \delta_i^{1-\varepsilon_j} \right)^{\frac{\alpha_j}{1-\varepsilon_j}} (w^e)^{1-\alpha_j} . \tag{22}$$

This equation is essentially the same as (6), which pins down labor productivity $q$ in the efficient allocation. This implies that the equilibrium pricing mechanisms do not introduce wedges that would distort firm-to-firm transactions. As a result, the equilibrium decisions $(c, l, x)$ coincide with the efficient allocation. The following proposition makes this point formally.

**Lemma 4.** *For a given entry decision vector $\theta$, distorted and undistorted equilibria are efficient. Furthermore, the equilibrium prices $w^e$ and $p_{jk}^x$ are equal (up to a choice of numeraire) to the planner's Lagrange multipliers $w$ and $\lambda_j$.*

This proposition establishes a connection between an equilibrium and the efficient allocation conditional on a given $\theta$. It follows that for the whole equilibrium allocation, including the vector $\theta$, to be efficient, entry decisions in the equilibrium and the efficient allocation must coincide. We now move on to characterize these decisions.

[20]See the proof of Lemma 4 for a derivation of this result.

**Entry in an undistorted equilibrium**

Combining the pricing rules with (19), we can write $j$'s profit in an undistorted equilibrium as

$$\pi_j^{\text{undist}} = \frac{\sigma}{\sigma-1}\delta_j c_j + \sum_{k=1}^{n} \frac{\varepsilon_k}{\varepsilon_k - 1}\delta_j x_{jk} - \sum_{i=1}^{n} \delta_i x_{ij} - w^e l_j - w^e f_j \theta_j. \tag{23}$$

We can compare this equation with its equivalent in the efficient allocation. From (12), it is straightforward to show (see Appendix A) that in the reshaped planner's problem we can write the first-order condition for $\theta_j$ as

$$(1+a_j)\,\lambda_j c_j + \sum_{k=1}^{n} (1+a_j+b_{jk})\,\lambda_j x_{jk} - \sum_{i=1}^{n} \lambda_i x_{ij} - w l_j - w\theta_j f_j = \theta_j \Delta\mu_j. \tag{24}$$

Comparing this last equation with (23) we see that both are essentially the same under the shape parameters given by ($\star$) (recall that $\delta_j = \lambda_j$ and $w^e = w$ by Lemma 4). This suggests that entry decisions in the undistorted equilibrium coincide with those of the reshaped planner's problem. The following proposition establishes this result.

**Proposition 4.** *A vector $\theta = \{0,1\}^n$ that satisfies the first-order conditions of the reshaped problem $\mathcal{R}$ with shape parameters ($\star$) is an undistorted equilibrium.*

Three consequences of this proposition are worth highlighting. First, whenever a solution $\theta = \{0,1\}^n$ of the reshaped problem coincides with the efficient allocation, that allocation is also an undistorted equilibrium. As a result, we can think of the efficient allocation as arising from market forces and individual agents interacting through decentralized markets as long as there are corrective subsidies in place. Second, if we solve the reshaped problem numerically and find a point $\theta = \{0,1\}^n$ that satisfies the first-order conditions, this point corresponds to an undistorted equilibrium, even though it might not coincide with the efficient allocation. This is reassuring from a practical point of view: even in a complicated economy that might not satisfy the conditions of Propositions 2 or 3, we can be sure that the allocation provided by the reshaping method is economically meaningful. Third, whenever there is a unique point $\theta^* \in \{0,1\}^n$ that satisfies $\mathcal{R}$'s first-order conditions, there must be a unique undistorted equilibrium. This is the case, for instance, when the conditions of Propositions 2 or 3 are satisfied. Recall that these conditions rule out isolated groups of firms in $\Omega$. To see why such groups can lead to multiple equilibria, consider as an example a large economy with a pair of firms (firms 1 and 2), such that 1's only potential input is 2, and vice versa. If $f$ is small enough, there are two equilibria in this setup: one in which both firms operate, and one in which both firms do not. Indeed, if both firms are not operating, there are no incentives for either of them to deviate and operate. The deviating firm would have access to no input and would not be able to produce. The conditions on $\Omega$ imposed by Propositions 2 and 3 rule out these isolated groups of firms, such that a unique equilibrium remains.

### Entry in a distorted equilibrium

The reshaping methodology can also be used to characterize an equilibrium that is distorted away from the efficient allocation. To see this, we can once again combine the firm's profit (21) together with the pricing rules to write profit in a distorted equilibrium as

$$\pi_j^{\text{dist}} = \frac{\sigma}{\sigma-1}\delta_j c_j + \sum_{k=1}^{n} \delta_j x_{jk} - \sum_{i=1}^{n} \delta_i x_{ij} - w^e l_j - w^e f_j \theta_j. \tag{25}$$

This equation is similar to the profit of the firm in the undistorted equilibrium, given by (23), with the exception that selling one unit of good to another firm brings in only the marginal cost $\delta_j$, instead of $\frac{\varepsilon_k}{\varepsilon_k-1}\delta_j$. These sales therefore generate no profit, and so losing a customer has no direct impact on entry decisions. This has important implications for the propagation of cascades, as we will see in the next section. The absence of a markup in firm-to-firm transactions also implies less profit than in the undistorted equilibrium and weaker incentives to operate. As a result, entry decisions are in general inefficient, but the reshaping method can still be used to characterize the equilibrium. Comparing (25) with the planner's equivalent condition (24) suggests to use the shape parameters

$$a_j^d = \frac{1}{\sigma-1} \quad \text{and} \quad b_{ij}^d = -\frac{1}{\sigma-1} \tag{26}$$

for that purpose instead of those given by ($\star$). The following proposition establishes that result.

**Proposition 5.** *A vector $\theta = \{0,1\}^n$ that satisfies the first-order conditions of the reshaped problem $\mathcal{R}$ with shape parameters (26) is a distorted equilibrium.*

This result shows that reshaping the planner's problem can be a useful tool to find not only the efficient allocation but also equilibrium allocations that are distorted. It also emphasizes the importance of the shape parameters $a_j$ and $b_{ij}$ for that purpose.

Finally, since $b_{ij}^\star \neq b_{ij}^d$ one might wonder if $\mathcal{R}$'s first-order conditions imply that any solution to the reshaped planner's problem belongs to $\{0,1\}^n$ under (26). To answer this question we can go back to Lemma 3 which described the marginal benefit and the marginal cost of increasing $\theta_j$ in the reshaped problem. Since $a_j^\star = a_j^d$, the first term in the first-order condition (12) is still independent of $\theta_j$, and since $a_j^d + b_{jk}^d = 0$ the second term is simply zero. It follows that, as with ($\star$), the first-order conditions will be satisfied at a corner $\{0,1\}$ as long as aggregate quantities are independent of $\theta_j$.

### Some remarks about the equilibrium definitions

Before exploring the forces at work in the economy, some comments about the equilibrium definitions are in order. First, I have assumed that firms behave atomistically when making decisions, in the sense that they take equilibrium quantities and the behavior of other producers as given. This assumption is common in the literature, and might be reasonable for large economies, but with only a small number of firms richer strategic interactions might play a more important role, and deviations from efficiency might occur.[21] Second, I have only considered one source of inefficiency: market power in firm-to-firm transactions. While this is arguably one of the main frictions in these markets, other distortions or market imperfections might certainly be at work in reality. The goal here is not to provide an exhaustive study of these frictions, but rather to investigate how one plausible source of distortions might affect operation decisions. Third, Propositions 4 and 5 cast the complicated problem of finding an equilibrium in an economy with discrete decisions as that of finding a stable point in a continuous optimization problem, which is computationally straightforward. Fourth, under both the undistorted and distorted versions of the equilibrium multiple equilibria can arise. While this multiplicity might be interesting on its own, it also raises the question of how an equilibrium is selected. In contrast, the efficient allocation is in general unique, which makes it a natural benchmark to study. As we have seen, the efficient allocation can also be thought of as arising from market forces, as either an undistorted equilibrium or a stable equilibrium. In what follows, I will therefore emphasize the efficient allocation, but I also highlight some discrepancies with the distorted equilibrium. In the quantitative section of the paper, we will see that both the efficient allocation and the distorted equilibrium behave in comparable ways, which suggests that the dominant economic forces are similar in both allocations.

# 5 Complementarities, cascades and clustering

We now explore how economic forces at work in the environment influence the structure of the production network and the propagation of shocks. Complementarities between the operating decisions of neighboring firms play an important role here. They lead to clustering of economic activity, cascades of firm shutdowns, and are responsible for large reorganizations of the production network in response to small shocks. In what follows, I first describe the origin of these complementarities and then turn to their impact on the economy. Some mechanisms are described in terms of the equilibrium allocation and others in terms of the planner's problem, depending on which perspective is more convenient. In the latter case, the same forces also operate in an undistorted equilibrium (Proposition 4) and in a stable equilibrium (Appendix C).

## 5.1 Upstream and downstream complementarities

In the model, neighboring firms tend to operate, or not, together. To highlight the origin of these complementarities, it is useful to consider how the profit of a firm is affected by the operating decisions of its suppliers and customers in partial equilibrium. In an undistorted equilibrium, the

[21] Note, however, that the stable equilibrium of Appendix C has rich firm-to-firm interactions and is efficient.

operating profit (19) of a firm $j$ under a given vector $\theta$ can be written as

$$\pi_j^{\text{undist}}(\theta) = \frac{1}{\sigma - 1}\delta_j(\theta)\, c_j(\delta_j(\theta)) + \sum_{k=1}^{n} \frac{1}{\varepsilon_k - 1}\theta_k \delta_j(\theta)\, x_{jk}(\delta_j(\theta)) - w^e f_j, \tag{27}$$

where the demand curves are defined, as before, as

$$c_j(\delta_j) = \beta_j C \left(\frac{\sigma}{\sigma - 1}\frac{\delta_j}{P^c}\right)^{-\sigma} \text{ and } x_{jk}(\delta_j) = \Omega_{jk} X_k \left(\frac{\delta_j}{P_k^x}\right)^{-\varepsilon_k}, \tag{28}$$

and where the unit cost $\delta_j(\theta)$ is given by (22). The terms $\delta_j c_j$ and $\delta_j x_{jk}$ in (27) are the cost of goods sold to the household and to other intermediate producers. When adjusted for markups, they contribute to the firm's profit.

We can use these equations to describe how the operating decision of a neighboring firm $i$ affects $j$'s own incentives to operate. In principle, a change from $\theta_i = 0$ to $\theta_i = 1$ would affect several equilibrium quantities such as the wage rate $w^e$ and aggregate consumption $C$. In practice, a single element of $\theta$ will have a negligible impact on these objects in an economy with a large number of firms, and we therefore take them as constant for now. This partial equilibrium analysis allows us to sharply characterize the main forces that affect a firm's operating decisions. The following definition and the notation in (27) and (28) specify which quantities are kept fixed in this exercise.

**Definition 2.** When studying the problem of firm $j$ in *partial equilibrium*, the following quantities are kept fixed: 1) the wage level $w^e$, 2) the demand shifters when selling to the household (aggregate consumption $C$ and the price index $P^c$), 3) the demand shifters when selling to any customer $k$ of $j$ (the input bundle $X_k$ and the price index $P_k^x$), and 4) the unit cost $\delta_i$ of other firms $i \neq j$ conditional on operating.

Equations (27) and (28) capture two channels through which the operating status of a neighbor affects a firm's own operation decision. The first channel operates downstream, from suppliers to customers, and involves the first two terms in (27). Suppose that a firm $i$ that is directly upstream from $j$ starts operating. Because of (22), this additional supplier implies that $j$ is able to produce at a lower unit cost $\delta_j$ and to sell at a lower price. Its good thus becomes more attractive and it sells more units. This, in turn, leads to a higher profit (both $\delta_j c_j$ and $\delta_j x_{jk}$ in (27) are decreasing functions of $\delta_j$) and $j$ is more likely to operate as a result.

There is also a second channel that operates upstream, from buyers to suppliers. The second term in (27) captures the importance for $j$'s profit of the demand from other producers. If a firm $k$ that is directly downstream from $j$ begins operating ($\theta_k = 1$) this creates more demand for $j$'s goods which results in a higher profit and, as a result, more incentives to operate.

These complementarities work differently in the distorted equilibrium. In that case, we can write

$j$'s profits as

$$\pi_j^{\text{dist}}(\theta) = \frac{1}{\sigma - 1}\delta_j(\theta) c_j(\delta_j(\theta)) - w^e f_j, \tag{29}$$

where the key difference is that, since the price in firm-to-firm transactions is equal to the marginal cost of production, the firm does not receive any profit from selling to other intermediate producers. This has important implications for the complementarities in operating decisions. Indeed, the second channel mentioned above, which works through the demand of intermediate producers, is absent here, and so the operating status of a firm $k$ that is downstream from $j$ has no direct impact on $j$'s decision to operate. Only $j$'s suppliers, through their impact on $\delta_j$, have any direct effect.

The following proposition formalizes this discussion.

**Proposition 6.** *In the partial equilibrium analysis of firm $j$ (Definition 2) operating a firm that is directly upstream or downstream from $j$ increases $\pi_j^{undist}$, and operating a firm that is directly upstream from $j$ increases $\pi_j^{dist}$.*

The complementarities described in this proposition suggest that immediate neighbors (customer-supplier pairs) tend to operate (or not) together with, as we will see, implications for the shape of the production network and the emergence of cascades of firm shutdowns. In addition to these complementarities, the model also features some standard substitution forces that apply between two suppliers of the same firm. Indeed, since the elasticity of substitution $\varepsilon_j$ between inputs is greater than one, the exit of one of firm $j$'s suppliers can increase the incentives for another of $j$'s suppliers to operate. These forces tend, however, to be somewhat more diffuse than the complementarities since they only apply to second neighbors in the production network. Intuitively, a typical firm has only a few direct neighbors but many second neighbors, so changes affecting a firm's second neighbors generally have a more limited impact on the firm itself. Nonetheless, substitution forces do influence the mechanisms of the model, as we explore below.

## 5.2 The structure of the production network

In the efficient allocation, the productivity vector $z$ interacts with the complementarities to influence the structure of the production network. In this section, I first focus on the role played by $z$ and then explore how the complementarities can lead to clustering of economic activity. Finally, I describe how the elasticities of substitution influence these mechanisms.

### Productivity, connections and operating status

The following proposition describes how the productivity of a firm and the set of its potential connections influences its operating status.

**Proposition 7.** *In the efficient allocation, the following holds.*

1. The operating decision $\theta_j(z)$ of firm $j$ is weakly increasing in $z_j$.

2. Denote by $\Omega^-$ a network of potential connections and let $\Omega^+ = \Omega^- + \Delta\Omega$, where $\Delta\Omega$ is a nonnegative matrix with potentially positive entries only in its $j$th row and $j$th column. Then, $\theta_{\Omega^+,j}(z) \geq \theta_{\Omega^-,j}(z)$ for all $z$, where $\theta_{\Omega,j}(z)$ denotes the operating decision of firm $j$ under $\Omega$.

The first part of the proposition shows that, unsurprisingly, a firm $j$ is more likely to operate when $z_j$ is higher. Intuitively, for $z_j$ high enough the output that operating $j$ generates more than compensates for the fixed cost $f_j$. The second part shows that the productivity threshold at which operating $j$ is efficient is lower when $j$ is better connected with its neighbors in the matrix $\Omega$. Additional downstream connections imply that operating $j$ can benefit the productivity $q$ of more customers, while if $j$ has access to additional suppliers, it would have a higher productivity $q$ upon operating. In both cases, the benefit of operating $j$ is larger.

**Clustering of economic activity**

The complementarities described in Proposition 6 also influence which groups of firms operate in the efficient allocation.

**Proposition 8.** *Let $\mathcal{J} \subset \mathcal{N}$ be a group of firms. Denote by $\theta^+, \theta^- \in \{0,1\}^n$ two operating vectors such that $\theta_j^+ = 1$ and $\theta_j^- = 0$ for $j \in \mathcal{J}$, and $\theta_j^+ = \theta_j^-$ for $j \notin \mathcal{J}$. Denote by $\Omega^-$ a network of potential connections and let $\Omega^+ = \Omega^- + \Delta\Omega$ where $\Delta\Omega$ is a matrix full of zeros except that $\Delta\Omega_{kl} > 0$ for some $k, l \in \mathcal{J}$ . Then*

$$C_{\Omega^+}\left(\theta^+\right) - C_{\Omega^+}\left(\theta^-\right) \geq C_{\Omega^-}\left(\theta^+\right) - C_{\Omega^-}\left(\theta^-\right),$$

*where $C_\Omega(\theta)$ denotes consumption in the efficient allocation under $\Omega$ and $\theta$.*

This proposition shows that the welfare benefit of operating a given group of firms is greater when there are more potential connections between them. Additional connections imply more ways for the complementarities to improve productivity, leading to higher consumption.

One implication of Proposition 8 is that the efficient allocation features clustering of economic activity, meaning that the planner prefers, all else equal, to operate firms that are better connected in the $\Omega$ matrix. An example is helpful to understand how this mechanism works.

**Example.** Consider an economy in which firms are located on a grid, as shown in Figure 5. Firms are identical except for their position in the $\Omega$ network and that the red firm is slightly more productive to break the symmetry.[22] In this case, it is never efficient to have more than one cluster of active firms, as in panel (a). Indeed, by grouping the two clusters together (panel b), the planner provides additional suppliers to some producers, which increases their labor productivity $q$ and, as a result, GDP. Clustering activity is therefore efficient.

[22] The slightly more productive firm is needed to position the cluster of operating firms. Without it, the configuration of panel (b) is still optimal but other solutions exist with the cluster translated on the grid.

In panels (a) and (b), firms are essentially identical but the tendency to cluster remains even when firms differ, for instance in terms of their productivity $z$. Panels (d) and (e) show the efficient allocation in the same economy but with idiosyncratic shocks $z$. In this case, the planner tends to cluster activity around the most productive firm, as in panel (d). If two distant firms have high productivity, organizing multiple clusters might be the optimal way of organizing production, as in panel (e).

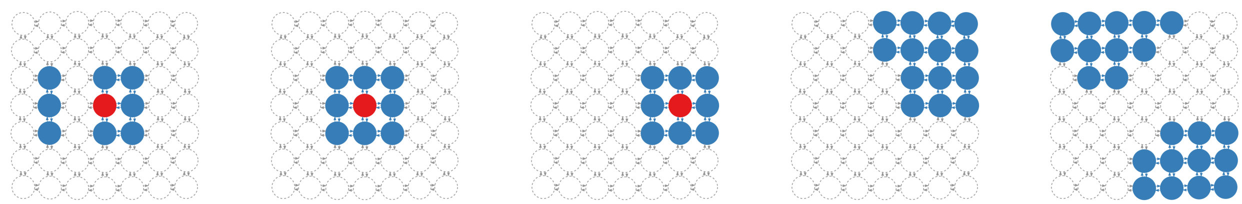

(a) Multiple clusters are inefficient | (b) Efficient clustering | (c) Small change, large reorganization | (d) Random $z$ | (e) Two high-$z$ firms

Figure 5: Clustering in a grid network

**The role of the elasticities of substitution**

The elasticities of substitution in the aggregators for final consumption ($\sigma$) and intermediate inputs ($\varepsilon$) play an important role in shaping the production network. Figure 6 shows the efficient network in four economies that differ only in terms of $\sigma$ and $\varepsilon$. In panel (a) both elasticities are large. Since firms are essentially producing the same good, the planner prefers to concentrate production in the hands of a small group of very productive producers (firms 1 and 2). In panel (b), instead, goods are poor substitutes (small $\varepsilon$) when they serve as intermediate inputs and additional suppliers are more valuable. The planner therefore provides additional inputs to firm 1 to increase its labor productivity $q$. If, instead, $\varepsilon$ remains large but $\sigma$ is small, as in panel (c), goods are poor substitutes in the consumption aggregator. The household prefers a wider variety of products and, as a result, the planner operates producers that are downstream from firm 1. These firms can take advantage of 1's high labor productivity to provide the household with cheap additional goods. When both elasticities are small, as in panel (d), the planner moves on both margins to operate some additional downstream and upstream producers.

## 5.3 The impact of shocks

The operating decisions matter for the structure of the network, but they also affect how productivity shocks propagate through the production network and impact GDP. In this section, we first explore how these shocks can trigger cascades of firm shutdowns, and how even a small shock can lead to an important reorganization of the production network. Finally, we look at how the endogenous reorganization of the network affects aggregate fluctuations.

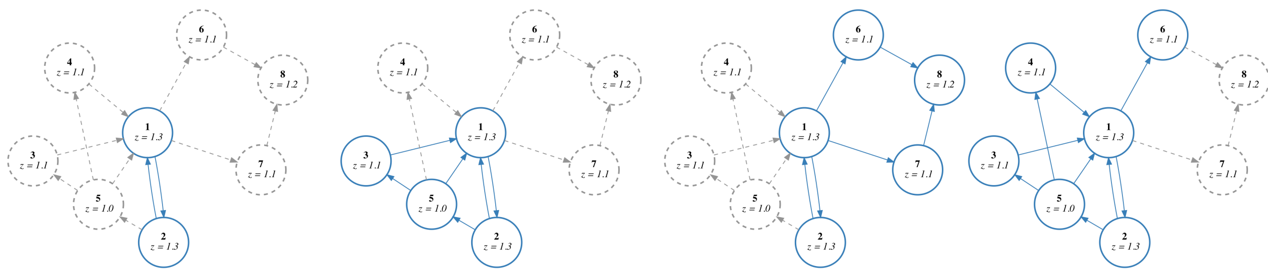

(a) Large $\varepsilon$, Large $\sigma$ (b) Small $\varepsilon$, Large $\sigma$ (c) Large $\varepsilon$, Small $\sigma$ (d) Small $\varepsilon$, Small $\sigma$

Figure 6: The impact of $\sigma$ and $\varepsilon$ on the production network

**Cascades of firm shutdowns**

One immediate consequence of the complementarities in operating decisions is that operating the neighbors of a firm is more beneficial when that firm operates.

