Endogenous Production Networks under Supply Chain Uncertainty^{*}

Alexandr Kopytov University of Hong Kong Bineet Mishra Cornell University

Kristoffer Nimark Cornell University Mathieu Taschereau-Dumouchel Cornell University

November 18, 2022

Abstract

Supply chain disturbances can lead to substantial increases in production costs. To mitigate these risks, firms may take steps to reduce their reliance on volatile suppliers. We construct a model of endogenous network formation to investigate how these decisions affect the structure of the production network and the level and volatility of macroeconomic aggregates. When uncertainty increases in the model, producers prefer to purchase from more stable suppliers, even though they might sell at higher prices. The resulting reorganization of the network leads to less macroeconomic volatility, but at the cost of a decline in aggregate output. The model also predicts that more productive and stable firms have higher Domar weights—a measure of their importance as suppliers—in the equilibrium network. We calibrate the model to U.S. data and find that the mechanism can account for a sizable decline in expected GDP during periods of high uncertainty like the Great Recession.

JEL Classifications: E32, C67, D57, D80, D85

^{*}We are grateful to the editor (Chad Jones), three anonymous referees, Ivan Alfaro, Matteo Bizzarri, Hafedh Bouakez, Ryan Chahrour, Roger Farmer, Simon Gilchrist, Yan Ji, Pablo Kurlat, Yueyuan Ma, Veronica Rappoport, Pascual Restrepo, David Zeke, and participants at various seminars and conferences for helpful comments and suggestions. This paper was previously circulated under the title "Endogenous Production Networks under Uncertainty". We gratefully acknowledge financial support from the Cornell Center for Social Sciences.

1 Introduction

Firms rely on complex supply chains to provide the intermediate inputs that they need for production. These chains can be disrupted by natural disasters, wars, trade barriers, changes in regulations, congestion in transportation links, etc. These shocks can then propagate to the rest of the economy through input-output linkages, resulting in aggregate fluctuations. However, individual firms may also take steps that mitigate such propagation by reducing their reliance on risky suppliers. In this paper, we study how this kind of mitigating behavior affects an economy's production network and, through that channel, macroeconomic aggregates.

Supply chain disruptions are one of the key challenges that business executives face, and firms devote substantial resources to reduce these risks (Ho et al., 2015). The COVID-19 pandemic provides a good example of how uncertainty can affect supply relationships. After the onset of the pandemic, many firms realized that their supply chains were exposed to substantially more risk than they previously thought. In a recent survey of firm managers, seventy percent agreed that the pandemic pushed companies to favor higher supply chain resiliency instead of simply purchasing from the lowest-cost supplier. Many also reported that they plan to diversify their supply chains across suppliers and geographies.¹

To study how supply chain uncertainty affects firms' sourcing decisions and how, in turn, these decisions affect the macroeconomy, we construct a macroeconomic model of endogenous network formation. In the model, firms produce differentiated goods that can be consumed by a risk-averse representative household or used as intermediate inputs by other producers. Firms can produce their goods in different ways, which we refer to as production techniques. A technique is a production function that specifies which intermediate inputs to use and how these inputs are to be combined. Techniques can also differ in terms of productivity. When choosing a production technique, a firm can marginally adjust the importance of a supplier or drop that supplier altogether. As a result, these decisions, when aggregated, lead to changes in the production network along both the intensive and extensive margins.

After production techniques have been chosen, firms are subject to random productivity shocks. They can then adjust how much they produce and the quantity of inputs that they use, subject to the constraints imposed by their selected technique. Competitive pressure between producers implies that the productivity shocks, as they affect production costs, are reflected in prices.