**Lemma 5.** *Let $\mathcal{G} \subset \mathcal{N}$ denote the potential neighbors of $j$, that is all firms $i \neq j$ such that $\Omega_{ij} > 0$ and/or $\Omega_{ji} > 0$. There exists a threshold $\bar{f} \geq 0$ such that if $f_j \leq \bar{f}$, then the consumption gain in the efficient allocation from operating $\mathcal{G}$ is larger when $j$ is operating, that is*

$$C\left(\theta_{\mathcal{G}} = 1, \theta_j = 1\right) - C\left(\theta_{\mathcal{G}} = 0, \theta_j = 1\right) \geq C\left(\theta_{\mathcal{G}} = 1, \theta_j = 0\right) - C\left(\theta_{\mathcal{G}} = 0, \theta_j = 0\right), \tag{30}$$

*where $C$ is computed keeping fixed $\theta_i$ for all $i \notin \{j \cup \mathcal{G}\}$.*

This lemma implies that cascades of firm shutdowns can occur in the efficient allocation.[23] Intuitively, if a severe $z_j$ shock pushes $j$ to shut down, its first neighbors, having lost a supplier or a customer, see their operating profit decline and are then more likely to shut down as well. In this case, $j$'s second neighbors are also losing a neighbor and are at a greater risk of shutting down themselves. Since the same logic applies to further neighbors of $j$, the initial $z$ shock can trigger a wave of shutdowns that spreads through the production network, both upstream and downstream.

The magnitude of these cascades depends on the strategic importance of firm $j$ in the production network, as the following lemma shows.

**Lemma 6.** *Let $\Omega^-$ be a network of potential connections and let $\mathcal{G} \subset \mathcal{N} \setminus j$ denote a subset of $j$'s neighbors in $\Omega^-$, that is $\Omega^-_{ij} > 0$ and/or $\Omega^-_{ji} > 0$ for all $i \in \mathcal{G}$. Let $\Omega^+ = \Omega^- + \Delta\Omega$ where $\Delta\Omega$ is a matrix of zeros except for $\Delta\Omega_{ij} > 0$ or $\Delta\Omega_{ji} > 0$ for some $i \in \mathcal{G}$. Denote by $\Delta^{\mathcal{G}}_{\Omega,j} C\left(\tilde{\theta}\right) = C_\Omega\left(\theta_{\mathcal{G}} = 1, \theta_j = \tilde{\theta}\right) - C_\Omega\left(\theta_{\mathcal{G}} = 0, \theta_j = \tilde{\theta}\right)$ the consumption gain in the efficient allocation from operating $\mathcal{G}$ when $\theta_j = \tilde{\theta}$ under $\Omega$, keeping the operating status of all other firms the same. Then the increase in consumption gain from operating $\mathcal{G}$ when $\theta_j = 1$ compared to when $\theta_j = 0$ is*

[23] The restriction $f_j \leq \bar{f}$ is needed to prevent a potential substitution effect between $j$ and $\mathcal{G}$: If $j$ starts operating and $f_j$ is large, the pressure on the labor market might be strong enough that it is welfare-improving to set $\theta_{\mathcal{G}} = 0$.

*greater under* $\Omega^+$ *than under* $\Omega^-$*, that is*

$$\Delta^{\mathcal{G}}_{\Omega^+,j} C(1) - \Delta^{\mathcal{G}}_{\Omega^+,j} C(0) \geq \Delta^{\mathcal{G}}_{\Omega^-,j} C(1) - \Delta^{\mathcal{G}}_{\Omega^-,j} C(0).$$

This result shows that the additional gain from operating $j$'s neighbors when $j$ operates is larger when $j$ is well connected with those neighbors. This suggests that well-connected firms can trigger larger cascades upon exit. Intuitively, if $j$ purchases from many firms or has many customers, its exit affects the profits of many neighbors and is likely to trigger multiple shutdowns. At the same time, Proposition 6 makes clear that firms with multiple neighbors have, all else equal, higher profit and would therefore remain in operation even after large adverse shocks. The model therefore predicts a negative correlation between the likelihood of a firm shutting down and the magnitude of the cascade it triggers upon exit. We will see in the next section that this correlation is visible in U.S. data.

Cascades of shutdowns can also arise in a distorted equilibrium but they propagate differently. In that case, complementarities operate only from supplier to customer (Proposition 6), which implies that the exit of a firm might trigger the exit of its customers, but not the other way around.

Finally, while cascades are driven by the complementarities between suppliers and customers, the substitution forces between a firm's suppliers can provide a stabilizing effect. Indeed, when a cascade leads a firm's supplier to exit, the incentives for other suppliers of that firm to operate increase. As a result, these substitution effects can dampen cascades and limit their reach. As expected, the strength of that stabilizing effect grows stronger as the elasticity of substitution $\varepsilon_j$ increases.

**Small shocks can lead to large reorganizations**

One perhaps unusual feature of the model is that a small change in the environment can trigger a large reorganization of the network. When designing the network, the planner compares the $2^n$ vectors $\theta$ in the set $\{0,1\}^n$ and selects the one providing the highest welfare. As, say, a firm's TFP $z_j$ declines, there is a point at which the planner shuts that firm down. But because of the complementarities between neighbors, it might be better to shut down the whole cluster around that firm and to move production elsewhere. In this case, the network might go through a large reorganization. Figure 5 provides an example. Recall that in panel (b), all firms are identical, except for their potential links in the $\Omega$ network and the fact that the red firm is slightly more productive. In panel (c), another firm (in red) becomes slightly more productive than its peers. While the change in $z$ is negligible, it triggers a large reorganization of the network, with several firms becoming active or inactive, and potentially large variations in firm-level distributions.

**The reorganization of the network and GDP**

The operating decisions of the firms can also influence how the productivity vector $z$ affects GDP $C(z)$. The following result shows that even though $\theta$ can change abruptly in response to a marginal variation in $z$, this response does not translate into an abrupt change in GDP.

**Proposition 9.** *In the efficient allocation, GDP $C(z)$ is a continuous function of $z$.*

Under a fixed $\theta$, the economy is standard and GDP is continuous. It follows that any discontinuity must come from changes in $\theta$. But if a marginal shock in $z$ triggers an abrupt drop in welfare that drop could be avoided by simply keeping $\theta$ constant after the shock. It follows that a discontinuous $C(z)$ cannot be a feature of the efficient allocation.

The previous result implies that a marginal shock in $z$ has a marginal impact on GDP. The following result quantifies the magnitude of that impact.

**Proposition 10.** *In the efficient allocation, at almost all $z$ the marginal impact of $z_j$ on GDP is given by*

$$\frac{d\log C}{d\log z_j} = \frac{\lambda_j y_j}{C}.$$

This proposition implies that at almost every point in the space of $z$ vectors, Hulten's (1978) theorem applies. At those points, the optimal operating decisions $\theta$ are unaffected by shocks to $z$, and the economy behaves as a standard CES production network economy. There are however points at which a marginal change in $z$ leads to a discrete jump in $\theta$. Hulten's theorem does not apply at those points, but their set has measure zero.

Importantly, Proposition 10 applies only to marginal changes in $z$. For larger shocks, the reorganization of the network does matter for GDP, as the next proposition shows.

**Proposition 11.** *Let $\theta^*(z)$ be the efficient allocation under $z$ and let $C(\theta, z)$ be consumption under $(\theta, z)$. Then the response of consumption after a change in productivity from $z$ to $z'$ is such that*

$$\underbrace{C\left(\theta^*\left(z'\right), z'\right) - C\left(\theta^*(z), z\right)}_{\textit{Change in consumption under a flexible network}} \geq \underbrace{C\left(\theta^*(z), z'\right) - C\left(\theta^*(z), z\right).}_{\textit{Change in consumption under a fixed network}}$$

This result shows that the endogenous reorganization of the network amplifies the impact of positive shocks and mitigates the impact of negative shocks. Under a flexible network, the planner is free to reorganize the production network to take advantage of the new productivity vector $z'$. For instance, clusters of firms that were built around formerly productive firms can be dismantled and the freed resources can be reallocated to producers that are now more productive.

# 6 Quantitative exploration

This section provides a basic calibration of the model and shows that it captures salient features of the data such as cascades of firm shutdowns and movements in the structure of the production network over the business cycle. The focus is on the efficient allocation, but I also consider a distorted equilibrium to see how inefficiencies affect the structure of the production network and the propagation of cascades.

## 6.1 Data

The model is calibrated using detailed U.S. data from the Factset Revere Supply Chain Relationships dataset, which provides annual firm-level input-output data from 2003 to 2016. This data is gathered by analysts from 10-Q and 10-K filings, annual reports, investor presentations, websites, press releases, etc. In an average year, the sample includes almost 13,000 private and public firms and more than 40,000 relationships. In that data, about 40% of all link destructions occur at the same time as either the supplier or the customer (or both) stops producing.[24] Figure 7 shows the Factset production network in a typical year.

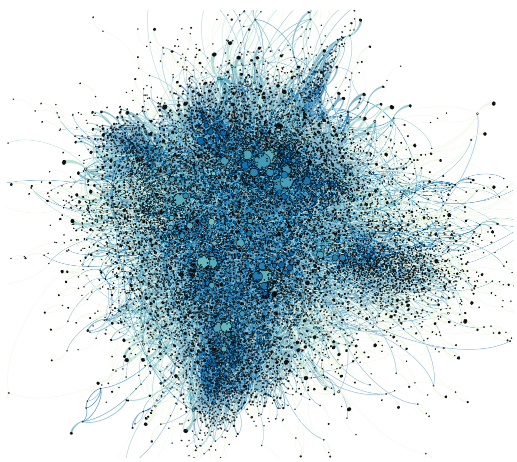

Notes: Vector image; zoom in. Each circle is a firm, and the size of the circle reflects the number of connections. Darker colors denote higher local clustering coefficients. In 2016, Walmart had the largest indegree (448) and Microsoft had the largest outdegree (332). Image generated using Gephi with the Yifan Hu layout. There are 20,702 firms and 62,474 links.

Figure 7: 2016 Factset Revere U.S. firm-level production network

To verify the robustness of some empirical patterns, I also rely on Compustat as another source of annual data. Compustat gathers information from financial statements about a firm's customers that purchase more than 10% of its sales. Since firms are not required to report less important customers, I rarely see a producer supplying to more than 10 clients in the data, and the tail of the outdegree distribution is likely to be artificially thinner as a result. Another limitation of that data is that the names of the customers are self-reported, so General Motors might enter the database as "General Motors", "GM", "General Mtrs", etc. To address this issue, Cohen and Frazzini (2008) (CF) and Atalay et al. (2011) (AHRS), use a combination of automatic algorithms and manual matching to identify each firm and to construct annual production networks. Their samples cover longer periods than Factset (1980 to 2004 and 1979 to 2008, respectively) and might therefore provide a more accurate picture of the evolution of the production network over the business cycle. On the other hand, they also cover fewer firms—about 1,300 firms and 1,500 relationships in an average year.

[24]The analogous exercise for link creations finds a similar number. To remove high-frequency gaps in the data, I assume that a link is created during the first year it appears in the dataset and is destroyed during its last. A firm is considered as shutting down during the last year that it is in the sample.

## 6.2 Parametrization

The model is static but I introduce a time dimension by allowing the idiosyncratic productivity shocks $z_{it}$ to be persistent across periods. Each $z_{it}$ is drawn from an iid AR(1) process with persistence $\rho_z$ and a standard deviation $\sigma_z$ for the ergodic distribution. Foster et al. (2008) and Bartelsman et al. (2013) find that firm-level physical productivity in the United States has a standard deviation of 0.39 and a persistence of 0.81. I set $\sigma_z = 0.44$ and $\rho_z = 0.935$ so that measured TFP in the model matches these targets. Aggregate TFP is normalized to $A = 1$.

There is no consensus in the literature about the cost of overhead labor $f$. Since employment in management occupations is about 5% of total employment, I set $f$ so that $f \times n = 5\%$. For the number of firms, I set $n = 1,000$ as a good trade-off between realism and computation time.[25] Since $n$ is finite, the idiosyncratic productivity shocks generate fluctuations in aggregate variables.

For the share of intermediate goods, I follow Jorgenson et al. (1987) and Jones (2011) and set $\alpha = 0.5$. The empirical literature provides little guidance about the elasticity of substitution between intermediate inputs at the firm level. I therefore rely on Broda and Weinstein (2006) who estimate an elasticity of substitution between product varieties using import data. As these data do not differentiate between items used for consumption and as intermediate inputs, their estimates capture a mix of $\sigma$ and $\varepsilon$. I set $\sigma = \varepsilon = 5$ as an average of their estimates and describe below how changes in these parameters affect the results.

I construct $\Omega$ by assuming that the number of *potential* incoming and outgoing connections, for any given firm, is drawn from a bivariate power law of the first kind. This family of distributions is entirely described by a single shape parameter $\xi$. I set $\xi = 1.78$ so that the distribution of *active* incoming connections generated by the model is close to its empirical counterpart in the Factset data.[26] These two distributions are well approximated by power laws, with an exponent parameter of 0.97 for the empirical distribution (see Section 6.3 below). I therefore target that moment in the calibration. This indirect inference approach ensures that the calibrated economy is consistent with a key feature of the empirical production network. To ensure that the results do not hinge on one particular matrix $\Omega$, I randomly draw 20 different $\Omega$'s and, for each of them, simulate the economy for 100 periods, each involving a different draw from the $z$ shock process.[27] The reported results are averages over these simulations.

[25]See Appendix D.4 for simulations with $n = 20,000$ firms and aggregate shocks. The results are similar.

[26]The probability that a firm has $x_{in}$ and $x_{out}$ inbound and outbound links in $\Omega$ is $\xi(\xi-1)(x_{in}+x_{out}-1)^{-\xi-1}$. The algorithm to construct $\Omega$ is in Appendix E.5. I target moments from Factset, instead of Compustat, as it is the most comprehensive data source for linkages. I also target the indegree distribution as it is less affected than the outdegree distribution by the 10% reporting threshold described in Section 6.1. Since Factset also relies on firms' financial disclosure, it is also affected by the 10% truncation, albeit to a lesser degree than Compustat.

[27]I discard and redraw simulations for which iterating on the first-order conditions does not converge to a point $\theta$ in $\{0,1\}^n$. This rarely happens and, overall, the rejected networks do not look different. Keeping all the simulations in the sample yields very similar results.

### 6.3 Calibrated economy

Table 3 shows how the calibrated network compares to U.S. data. I focus on six key moments to describe the overall structure of the network. The first four moments are the indegree, outdegree eigenvector centrality and sales distributions.[28] In the model and in the data, these distributions are close to power laws so that their exponent parameters provide a good description of the full distributions. These exponents have an important influence on the aggregate impact of idiosyncratic shocks (Acemoglu et al., 2012). The fifth and sixth moments are the global clustering coefficient and the average distance between firms. Both measures describe how tightly firms are connected with each other—a key metric given the importance of clustering for productivity.[29]

We see from Table 3 that the calibrated economy (column 1), despite its simplicity, fits the Factset data (column 4) relatively well, but there are some discrepancies with the Compustat datasets (columns 5 and 6), which is not surprising given their coverage. These discrepancies are particularly large when looking at the outdegree distribution and the clustering coefficient—a likely consequence of the 10% truncation threshold described above.[30]

Table 3: Production network in the calibrated economy and in the data

| | Model | | | Dataset | | |
|---|---|---|---|---|---|---|
| | Calibrated | Random | Inefficient | Factset | Compustat | |
| | | | | | CF | AHRS |
| Power law exponents | | | | | | |
| Indegree distribution | 0.97 | 1.18 | 1.02 | 0.97 | 1.32 | 1.13 |
| Outdegree distribution | 0.92 | 1.15 | 0.95 | 0.83 | 2.22 | 2.24 |
| Centrality distribution | 1.16 | 1.26 | 1.22 | 0.59 | 0.06 | 0.08 |
| Sales distribution | 0.80 | 0.60 | 0.80 | – | 0.54 | 0.54 |
| Global clustering coefficient | 3.45 | 2.08 | 2.99 | 3.46 | 0.09 | 0.08 |
| Average distance | 2.64 | 3.04 | 2.66 | 4.81 | 1.04 | 1.06 |

Notes: To focus on the right tail, the eigenvector centrality and the sales distribution are truncated below the first quartile. Global clustering coefficients are computed on the undirected graph and multiplied by the square roots of the number of nodes. See footnote 29 for details. The average distance and the eigenvector centrality are computed on the undirected graph. "Inefficient" refers to the distorted equilibrium of Section 4. All power law exponents are computed using the estimator of Gabaix and Ibragimov (2011). Since a large fraction of firms in Factset are private, its sales data is sparse and is not reported.

The top panels of Figure 8 show the degree distributions in the model and in the Factset data for 2016, the most recent year in the sample. To highlight the shape of these distributions, the figure

[28] Sales in the efficient allocation are computed as $\lambda_j y_j$.

[29] The global clustering coefficient is computed on the undirected graph. It equals three times the number of triangles (three fully connected nodes) divided by the number of triplets (three connected nodes). In power law graphs, that coefficient declines naturally with $n$. Following Ostroumova Prokhorenkova and Samosvat (2014), I therefore normalize the means of the coefficients by multiplying them by the square root of the number of nodes. This normalization allows for a better comparison of networks across datasets.

[30] For a better comparison with Compustat, we can truncate the model-generated data using the 10% threshold. In that case, the exponent of the indegree, outdegree, centrality and sales distributions are 1.06, 1.08, 1.15 and 0.80. The clustering coefficient is 1.41 and the average distance is 3.14. Recall also that the Compustat data only covers public firms, which tend to be bigger, and so the sales distribution exponent might be artificially smaller.

uses a log-log scale and plots the complementary cumulative distributions (CCDF) on the vertical axis. The roughly linear shapes confirm that they are close to power laws. As we can see, the model fits both distributions reasonably well. The bottom left panel shows that, compared to the data, the model features somewhat too many firms with low centrality and too few firms with high centrality. Since a firm's centrality depends on its indirect connections, it might be possible to improve the fit here by using a more complicated distribution for $\Omega$. Finally, the bottom right panel of Figure 8 shows that the model is also able to capture the broad shape of the sales distribution.

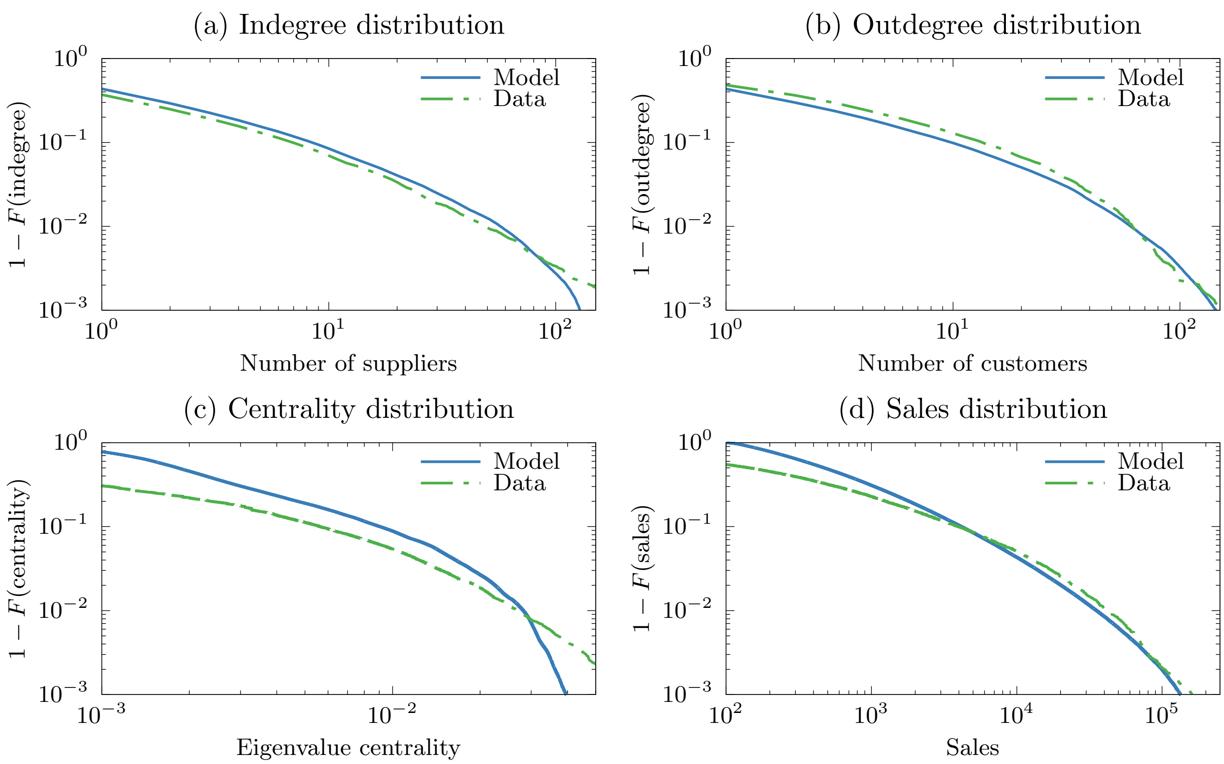


Notes: Efficient allocation compared to the data. The data indegree, outdegree and eigenvector centrality distributions are from Factset in 2016 (last year in Factset). The data sales distribution is from Compustat in 2004 (last year in CF). Eigenvector centralities from the undirected graphs. Model sales and centralities are normalized to match their data averages.

Figure 8: Distributions in the model and the data

### 6.4 Comparison with a random network

To highlight which features of a network are desirable for efficiency, I compare the calibrated network, which has been designed optimally by the planner, to a *random* benchmark built by operating each firm with some probability $p > 0$, where $p$ is set so that both networks have the same number of active firms. All draws are independent, and all quantities, except for the network, are chosen optimally by the planner. Since it is completely random, any discrepancies between this benchmark and the efficient network are design decisions taken by the planner to improve efficiency.

The first two columns of Table 3 show how both networks differ. The power law exponents of the indegree and outdegree distributions are smaller in the efficient network, indicating thicker tails than in the random benchmark. The efficient network therefore features a larger share of highly connected suppliers and customers. For instance, the likelihood that a random firm has more than 25 suppliers in the efficient allocation is 3.2% while the same number in the random network economy is 1.6%. Similarly, the clustering coefficient is also larger, and the average distance is smaller, in the efficient network. These moments highlight that the planner prefers to cluster firms to take advantage of the gains from input variety.

The exponents of the sales distributions suggest, however, a *thinner* tail in the efficient allocation. This is because the random network features many firms that are poorly located in the network. As a result, the planner moves resources, such as labor, away from those firms and toward the few firms that are well connected. It follows that the biggest producers have *relatively* large sales in comparison to their smallest counterparts, which explains why the power law exponent is smaller in the random network economy. For any given threshold in sales, however, the efficient economy features a greater fraction of firms above that threshold.

### 6.5 Cascades of firm shutdowns

We can use the calibrated model to evaluate how cascades of firm shutdowns arise and propagate through the network. To do so, I set the productivity $z$ of a random active firm to zero so that it stops production. I then compute the new efficient allocation and count how many of that firm's neighbors also shut down. Figure 9 shows the outcome of this exercise. The left panel looks at the firm's downstream neighbors, and the vertical axis shows the cumulative number of shutdowns as we move away from the shuttered firm. The right panel provides the same information but for upstream neighbors. The figure also differentiates between cascades originating from an average firm, and from firms with a high number of neighbors (above 99th percentile).

We see that the shutdown of an average firm is likely to only create a small cascade: about 0.6 of its downstream neighbors, and even fewer of its upstream neighbors, shut down. But as we move to high-degree firms, the cascades become larger. For important suppliers about 24 downstream neighbors are wiped by the cascade and the production network is extensively reorganized. Figure 9 also shows that cascades mostly propagate downstream, from supplier to customer, instead of upstream. This is a consequence of the gains from input variety, embedded in equation (6), which make losing a supplier particularly costly in terms of productivity.

The TFP shocks that trigger cascades also lead to declines in GDP. The forces of the model imply a positive correlation between the size of that decline and the size of the cascade. While GDP barely moves after the exit of an average firm, a cascade that originates from a high-degree firm is associated with a 2.4 percent drop in GDP on average. Firms with high outdegrees—the star suppliers—have a disproportionate impact on aggregate output upon shutting down. Since they help to improve the productivity of many producers, their exit lowers the aggregate productivity of the

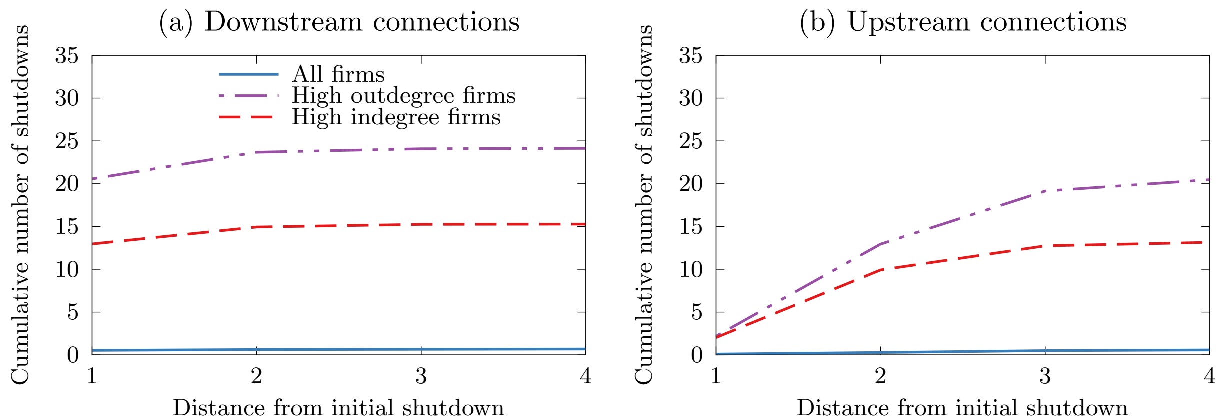


Notes: Cumulative number of exits at different distances from the shuttered firm. "High degree" refers to the firms above the 99th percentile. Simulations of 100 randomly drawn matrices $\Omega$, for each of which 1,000 cascades are created.

Figure 9: Cumulative cascades by degree of originator

network substantially.

Cascades in the model are simply the manifestation of the optimal reorganization of the network after a shock. Preventing cascades (by keeping $\theta$ fixed) must therefore lead to a larger decline in GDP. Appendix D.1 shows that this is indeed the case and that the decline in GDP is larger when inputs are poor substitutes (low $\varepsilon$). Appendix D.3 considers cascades in the distorted equilibrium of Section 4 and shows that they propagate relatively more downstream compared to the efficient allocation. This is because of the difference in complementarities described in Proposition 6. Cascades in the distorted equilibrium are also associated with larger GDP losses. Finally, Appendix D.2 shows that while cascades generally result in a net exit of firms, the substitution forces at work in the model also lead to the entry of some producers amidst these disruptions.

**Cascades in the model and the data**

In the data, many firms are simultaneously hit by (unobserved) shocks and multiple cascades might overlap, so that there is no straightforward way to use the exercise above to evaluate the fit of the model. We can, however, use simple regressions to capture the impact of an exiting firm on its neighbors. Specifically, I compute the fraction of each firm $j$'s neighbors that exit in a given period and regress that number on whether $j$ itself shuts down. I run separate regressions for upstream and downstream neighbors at various distances from $j$. To be precise, denote by $\mathrm{DX}_{jdt}$ and $\mathrm{UX}_{jdt}$ the fraction of firm $j$'s downstream and upstream neighbors located at a distance $d$ that exit between $t$ and $t+1$. I regress

$$\mathrm{DX}_{jdt} = \alpha^D + \beta_d^D \mathrm{Exit}_{jt} + \mathrm{Controls}_{jt} + \varepsilon_{jdt}, \tag{31}$$

and

$$\mathrm{UX}_{jdt} = \alpha^U + \beta_d^U \mathrm{Exit}_{jt} + \mathrm{Controls}_{jt} + \varepsilon_{jdt}, \tag{32}$$

where $\text{Exit}_{jt}$ equals 1 if $j$ exits between $t$ and $t+1$ and 0 otherwise.[31] The coefficients $\beta_d^D$ and $\beta_d^U$ capture the increase in shutdown probability associated with the exit of a neighboring firm located at a distance $d$.

Figure 10 shows the coefficients estimated from the Factset data (green dashed lines). We see that the shutdown of a firm is associated with about a 10% increase in the probability that one of its direct suppliers or customers also exits. This number falls to about 2% for the second neighbors and keeps declining afterward. The model (solid blue lines) is roughly able to match these patterns, suggesting that it broadly captures the joint operating decisions of nearby firms.[32]

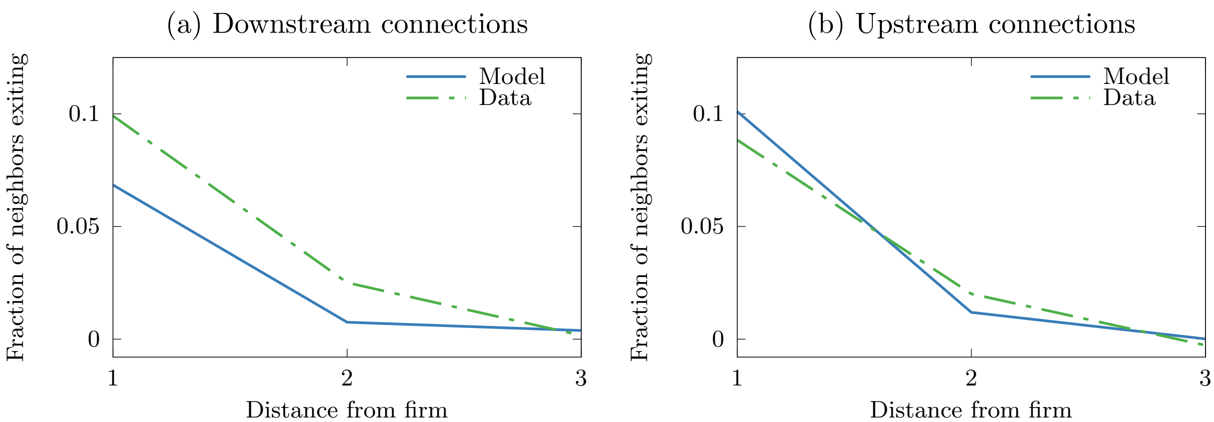


Notes: Factset data. Estimated coefficients from regressing the fraction of exiting neighbors on whether a firm exits. Indegree and outdegree controls are included. The distance is the smallest number of connections between two firms.

Figure 10: Cascades of firm shutdowns in the model and in the data

While the shutdown of any firm can push a neighbor out of business, the exit of well-connected producers generally trigger larger cascades. To see this, we can first measure the size of a cascade as the total number of shutdowns, summed up to the second neighbors, associated with the exit of a firm, and then compare this statistic across firms with different numbers of neighbors.

The first column of Table 4 shows that, in the data, firms that are above the 90th percentile of the degree distribution are associated with cascades that are about three times larger than those associated with the average firm. High-degree firms are, however, less likely to actually shut down in response to shocks, as the third column shows. In the data, an average firm has a 11.8% chance of exiting in a given year, while this number drops to 2.5% for a high-degree firm.

The model does well in terms of the size of the cascades and is also able to roughly replicate the exit probabilities. In the model, high-degree firms are particularly valuable to the planner and are therefore kept in operation even after severe shocks. When they do shut down, however, the

[31] In the data, I consider that a firm shuts down during its last year in the sample. I use SDC Platinum to exclude mergers and acquisitions. The controls in (31) and (32) include the in- and outdegree of firm $j$.

[32] One possibility is that the regressions (31) and (32) capture common shocks across firms instead of the propagation over the network. For instance, since trading partners are likely to be geographically close to each other, a local shock could directly affect both of them at the same time. To alleviate this concern, I run the same regressions on supplier/customer pairs located in different zip codes. Reassuringly, the results are essentially the same.

planner reorganizes the whole cluster of producers that was built around them, which explains the large cascades that they trigger.

Table 4: High-degree firms are more resilient but create larger cascades

| | Size of cascades | | Probability of exit | |
|---|---|---|---|---|
| | Data | Model | Data | Model |
| Average firm | 0.9 | 1.1 | 11.8% | 11.3% |
| High-degree firm | 3.0 | 4.3 | 2.5% | 1.7% |

Notes: "High degree firms" are above the 90th percentile of the degree distribution. "Size of cascades" is the sum of exiting firms up to the second neighbors downstream and upstream, computed by multiplying the regression coefficients in Figure 10 by the number of neighbors at the corresponding distance.

### 6.6 Aggregate fluctuations

There is a finite number of firms in the economy, and as a result, firm-level productivity shocks create aggregate fluctuations. Since these shocks also affect the production network, aggregate output is endogenously correlated with the structure of the network. I investigate that correlation in this section, and I also consider how the endogenous reorganization of the network amplifies or dampens fluctuations in macroeconomic aggregates.[33]

#### Comovements

Table 5 shows the correlation between GDP and the structure of the network in the calibrated economy and the data. In the model the exponent of the degree, eigenvector centrality and sales distributions are negatively correlated with output, which indicates thicker right tails, and thus an abundance of well-connected and high-centrality firms, during expansions. The economy also features more clustering and smaller average distances during booms. These correlations are similar in the data, although there are some discrepancies across datasets. The model is closest to the Factset data, which provides the most comprehensive link coverage.[34]

These patterns can be explained through the lens of the model. When well-positioned firms receive good shocks, the planner builds highly-connected clusters around them. As discussed before, these clusters are particularly productive, which generates the observed correlations between output, clustering and the degree distributions. Inversely, during recessions it might be too costly to organize these productive clusters—perhaps because a few critical firms face low $z$ shocks. As a result, production is more dispersed and output is lower.

[33] One can show that aggregate productivity shocks would have no effect on the network. We can therefore abstract from them when exploring the interaction between GDP and the shape of the network.

[34] If we truncate the model-generated data for a better comparison with Compustat, we find that the correlations between output and the indegree, outdegree, centrality and sales distribution exponents are $-0.48$, $-0.56$, $-0.03$ and $-0.44$. The correlations with the clustering coefficient and the average distance are 0.80 and $-0.76$.

Table 5: Correlation between network moments and GDP

| | Model | | Dataset | | |
|---|---|---|---|---|---|
| | Calibrated | Inefficient | Factset | Compustat | |
| | | | | CF | AHRS |
| Power law exponents | | | | | |
| Indegree distribution | −0.53 | −0.18 | −0.87 | −0.12 | −0.35 |
| Outdegree distribution | −0.63 | −0.46 | −0.97 | −0.11 | −0.31 |
| Centrality distribution | −0.10 | 0.03 | −0.15 | −0.37 | 0.29 |
| Sales distribution | −0.44 | −0.38 | – | −0.24 | 0.04 |
| Global clustering coefficient | 0.60 | 0.40 | 0.76 | 0.11 | 0.18 |
| Average distance | −0.82 | −0.65 | −0.69 | 0.00 | 0.18 |

Notes: All time series are in logs. In the data, output is annual real GDP, detrended linearly in sample. Since there are only 13 years in the Factset data we use the CBO 10-year projection for real GDP growth at the beginning of the sample in 2003 (2.58%) to detrend the series. To focus on the right tail, the eigenvector centrality and sales distributions are truncated below the first quartile. The global clustering coefficient, the average distance and the eigenvector centrality are computed on the undirected graph. Since Factset covers many private firms, its sales data is sparse and is not reported. "Inefficient" refers to the distorted equilibrium of Section 4.

**Level and volatility of output**

The endogenous formation of the network also matters for the level and the volatility of aggregate output. To see how, it is useful to compare the efficient allocation, in which the network is constantly reorganized in response to shocks, to an alternative economy in which the network is designed efficiently in the first period but then kept fixed afterward. The differences between these two economies capture the role played by the endogenous response of the network to shocks.

There are large differences between these two economies. First, aggregate output is 11% lower when the network is kept fixed, which suggests that frictions that might impede the reorganization of the network can have large welfare consequences. Second, aggregate output is 17% more volatile when the network is fixed, which shows the importance of the endogenous network for the aggregation of firm-level shocks into macroeconomic fluctuations.[35,36]

To understand how the reorganization of the network dampens fluctuations, it helps to think of the planner as choosing, for each productivity vector $z$, the best network $\theta$ out of $2^n$ possibilities. GDP $C(z)$ can therefore be written as $C(z) = \max_{k\in\{1,\dots,2^n\}} C_k(z)$ ,where $C_k(z)$ is GDP under the $k$th network. By itself, each network $k$ is associated with a probability distribution for $C_k$ where the randomness comes from the underlying shocks $z$. The mean and the variance of these distributions vary with $k$, but for the networks that are actually selected by the planner the differences are limited and the distributions overlap substantially. Figure 11 provides an example. Each curve represents

[35]These numbers are similar in an economy with a larger number of firms and aggregate shocks (Appendix D.4). This suggests that the network adjustment margin (or the discrete margin) remains important even in large economies. The gaps in $E[C]$ and $V[C]$ between the fixed and flexible networks are larger under higher elasticities $\sigma$ and $\varepsilon$.

[36]When, in addition to the network itself, all the other inputs of the firms are kept fixed, GDP volatility doubles compared to the flexible network benchmark. On its own, the endogenous formation of the network is therefore able to explain about one fifth of the reduction in volatility generated by all the adjustment margins together.

the PDF of GDP under a fixed network $k$. The figure also shows the output produced by each network under four different productivity vectors $z$, indexed by symbols. For instance, we see that network $k = 1$ performs poorly under the ■ shock, while network $k = 5$ performs well. The symbols in blue indicate which network performs the best under a given $z$ and is therefore chosen by the planner. We see that, for any fixed network, the PDFs are spread out and the variance of output is relatively large. In contrast, the output produced by the best network—the blue symbols in the right tails of the distributions—are close to each other indicating that the variance of output $C_{k^*(z)}(z)$ under the efficient network $k^*(z) = \arg\max_k C_k(z)$ is relatively small.[37]

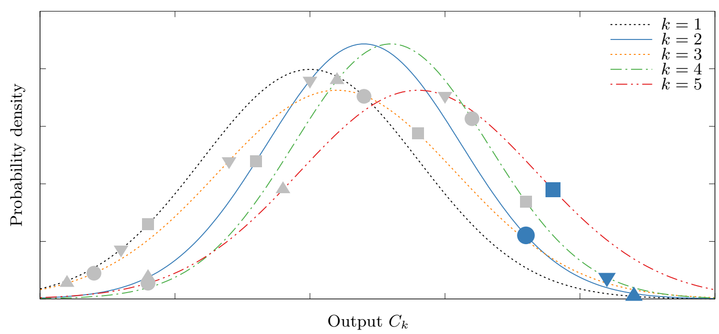


Figure 11: Distribution of output $C_k$ under different networks $k$

## 6.7 The inefficient equilibrium and the data

So far, the quantitative exploration of the model has focused on the efficient allocation, but in reality market power might play a role in shaping the production network. To have a better sense of the quantitative importance of this friction, I now solve for the distorted equilibrium introduced in Section 4. Remember that in that equilibrium entry decisions and the overall structure of the production network are distorted away from the efficient allocation because of market power.

Column "inefficient" in Table 3 shows that the power law exponents of the distributions are roughly similar in the efficient and distorted allocations. It is notable, however, that these exponents are all smaller in the efficient allocation, such that the distorted equilibrium features fewer highly-connected firms. Similarly, firms in the distorted equilibrium are less clustered together and are further away from each other. Given the role of clustering for productivity, this might explain part of the decline in welfare compared to the efficient allocation.

Table 5 describes how the network in the inefficient equilibrium evolves over the business cycle. As we can see, the correlations are similar to those in the efficient allocation. They are however

[37] This intuition is reminiscent of results from extreme value theory that show that the variance of the maximum of $m$ independent normal variables declines with $m$. This result does not hold when the underlying random variables have fat tails. In a network economy in which GDP can have fat tails (Acemoglu et al., 2015; Baqaee and Farhi, 2017a), the flexible network economy might therefore be more volatile than its fixed-network counterparts.

smaller (in absolute value) in the distorted equilibrium, indicating that this network is more rigid and less able to adapt to take advantage of changing economic conditions.

Overall, these tables show that while the efficient and the distorted networks are not identical, the differences are somewhat limited. Both allocations are also similar in terms of macroeconomic aggregates: GDP is 0.9% lower and 0.2% more volatile in the distorted equilibrium. This suggests that the key forces that shape the efficient allocation are at work in the distorted equilibrium—not only qualitatively but also quantitatively. While other distortions might lead to more significant departures from the efficient allocation, the findings of this section suggest that the efficient allocation might be able to provide a reasonable approximation to some distorted economies.

## 7 Conclusion

This paper proposed a theory of network formation that operates through the firms' extensive margin of production. The focus has been on the efficient allocation but the paper also considered the role of market power. Clearly, many other types of externalities, coordination problems or market frictions might be at work in reality. Studying the importance of these distortions and figuring out which ones, if any, matter quantitatively is an important topic for future research. In a framework with richer inefficiencies, it would also be interesting to evaluate which types of policies, such as bailouts or subsidies, could be beneficial.

## References


Acemoglu, D., and P. D. Azar (2018): "Endogenous Production Networks," Working Paper 24116, National Bureau of Economic Research.

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012): "The Network Origins of Aggregate Fluctuations," *Econometrica*, 80, 1977–2016.

Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi (2015): "Microeconomic Origins of Macroeconomic Tail Risks," Working Paper 20865, National Bureau of Economic Research.

Acemoglu, D., and A. Tahbaz-Salehi (2020): "Firms, Failures, and Fluctuations: The Macroeconomics of Supply Chain Disruptions," *Working Paper*.

Atalay, E. (2017): "How Important Are Sectoral Shocks?" *American Economic Journal: Macroeconomics*, 9, 254–80.

Atalay, E., A. Hortaçsu, J. Roberts, and C. Syverson (2011): "Network structure of production," *Proceedings of the National Academy of Sciences*, 108, 5199–5202.

Baqaee, D. R. (2018): "Cascading Failures in Production Networks," *Econometrica*, 86, 1819–1838.

Baqaee, D. R., and E. Farhi (2017a): "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem," Working Paper 23145, National Bureau of Economic Research.

——— (2017b): "Productivity and Misallocation in General Equilibrium.," Working Paper 24007, National Bureau of Economic Research.

Barrot, J.-N., and J. Sauvagnat (2016): "Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks," *The Quarterly Journal of Economics*, 131, 1543–1592.

Bartelsman, E., J. Haltiwanger, and S. Scarpetta (2013): "Cross-Country Differences in Productivity: The Role of Allocation and Selection," *American Economic Review*, 103, 305–34.

Bigio, S., and J. La'O (2016): "Financial Frictions in Production Networks," Working Paper 22212, National Bureau of Economic Research.

Bloom, P. N., and V. G. Perry (2001): "Retailer power and supplier welfare: The case of Wal-Mart," *Journal of Retailing*, 77, 379–396.

Boehm, J., and E. Oberfield (2018): "Misallocation in the Market for Inputs: Enforcement and the Organization of Production," Working Paper 24937, National Bureau of Economic Research.

Bresnahan, T. F., and V. A. Ramey (1994): "Output Fluctuations at the Plant Level," *The Quarterly Journal of Economics*, 593–624.

Broda, C., and D. E. Weinstein (2006): "Globalization and the Gains From Variety," *The Quarterly Journal of Economics*, 121, 541–585.

Caliendo, L., F. Parro, E. Rossi-Hansberg, and P.-D. Sarte (2017a): "The Impact of Regional and Sectoral Productivity Changes on the U.S. Economy," working paper.

Caliendo, L., F. Parro, and A. Tsyvinski (2017b): "Distortions and the Structure of the World Economy," Working Paper 23332, National Bureau of Economic Research.

Carvalho, V. M., , M. Nirei, and Y. Saito (2014): "Supply Chain Disruptions: Evidence from the Great East Japan Earthquake," Working Paper 14-E-035, The Research Institute of Economy, Trade and Industry.

Carvalho, V. M., and A. Tahbaz-Salehi (2018): "Production networks: A primer."

Chahrour, R., K. Nimark, and S. Pitschner (2019): "Sectoral Media Focus and Aggregate Fluctuations," *Available at SSRN 3477432*.

Cohen, L., and A. Frazzini (2008): "Economic Links and Predictable Returns," *The Journal of Finance*, 63, 1977–2011.

Elliott, M., B. Golub, and M. V. Leduc (2020): "Supply network formation and fragility," *Available at SSRN 3525459*.

Foster, L., J. Haltiwanger, and C. Syverson (2008): "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *The American Economic Review*, 98, 394–425.