Importantly, firms' beliefs about the distribution of sectoral productivities can influence their choice of production technique and, thus, the structure of the network. Since firms are owned by the representative household, they compare profits across different states of the world using its

¹Survey by Foley & Lardner LLP, available online at https://www.foley.com/-/media/files/insights/ publications/2020/09/foley-2020-supply-chain-survey-report-1.pdf. See also Wagner and Bode (2008) and Zurich Insurance Group (2015) for other surveys documenting the importance of supply chain risks. Alessandria et al. (2022) investigate the impact of supply chain disturbances in the context of the COVID pandemic.

stochastic discount factor. As such they inherit the household's attitude toward risk. Consequently, while a firm would generally prefer to purchase from a more productive firm, it might decide not to do so if this firm is also more risky. Such a firm would sell at a lower price on average, but it is also more likely to suffer from a large negative productivity shock, in which case the price of its good would rise substantially. Potential customers take this possibility into account and balance concerns about average productivity and stability when choosing a production technique.

As an example, consider a car manufacturer that must decide what materials to use as inputs. If carbon fiber prices are expected to increase or to be more volatile, it may instead use steel for some components. If the change is large enough, it may switch away from using carbon fiber altogether, in which case the link between the car manufacturer and its carbon fiber supplier would disappear.

We prove that there always exists an efficient equilibrium in this environment, so that the implied equilibrium production network can be understood as resulting from a social planner maximizing the utility of the representative household.² That network thus optimally balances a higher level of expected GDP against a lower variance, with the relative importance of these two objectives being determined by the household's risk aversion. We further show that in the efficient equilibrium the importance of a producer (as measured by its sales share or Domar weight) increases in response to (1) an increase in the expected value of its productivity, or (2) a decrease in the variance of its productivity.

The model features a novel mechanism through which uncertainty can lower expected aggregate output. In the presence of uncertainty, firms prefer stable input prices and, as a result, move away from suppliers that are expected to be the most productive in favor of producers that are less susceptible to risk. This flight to safety implies that less productive producers gain in importance, and aggregate productivity and GDP fall as a result. On the flip side, this supply chain reshuffling leads to a more resilient network that dampens the effect of shocks and reduces aggregate fluctuations.

In some circumstances, our model mechanisms can have counterintuitive implications for how the productivity process affects aggregate quantities. While an increase in expected productivity or a decline in volatility always have a positive effect on welfare, their impact on expected GDP can be the opposite of what one would expect. For instance, an *increase* in expected productivity can lead to a *decline* in expected GDP, so that Hulten's (1978) theorem is not a good guide to understanding changes in GDP, even as a first-order approximation. To understand why, consider a producer with (on average) low but stable productivity. Its high output price makes it unattractive as a supplier. But if its expected productivity increases, its risk-reward profile improves, and other producers might begin to purchase from it. Doing so, they might move away from more productive—but also riskier—producers and, as a result, expected GDP might fall. We show that a similar mechanism is also at work for the variance of shocks, such that an increase in the volatility of a firm's productivity

 $^{^{2}}$ While we have been unable to find multiple equilibria in the model, we do not have a proof of uniqueness. We focus throughout on the efficient equilibrium which is generically unique.

can lead to a decline in the variance of aggregate output.

To evaluate the quantitative importance of allowing firms to adjust their production techniques in response to changes in beliefs, we calibrate the model using sectoral data for the United States. The model matches salient properties of the U.S. input-output structure such as the average and the standard deviation of sectoral Domar weights reasonably well. The calibrated economy is also able to replicate key features of the data that speak to the importance of beliefs for the structure of the production network. In particular, the Domar weight of a sector is positively correlated with its expected productivity and negatively correlated with its volatility. This evidence suggests that firms move away from uncertain suppliers in the data, as is predicted by the model.

We then use the calibrated model to evaluate the importance of the changing structure of the production network for macroeconomic aggregates. For this exercise, we first compare our baseline calibration to an alternative economy in which the production network is kept fixed, so that firms cannot move away from suppliers that become unproductive or volatile. We find that aggregate output is about 2.1% lower in this case, so that the endogenous response of the network can have a significant impact on welfare. This finding suggests that policies that impede the reorganization of the network (for instance, trade barriers) might have a sizable adverse effect.