Gabaix, X. (2011): "The Granular Origins of Aggregate Fluctuations," *Econometrica*, 79, 733–772.

Gabaix, X., and R. Ibragimov (2011): "Rank - 1/2: A Simple Way to Improve the OLS Estimation of Tail Exponents," *Journal of Business and Economic Statistics*, 29, 24–39.

Garey, M. R., and D. S. Johnson (1990): *Computers and Intractability; A Guide to the Theory of NP-Completeness*, New York, NY, USA: W. H. Freeman and Co.

di Giovanni, J., A. A. Levchenko, and I. Mejean (2014): "Firms, Destinations, and Aggregate Fluctuations," *Econometrica*, 82, 1303–1340.

Hatfield, J. W., S. D. Kominers, A. Nichifor, M. Ostrovsky, and A. Westkamp (2013): "Stability and Competitive Equilibrium in Trading Networks," *Journal of Political Economy*, 121, 966–1005.

Hendricks, K. B., and V. R. Singhal (2005): "An Empirical Analysis of the Effect of Supply Chain Disruptions on Long-Run Stock Price Performance and Equity Risk of the Firm," *Production and Operations Management*, 14, 35–52.

Hulten, C. R. (1978): "Growth Accounting with Intermediate Inputs," *The Review of Economic Studies*, 45, 511–518.

Jackson, M. O., and A. Wolinsky (1996): "A Strategic Model of Social and Economic Networks," *Journal of Economic Theory*, 71, 44 – 74.

Jones, C. I. (2011): "Intermediate Goods and Weak Links in the Theory of Economic Development," *American Economic Journal: Macroeconomics*, 3, 1–28.

Jorgenson, D. W., F. Gollop, and B. Fraumeni (1987): *Productivity and U.S. Economic Growth*: Harvard University Press, Reprinted- Universe, 1999.

Kennan, J. (2001): "Uniqueness of Positive Fixed Points for Increasing Concave Functions on $R^n$: An Elementary Result," *Review of Economic Dynamics*, 4, 893 – 899.

Kopytov, A., K. Nimark, B. Mishra, and M. Taschereau-Dumouchel (2022): "Endogenous Production Networks under Supply Chain Uncertainty," *Available at SSRN 3936969*.

Li, D., and X. Sun (2006): *Nonlinear integer programming* Volume 84: Springer Science.

Lim, K. (2018): "Endogenous Production Networks and the Business Cycle," working paper.

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

Murray, W., and K.-M. Ng (2010): "An algorithm for nonlinear optimization problems with binary variables," *Computational Optimization and Applications*, 47, 257–288.

Oberfield, E. (2018): "A Theory of Input-Output Architecture," *Econometrica*, 86, 559–589.

Ostroumova Prokhorenkova, L., and E. Samosvat (2014): *Global Clustering Coefficient in Scale-Free Networks*, 47–58, Cham: Springer International Publishing.

Ozdagli, A., and M. Weber (2017): "Monetary Policy through Production Networks: Evidence from the Stock Market," working paper.

Tintelnot, F., A. K. Kikkawa, M. Mogstad, and E. Dhyne (2018): "Trade and Domestic Production Networks," Working Paper 25120, National Bureau of Economic Research.

Wagner, S. M., and C. Bode (2008): "An Empirical Examination of Supply Chain Performance Along Several Dimensions of Risk," *Journal of Business Logistics*, 29, 307–325.

# Online Appendix

## A Alternative first-order condition

The following lemma proposes a different version of the reshape planner's first-order conditions.

**Lemma 7.** *The first-order condition of problem $\mathcal{R}$ with respect to $\theta_j$ can be written as*

$$(1+a_j)\,\lambda_j c_j + \sum_{k=1}^{n} (1+a_j+b_{jk})\,\lambda_j x_{jk} - \sum_{i=1}^{n} \lambda_i x_{ij} - w l_j - w\theta_j f_j = \theta_j \Delta\mu_j. \tag{24}$$

This equation provides the accounting of the resources that go into the decision to operate firm $j$. The first two terms in (24) capture the value of the goods produced by firm $j$ that go to the household and other producers. The last three terms on the left-hand side correspond to the intermediate inputs and the amount of labor that are needed to operate $j$. It is clear from (24) that $a_j$ and $b_{ij}$ are essentially changing the value of good $j$ in the first-order conditions.

## B Additional numerical tests

This section provides the details of the numerical simulations of Section 3.4 as well as several additional exercises to show the robustness of the solution approach.

### B.1 Details of the simulations of Table 1

The numerical simulations of Table 1 involve a large number of economies that are generated randomly from a broad set of parameters.

**Aggregate parameters.** The aggregate parameters are selected from: $n \in \{4, 6, 8, 10, 12, 14\}$ for the number of firms and $\sigma \in \{4, 6, 8\}$ for the elasticity of the consumption aggregator. The matrix $\Omega$ is such that each firm has on average 3, 4, ... up to $n$ *potential* incoming connections (non-zero $\Omega_{ij}$).[38] I restrict $\Omega$ to have empty diagonals, as in the data. Each non-zero element in $\Omega$ is drawn from $\Omega_{ij} \sim$ iid $U([0,1])$. Appendix E.1 describes the algorithm to build $\Omega$.

**Firm-level parameters.** The firm-level parameters are drawn from: $\log(z_k) \sim$ iid $\mathcal{N}\left(0, 0.25^2\right)$ for the productivities, $f_j \sim$ iid $U([0, 0.2/n])$ for the fixed costs, $\alpha_j \sim$ iid $U([0.25, 0.75])$ for the intermediate input shares, $\varepsilon_j \sim$ iid $U([4, 8])$ for the elasticities between intermediate inputs, and $\beta_j \sim$ iid $U([0,1])$ for the demand shifters.

[38] The corresponding average numbers of *active* incoming connections are 2.1, 3.0, 3.8, 4.5, 5.3 and 5.8, respectively. See Appendix B.3 below for tests on very sparse matrices $\Omega$.

**Procedure.** For every combination of the aggregate parameters, I simulate 500 economies. In each case, the matrix $\Omega$ and the individual characteristics of the firms are drawn from the distributions described above. I then use the exhaustive search algorithm of Appendix E.2 to compute the true solution to $\mathcal{P}$. I also use the algorithm of Appendix E.4 to compute solutions to the reshaped and non-reshaped versions of the planner's problem. These two solutions are then compared to the true solution and the results are reported in Table 1. I exclude from the simulations pathological cases in which the algorithms find an aggregate consumption of 0.[39] For the benchmark tests, an economy is kept in the sample only if the first-order conditions of the reshaped problem yield a solution in $\{0,1\}^n$. Appendix B.2 shows that the algorithm performs well when these simulations are kept in the sample.

## B.2 When the solution to the reshaped problem is not in $\{0,1\}^n$

The results presented in Table 1 exclude economies in which $\mathcal{R}$'s first-order conditions are such that $0 < \theta_j < 1$ for at least one firm $j$, which happens in less than a tenth of the simulations. But even when these cases are not excluded, the solution approach performs well. To see this, Table 6 shows the outcome of the same simulations as Table 1 but without excluding any economies. On average the error in $C$ is less than 0.008% and about 99.7% of firms are assigned the correct status. In contrast, without reshaping the average error in $C$ is 0.867%—108 times more.

Table 6: Testing the reshaping on small networks without exclusions

| | With reshaping | | Without reshaping | |
|---|---|---|---|---|
| $n$ | Correct $\theta$ | Error in $C$ | Correct $\theta$ | Error in $C$ |
| 4 | 99.9% | 0.000% | 91.5% | 0.499% |
| 6 | 99.9% | 0.007% | 88.2% | 0.689% |
| 8 | 99.7% | 0.011% | 86.5% | 0.789% |
| 10 | 99.7% | 0.006% | 85.3% | 0.852% |
| 12 | 99.7% | 0.007% | 84.6% | 0.901% |
| 14 | 99.7% | 0.007% | 84.0% | 0.926% |

Notes: Same simulations as Table 1 but without excluding economies such that $\theta \notin \{0,1\}^n$. See Section B.1 in the Appendix for details.

A similar exercise can be done for economies with a large number of firms. This exercise is analogous to that of Table 2 and is presented in Table 7. We see that even when the first-order conditions of $\mathcal{R}$ yield a solution $\theta \notin \{0,1\}^n$, the error in aggregate output is negligible.

## B.3 Performance with very sparse matrices $\Omega$

Table 8 shows the same simulations as Table 1 but with matrices $\Omega$ that are drawn so that firms have only 1 or 2 *potential* incoming connections on average. As a result, the networks described by

[39] This happens, for instance, when $\Omega$ is so sparse that a closed-loop of suppliers does not exist.

Table 7: Testing the reshaping on large networks without exclusions

| | With reshaping | | Without reshaping | |
|---|---|---|---|---|
| $n$ | Correct $\theta$ | Error in $C$ | Correct $\theta$ | Error in $C$ |
| 1000 | 99.9% | $< 0.001\%$ | 66.5% | 0.56% |

Notes: Same as Table 2 except all simulations are kept in the sample. $x < 0.001\%$ indicates that $x > 0$ but that proper rounding would yield 0, and similarly for $x > 99.9\%$.

the matrices $\Omega$ are extremely sparse. The algorithm still performs well, with an average error in aggregate output that is 66 times smaller than when the problem is not reshaped.

Table 8: Testing the reshaping on sparse networks $\Omega$

| | With reshaping | | Without reshaping | |
|---|---|---|---|---|
| $n$ | Correct $\theta$ | Error in $C$ | Correct $\theta$ | Error in $C$ |
| 8 | 99.6% | 0.008% | 91.0% | 0.508% |
| 10 | 99.5% | 0.007% | 91.1% | 0.496% |
| 12 | 99.6% | 0.007% | 90.8% | 0.503% |
| 14 | 99.5% | 0.009% | 90.5% | 0.531% |

Notes: Same simulations as Table 1 but with matrices $\Omega$ in which firms have on average only 1 or 2 potential connections.

## B.4 Formation of the network link by link

This appendix considers two exercises that show that reshaping the planner's problem is also useful when the production network is constructed link by link instead of through the extensive margin of the firms. In both exercises, the economy contains $m$ real firms that are always active ($f_j = 0$). Any two of these real firms are connected to each other by a link: for any ordered pair of real firms $i, j$ with $i \neq j$, there exists a "link firm" $k$ such that $\Omega_{ik} > 0$ and $\Omega_{kj} > 0$). There are no other connections in $\Omega$. These link firms operate or not as a function of economic conditions.

### Individual link formation in small networks

When the number $m$ of real firms is small, we can use the same approach as in Section 3.4 and find the true solution to the planner's problem by comparing the welfare provided by each possible network $\theta$. There are at most $m\,(m-1)$ links in an economy, in which case the utility provided by $2^{m(m-1)}$ networks must be compared. Since this quantity grows rapidly with $m$, Table 9 shows the results of these tests when there are only $m \in \{3, 4, 5\}$ real firms. As before, the outcome of this exhaustive search is compared to the allocation found by reshaping the planner's problem.

We see from Table 9 that the reshaping algorithm works well. Over all the simulations, more than 99.7% of the links are assigned the proper operating status $\theta$ and the errors in aggregate output are small. Without reshaping, large fractions of the links are assigned the wrong operating status and the error in aggregate output can be sizable.

Table 9: Individual links formation with few firms

| Number of firms | | With reshaping | | Without reshaping | |
|---|---|---|---|---|---|
| Real firms $m$ | Link firms $n-m$ | Correct $\theta$ | Error in $C$ | Correct $\theta$ | Error in $C$ |
| 3 | 6 | 99.9% | 0.001% | 90.9% | 0.25% |
| 4 | 12 | 99.8% | 0.004% | 85.9% | 0.39% |
| 5 | 20 | 99.7% | 0.004% | 82.0% | 0.52% |

Notes: Real firms: $f_j = 0$, $\alpha_j = 0.5$, $\sigma = \varepsilon_j = 6$, $\sigma_z = 0.25$. Link firms: $f_{link} \sim$ iid $U([0.0, 0.1/n])$, $\alpha_{link} \sim$ iid $U([0.5, 1.0])$ and $\sigma_{z_{link}} = 0.25$. For simplicity all non-zero $\Omega_{ij}$ are set to 1. For each $m$, 500 economies are generated randomly and the algorithm of Section E.4 is used to solve the planner's problem. An economy is kept in the sample only if the first-order conditions converge to a point in $\{0,1\}^n$. More than 80% of the economies are kept in the sample.

**Individual link formation in large networks**

For economies with a large number of firms, the true solution to the planner's problem is unknown but we can check whether there exist welfare-improving deviations from the allocation found using the reshaped problem. The procedure is the same as in the Section 3.4. The parameters of the tests are the same as in Table 9 but the economies feature $m \in \{10, 25, 40\}$ real firms and $n \in \{100, 625, 1600\}$ total firms (real plus links). The results are presented in Table 10. Reshaping the planner's problem yields solutions with few welfare-improving deviations so that the vast majority of links are assigned the correct status and the errors in aggregate output are negligible. In contrast, a large fraction of the links are assigned the wrong status and the errors in aggregate output are significant when the problem is not reshaped.

Table 10: Individual links formation with a large number of firms

| Number of firms | | With reshaping | | Without reshaping | |
|---|---|---|---|---|---|
| Real firms $m$ | Link firms $n-m$ | Correct $\theta$ | Error in $C$ | Correct $\theta$ | Error in $C$ |
| 10 | 90 | 99.9% | 0.002% | 76.8% | 0.66% |
| 25 | 600 | $>$ 99.9% | $<$ 0.001% | 74.0% | 0.73% |
| 40 | 1560 | $>$ 99.9% | $<$ 0.001% | 73.4% | 0.74% |

Notes: The parameters of these tests, except for $m$, are as in Table 9. An economy is kept in the sample only if the first-order conditions converge to a point in $\{0,1\}^n$. $x < 0.001\%$ indicates that $x > 0$ but that proper rounding would yield 0, and similarly for $x > 99.9\%$. For each $m$, 500 economies are generated randomly and the algorithm of Section E.4 is used to solve the planner's problem.

One potential concern of using the reshaping method in this context is that the first-order conditions often converge on a vector $\theta$ such that $\theta_j \notin \{0,1\}$ for at least one firm.[40] There are two reasons

[40] In the simulations of Table 10, the first-order conditions converge to a point $\theta_j \in \{0,1\}$ *for all* $j$ in 43% of the simulations for $m = 10$, 12% for $m = 25$ and 5% for $m = 40$.

for this. First, as the total number of firms increases (up to $n = 1600$ for the economies with $m = 40$) it's more likely that at least one firm ends up with $\theta_j \notin \{0, 1\}$. Second, the matrices $\Omega$ considered here are extremely sparse. As a result, the forces pushing the first-order conditions to hit the bounds are weakened. In practice, however, these issues have limited implications. Only a small fraction of the links end up away from the $\{0, 1\}$ bounds, and their impact on aggregate output is minimal. Table 11 shows the outcome of the same simulations but without excluding any simulations. We see that the results are essentially unchanged and that the solution approach also performs well in these situations.

Table 11: Individual links formation with a large number of firms and without exclusions

| Number of firms | | With reshaping | | Without reshaping | |
|---|---|---|---|---|---|
| Real firms $m$ | Link firms $n - m$ | Correct $\theta$ | Error in $C$ | Correct $\theta$ | Error in $C$ |
| 10 | 90 | 99.8% | 0.004% | 76.9% | 0.65% |
| 25 | 600 | $> 99.9\%$ | $< 0.001\%$ | 74.3% | 0.71% |
| 40 | 1560 | $> 99.9\%$ | $< 0.001\%$ | 73.4% | 0.73% |

Notes: The parameters of these tests are the same as in Table 9. No economies are excluded from the sample. $x < 0.001\%$ indicates that $x > 0$ but that proper rounding would yield 0, and similarly for $x > 99.9\%$. For each $m$, 500 economies are generated randomly and the algorithm of Section E.4 is used to solve the planner's problem.

# C Stable equilibrium

In this section, I consider an environment in which firms are facing contractual obligations to purchase and deliver goods. An equilibrium, in that context, is an allocation in which there is no group of firms that want to change the terms of the contracts. This equilibrium concept has proven particularly convenient in network economies (Jackson and Wolinsky, 1996; Hatfield et al., 2013). The approach followed here is most closely related to Oberfield (2018). One of the result of this section is that the efficient allocation can be decentralized as a stable equilibrium.

I first describe the contractual environment. Define a *contract* between two firms $i$ and $j$ as a pair $\{x_{ij}, T_{ij}\}$ where $x_{ij}$ is a quantity shipped from $i$ to $j$, and $T_{ij}$ is a payment from $j$ to $i$. An *arrangement* is a collection of contracts between all possible pairs of firms $\{x_{ij}, T_{ij}\}_{i,j \in \mathcal{N}^2}$.

Under a given arrangement, a firm $j$ must supply and purchase the prescribed quantities, but it can decide on a price $p_j$ to charge the household, an amount $c_j$ to sell to the final good producer, how much labor $l_j$ to employ, and its operating status $\theta_j$. It makes these decisions to maximize profits

$$\pi_j = p_j c_j - w l_j + \sum_{i \in \mathcal{N}} T_{ji} - \sum_{i \in \mathcal{N}} T_{ij} - w \theta_j f_j, \tag{33}$$

where $w$ is the wage, and subject to a technology constraint

$$c_j + \sum_{k \in \mathcal{N}} x_{jk} \leq y_j, \tag{34}$$

where $y_j$ satisfies (2), and to the usual demand curve

$$c_j = \beta_j C \left(p_j / P\right)^{-\sigma} \tag{35}$$

where $P$ is a price index. When making decision each firm takes $w$, $C$ and $P$ as given. In equilibrium, the price index is given by $P = \left(\sum_j \beta_j P_j^{1-\sigma}\right)^{1/(1-\sigma)}$.

An allocation is *feasible* if all the technology constraints (34) and the labor resource constraint $\sum_j l_j + \sum_j \theta_j f_j \leq$ are satisfied.

A *coalition* is a set of firms $J$. All coalitions behave atomistically, in the sense that they take aggregate consumption $C$, the wage $w$ and the aggregate price level $P$ as given. A *deviation* for a given coalition $J$ consists of (i) dropping any contracts that involve at least one firm in $J$ and (ii) altering the terms of any contract involving a buyer and a supplier that are both members of the coalition. Finally, a *dominating deviation* for a given coalition is a deviation that delivers at least the same amount of profit to all members of the coalition and strictly greater profits to at least one member.

We can now define a stable equilibrium in this environment.

**Definition 3.** A stable equilibrium is an arrangement $\{x_{ij}, T_{ij}\}_{i,j \in \mathcal{N}^2}$, firms' choices $\{p_j, c_j, l_j, \theta_j\}_{j \in \mathcal{N}}$ and a wage $w$ such that (i) given the wage, total profit, and prices, the consumption choices $\{c_j\}_{j \in \mathcal{N}}$ maximize the utility of the representative household; (ii) for each $j \in \mathcal{N}$, $\{p_j, c_j, l_j, \theta_j\}$ maximizes the profits of $j$ given the arrangement, the wage, the household's demand and the technology constraint; (iii) labor and final goods markets clear; (iv) there are no dominating deviations available to any coalition; and (v) the equilibrium allocation is feasible.

The following proposition shows how equilibria and the efficient allocation are related.

**Proposition 12.** *Every stable equilibrium is efficient.*

This proposition shows that every equilibrium allocation is a solution to the planner's problem $\mathcal{P}$.[41] As a result, solutions to $\mathcal{P}$ implicitly characterize equilibrium outcomes in this economy.

[41] Since the household is not part of any coalition, the grand coalition does not seek to maximize social welfare.

# D Appendix for Section 6

## D.1 Cascades with $\varepsilon = 3$ and when the network is held fixed

We can use the model to evaluate how cascades propagate in counterfactual economies with different parameters. One parameter with an important role for the cascades is the elasticity of substitution between intermediate inputs, $\varepsilon$. Figure 12 shows the outcome of the same exercise as that of Figure 9 but in an economy in which $\varepsilon = 3$, such that intermediate inputs are worse substitutes than in the calibrated economy. By comparing the figures, we see that the lower elasticity affects the cascades in two important ways: 1) shocks to high degree firms now trigger larger cascades (notice the different scales) and, 2) these cascades have more substantial upstream propagation compared to those in the benchmark economy.

Why cascades are larger when they originate from high-degree firms is easy to understand. With $\varepsilon$ small, intermediate inputs are poor substitutes and losing a supplier has a larger negative impact on a firm's productivity, which leads to more shutdowns. To understand why it also makes cascades have more upstream propagation, it is useful to think about the planner's incentives to operate a firm in this economy. Since the elasticity of substitution in the consumption aggregator, $\sigma$, is relatively large, equation (7) implies that firms with high productivity $q$ are particularly valued by the household. But because $\varepsilon$ is small, these high-$q$ firms are likely to get their high productivity from a large number of suppliers. As a result, if one of the high-$q$ firm shuts down, its many suppliers are no longer useful (they don't contribute much to $Q$) and the planner is likely to shut them down as well, thereby triggering an upstream cascade. In contrast, in the benchmark economy, where $\varepsilon = \sigma = 5$, the planner puts a higher value on the direct contribution to final consumption of these many suppliers, and they are therefore more likely to remain if one of their large customers shuts down.

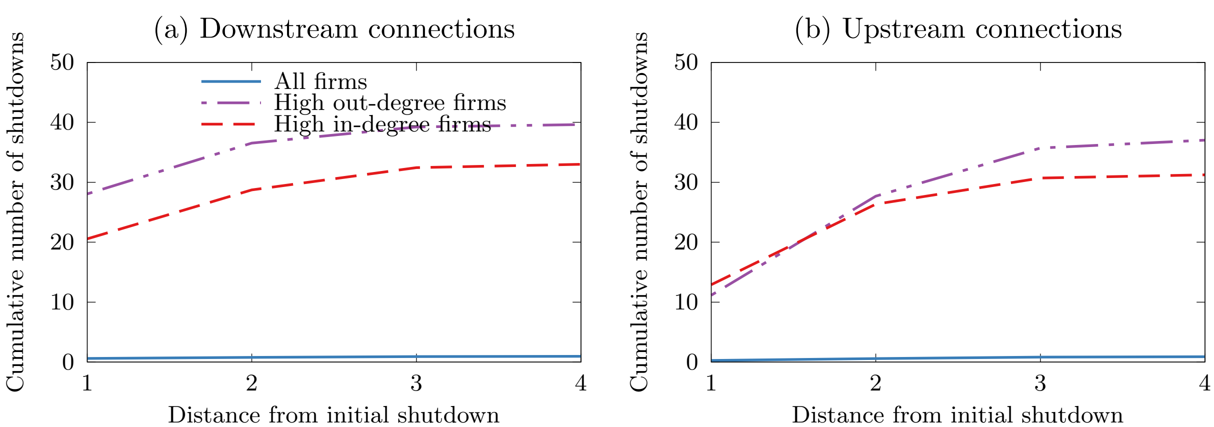


Notes: Cumulative number of exits at different distances from shuttered firm. "High degree" refers to the firms above the 99th percentile. Simulations of 100 randomly drawn matrices $\Omega$, for each of which 1000 cascades are created.

Figure 12: Cumulative cascades by degree of originator, $\varepsilon = 3$.

Table 12 shows that the decline in GDP associated with a cascade is larger when $\varepsilon = 3$. In this case, the loss of an input is more detrimental and leads to a larger decline in $q$. The same table also show that the decline in GDP associated with a shock is larger when the network is kept fixed.

| | Impact on GDP | | | |
|---|---|---|---|---|
| | Flexible network | | Fixed network | |
| | Benchmark | $\varepsilon = 3$ | Benchmark | $\varepsilon = 3$ |
| Average firm | −0.1% | −0.1% | −0.1% | −0.1% |
| High indegree | −1.8% | −4.7% | −1.9% | −4.8% |
| High outdegree | −2.5% | −4.9% | −2.6% | −5.1% |
| High degree | −2.4% | −5.1% | −2.5% | −5.3% |

Notes: "High degree" refers to firms above the 99th percentile. Simulations of 100 randomly drawn matrices $\Omega$, for each of which 1000 cascades are created.

Table 12: Correlation between output drop and firm degree

## D.2 Cascades and firm entry

In the main text, we investigated firm exit during cascades of shutdowns. But since intermediate inputs are substitutes in the production process, the exit of a producer can also trigger the entry of other firms. For instance, if firm $i$ loses one of its suppliers, another of its suppliers might decide to start operating to fill in the gap. To investigate this force, Figure 13 shows the entry of producers around an exiting firm in the calibrated economy. The exercise mirrors the analysis in Figure 9, but instead focuses on firm entry rather than exit and includes all firms around the shuttered one, rather than only those with direct upstream or downstream links. We see that as expected the exit of a producer triggers the entry of firms, with the shutdown of a high out-degree firm leading to the entry of about 13 producers. In contrast, Figure 9 showed that the same shock results in the exit of approximately 44 producers when accounting for both upstream and downstream connections. These results indicate that while substitution forces are clearly present, they are somewhat dominated by the complementarities at work between suppliers and customers.

## D.3 Cascades in the distorted equilibrium

In this section, I consider how cascades propagate in the distorted equilibrium. To do so, I proceed as in the exercise of Figure 9, where I looked at the downstream and upstream propagation that follows from the exit of a single producer. In Figure 14 I report the ratio of downstream to upstream cumulative shutdowns in the efficient allocation (solid blue line) and the inefficient equilibrium (dashed purple line). As we can see, cascades in the equilibrium propagate much more downstream than upstream. This difference is particularly important when we consider all cascades (panel a), but also when we restrict the sample to cascades originating from high-degree firms (panels

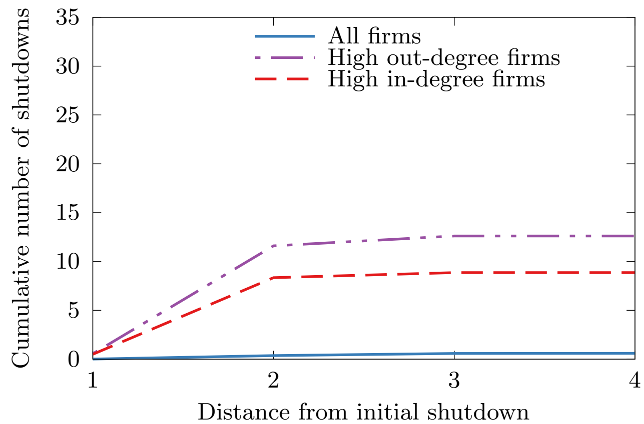


Notes: Cumulative number of entries at different distances from shuttered firm. "High degree" refers to the firms above the 99th percentile. Simulations of 100 randomly drawn matrices Ω, for each of which 1000 cascades are created.

Figure 13: Cumulative entries by degree of originator, all connections.

b and c). Additional downstream propagation is also visible in the efficient allocation, although the effect is less strong here, implying that there is significantly more upstream propagation in the efficient allocation. This is in line with Proposition 6 and the overall discussion in Section 5.3, which describe that because of the pricing distortions the equilibrium features less upstream propagation of shocks.

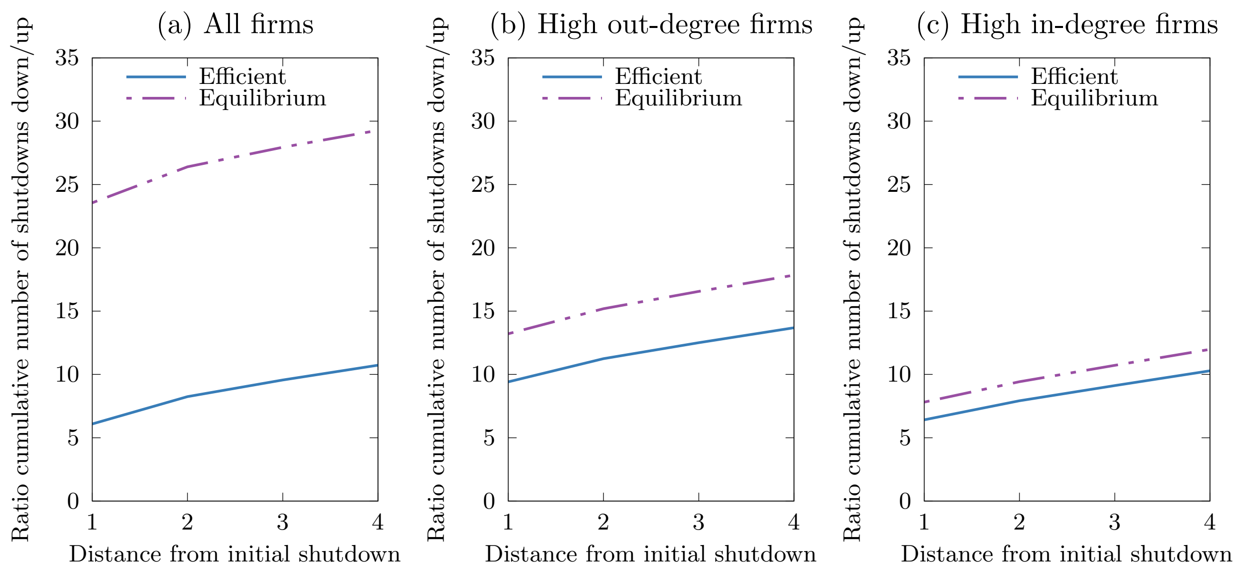


Notes: Cumulative number of exits at different distances from shuttered firm. The figure reports the downstream to upstream ratio of these numbers in the efficient allocation and the distorted equilibrium. "High degree" refers to the firms above the 99th percentile. Simulations of 100 randomly drawn matrices Ω, for each of which 1000 cascades are created.

Figure 14: Ratio of downstream to upstream cumulative cascades in the efficient allocation and the distorted equilibrium

We can also look at the welfare cost of the cascades in the efficient allocation. Table 13 shows the same exercise as Table 12 but for the distorted equilibrium. We see that cascades have a larger impact on output in the distorted equilibrium. In that case, the reorganization of the network after the shock is suboptimal and welfare is adversely affected as a result.

| | Impact on output | |
|---|---|---|
| | Efficient | Equilibrium |
| Average firm | −0.11% | −0.13% |
| High indegree | −1.79% | −1.83% |
| High outdegree | −2.47% | −2.56% |
| High degree | −2.36% | −2.47% |

Notes: "High degree" refers to firms above the 99th percentile. Simulations of 100 randomly drawn matrices $\Omega$, for each of which 1000 cascades are created.

Table 13: Correlation between output drop and firm degree in the efficient allocation and the distorted equilibrium

## D.4 Large number of firms and aggregate shocks

To investigate how the model behaves under a more realistic parametrization, I simulate the calibrated economy with $n = 20,000$ firms (roughly the number of firms in Factset) and with aggregate shocks to total factor productivity $A$. I assume that $\log(A_t)$ follows an AR(1) process with an autocorrelation of 0.9 and and a standard deviation parameter set to match empirical estimates about the impact of aggregate shocks on volatility.[42,43]

Table 14 shows the correlations between aggregate output and the shape of the network. We see that the numbers are broadly similar to those of the benchmark calibration. We can also compute the difference in output volatility between the flexible and fixed networks in this setting. I find that the flexible network economy is about 11% less volatile. Finally, aggregate output is also 13% larger under the flexible network, roughly the same number as in the benchmark economy. This suggest that the importance of the discrete margin of adjustment matters even in large economies.

[42] Atalay (2017) finds that aggregate shocks account for 17% of GDP volatility. I parametrize the stochastic process followed by $\log(A_t)$ to match that estimate.

[43] One can show that we can write the planner's problem as

$$\max_{\theta\in\{0,1\}^n} A^{\frac{1}{1-\alpha}} Q\left(1-\sum_{j\in\mathcal{N}}\theta_j f_j\right)L, \text{ where } q_j = z_j\theta_j\left(\sum_{i\in\mathcal{N}}\Omega_{ij}q_i^{\varepsilon_j-1}\right)^{\frac{\alpha}{\varepsilon_j-1}}, \forall j. \tag{36}$$

Since $A$ only enters as a multiplicative constant in the objective function, it has no impact on the optimal $\theta$ and thus on the production network. (36) also shows that it is straightforward to compute the variance of output with aggregate shocks. If the variance of log output is $x$ without fluctuations in $A$, then the overall variance of log output would simply be $x + (1-\alpha)^{-2}\, y$ if we added shocks with a variance of $y$ to $\log(A)$.

Table 14: Correlation between with aggregate output with $n = 20,000$ firms and aggregate shocks

| Network | Power law exponents | | Clustering coefficient |
|---|---|---|---|
| | indegree | outdegree | |
| Model with $n = 20,000$ firms and aggregate shocks | -0.75 | -0.83 | 0.79 |
| Benchmark model | -0.53 | -0.63 | 0.60 |

Notes: All time series are in logs. The parameters of the economy are as in the benchmark calibration except as mentioned in the text. Since these simulations are computationaly intensive, I simulate four economies instead of twenty in the benchmark exercises.

# E Algorithms

## E.1 Construction of the matrix $\Omega$ for the numerical tests.

This algorithm constructs the matrices $\Omega$ used in the numerical tests of Section 3.4 in the main text and of Sections B.1 and B.2 in the Appendix. Consider an economy with $n$ firms, each with $m$ incoming potential connections on average. Set $p = m/(n-1)$ and $\Omega_{ij} = 0$ for all $i, j \in \mathcal{N}^2$.[44] The steps are 1) draw $\Omega_{ij} \sim$ iid Bernoulli$(p)$, for all $i, j \in \mathcal{N}^2$, 2) for each $i, j \in \mathcal{N}^2$ such that $\Omega_{ij} = 1$, draw $\Omega_{ij} \sim$ iid $U[0, 1]$, and 3) set $\Omega_{ii} = 0$ for all $i \in \mathcal{N}$.

## E.2 Exhaustive search

This algorithm performs an exhaustive search of the $2^n$ vectors $\theta \in \{0, 1\}^n$. It is used in Section 3.4 in the main text as well as in Sections B.1, B.2 and B.4 in the Appendix.

1. Order in an arbitrary way all the possible $\theta \in \{0, 1\}^n$, from $\theta^1$ to $\theta^{2^n}$.
2. For each $p \in \{1, \ldots, 2^n\}$, use equations (6) and (8) to compute the aggregate consumption associated with $\theta^p$.
3. The vector $\theta$ that provides the highest aggregate consumption corresponds to the efficient allocation.

This algorithm is guaranteed to find the global maximum of $\mathcal{P}$ but it is infeasible for large $n$ given the speed at which the number of vectors in $\{0, 1\}^n$ grows with $n$.

## E.3 Deviation-free allocation

This algorithm starts from an allocation $\theta^0 \in \{0, 1\}^n$ and looks for welfare-improving deviations. It is used in Sections 3.4, B.2 and B.4. Note that this algorithm performs a local search and will not in general find the true solution of the optimization problem. Nonetheless we can use it to see if the

[44] The division by $n - 1$ is needed because the diagonal is forced to be empty.

reshaping procedure produces solution with "obvious" mistakes. It is also useful when comparing the reshaped planner's problem with the solution of the non-reshaped procedure, as in Table 2.

1. Initialize the 0-th iteration with $\theta^0$.

2. For the $p$-th iteration, define $\tilde{\theta} = \theta^p$ and set $j = 1$.

   (a) If $\theta_j^p = 0$, set $\tilde{\theta}_j = 1$. If, instead, $\theta_j^p = 1$, set $\tilde{\theta}_j = 0$.

   (b) Using equations (6) and (8) compute the welfare associated with $\tilde{\theta}$.

   (c) If the welfare under $\tilde{\theta}$ is larger than the welfare under $\theta^p$ set $\theta^p = \tilde{\theta}$.

   (d) Set $j = j + 1$, set $\tilde{\theta} = \theta^p$ and repeat steps (a) through (d) until $j = n$.

3. Repeat step 2 above until no welfare-improving deviations are found for some $\theta^p$.

### E.4 Iterating on the first-order conditions

A convenient way to solve the reshaped planner's problem is to iterate on the first-order conditions of the log of the objective function of $\mathcal{R}$ while treating (9) as an inequality constraint. In what follows $\zeta_k$ is the Lagrange multiplier on the $k$-th inequality constraint (9), and $\underline{\mu}_j$ and $\overline{\mu}_j$ are the Lagrange multipliers on the constraint $\theta_j \geq 0$ and $\theta_j \leq 1$. The algorithm is as follows:

1. Initialize the 0-th iteration with $\Delta\mu_k^0 = \underline{\mu}_j^0 - \bar{\mu}_j^0 = -1$ for all $k \in \mathcal{N}$.

2. For the $p$-th iteration:

   (a) Using the complementary slackness condition set $\theta_k^p = 1$ if $\Delta\mu_k^p \leq 0$ and $\theta_k^p = 0$ if $\Delta\mu_k^p > 0$.

   (b) With $\theta^p$, iterate on (9) until convergence to find the vector $q^p$.

   (c) For each $j$, compute $B_j = \left(\sum_{i=1}^n \Omega_{ij} q_i^{\varepsilon_j - 1}\right)^{\frac{1}{\varepsilon_j - 1}}$ and $\Lambda_j = \frac{\theta_j}{B_j^{\varepsilon_j - 1}}$ if $B_j > 0$ and $\Lambda_j = 0$ otherwise.

   (d) Find $\frac{\zeta_k^p q_k^p}{\theta_k^p}$ by solving the following system of linear equations derived from the first-order conditions:

$$\frac{\beta_k \left(A z_k B_k^{\alpha_k}\right)^{\sigma-1}}{\sum_{j\in\mathcal{N}} \beta_j q_j^{\sigma-1}} + \sum_{j\in\mathcal{N}} \left(A z_k B_k^{\alpha_k}\right)^{\varepsilon_j - 1} \Omega_{kj} \alpha_j \Lambda_j \frac{\zeta_j q_j}{\theta_j} = \frac{\zeta_k q_k}{\theta_k}$$

   for each $k$, and where $\frac{\beta_k \left(A z_k B_k^{\alpha_k}\right)^{\sigma-1}}{\sum_{j\in\mathcal{N}} \beta_j q_j^{\sigma-1}}$ should be set to 0 if $\sum_{j\in\mathcal{N}} \beta_j q_j^{\sigma-1} = 0$.

   (e) Compute $\Delta\mu_k$ using the following equation derived from the first-order conditions

$$\frac{f_k}{L - \sum_{j\in\mathcal{N}} f_j \theta_j} = a_k \frac{\zeta_k q_k}{\theta_k} + \sum_{j\in\mathcal{N}} b_{kj} \alpha_j \Omega_{kj} \Lambda_j \left(A z_k B_k^{\alpha_k}\right)^{\varepsilon_j - 1} \frac{\zeta_j q_j}{\theta_j} + \Delta\mu_k$$

for each $k \in \mathcal{N}$ and update to $\Delta\mu^{p+1} = \psi\Delta\mu + (1-\psi)\,\Delta\mu^{p}$ where $0 < \psi \leq 1$ is some parameter to control the speed of convergence.

3. Repeat step 2 above until convergence on $\Delta\mu$.

In practice, it is useful to slow down the updating rule by setting $\psi = 0.9$.

Notice that this algorithm imposes that $\theta \in \{0,1\}^n$ at every iteration. When the solution to $\mathcal{R}$ is not in $\{0,1\}^n$, the algorithm does not converge and the status $\theta$ of some firms keeps alternating between 0 and 1. In practice, I stop the algorithm when the distance between $\Delta\mu_k^{p+1}$ and $\Delta\mu_k^{p}$ starts to increase, which usually indicates that there will be no convergence. I then look at the set of firms for which $\theta$ keeps alternating (different sign for $\Delta\mu_k^{p+1}$ and $\Delta\mu_k^{p}$), and then pick the best $\theta \in \{0,1\}$ to maximize the planner's objective function.

### E.5 Construction of the matrix $\Omega$ in the calibrated economy

The matrix $\Omega$ is constructed by assuming that the number of potential incoming and outgoing connections $(x_{in}, x_{out})$, for any given firm, is drawn from a bivariate power law of the first kind $G$ for which the joint density over $(x_{in}, x_{out})$ is $g\,(x_{in}, x_{out}) = \xi\,(\xi-1)\,(x_{in}+x_{out}-1)^{-(\xi+1)}$. The full algorithm to construct the matrix is as follows:

1. Begin with $\Omega_{ij} = 0$ for all $i, j$.

2. For each firm $j$, draw from $G$ a pair $\left(x_{in}^{j}, x_{out}^{j}\right)$ for the number of incoming and outgoing connections for $j$. Redraw until $\sum_j x_{in}^{j} = \sum_j x_{out}^{j}$ so that the total number of incoming connections is equal to the total number of outgoing connections.

3. For each firm $j$, create $x_{in}^{j}$ incoming stubs and $x_{out}^{j}$ outgoing stubs.

4. Randomly match each incoming stub to an outgoing stub. An incoming stub has the same probability of being matched with any outgoing stub. Set $\Omega_{ij} = 1$ where $i$ is the firm associated with the outgoing stub and $j$ is the firm associated with the incoming stub.

5. Since there are no self-links in the data, set $\Omega_{ii} = 0$ for all $i$.

6. Verify that each firm has at least one potential input so that it can potentially produce, otherwise go back to step 1.

## F Proofs

### F.1 Preliminary results

This section contains preliminary definitions and results that are used in the proofs. The proof of Lemma 1 relies on the following definitions from Kennan (2001).

**Definition.** A function $g : \mathbb{R}^n \to \mathbb{R}^n$ is radially quasiconcave ("R-concave") if $g(x) = 0$ and $x > 0$ and $0 \leq \lambda \leq 1$ implies $g(\lambda x) \geq 0$. If (in addition) $0 < \lambda < 1$ implies $g(\lambda x) > 0$, then $g$ is strictly R-concave.

**Definition.** A function $g = (g_1, g_2, \ldots, g_n) : \mathbb{R}^n \to \mathbb{R}^n$ is quasi-increasing if $y_i = x_i$ and $y_j \geq x_j$ for all $j$ implies $g_i(y) \geq g_i(x)$.

The following Lemma is used as an intermediate step to prove Lemma 1

**Lemma 8.** *Denote by $\tilde{\mathcal{N}}$ any subset of $\mathcal{N}$ with $\tilde{n}$ firms and such that $\sum_{i\in\tilde{\mathcal{N}}}\Omega_{ij} > 0$ for all $j \in \tilde{\mathcal{N}}$. The function $g : \mathbb{R}^{\tilde{n}} \to \mathbb{R}^{\tilde{n}}$ defined, for all $j \in \tilde{\mathcal{N}}$, as*

$$g_j(p) = (z_j A)^{\varepsilon_j} \left(\sum_{i\in\tilde{\mathcal{N}}} \Omega_{ij} p_i^{\frac{\varepsilon_j - 1}{\varepsilon_j}}\right)^{\alpha_j \frac{\varepsilon_j}{\varepsilon_j - 1}} - p_j,$$

*is strictly R-concave.*

*Proof.* Suppose that there exists a $p^* > 0$ such that $g(p^*) = 0$. Then, for $0 \leq \lambda \leq 1$ and for all $j \in \tilde{\mathcal{N}}$

$$\begin{aligned} g_j(\lambda p^*) &= \lambda^{\alpha_j} (z_j A)^{\varepsilon_j} \left(\sum_{i\in\tilde{\mathcal{N}}} \Omega_{ij} (p_i^*)^{\frac{\varepsilon_j - 1}{\varepsilon_j}}\right)^{\alpha_j \frac{\varepsilon_j}{\varepsilon_j - 1}} - \lambda p_j^*, \\ &\geq \lambda (z_j A)^{\varepsilon_j} \left(\sum_{i\in\tilde{\mathcal{N}}} \Omega_{ij} (p_i^*)^{\frac{\varepsilon_j - 1}{\varepsilon_j}}\right)^{\alpha_j \frac{\varepsilon_j}{\varepsilon_j - 1}} - \lambda p_j^* \geq \lambda g_j(p^*) = 0, \end{aligned}$$

where the first inequality is strict for $0 < \lambda < 1$ since $0 < \alpha_j < 1$ and $\sum_{i\in\tilde{N}}\Omega_{ij} > 0$ by assumption. □

The proof of Proposition 3 relies on the following intermediary lemma.