To isolate the impact of uncertainty, we also compare our calibrated model to an alternative economy in which firms are unconcerned about risk when making sourcing decisions. While this economy is similar to the calibrated one during normal times, significant discrepancies appear during high-volatility periods, such as the Great Recession. During that episode, we find that firms respond to uncertainty by moving to safer but less productive suppliers. Taken together, these decisions lead to a 2.4% reduction in the volatility of GDP. However, the added stability comes at the cost of a 0.25% additional decline in expected GDP. Interestingly, this increase in resilience appears to have paid off ex post: According to our estimates, *realized* GDP in the baseline economy is 2.7% higher during the Great Recession compared to the economy in which firms did not adjust their techniques in response to the increased uncertainty. We also find a significant role for uncertainty when comparing our calibrated model to an alternative economy in which firms know the realization of the productivity shocks before deciding which production technique to adopt.

The model that we use for our quantitative analysis relies on some simplifying assumptions for tractability. To verify the robustness of our findings, we provide some empirical evidence that does not rely on the structure of the model. Taking advantage of rich firm-level data, we find that, as in the model, higher uncertainty leads to a decline in Domar weight, and that network connections involving riskier suppliers are more likely to break down. These results are robust to using different measures of uncertainty and instruments to tease out exogenous variation in uncertainty.

Our work is related to a large literature that investigates the impact of uncertainty on macroeconomic aggregates (Bloom, 2009, 2014; Bloom et al., 2018). In this paper, we propose a novel mechanism through which uncertainty can lower expected GDP. This mechanism operates through a flight to safety process in which firms facing higher uncertainty switch to safer but less productive suppliers, leading to lower but less volatile GDP. In a recent paper, David et al. (2022) argue that uncertainty may lead capital to flow to firms that are less exposed to aggregate risk, rather than to those firms where it would be most productive. In their model, as in ours, uncertainty leads to lower aggregate output and measured TFP.³

There is a large and growing literature that studies how shocks propagate through production networks, in the spirit of early contributions by Long and Plosser (1983), Dupor (1999) and Horvath (2000). Acemoglu et al. (2012) derive conditions on input-output networks under which idiosyncratic shocks result in aggregate fluctuations even when the number of producers is large.⁴ Acemoglu et al. (2017) and Baqaee and Farhi (2019a) describe conditions under which production networks can generate fat-tailed aggregate output. Foerster et al. (2011) and Atalay (2017) study the empirical contributions of sectoral shocks for aggregate fluctuations. The mechanisms studied in these papers are also present in our model. Carvalho and Gabaix (2013) argue that the reduction in aggregate volatility during the Great Moderation (and its potential recent undoing) can be explained by changes in the input-output network.⁵

In most of the literature, Hulten's (1978) theorem applies, so that sales shares are a sufficient statistic to predict the impact of microeconomic shocks on macroeconomic aggregates. In contrast, since firms can adjust production techniques ex ante, in our model Hulten's theorem is not a useful guide to how shocks affect expected GDP, even as a first-order approximation.⁶ An increase in expected sectoral productivity can even have a negative impact on expected GDP.

Our paper is not the first to study the endogenous formation of production networks. Closely related to our work are earlier papers by Oberfield (2018) and Acemoglu and Azar (2020) in which the input-output network emerges as the outcome of individual technique choice by firms. We adopt a similar network formation process but allow for uncertainty and beliefs to influence the structure of the network. Taschereau-Dumouchel (2020), Acemoglu and Tahbaz-Salehi (2020) and Elliott et al. (2022) study economies in which the firms' decisions to operate or not shape the production network. Lim (2018) and Huneeus (2018) evaluate the importance of endogenous changes in the production network for business cycle fluctuations. Dhyne et al. (2021) build a model of endogenous network formation and international trade. Boehm and Oberfield (2020) estimate a

³Fernández-Villaverde et al. (2011) investigate the real impact of interest rate volatility for emerging economies. Jurado et al. (2015) provide econometric estimates of time-varying macroeconomic uncertainty. Baker et al. (2016) measure economic policy uncertainty based on newspaper coverage. Nieuwerburgh and Veldkamp (2006) and Fajgelbaum et al. (2017) develop models in which uncertainty can have long-lasting impacts on economic aggregates.

⁴Production networks are one mechanism through which granular fluctuations can emerge (Gabaix, 2011).

⁵Other works have looked at the importance of production networks outside of the business cycle literature. Jones (2011) investigates their importance to explain the income difference between countries. Barrot and Sauvagnat (2016), Boehm et al. (2019) and Carvalho et al. (2021) study the propagation of shocks after natural disasters.

⁶Baqaee and Farhi (2019a) investigate departures from Hulten's theorem due to higher-order effects of shocks. Recent work that has studied production networks under distortions, where Hulten's theorem generally does not hold, include Baqaee (2018), Liu (2019), Baqaee and Farhi (2019b), Bigio and La'O (2020) and Caliendo et al. (2022).

network formation model using Indian micro data to study misallocation in input markets. Bernard et al. (2022) build a model of network formation to explain firm heterogeneity.⁷ A key distinguishing feature of our work is its focus on how uncertainty affects the structure of the production network.

Several papers in the network literature endow firms with CES production functions, so that the input-output matrix varies with factor prices. Our model generates endogenous changes in the production network through a different mechanism, which is closer to Oberfield (2018) and Acemoglu and Azar (2020). In contrast to the standard CES setup, our model allows links between sectors to be created or destroyed. In addition, the existing literature using CES production network models has not studied how uncertainty and beliefs shape production networks, and introducing such mechanisms while keeping the model tractable is not straightforward.

The next section introduces our model of network formation under uncertainty. In Section 3, we first characterize the equilibrium when the network is kept fixed. We then consider the full equilibrium with a flexible network in Section 4. In Section 5, we describe the mechanisms at work in the model. In Section 6, we calibrate the model to U.S. data. Section 7 provides additional empirical evidence in support of the mechanisms. The last section concludes. All proofs are in Appendix E.

2 A model of endogenous network formation under uncertainty

We study the formation of production networks under uncertainty in a multi-sector economy. Each sector is populated by a representative firm that produces a differentiated good that can be used either as an intermediate input or for consumption. To produce, each firm must choose a production technique, which specifies a set of inputs to use. Firms are owned by a risk-averse representative household and are subject to sector-specific productivity shocks. Since firms choose production techniques before these shocks are realized, the probability distribution of the shocks affects the input-output structure of the economy.

2.1 Firms and production functions

There are *n* sectors, indexed by $i \in \{1, ..., n\}$, each producing a differentiated good. In each sector, there is a representative firm that behaves competitively so that equilibrium profits are always zero. When this creates no confusion, we use sector *i*, product *i* and firm *i* interchangeably.

As in Oberfield (2018) and Acemoglu and Azar (2020), the representative firm in sector *i* has access to a set of production techniques \mathcal{A}_i . A technique $\alpha_i \in \mathcal{A}_i$ specifies the set of inputs that are used in production, how these inputs are to be combined, and a productivity shifter $A_i(\alpha_i)$.

⁷Atalay et al. (2011) show that a modified "preferential attachment" model can fit features of the U.S. firm-level production network. Carvalho and Voigtländer (2014) build a rule-based model of network formation to study the adoption and diffusion of intermediate inputs. Kopytov (2018) constructs a model to study financial interconnected-ness and systemic risk under uncertainty.