**Lemma 9.** *Let $F = A - fB$ where $f > 0$, $A$ is the all-one $n \times n$ matrix and $B$ is an $n \times n$ matrix. If $B$ is negative definite on the subspace $S : \sum_{i=1}^n x_i = 0$ then $F$ is positive definite for $f > 0$ small enough.*

*Proof.* The negative definiteness of $B$ on $S$ implies that $x'Bx \leq -d\|x\|^2$ for $x \in S$ and some $d > 0$. We can write any vector $z$ as $z = x + y$ where $x \in S$ and $y \perp S$. Then,

$$\begin{aligned} z'(A - fB)z &= n\|y\|^2 - fx'Bx - fy'By - 2fy'Bx \\ &\geq (n - 1/2)\|y\|^2 + df\|x\|^2 - 2f\|B\|\|x\|\|y\| \end{aligned}$$

for $f$ small enough. For $f$ small enough, this last expression is strictly convex in $(\|x\|, \|y\|)$ with a minimum of 0 at $(0,0)$ or, equivalently, at $z=0$. Since $z'(A-fB)z>0$ for any $z\neq 0$, it follows that $F$ is positive definite. □

The following lemma is useful to show that every stable equilibrium is efficient.

**Lemma 10.** *Let $X, Y \subset \mathbb{R}^n$. Define Problem A as*

$$\sup_{x\in X, y\in Y} f(x,y) \text{ subject to } \sum_j y_j \leq 0$$

*and Problem B as*

$$\sup_{x\in X, y\in Y} g(f(x,y)) - \lambda\left(\sum_j y_j\right)$$

*where $g$ is a strictly increasing function and where $\lambda$ is such that $\sum_j y_j = 0$ at any solution. Suppose that for any solution to Problem A the constraint binds, then Problems A and B have the same solutions.*

*Proof.* Take a point $(x^A, y^A)$ that solves Problem A and such that, since the constraint binds, $\sum_j y_j^A = 0$. Towards a contradiction, suppose $(x^A, y^A)$ does not solve Problem B. Then there is another point $(\tilde{x}, \tilde{y})$ such that $\sum_j \tilde{y}_j = 0$ (by the definition of $\lambda$) and such that $g(f(\tilde{x},\tilde{y})) - \lambda\left(\sum_j \tilde{y}_j\right) > g\left(f\left(x^A, y^A\right)\right) - \lambda\left(\sum_j y_j^A\right)$. Since $g$ is strictly increasing this implies that $f(\tilde{x},\tilde{y}) > f\left(x^A, y^A\right)$ but, since $(\tilde{x},\tilde{y})$ is in the feasible set of Problem A, this implies that $(x^A, y^A)$ was not a solution to Problem A, which is a contradiction. Conversely, take a point $(x^B, y^B)$ that solves Problem B. Then by the definition of $\lambda$ it must be that $\sum_j y_j^B = 0$. Towards a contradiction, suppose $(x^B, y^B)$ does not solve Problem A. Then there is another point $(\tilde{x},\tilde{y})$ such that $\sum_j \tilde{y}_j = 0$ (since the constraint in Problem A binds at the optimum) and such that $f(\tilde{x},\tilde{y}) > f\left(x^B, y^B\right)$. Since $g$ is strictly increasing this implies that $g(f(\tilde{x},\tilde{y})) - \lambda\left(\sum_j \tilde{y}_j\right) > g\left(f\left(x^B, y^B\right)\right) - \lambda\left(\sum_j y_j^B\right)$ so that $(x^B, y^B)$ is not a solution to Problem B, which is a contradiction. □

## F.2 Main results

**Lemma 1.** *In the efficient allocation, the labor productivity vector $q$ satisfies*

$$q_j = z_j \theta_j A \left(\sum_{i\in\mathcal{N}} \Omega_{ij} q_i^{\varepsilon_j - 1}\right)^{\frac{\alpha_j}{\varepsilon_j - 1}}, \tag{6}$$

*for all $j$. Furthermore, there is a unique $q$ that solves (6) such that $q_j > 0$ if $j$ operates and has access to an operating cycle, and $q_j = 0$ otherwise.*

*Proof.* We only need to consider the firms that operate and have access to an operating cycle, as defined in footnote 7, since the problem of the other firms is trivial. The first-order conditions of $\mathcal{P}$

with respect to $l_j$ and $x_{ij}$ are

$$wl_j = \lambda_j (1 - \alpha_j) y_j, \tag{37}$$

$$\lambda_i x_{ij} = \alpha_j \Omega_{ij}^{\frac{1}{\varepsilon_j}} \left(\frac{x_{ij}}{X_j}\right)^{\frac{\varepsilon_j - 1}{\varepsilon_j}} \lambda_j y_j, \tag{38}$$

where $X_j = \left(\sum_k \Omega_{kj}^{\frac{1}{\varepsilon_j}} x_{kj}^{\frac{\varepsilon_j-1}{\varepsilon_j}}\right)^{\frac{\varepsilon_j}{\varepsilon_j-1}}$ is an aggregator of intermediate goods. Notice that these expressions imply that

$$\lambda_j y_j = wl_j + \sum_i \lambda_i x_{ij}. \tag{39}$$

Combining the second first-order condition with the definition of $X_j$, we find that

$$\Lambda_j X_j = \alpha_j \lambda_j y_j. \tag{40}$$

where $\Lambda_j = \left(\sum_k \Omega_{kj} \lambda_k^{1-\varepsilon_j}\right)^{\frac{1}{1-\varepsilon_j}}$ is the shadow price of the bundle $X_j$. The left-hand side corresponds to the social value of $X_j$ while the right-hand side is the social value of $y_j$ times the share of intermediate input $\alpha_j$ in production. Combining (37) and (40) with (2) yields (6).

For future reference, note that we can combine this (40) with (38) to find

$$x_{ij} = \Omega_{ij} X_j \left(\frac{\lambda_i}{\Lambda_j}\right)^{-\varepsilon_j}. \tag{41}$$

Similarly, we can write the social cost share of input $x_{ij}$ in $j$'s input bundle as

$$\frac{\lambda_i x_{ij}}{\Lambda_j X_j} = \frac{\alpha_j \Omega_{ij}^{\frac{1}{\varepsilon_j}} \left(\frac{x_{ij}}{X_j}\right)^{\frac{\varepsilon_j-1}{\varepsilon_j}} \lambda_j y_j}{\alpha_j \lambda_j y_j} = \Omega_{ij}^{\frac{1}{\varepsilon_j}} \left(\frac{x_{ij}}{X_j}\right)^{\frac{\varepsilon_j-1}{\varepsilon_j}}. \tag{42}$$

I follow Kennan (2001) to prove the uniqueness of $q$. Consider the change of variable $p_j = q_j^{\varepsilon_j}$, and let $\tilde{\mathcal{N}}$ be the set of firms that operate and have access to an operating cycle. Let $\tilde{n}$ be the the number of such firms. Clearly, $p_j = 0$ for $j \notin \tilde{\mathcal{N}}$. We can rewrite (6) as a mapping from $\mathbb{R}^{\tilde{n}}$ to $\mathbb{R}^{\tilde{n}}$:

$$p_j = (z_j A)^{\varepsilon_j} \left(\sum_{i \in \tilde{\mathcal{N}}} \Omega_{ij} p_i^{\frac{\varepsilon_j - 1}{\varepsilon_j}}\right)^{\alpha_j \frac{\varepsilon_j}{\varepsilon_j - 1}}, \tag{43}$$

for all $j \in \tilde{\mathcal{N}}$. Denote the right-hand side of (43) by $f_j(p)$ and define $g : \mathbb{R}^{\tilde{n}} \to \mathbb{R}^{\tilde{n}}$ as $g(p) = f(p) - p$. A solution to (43) is therefore a vector $p$ such that $g(p) = 0$. By Lemma 8, $g$ is strictly R-concave. Note also that $g$ is quasi-increasing.

Consider the mapping $h : \mathbb{R} \to \mathbb{R}^{\tilde{n}}$ defined as $h(s) = f(\mathbb{1}_{\tilde{n}} s)$ where $\mathbb{1}_{\tilde{n}}$ is the all-one vector of size

$\tilde{n}$. Then $h(s)$ is strictly concave, strictly increasing and differentiable with $h(0) = 0$, $\lim_{s\to 0} h'(s) = \infty$ and $\lim_{s\to\infty} h'(s) = 0$, in all dimensions. As a result, there exist constants $\bar{p} > \underline{p} > 0$ such that $h(\underline{p}) > \underline{p}$ and $h(\bar{p}) < \bar{p}$. Then, Theorems 3.1 and 3.2 in Kennan (2001) apply: (43) has a unique positive fixed point $p^*$ and there is therefore a unique positive $q^*$ that satisfies (6). It is such that $q_j^* = \left(p_j^*\right)^{\frac{1}{\varepsilon_j}}$ if $j$ operates and has access to an operating cycle, and $q_j^* = 0$ otherwise. Kennan (2001) also shows that the fixed point can be found by iterating on (6). Note also that the proof is essentially unchanged if we use the reshaped equation (9) instead of (6). □

**Lemma 2.** *In the efficient allocation, GDP is given by*

$$C = Q\left(1 - \sum_{j\in\mathcal{N}} \theta_j f_j\right). \tag{8}$$

*Proof.* The first-order condition of $\mathcal{P}$ with respect to $c_j$, where $j$ is an operating firm is is

$$c_j = \beta_j \left(\frac{q_j}{w}\right)^{\sigma} C. \tag{44}$$

Raising both sides to the power $\frac{\sigma-1}{\sigma}$, multiplying by $\beta_j^{\frac{1}{\sigma}}$ and summing across $j$'s we find that

$$w = Q. \tag{45}$$

such that the shadow value of labor is equal to aggregate labor productivity. Next, note that from the resource constraint (4) we can write

$$y_j = c_j + \sum_{k=1}^{n} x_{jk} = \beta_j \left(\frac{q_j}{w}\right)^{\sigma} C + \sum_{k\in\tilde{\mathcal{N}}} \Omega_{jk} X_k \left(\frac{\lambda_j}{\left(\sum_l \Omega_{lk}\lambda_l^{1-\varepsilon_k}\right)^{\frac{1}{1-\varepsilon_k}}}\right)^{-\varepsilon_k},$$

where I have used (44) and (41) for the second line. As before, $\tilde{\mathcal{N}}$ is the set of firms that operate and have access to an operating cycle. Combining with (37), and (40), this expression becomes

$$\beta_j \left(\frac{q_j}{w}\right)^{\sigma} C + \sum_{k\in\tilde{\mathcal{N}}} \frac{\alpha_k}{1-\alpha_k} \frac{\Omega_{jk} q_j^{\varepsilon_k}}{\sum_l \Omega_{lk} q_l^{\varepsilon_k - 1}} l_k = \frac{q_j l_j}{1-\alpha_j}.$$

For $q_j = 0$, both sides of this expression are equal to zero. For $j$ such that $q_j > 0$, we can write

$$0 = \beta_j \left(\frac{q_j}{Q}\right)^{\sigma-1} \frac{C}{Q} + \sum_{k\in\tilde{\mathcal{N}}} \alpha_k \frac{\Omega_{jk} q_j^{\varepsilon_k - 1}}{\sum_{i\in\mathcal{N}} \Omega_{ik} q_i^{\varepsilon_k - 1}} \frac{l_k}{1-\alpha_k} - \frac{l_j}{1-\alpha_j}, \tag{46}$$

for all $j \in \mathcal{N}$. Summing across $j$'s and simplifying yields (8). Note that once $q$ is known we can find $l$ by inverting (46). We can then find $y$ and $x$ using the first-order conditions (37) and (40). For future use, we can write a version of expression (46) for the reshaped problem $\mathcal{R}$, defined in the right-column of Figure 2, through the transformation $\Omega_{ij} \to \Omega_{ij}\theta_i^{b_{ij}(\varepsilon_j-1)}$. We find

$$0 = \beta_j \left(\frac{q_j}{Q}\right)^{\sigma-1} \frac{C}{Q} + \sum_{k\in\tilde{\mathcal{N}}} \alpha_k \frac{\Omega_{jk}\theta_j^{b_{jk}(\varepsilon_k-1)} q_j^{\varepsilon_k-1}}{\sum_{i\in\mathcal{N}} \Omega_{ik}\theta_i^{b_{ik}(\varepsilon_k-1)} q_i^{\varepsilon_k-1}} \frac{l_k}{1-\alpha_k} - \frac{l_j}{1-\alpha_j}. \tag{47}$$

□

**Proposition 1.** *If $\theta^* \in \{0,1\}^n$ solves $\mathcal{R}$, then $\theta^*$ also solves $\mathcal{P}$.*

*Proof.* By construction, the objective function $V_{\mathcal{R}}$ of $\mathcal{R}$ and the objective function $V_{\mathcal{P}}$ of $\mathcal{P}$ coincide over $\{0,1\}^n$. Therefore $V_{\mathcal{R}}(\theta^*) = V_{\mathcal{P}}(\theta^*)$. Since the feasible set of $\mathcal{R}$, $[0,1]^n$, contains the feasible set of $\mathcal{P}$, $\{0,1\}^n$, it must be that $V_{\mathcal{P}}(\theta^*) \geq V_{\mathcal{P}}(\theta)$ for $\theta \in \{0,1\}^n$, otherwise $\theta^*$ would not be a solution to $\mathcal{R}$. $\theta^*$ therefore solves $\mathcal{P}$. □

**Lemma 3.2.** *The first-order condition of problem $\mathcal{R}$ with respect to $\theta_j$ can be written as*

$$a_j\lambda_j\theta_j^{-1}y_j + \sum_k b_{jk}\lambda_k\theta_j^{-1}x_{jk} - wf_j = \Delta\mu_j, \tag{11}$$

*where $\Delta\mu_j = \overline{\mu}_j - \underline{\mu}_j$ is the difference between the Lagrange multipliers on the constraints $\theta_j \leq 1$ and $\theta_j \geq 0$, respectively.*

*Proof.* We can recast the reshaped problem into the original formulation of (3)–(5) and (2) by noticing that the reshaping procedure amounts to transforming the production function as $z_j\theta_j \to z_j\theta_j^{a_j}$ and $\Omega_{ij} \to \Omega_{ij}\left(\theta_i^{b_{ij}}\right)^{\varepsilon_j-1}$. We can then write the Lagrangian of that problem as

$$\mathcal{L} = \left(\sum_{j\in\mathcal{N}} \beta_j^{\frac{1}{\sigma}} c_j^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}} - \sum_j \lambda_j \left(c_j + \sum_{k\in\mathcal{N}} x_{jk} - \frac{A}{\alpha_j^{\alpha_j}(1-\alpha_j)^{1-\alpha_j}} z_j\theta_j^{a_j} X_j^{\alpha_j} l_j^{1-\alpha_j}\right)$$
$$- w\left(\sum_{j\in\mathcal{N}} l_j + \sum_{j\in\mathcal{N}} f_j\theta_j - 1\right) - \sum_{j\in\mathcal{N}} \bar{\mu}_j(\theta_j - 1) + \sum_{j\in\mathcal{N}} \underline{\mu}_j\theta_j,$$

where $X_j = \left(\sum_{i\in\mathcal{N}} \theta_i^{b_{ij}\frac{\varepsilon_j-1}{\varepsilon_j}} \Omega_{ij}^{\frac{1}{\varepsilon_j}} x_{ij}^{\frac{\varepsilon_j-1}{\varepsilon_j}}\right)^{\frac{\varepsilon_j}{\varepsilon_j-1}}$ is the reshaped intermediate input bundle of firm $j$. The first-order condition with respect to $\theta_l$ is

$$\lambda_l a_l \theta_l^{-1} y_l + \sum_j \lambda_j b_{lj} y_j \alpha_j \theta_l^{-1} \theta_l^{b_{lj}\frac{\varepsilon_j-1}{\varepsilon_j}} \Omega_{lj}^{\frac{1}{\varepsilon_j}} \left(\frac{x_{lj}}{X_j}\right)^{\frac{\varepsilon_j-1}{\varepsilon_j}} - wf_l - \bar{\mu}_l + \underline{\mu}_l = 0.$$

We get (11) by noticing from (42) that $\theta_l^{b_{lj}\frac{\varepsilon_j-1}{\varepsilon_j}}\Omega_{lj}^{\frac{1}{\varepsilon_j}}\left(\frac{x_{lj}}{X_j}\right)^{\frac{\varepsilon_j-1}{\varepsilon_j}}$ is the reshaped version of $\frac{\lambda_l x_{lj}}{\Lambda_j X_j}$, and using (40). □

**Lemma 3.** *The first-order condition of problem $\mathcal{R}$ with respect to $\theta_j$ can be written as*

$$a_j\frac{\lambda_j}{\theta_j}\underbrace{\beta_j\lambda_j^{-\sigma}C}_{c_j}+\sum_k\left(a_j+b_{jk}\right)\frac{\lambda_j}{\theta_j}\underbrace{\theta_j^{b_{jk}(\varepsilon_k-1)}\Omega_{jk}X_k\left(\frac{\lambda_j}{\Lambda_k}\right)^{-\varepsilon_k}}_{x_{jk}}-wf_j=\Delta\mu_j, \tag{12}$$

*where*

$$X_j=\left(\sum_k\theta_k^{b_{kj}\frac{\varepsilon_j-1}{\varepsilon_j}}\Omega_{kj}^{\frac{1}{\varepsilon_j}}x_{kj}^{\frac{\varepsilon_j-1}{\varepsilon_j}}\right)^{\frac{\varepsilon_j}{\varepsilon_j-1}}\text{ and }\Lambda_j=\left(\sum_k\theta_k^{b_{kj}(\varepsilon_j-1)}\Omega_{kj}\lambda_k^{1-\varepsilon_j}\right)^{\frac{1}{1-\varepsilon_j}}, \tag{13}$$

*are, respectively, the reshaped intermediate input bundle of firm $j$ and the social value of that bundle, and where*

$$\lambda_j=\frac{1}{z_j\theta_j^{a_j}A}\Lambda_j^{\alpha_j}w^{1-\alpha_j}, \tag{14}$$

*is the social value of a unit of good $j$.*

*Proof.* Combining (11) with the resource constraint (4) yields

$$\frac{\lambda_l}{\theta_l}a_lc_l+\sum_j\left(a_l+b_{lj}\right)\frac{\lambda_l}{\theta_l}x_{lj}-wf_l=\Delta\mu_l.$$

With the reshaped version of the first-order conditions (41) and (44), this equation becomes (12). Finally, note that using the definition $q_j:=w/\lambda_j$ in (9) yields (14). □

**Lemma 7.** *The first-order condition of problem $\mathcal{R}$ with respect to $\theta_j$ can be written as*

$$\left(1+a_j\right)\lambda_jc_j+\sum_{k=1}^n\left(1+a_j+b_{jk}\right)\lambda_jx_{jk}-\sum_{i=1}^n\lambda_ix_{ij}-wl_j-w\theta_jf_j=\theta_j\Delta\mu_j. \tag{11}$$

*Proof.* Equation (12) can be written as

$$\lambda_ja_jc_j+\sum_k\left(a_j+b_{jk}\right)\lambda_jx_{jk}-wf_j=\theta_j\Delta\mu_j.$$

Adding 39 and using the resource constraint (4), it becomes

$$\lambda_ja_jc_j+\lambda_j\left(c_j+\sum_{k=1}^nx_{jk}\right)+\sum_k\left(a_j+b_{jk}\right)\lambda_jx_{jk}-wl_j-\sum_i\lambda_ix_{ij}-wf_j=\theta_j\Delta\mu_j.$$

Grouping terms in this expression yields the result. □

**Proposition 2.** *Let $\varepsilon_j = \varepsilon$ and $\alpha_j = \alpha$ for all $j$. If $\Omega_{ij} = d_i e_j$ for some vectors $d$ and $e$ then the Karush-Kuhn-Tucker conditions are necessary and sufficient to characterize a solution to $\mathcal{R}$.*

*Proof.* Raise both sides of (9) to the power $\varepsilon - 1$, multiply by $d_j \theta_j^{b(\varepsilon-1)}$ and sum across $j$'s to find

$$\sum_{j\in\mathcal{N}} d_j \left(\theta_j^b q_j(\theta)\right)^{\varepsilon-1} = \left(\sum_{j\in\mathcal{N}} d_j e_j^\alpha (Az_j)^{\varepsilon-1} \theta_j^{a(\varepsilon-1)+b(\varepsilon-1)}\right)^{\frac{1}{1-\alpha}},$$

so that, once combined with (9), we find

$$q_j(\theta) = Az_j \theta_j^a d_j^{\frac{\alpha}{\varepsilon-1}} \left(\sum_{i\in\mathcal{N}} d_i e_i^\alpha (Az_i)^{\varepsilon-1} \theta_i^{a(\varepsilon-1)+b(\varepsilon-1)}\right)^{\frac{\alpha}{1-\alpha}\frac{1}{\varepsilon-1}}.$$

Computing the log of $Q$, we get

$$\begin{aligned}\log(Q) =& \frac{1}{\sigma-1}\log\left(\left(\sum_{j\in\mathcal{N}}\beta_j\left(z_j\theta_j^a A d_j^{\frac{\alpha}{\varepsilon-1}}\right)^{\sigma-1}\right)\left(\sum_{i\in\mathcal{N}} d_i e_i^\alpha (Az_i)^{\varepsilon-1}\theta_i^{a(\varepsilon-1)+b(\varepsilon-1)}\right)^{\frac{\alpha}{1-\alpha}\frac{\sigma-1}{\varepsilon-1}}\right)\\ =&\frac{1}{\sigma-1}\log\left(\sum_{j\in\mathcal{N}}\beta_j\left(z_j\theta_j^a A d_j^{\frac{\alpha}{\varepsilon-1}}\right)^{\sigma-1}\right)+\frac{1}{\varepsilon-1}\frac{\alpha}{1-\alpha}\log\left(\sum_{i\in\mathcal{N}} d_i e_i^\alpha (Az_i)^{\varepsilon-1}\theta_i^{a(\varepsilon-1)+b(\varepsilon-1)}\right).\end{aligned} \tag{48}$$

If $0 < a \leq (\sigma-1)^{-1}$ and $-a \leq b \leq (\varepsilon-1)^{-1} - a$ (and in particular if ($\star$) holds) the exponents on $\theta$ are all between 0 and 1 so that the summations in $\log(Q)$ are concave functions of $\theta$. The log of a concave function is concave so $\log(Q)$ is also concave. Moving toward the full objective function, the term $1 - \sum_{j\in\mathcal{N}} \theta_j f_j$ is concave and so is $\log V_R$. Since, in addition, the constraint set $\theta \in [0,1]^n$ is convex and the Slater's qualification condition is obviously satisfied, the Karush-Kuhn-Tucker conditions are necessary and sufficient to characterize an optimal solution to the maximization of $\log(V_R(\theta))$ on the set $\theta \in [0,1]^n$. Since log is an increasing transformation, a solution to this problem also solves $\mathcal{R}$. □

**Proposition 3.** *Let $\sigma = \varepsilon_j$ for all $j$. Suppose that the $\{\beta_j\}_{j\in\mathcal{N}}$ are not too far from each other and that the matrix $\Omega$ is close enough to $\bar{\Omega}$. Then there exists a threshold $\bar{f} > 0$ such that if $f_j < \bar{f}$ for all $j$ the Karush-Kuhn-Tucker conditions are necessary and sufficient to characterize a solution to $\mathcal{R}$.*

*Proof.* To simplify the notation, define $p_j = q_j^{\sigma-1}$ and let $g^j = p_j\left(z_j^{\sigma-1}\left(\sum_i \Omega_{ij} p_i\right)^{\alpha_j}\right)^{-1}$. $\mathcal{R}$ can then be written as

$$\min_{p\in P} -\frac{1}{\sigma-1}\log\left(\sum_{j\in\mathcal{N}}\beta_j p_j\right) - \log\left(1 - \sum_{j\in\mathcal{N}} f_j g^j(p)\right)$$

where $P=\left\{p\in\mathbb{R}^n_{\geq 0}:p_j\leq z_j^{\sigma-1}\left(\sum_{i\in\mathcal{N}}\Omega_{ij}p_i\right)^{\alpha_j},\forall j\right\}$.

Denote the objective function by $\Lambda$. Its Hessian matrix has typical element

$$\frac{\partial^2\Lambda}{\partial p_k\partial p_l}=\underbrace{\frac{1}{\sigma-1}\beta_k\beta_l\left(\sum_j\beta_j p_j\right)^{-2}}_{A_{kl}}+\underbrace{\frac{\sum_j f_j g^j_{kl}(p)}{1-\sum_j f_j g^j(p)}}_{B_{kl}}+\underbrace{\frac{\left(\sum_j f_j g^j_k(p)\right)\left(\sum_j f_j g^j_l(p)\right)}{\left(1-\sum_j f_j g^j(p)\right)^2}}_{C_{kl}}, \tag{49}$$

and define $A$, $B$ and $C$ as the matrices with typical elements $A_{kl}$, $B_{kl}$ and $C_{kl}$.

I will show that in the limit as $\Omega\to\bar{\Omega}$ and $\beta_j\to\bar{\beta}$ for all $j$ the Hessian is positive definite when the largest fixed cost $\max_j f_j$ is small enough. To do so, I will rely on Lemma 9 above. For that purpose, notice that, in the limit, $A$ is a positive multiple of the all-one matrix.

Pick $\bar{f}>0$ and $\tilde{f}_j\in[0,1]^n$ so that $f_j=\bar{f}\tilde{f}_j$ for all $j$. Taking the derivatives of $g$, we find

$$B_{kl}=\frac{1}{L-\sum_j f_j g^j(p)}\left(\underbrace{\sum_{j\in\mathcal{N}}f_j\frac{\alpha_j(\alpha_j+1)p_j\Omega_{kj}\Omega_{lj}}{z_j^{\sigma-1}\left(\sum_i\Omega_{ij}p_i\right)^{\alpha+2}}}_{B^1_{kl}}-\bar{f}\underbrace{\left(\frac{\tilde{f}_k\alpha_k\Omega_{lk}}{z_k^{\sigma-1}\left(\sum_i\Omega_{ik}p_i\right)^{\alpha_k+1}}+\frac{\tilde{f}_l\alpha_l\Omega_{kl}}{z_l^{\sigma-1}\left(\sum_i\Omega_{il}p_i\right)^{\alpha_l+1}}\right)}_{B^2_{kl}}\right). \tag{50}$$

$B^1$ is a Gramian matrix where $B^1_{kl}$ is the scalar product of a pair of vectors $v_k$ and $v_l$ defined as

$$v_m=\left[\begin{array}{ccccc}\sqrt{\frac{f_1\alpha_1(\alpha_1+1)p_1}{z_1^{\sigma-1}\left(\sum_i\Omega_{i1}p_i\right)^{\alpha_1+2}}}\Omega_{m1} & \dots & \sqrt{\frac{f_j\alpha_j(\alpha_j+1)p_j}{z_j^{\sigma-1}\left(\sum_i\Omega_{ij}p_i\right)^{\alpha_j+2}}}\Omega_{mj} & \dots & \sqrt{\frac{f_n\alpha_n(\alpha_n+1)p_n}{z_n^{\sigma-1}\left(\sum_i\Omega_{in}p_i\right)^{\alpha_n+2}}}\Omega_{mn}\end{array}\right]'.$$

Since Gramian matrices are positive semi-definite, so is $B^1$. For $B^2$, I will show that, in the limit, it is negative definite on the subspace $S:\sum_{i=1}^n x_i=0$. Define $b$ as a vector with typical element $b_j=\tilde{f}_j\alpha_j\left[z_j^{\sigma-1}\left(\sum_i\Omega_{ij}p_i\right)^{\alpha_j+1}\right]^{-1}$. We can write $B^2=\Omega\text{diag}(b)+(\Omega\text{diag}(b))'$. Take any vector $x\in S$, then in the limit,

$$\begin{aligned}x'B^2x&=x'\left[\omega(O_n-I_n)\,\text{diag}(b)+(\omega(O_n-I_n)\,\text{diag}(b))'\right]x\\&=x'\left[-2\omega I_n\text{diag}(b)\right]x<0\end{aligned}$$

for any $x\neq 0$. The matrix $B^2$ is therefore negative definite on $S$. Using Lemma 9, $A-\bar{f}B^2$ is therefore positive definite for $\bar{f}>0$ small enough. Finally, the matrix $C$ in (49), is also a Gramian matrix and its contribution to the Hessian is thus positive semi-definite.

Putting the pieces together, we have shown that the Hessian of the objective function $\Lambda$ is positive definite for $\max_j f_j$ small enough when $\Omega=\bar{\Omega}$ and $\beta=\bar{\beta}$. Now, each element of the Hessian is also a continuous function of $(\Omega, \beta)$ in a neighborhood of $(\bar{\Omega}, \bar{\beta})$.[45] Since the eigenvalues are continuous functions of the elements of a matrix, they are also continuous functions of $\Omega$ and $\beta$. There is therefore a ball $\mathcal{B} = \left\{(\Omega, \beta) : \left\|(\Omega, \beta) - (\bar{\Omega}, \bar{\beta})\right\| < \nu\right\}$ for some $\nu > 0$ such that the Hessian is also positive definite for $(\Omega, \beta) \in \mathcal{B}$.[46] □

**Lemma 4.** *For a given entry decision vector $\theta$, distorted and undistorted equilibria are efficient. Furthermore, the equilibrium prices $W$ and $Q_{jk}$ are equal (up to a choice of numeraire) to the planner's Lagrange multipliers $w$ and $\lambda_j$.*

*Proof.* We begin by computing the first-order conditions of the firms in equilibrium. Under both equilibrium definitions, the cost-minimization problem of a firm $j$ is

$$\min_{x,l} \sum_{i=1}^{n} p^x_{ij} x_{ij} + w^e l_j, \quad \text{subject to} \quad \frac{A}{\alpha_j^{\alpha_j} (1-\alpha_j)^{1-\alpha_j}} z_j \theta_j \left(\sum_{i=1}^{n} \Omega_{ij}^{\frac{1}{\varepsilon_j}} x_{ij}^{\frac{\varepsilon_j - 1}{\varepsilon_j}}\right)^{\frac{\varepsilon_j}{\varepsilon_j - 1}\alpha_j} l_j^{1-\alpha_j} \geq y_j,$$

for some fixed output level $y_j$. Denote $\delta_j$ the Lagrange multiplier on the constraint, which also corresponds to the marginal cost of production. The first-order conditions with respect to $l_j$ and $x_{kj}$ are

$$(1-\alpha_j)\,\delta_j y_j = w^e l_j \tag{51}$$

$$\delta_j \frac{A}{\alpha_j^{\alpha_j} (1-\alpha_j)^{1-\alpha_j}} \theta_j z_j \alpha_j \left(\sum_{i=1}^{n} \Omega_{ij}^{\frac{1}{\varepsilon_j}} x_{ij}^{\frac{\varepsilon_j - 1}{\varepsilon_j}}\right)^{\alpha_j \frac{\varepsilon_j}{\varepsilon_j - 1} - 1} \Omega_{kj}^{\frac{1}{\varepsilon_j}} x_{kj}^{\frac{\varepsilon_j - 1}{\varepsilon_j}} l_j^{1-\alpha_j} = p^x_{kj} x_{kj}. \tag{52}$$

I now show that these equations imply that the ratio $w^e/\delta_j$ in the equilibrium is equal to $q_j$ in the efficient allocation. In the undistorted equilibrium, the standard pricing equation applies such that $p^x_{kj} = \frac{\varepsilon_j}{\varepsilon_j - 1} \frac{1}{1+s^x_{kj}} \delta_k = \delta_k$, under the assumed subsidy. In the distorted equilibrium, the price $p^x_{kj}$ is equal to the marginal cost of production $\delta_k$. As a results, the equations (51) and (52) coincide with their efficient allocation counterparts, (37) and (38), if we replace $\delta_j$ with $\lambda_j$ and $W$ with $w$. We can therefore follow the same steps and find that the equation

$$\frac{w^e}{\delta_j} = z_j \theta_j A \left(\sum_{i=1}^{n} \Omega_{ij} \left(\frac{w^e}{\delta_i}\right)^{\varepsilon_j - 1}\right)^{\frac{\alpha_j}{\varepsilon_j - 1}}, \tag{53}$$

holds. Since this equation is the same as (6), it follows that $\frac{w^e}{\delta_j} = q_j = \frac{w}{\lambda_j}$ for all $j$.

Next we turn to consumption $c_j$. Under both equilibrium definitions, the usual markup over

[45] Elements of the Hessian become infinite at the boundary of $P$ where $\sum_i \Omega_{ij} p_i = 0$ for some $j$. However, these points are not relevant to the planner under our assumptions. When the fixed costs are small enough and when $\Omega$ is close enough to $\bar{\Omega}$, each firm is connected to at least one producing firm at the optimum. Therefore, we can exclude these points easily adding the constraints $\sum_i \Omega_{ij} p_i \geq D$ for some $D > 0$ small. These constraints will never bind and the solution to the planner's problem is therefore unchanged.

[46] Since all norms are equivalent in a finite dimensional space, there is no need to specify one here.

marginal cost pricing rule applies so that $p_j^c = \frac{\sigma}{\sigma-1}\delta_j$. and

$$c_j = \beta_j C \left(\frac{p_j^c}{P^c}\right)^{-\sigma} = \beta_j C \left(\frac{\frac{\sigma}{\sigma-1}\delta_j}{P^c}\right)^{-\sigma} \text{ and } P^c = \left(\sum_{j=1}^n \beta_j \left(p_j^c\right)^{1-\sigma}\right)^{\frac{1}{1-\sigma}} = \frac{\sigma}{\sigma-1}\left(\sum_{j=1}^n \beta_j \delta_j^{1-\sigma}\right)^{\frac{1}{1-\sigma}}.$$

It will be convenient to normalize the aggregate price level so that $P^c = \frac{\sigma}{\sigma-1}$ and $\left(\sum_j \beta_j \delta_j^{1-\sigma}\right)^{\frac{1}{1-\sigma}} = 1$, and we can write

$$c_j = \beta_j C \left(\delta_j\right)^{-\sigma}, \tag{54}$$

which is exactly the first-order condition (44) of the planner.

I now show that the equilibrium prices correspond to the planner's Lagrange multipliers. Since $\frac{w^e}{\delta_j} = q_j$ we can write

$$1 = \left(\sum_{j=1}^n \beta_j \delta_j^{1-\sigma}\right)^{\frac{1}{1-\sigma}} = \left(\sum_{j=1}^n \beta_j \left(\frac{w^e}{q_j}\right)^{1-\sigma}\right)^{\frac{1}{1-\sigma}} = w^e \times Q^{-1},$$

and since $Q = w$ by (45) we find that $w^e = w$. It follows immediately that $\delta_j = \lambda_j$.

I have shown that all first-order conditions are the same in the equilibria and the planner's problem, and that all prices are equal to their respective shadow value in the efficient allocation.[47] Since the resources constraints are also the same, the distorted and undistorted equilibria coincide with the efficient allocation, for an exogenously given entry vector $\theta$. □

**Proposition 4.** *A vector $\theta = \{0,1\}^n$ that satisfies the first-order conditions of the reshaped problem $\mathcal{R}$ with reshaping parameters ($\star$) is an undistorted equilibrium.*

*Proof.* Because of Lemma 4, we just need to show that the vector $\theta$ that solves the reshaped problem's first-order conditions is also an undistorted equilibrium.

Recall from (24) that the planner entry decision for firm $k$ can be written as

$$\left(1 + a_k\right)\lambda_k c_k + \sum_{j=1}^n \left(1 + a_k + b_{kj}\right) x_{kj} \lambda_k - \sum_{j=1}^n \lambda_j x_{jk} - w l_k - w\theta_k f_k = \theta_k \Delta\mu_k,$$

which, using the first-order conditions (37) and (38) and the resource constraint (4), boils down to

$$a_k \lambda_k c_k + \sum_{j\in\mathcal{N}} \left(a_k + b_{kj}\right) x_{kj} \lambda_k - w\theta_k f_k = \theta_k \Delta\mu_k.$$

We can combine this equation with the planner's consumption decision (44) and the previously-

[47] Recall that for a given $\theta$ the problem of the social planner is convex such that first-order conditions are sufficient to characterize the efficient allocation.

derived equation for $x_{kj}$ (41) to write

$$a_k \lambda_k \beta_k \lambda_k^{-\sigma} C + \sum_{j \in \mathcal{N}} (a_k + b_{kj}) \lambda_k \Omega_{kj} X_j \left( \frac{\lambda_k}{\left(\sum_i \Omega_{ij} \lambda_i^{1-\varepsilon_j}\right)^{\frac{1}{1-\varepsilon_j}}} \right)^{-\varepsilon_j} - w\theta_k f_k = \theta_k \Delta\mu_k.$$

Now suppose that $\theta_k > 0$. We can divide this equation by $\theta_k$, such that

$$a_k \beta_k \frac{\lambda_k^{1-\sigma}}{\theta_k} C + \sum_{j \in \mathcal{N}} (a_k + b_{kj}) \frac{\lambda_k^{1-\varepsilon_j}}{\theta_k} \Omega_{kj} X_j \left( \frac{1}{\left(\sum_i \Omega_{ij} \lambda_i^{1-\varepsilon_j}\right)^{\frac{1}{1-\varepsilon_j}}} \right)^{-\varepsilon_j} - w f_k = \Delta\mu_k. \tag{55}$$

It turns out that this equation is also valid for $\theta_k = 0$. To see this, notice that in this case $\lambda_k \to \infty$ (since $q_k = 0$), and that the ratios $\frac{\lambda_k^{1-\sigma}}{\theta_k}$ and $\frac{\lambda_k^{1-\varepsilon_j}}{\theta_k}$ remain finite. Indeed, from (6) we can write

$$\begin{aligned} \frac{1}{\theta_j} \left( \frac{w}{\lambda_j} \right)^{\sigma-1} &= \frac{1}{\theta_j} \left( z_j \theta_j^{a_j} A \left( \sum_{i \in \mathcal{N}} \Omega_{ij} \left( \theta_i^{b_{ij}} q_i \right)^{\varepsilon_j - 1} \right)^{\frac{\alpha_j}{\varepsilon_j - 1}} \right)^{\sigma-1} \\ &= \left( z_j A \left( \sum_{i \in \mathcal{N}} \Omega_{ij} \left( \theta_i^{b_{ij}} q_i \right)^{\varepsilon_j - 1} \right)^{\frac{\alpha_j}{\varepsilon_j - 1}} \right)^{\sigma-1} \end{aligned} \tag{56}$$

such that $\frac{\lambda_k^{1-\sigma}}{\theta_k}$ is indeed well-defined when $\theta_k = 0$ under ($\star$). For $\frac{\lambda_k^{1-\varepsilon_j}}{\theta_k}$, note that since $\theta \in \{0,1\}^n$, we can write (6) as $q_j = z_j \theta_j^{\frac{1}{\varepsilon_j - 1}} A \left( \sum_{i \in \mathcal{N}} \Omega_{ij} q_i^{\varepsilon_j - 1} \right)^{\frac{\alpha_j}{\varepsilon_j - 1}}$ since $\varepsilon_j > 1$ for all $j$. Then this equation can be reorganized as

$$\frac{1}{\theta_j} \left( \frac{w}{\lambda_j} \right)^{\varepsilon_j - 1} = \left( z_j A \left( \sum_{i \in \mathcal{N}} \Omega_{ij} \left( \frac{w}{\lambda_i} \right)^{\varepsilon_j - 1} \right)^{\frac{\alpha_j}{\varepsilon_j - 1}} \right)^{\varepsilon_j - 1}, \tag{57}$$

and so $\frac{\lambda_k^{1-\varepsilon_j}}{\theta_k}$ is also well-defined when $\theta_k = 0$.

I now compare (55), to its equivalent in the undistorted equilibrium. Combining the profits (23) with the resource constraint (20) and the demand curves (54) and (18), we can write the entry decision of a firm $j$ as

$$\frac{1}{\sigma - 1} \delta_j \beta_j C \frac{\delta_j^{1-\sigma}}{\theta_j} + \sum_{i \in \mathcal{N}} \frac{1}{\varepsilon_i - 1} \frac{\delta_j^{1-\varepsilon_j}}{\theta_j} \Omega_{ji} X_i \left( \frac{1}{\left(\sum_k \Omega_{ki} \delta_k^{1-\varepsilon_i}\right)^{\frac{1}{1-\varepsilon_i}}} \right)^{-\varepsilon_i} - w^e f_j \geq 0 \tag{58}$$

where I have used the fact that $p_{ji}^x = \delta_j$. Here, notice that if $\theta_j = 1$ the expressions $\frac{\delta_j^{1-\sigma}}{\theta_j} = \delta_j^{1-\sigma}$

and $\frac{\delta_j^{1-\varepsilon_j}}{\theta_j} = \delta_j^{1-\varepsilon_j}$ correspond to the actual prices of the firm raised to the powers $1-\sigma$ or $1-\varepsilon_j$. If instead $\theta_j = 0$, these expressions correspond to the *counterfactual* prices of the firm (raised to the appropriate powers), if it were to enter. This can be seen from the right-hand sides of (56) and (57), and by remembering from the proof of Lemma 4 that the same equations hold with $\delta_j$ and $W$ instead of $\lambda_j$ and $w$. As a result, (58) is the correct equation to think of entry given our equilibrium concept.

Now, suppose that in the reshaped planner's problem $\theta_k = 1$, then it must be that $\Delta\mu_k \geq 0$ by the slackness condition, and so the left-hand side of (55) is positive (under the reshaping parameters $\star$). Then the left-hand side of (58) is also positive and the firm enters in the undistorted equilibrium. Similarly, if $\theta_k = 0$, $\Delta\mu_k < 0$ by the complementary slackness conditions and the firm does not enter in equilibrium. □

**Proposition 5.** *A vector $\theta = \{0,1\}^n$ that satisfies the first-order conditions of the reshaped problem $\mathcal{R}$ with parameters (26) is a distorted equilibrium.*

*Proof.* The steps are essentially the same as for Proposition 4, except that I use the reshaping parameters (26) instead and the entry equation (25) for the equilibrium. □

**Proposition 6.** *In the partial equilibrium analysis of firm $j$ (Definition 2) operating a firm that is directly upstream or downstream from $j$ increases $\pi_j^{undist}$, and operating a firm that is directly upstream from $j$ increases $\pi_j^{dist}$.*

*Proof.* Consider a firm $i$ that is directly upstream from $j$ so that $\Omega_{ij} > 0$. A transition from $\theta_i = 0$ to $\theta_i = 1$ results in a decline in $\delta_j$ through the recursive structure of (22). This in turn, leads to an increase in $\delta_j c_j$ and $\delta_j x_{jk}$ for all $k$ by (28). It follows that both $\pi_j^{\text{undist}}$ and $\pi_j^{\text{dist}}$ increase as a result of the change in $\theta_i$. Now consider a firm $k$ that is directly downstream from $j$ so that $\Omega_{jk} > 0$. The transition from $\theta_k = 0$ to $\theta_k = 1$ leads to an increase in $k$'s purchases of $j$'s good (second term in (27)). Since in the undistorted equilibrium $j$ earns a markup above its marginal cost for each unit sold, $\pi^{\text{undist}}$ increases. We see from (29) that this channel is absent from $\pi_j^{\text{dist}}$ and so the operating decisions of upstream firms have no direct impact on $j$'s operating decision. □

**Proposition 7.** *In the efficient allocation, the following holds.*

1. *The operating decision $\theta_j(z)$ of firm $j$ is weakly increasing in $z_j$.*

2. *Denote by $\Omega^-$ a network of potential connections and let $\Omega^+ = \Omega^- + \Delta\Omega$, where $\Delta\Omega$ is a nonnegative matrix with potentially positive entries only in its $j$th row and $j$th column. Then, $\theta_{\Omega^+,j}(z) \geq \theta_{\Omega^-,j}(z)$ for all $z$, where $\theta_{\Omega,j}(z)$ denotes the operating decision of firm $j$ under $\Omega$.*

*Proof.* I first prove part 1. Take two productivity vectors $z'$ and $z$ such that $z'_i = z_i$ for all $i \neq j$ and $z'_j > z_j$. I will show that if $(\theta_j(z'), \theta_{-j}(z')) = (0, \theta_{-j}(z'))$ is optimal at $z'$ then it must also be optimal at $z$, which would imply that $\theta_j(z)$ cannot be strictly decreasing in $z_j$. Suppose not, then there must be another allocation $\tilde{\theta}$ such that

$$C\left(z,\tilde{\theta}_j,\tilde{\theta}_{-j}\right) > C\left(z,0,\theta_{-j}\left(z'\right)\right). \tag{59}$$

But then

$$C\left(z',0,\theta_{-j}\left(z'\right)\right) \geq C\left(z',\tilde{\theta}_j,\tilde{\theta}_{-j}\right) \geq C\left(z,\tilde{\theta}_j,\tilde{\theta}_{-j}\right) > C\left(z,0,\theta_{-j}\left(z'\right)\right) = C\left(z',0,\theta_{-j}\left(z'\right)\right), \tag{60}$$

where the first inequality comes from the fact that $(0,\theta_{-j}(z'))$ is optimal at $z'$, the second inequality is because a higher $z$ is always weakly preferable for a fixed $\theta$, the third inequality comes from (59), and the equality comes from the fact that $z_j$ cannot influence consumption since $\theta_j = 0$. But (60) implies that $C(z',0,\theta_{-j}(z')) > C(z',0,\theta_{-j}(z'))$ which is contradiction. It follows that if $\theta_j(z') = 0$ it must be $\theta_j(z) = 0$, and $\theta_j(z)$ cannot be decreasing in $z_j$.

I now turn to the proof of part 2. Denote $C_\Omega(z,\theta)$ consumption under $(\Omega,z,\theta)$. I first show that $C_\Omega(z,\theta)$ is increasing in $\Omega$ in the sense that additional connections are weakly welfare improving, for a fixed $z$ and $\theta$. Denote by $\Omega^-$ a network of potential connections and let $\Omega^+$ be identical to $\Omega^-$ except that it has an additional potential connection between two firms. From (7), (8) and (6) we can write

$$C_\Omega(z,\theta) = \left(\sum_{j=1}^{n} \beta_j q_j^{\sigma-1}\right)^{\frac{1}{\sigma-1}} \left(1 - \sum_{j=1}^{n} \theta_j f_j\right) \tag{61}$$

where $q_j$ is the (unique) fixed point of (6). When comparing $C_{\Omega^+}(z,\theta)$ and $C_{\Omega^-}(z,\theta)$ the vector $\theta$ is the same and, therefore, so is $\sum_{j=1}^{n}\theta_j f_j$. The only difference in the objective function comes from the vector $q$. I now show that $q_{\Omega^+}(z,\theta) \geq q_{\Omega^-}(z,\theta)$. Recall from the proof of Lemma 1 that we can find the fixed point $q$ by iterating on (6). Further, as explained in Kennan (2001), if we start iterating from the vector $\underline{p}$ (defined in the proof of Lemma 1) each iteration is larger than the previous one until convergence. Denote by$F^+(q)$ the vectorized right-hand side of (6) for all $j$ under $\Omega^+$, and $F^-(q)$ the same object but under $\Omega^-$. Notice that $F^+(q) \geq F^-(q)$ for all $q$. Next, consider the sequences $q^{(t+1),+} = F^+\left(q^{(t),+}\right)$ and $q^{(t+1),-} = F^+\left(q^{(t),-}\right)$ where $q^{(0),+} = q^{(0),-} = \underline{p}$. I will show by induction that $q^{(t),+} \geq q^{(t),-}$ for all $t$. First, notice that $q^{(1),+} = F^+\left(\underline{p}\right) \geq F^-\left(\underline{p}\right) = q^{(1),-}$. Next, suppose that $q^{(t),+} \geq q^{(t),-}$ for some $t$. Then,

$$q^{(t+1),+} = F^+\left(q^{(t),+}\right) \geq F^+\left(q^{(t),-}\right) \geq F^-\left(q^{(t),-}\right) = q^{(t+1),-},$$

where the first inequality comes from the fact that $F^+$ is weakly increasing in $q$, and the second inequality comes from the fact that $F^+(q) \geq F^-(q)$ for all $q$. By induction, it follows that $q^{(t),+} \geq q^{(t),-}$ for all $t$. Since both sequences converge, we have that $q^+ \geq q^-$. Since $\sum_{j=1}^{n}\beta_j q_j^{\sigma-1}$ is increasing in $q$, it is then immediate from (61) that $C_{\Omega^+}(z,\theta) \geq C_{\Omega^-}(z,\theta)$.[48]

I now turn to part 2 of the lemma. I will show that if $\left(\theta_j^{\Omega^+}(z), \theta_{-j}^{\Omega^+}(z)\right) = \left(0, \theta_{-j}^{\Omega^+}(z)\right)$ is optimal under $\Omega^+$ then it must also be optimal under $\Omega^-$, which would immediately imply that $\theta_{\Omega^+,j}(z) \geq \theta_{\Omega^-,j}(z)$ since $\theta$ is binary. Suppose not, then there must be another allocation $\tilde{\theta}$ such that

$$C^{\Omega^-}\left(\tilde{\theta}_j, \tilde{\theta}_{-j}\right) > C^{\Omega^-}\left(0, \theta_{-j}^{\Omega^+}(z)\right). \tag{62}$$

But then

$$C^{\Omega^+}\left(0, \theta_{-j}^{\Omega^+}(z)\right) \geq C^{\Omega^+}\left(\tilde{\theta}_j, \tilde{\theta}_{-j}\right) \geq C^{\Omega^-}\left(\tilde{\theta}_j, \tilde{\theta}_{-j}\right) > C^{\Omega^-}\left(0, \theta_{-j}^{\Omega^+}(z)\right) = C^{\Omega^+}\left(0, \theta_{-j}^{\Omega^+}(z)\right), \tag{63}$$

where the first inequality follows since $\left(0, \theta_{-j}^{\Omega^+}(z)\right)$ is optimal under $\Omega^+$, the second inequality follows from the intermediary result we just proved, the strict inequality follows from (62), and the equality follows since the extra link in $\Omega^+$ does not matter for consumption since $\theta_j = 0$. But (63) implies that $C^{\Omega^+}\left(0, \theta_{-j}^{\Omega^+}(z)\right) > C^{\Omega^+}\left(0, \theta_{-j}^{\Omega^+}(z)\right)$ which is contradiction. It follows that if $\theta_j^{\Omega^+} = 0$ it must be $\theta_j^{\Omega^-} = 0$, and the result follows.

**Proposition 8.** *Let $\mathcal{J} \subset \mathcal{N}$ be a group of firms. Denote by $\theta^+, \theta^- \in \{0,1\}^n$ two operating vectors such that $\theta_j^+ = 1$ and $\theta_j^- = 0$ for $j \in \mathcal{J}$, and $\theta_j^+ = \theta_j^-$ for $j \notin \mathcal{J}$. Denote by $\Omega^-$ a network of potential connections and let $\Omega^+ = \Omega^- + \Delta\Omega$ where $\Delta\Omega$ is a matrix full of zeros except that $\Delta\Omega_{kl} > 0$ for some $k, l \in \mathcal{J}$ . Then $C_{\Omega^+}(\theta^+) - C_{\Omega^+}(\theta^-) \geq C_{\Omega^-}(\theta^+) - C_{\Omega^-}(\theta^-)$, where $C_\Omega(\theta)$ denotes consumption in the efficient allocation under $\Omega$ and $\theta$.*

First note that $C_{\Omega^+}(\theta^-) = C_{\Omega^-}(\theta^-)$ since in either case the firms in $\mathcal{J}$ are not operating and so the extra link in $\Omega^+$ is irrelevant. I therefore only have to show that $C_{\Omega^+}(\theta^+) \geq C_{\Omega^-}(\theta^+)$, but this was proven as part of the proof of Proposition 7. □

**Lemma 5.** *Let $\mathcal{G} \subset \mathcal{N}$ denote the potential neighbors of $j$, that is all firms $i \neq j$ such that $\Omega_{ij} > 0$ and/or $\Omega_{ji} > 0$. There exists a threshold $\bar{f} \geq 0$ such that if $f_j \leq \bar{f}$, then the consumption gain in the efficient allocation from operating $\mathcal{G}$ is larger when $j$ is operating, that is*

$$C(\theta_\mathcal{G} = 1, \theta_j = 1) - C(\theta_\mathcal{G} = 0, \theta_j = 1) \geq C(\theta_\mathcal{G} = 1, \theta_j = 0) - C(\theta_\mathcal{G} = 0, \theta_j = 0), \tag{30}$$

*where $C$ is computed keeping fixed $\theta_i$ for all $i \notin \{j \cup \mathcal{G}\}$.*

*Proof.* The second and the fourth term in (30) cancel out since in both cases $j$ has no supplier and cannot produce. The result follows since $C(\theta_\mathcal{G} = 1, \theta_j = 1) \geq C(\theta_\mathcal{G} = 1, \theta_j = 0)$ is always true when $f_j$ is small enough. Indeed, in the limit as $f_j \to 0$ there is no fixed cost of operating $\theta_j$ and so setting $\theta_j = 1$ only benefits consumption through its impact on $Q$. □

[48] Note that this result holds for more general matrices $\Omega^+$ and $\Omega^-$. Specifically, it holds for any $\Omega^+ = \Omega^- + \Delta\Omega$ where $\Delta\Omega$ is a nonnegative matrix. I use this more general result in other proofs below.

**Lemma 6.** *Let $\Omega^-$ be a network of potential connections and let $\mathcal{G} \subset \mathcal{N} \setminus j$ denote a subset of j's neighbors in $\Omega^-$, that is $\Omega^-_{ij} > 0$ and/or $\Omega^-_{ji} > 0$ for all $i \in \mathcal{G}$. Let $\Omega^+ = \Omega^- + \Delta\Omega$ where $\Delta\Omega$ is a matrix of zeros except for $\Delta\Omega_{ij} > 0$ or $\Delta\Omega_{ji} > 0$ for some $i \in \mathcal{G}$. Denote by $\Delta^{\mathcal{G}}_{\Omega,j} C\left(\tilde{\theta}\right) = C_\Omega\left(\theta_{\mathcal{G}} = 1, \theta_j = \tilde{\theta}\right) - C_\Omega\left(\theta_{\mathcal{G}} = 0, \theta_j = \tilde{\theta}\right)$ the consumption gain in the efficient allocation from operating $\mathcal{G}$ when $\theta_j = \tilde{\theta}$ under $\Omega$, keeping the operating status of all other firms the same. Then the increase in consumption gain from operating $\mathcal{G}$ when $\theta_j = 1$ compared to when $\theta_j = 0$ is greater under $\Omega^+$ than under $\Omega^-$, that is $\Delta^{\mathcal{G}}_{\Omega^+,j} C(1) - \Delta^{\mathcal{G}}_{\Omega^+,j} C(0) \geq \Delta^{\mathcal{G}}_{\Omega^-,j} C(1) - \Delta^{\mathcal{G}}_{\Omega^-,j} C(0)$.*

*Proof.* Using the definition of $\Delta^{\mathcal{G}}_{\Omega,j} C$ we can write the inequality as

$$\begin{aligned} &C_{\Omega^+}\left(\theta^{\mathcal{G}} = 1, \theta_j = 1\right) - C_{\Omega^+}\left(\theta^{\mathcal{G}} = 0, \theta_j = 1\right) - C_{\Omega^+}\left(\theta^{\mathcal{G}} = 1, \theta_j = 0\right) + C_{\Omega^+}\left(\theta^{\mathcal{G}} = 0, \theta_j = 0\right) \\ &\geq C_{\Omega^-}\left(\theta^{\mathcal{G}} = 1, \theta_j = 1\right) - C_{\Omega^-}\left(\theta^{\mathcal{G}} = 0, \theta_j = 1\right) - C_{\Omega^-}\left(\theta^{\mathcal{G}} = 1, \theta_j = 0\right) + C_{\Omega^-}\left(\theta^{\mathcal{G}} = 0, \theta_j = 0\right). \end{aligned}$$

The second and sixth terms cancel out as the extra link in $\Omega^+$ does not matter for consumption since $\theta_{\mathcal{G}} = 0$. The third and seventh terms cancel out as the extra link in $\Omega^+$ does not matter for consumption since $\theta_j = 0$. The same logic implies that the fourth and the eight terms also cancel out. The inequality therefore boils down to $C_{\Omega^+}\left(\theta^{\mathcal{G}} = 1, \theta_j = 1\right) \geq C_{\Omega^-}\left(\theta^{\mathcal{G}} = 1, \theta_j = 1\right)$, which we have already proven in the proof of Proposition 7. □

**Proposition 9.** *In the efficient allocation, GDP $C(z)$ is a continuous function of $z$.*

*Proof.* I proceed by contradiction. Suppose that $C(z)$ is discontinuous at a point $z$ and, without loss of generality, suppose that this discontinuity happens in direction $i$: there exists a $\delta > 0$ such that

$$C(z) - \lim_{\epsilon \to 0^+} C(z - \epsilon 1_i) > \delta, \tag{64}$$

where $1_i$ is a vector full of zeros except for a 1 in the $i$th row. For a fixed $\theta$, the economy is a standard CES production network economy such that Hulten's theorem applies and $C(z, \theta)$ is continuous in $z$. For $C$ to be discontinuous at $z$, it must therefore be that $\theta(z) \neq \lim_{\epsilon \to 0^+} \theta(z - \epsilon 1_i)$. But since $C(z, \theta)$ is continuous for a fixed $\theta$, it must be that

$$C(z, \theta(z)) = \lim_{\epsilon \to 0^+} C(z - \epsilon 1_i, \theta(z)) > \lim_{\epsilon \to 0^+} C(z - \epsilon 1_i, \theta(z - \epsilon 1_i)),$$

where the inequality comes from (64). It follows that $\theta(z - \epsilon 1_i)$ is not optimal at $z - \epsilon 1_i$ for $\epsilon \to 0^+$, which is a contradiction. □

**Proposition 10.** *In the efficient allocation, at almost all $z$ the marginal impact of $z_j$ on GDP is given by $\frac{d \log C}{d \log z_j} = \frac{\lambda_j y_j}{C}$.*

*Proof.* Suppose that $\theta$ is fixed. In this case, the allocation solves the planning problem

$$C(z) = \max_{c,x,l} \left( \sum_{j=1}^{n} \beta_j^{\frac{1}{\sigma}} c_j^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} - \sum_{j=1}^{n} \lambda_j \left( c_j + \sum_{k=1}^{n} x_{jk} - y_j \right) - w \left( \sum_{j=1}^{n} l_j + \sum_{j=1}^{n} \theta_j f_j - 1 \right),$$

where as before $\lambda_j$ and $w$ are Lagrange multipliers. The envelope theorem implies that $\frac{dC}{dz_j} = \frac{\lambda_j y_j}{z_j}$,and so Hulten's theorem applies when $\theta$ is fixed. It follows that Hulten's theorem might not apply only at a point $z$ if a marginal shock to $z$ can lead to a change in $\theta$. To prove the proposition, I begin by measuring the set of points $z$ for which two vectors $\theta$ are efficient.

**Lemma.** *The efficient operating status vector $\theta(z)$ is unique for almost all $z$.*

*Proof.* Denote by $\mathcal{Z}$ the set of $z$'s such that two vectors $\theta$ are efficient, and by $\mathcal{Z}(\theta^*, \theta^{**}) = \{z : \theta^* \text{ and } \theta^{**} \text{ are efficient}\}$. Since there is only a finite set of potential vectors $\theta$, we can write $\mathcal{Z}$ as the finite union $\mathcal{Z} = \bigcup_{\theta^*,\theta^{**}} \mathcal{Z}(\theta^*, \theta^{**})$. Thus if $\mathcal{Z}(\theta^*, \theta^{**})$ has measure zero for all pairs $(\theta^*, \theta^{**})$ then $\mathcal{Z}$ also has measure zero. Define $\Delta(z, \theta^*, \theta^{**}) = C(z, \theta^*) - C(z, \theta^{**})$ and notice that for all $z \in \mathcal{Z}(\theta^*, \theta^{**})$ we have $\Delta(z', \theta^*, \theta^{**}) = 0$. Since $\theta^* \neq \theta^{**}$ there is at least one index $i$ such that $\theta_i^* \neq \theta_i^{**}$. Without loss of generality, assume that $\theta_i^* = 1$ and $\theta_i^{**} = 0$. This implies that $C(z_i, z_{-i}, \theta^*)$ is strictly increasing in $z_i$ (recall that Hulten's theorem applies in this case) and that $C(z_i, z_{-i}, \theta^{**})$ does not vary with $z_i$. It follows that there is a unique $z_i$ such that $\Delta(z_i, z_{-i}, \theta^*, \theta^{**}) = 0$, and that $\mathcal{Z}(\theta^*, \theta^{**})$ has measure zero. The finite union $\mathcal{Z} = \bigcup_{\theta^*,\theta^{**}} \mathcal{Z}(\theta^*, \theta^{**})$ thus also has measure zero, which is the result. □

With that lemma in hand, I now show that any point $z$ for which a marginal increase in $z_i$ leads to a transition in $\theta$ implies that two $\theta$'s are optimal at $z$. More precisely, if there is a point $z$ and an index $i$ such that $\theta^* = \lim_{z_i' \to z_i^-} \theta(z_i', z_{-i}) \neq \lim_{z_i' \to z_i^+} \theta(z_i', z_{-i}) = \theta^{**}$, then $C(z_i, z_{-i}, \theta^*) = C(z_i, z_{-i}, \theta^{**})$. Suppose not, and assume without loss of generality that $C(z_i, z_{-i}, \theta^*) - C(z_i, z_{-i}, \theta^{**}) > \delta$ for some $\delta > 0$. Since $C$ is continuous in $z$ this implies that $\theta^{**}$ is not optimal for $z_i' \to z_i^+$, which is a contradiction. It follows that for Hulten's theorem to not apply at a point $z$ in a direction $i$ it must be that two vectors $\theta$ are optimal at $z$, but we know from the preceding lemma that the set of such points has measure zero.

With that lemma in hand, I now show that any point $z$ for which a marginal increase in $z_i$ leads to a transition in $\theta$ implies that two $\theta$'s are optimal at $z$. More precisely, if there is a point $z$ and an index $i$ such that $\theta^* = \lim_{z_i' \to z_i} \theta(z_i', z_{-i}) \neq \theta(z_i, z_{-i}) = \theta^{**}$, then $C(z_i, z_{-i}, \theta^*) = C(z_i, z_{-i}, \theta^{**})$. Suppose not, and assume without loss of generality that $C(z_i, z_{-i}, \theta^*) - C(z_i, z_{-i}, \theta^{**}) > \delta$ for some $\delta > 0$. Since $C$ is continuous in $z$ this implies that $\theta^{**}$ is not optimal for $z_i' \to z_i$, which is a contradiction. It follows that for Hulten's theorem to not apply at a point $z$ in a direction $i$ it must be that two vectors $\theta$ are optimal at $z$, but we know from the preceding lemma that the set of such points has measure zero. □

**Proposition 11.** *Let $\theta^*(z)$ be the efficient allocation under $z$ and let $C(\theta, z)$ be consumption under $(\theta, z)$. Then the response of consumption after a change in productivity from $z$ to $z'$ is such that*

$$\underbrace{C\left(\theta^*\left(z'\right), z'\right) - C\left(\theta^*\left(z\right), z\right)}_{\text{Change in consumption under a flexible network}} \geq \underbrace{C\left(\theta^*\left(z\right), z'\right) - C\left(\theta^*\left(z\right), z\right)}_{\text{Change in consumption under a fixed network}}.$$

*Proof.* By definition $\theta^*(z')$ maximizes welfare under $z'$. This implies that $C(\theta^*(z'), z') \geq C(\theta^*(z), z')$. Subtracting $C(\theta^*(z), z)$ from both sides yields the result. □

**Proposition 12.** *Every stable equilibrium is efficient.*

*Proof.* The proof proceeds by establishing restrictions that any stable equilibrium must satisfy. It then shows that any allocation that satisfies these restrictions must be efficient.

Consider a coalition made of all the firms in the economy. For the equilibrium to be stable there cannot be an alternative arrangement that would yield larger aggregate profits. Otherwise, transfers could be designed to make one firm better off while keeping the other firms at the same profit level. The arrangement $\{x_{ij}, T_{ij}\}_{i,j}$ must therefore maximize $\sum_{j\in\mathcal{N}} \pi_j$. But, by the definition of an equilibrium, this maximization is subject to the behavior of the firms. Any equilibrium allocation therefore solves

$$\max_{\{x_{ij}, T_{ij}\}_{i,j}} \sum_{j\in\mathcal{N}} \left\{ \max_{\{p_j, c_j, l_j, \theta_j\}} \pi_j \left(p_j, c_j, l_j, \theta_j, \{x_{ij}\}_{ij}\right) \text{ s.t. (34) and (35)} \right\}. \tag{65}$$

It is, however, equivalent to let the coalition itself directly optimize over $\{p_j, c_j, l_j, \theta_j\}_j$. To see this, notice that, conditional on the arrangement, the inner maximization problems in (65) are all independent from each other. In other words, the decisions of a firm $i$ have no effect on the profit of a firm $j$ as long as the contracts specified by the arrangement are fulfilled. As a result, we can write (65) as $\max_{\{x_{ij}\}_{ij}, \{c_j, l_j, \theta_j\}_j} \sum_{j\in\mathcal{N}} \pi_j$ subject to the constraints (34) and (35) for all firms. By including the household's demand curves directly in the objective function, and by using the definition of $\pi_j$, the absence of dominating deviations therefore implies that the allocation must solve

$$\max_{\{x_{ij}\}_{ij}, \{c_j, l_j, \theta_j\}_j} C^{\frac{1}{\sigma}} P \sum_{j\in\mathcal{N}} \beta_j^{\frac{1}{\sigma}} c_j^{\frac{\sigma-1}{\sigma}} - w \sum_{j\in\mathcal{N}} (l_j + \theta_j f_j)$$

subject to (34) for all $j \in \mathcal{N}$, and where $C$ and $P$ are taken as given. Now, by Lemma 10 this problem is equivalent to an alternative problem in which the coalition maximizes $\left(\sum_{j\in\mathcal{N}} \beta_j^{\frac{1}{\sigma}} c_j^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}$ subject to $\sum_j l_j + \theta_j f_j \leq 1$ in addition to the other constraints.[49] This reformulated problem is identical to the problem $\mathcal{P}$ of the social planner such that any stable equilibrium must be efficient. □

[49] The corresponding function $g$ is $g(x) = x^{\frac{\sigma}{\sigma-1}}$. The constraints (34) can be included directly in the function $f$ in Lemma 10 by setting $f = -\infty$ for points outside the constraint set.