We model these techniques as Cobb-Douglas technologies that can vary in terms of factor shares and total factor productivity. It is therefore convenient to identify a technique $\alpha_i \in \mathcal{A}_i$ with the intermediate input shares associated with that technique, $\alpha_i = (\alpha_{i1}, \ldots, \alpha_{in})$, and to write the corresponding production function as

$$F(\alpha_i, L_i, X_i) = e^{\varepsilon_i} A_i(\alpha_i) \zeta(\alpha_i) L_i^{1-\sum_{j=1}^n \alpha_{ij}} \prod_{j=1}^n X_{ij}^{\alpha_{ij}}, \qquad (1)$$

where L_i is labor and $X_i = (X_{i1}, \ldots, X_{in})$ is a vector of intermediate inputs. The term ε_i is the stochastic component of sector *i*'s total factor productivity. Finally, $\zeta(\alpha_i)$ is a normalization to simplify future expressions.⁸

Since a technique α_i corresponds to a vector of factor shares, we define the set of feasible production techniques \mathcal{A}_i for sector i as $\mathcal{A}_i = \left\{ \alpha_i \in [0,1]^n : \sum_{j=1}^n \alpha_{ij} \leq \overline{\alpha}_i \right\}$, where $0 < 1 - \overline{\alpha}_i < 1$ provides a lower bound on the share of labor in the production of good i. We denote by $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$ the Cartesian product of the sets $\{\mathcal{A}_1, \ldots, \mathcal{A}_n\}$, such that an element $\alpha \in \mathcal{A}$ corresponds to a set of input shares for each sector. As such, it fully characterizes the production network and firms can influence the structure of this network through their choice of techniques. Importantly, the set \mathcal{A} allows firms to adjust the importance of a supplier at the margin or to not use a particular input at all by setting the corresponding share to zero. The model is therefore able to capture network adjustments along both the intensive and extensive margins.⁹

The choice of technique influences the total factor productivity of sector *i* through the term $A_i(\alpha_i)$ in (1). This term is given by nature and represents how effective some combinations of inputs are at producing a given good. For instance, beach towels and flowers are not very useful when making a car, and a technique that relies only on these inputs would have a low A_i . In contrast, a technique that would use aluminum, steel, car engines, etc. would be associated with a higher productivity. When deciding on its optimal production technique a firm will take A_i into account, but it will also evaluate the expected level and volatility of each input price.

We impose the following structure on $A_i(\alpha_i)$.

Assumption 1. $A_i(\alpha_i)$ is smooth and strictly log-concave.

This assumption is both technical and substantial in nature. The strict log-concavity ensures that there exists a unique technique that solves the optimization problem of the firm. It also implies that, for each sector *i*, there is a set of *ideal* input shares α_{ij}° that maximize A_i and that represent the most efficient way to combine intermediate inputs to produce good *i*.

⁸Namely, $\zeta(\alpha_i) = \left[\left(1 - \sum_{j=1}^n \alpha_{ij} \right)^{1 - \sum_{j=1}^n \alpha_{ij}} \prod_{j=1}^n \alpha_{ij}^{\alpha_{ij}} \right]^{-1}$. This normalization is useful to simplify the unit cost expression, given by (8) below. $\zeta(\alpha_i)$ could instead be included in $A_i(\alpha_i)$ without any impact on the model.

⁹This is in contrast to standard network models with CES production. In those models, the share of an input can fluctuate but can never reach zero. As a result, these models cannot generate the destruction or creation of links observed in the sectoral data that we use in Section 6 and in the firm-level network data studied in Section 7.

Example. One example of a function $A_i(\alpha_i)$ that satisfies Assumption 1 and that we will use in the quantitative part of the paper is the quadratic form

$$\log A_i\left(\alpha_i\right) = a_i^\circ - \sum_{j=1}^n \kappa_{ij} \left(\alpha_{ij} - \alpha_{ij}^\circ\right)^2 - \kappa_{i0} \left(\sum_{j=1}^n \alpha_{ij} - \sum_{j=1}^n \alpha_{ij}^\circ\right)^2,\tag{2}$$

where $\alpha_i^{\circ} = (\alpha_{i1}^{\circ}, \dots, \alpha_{in}^{\circ})$ is the vector of ideal TFP-maximizing input shares and α_i° is log TFP at the ideal shares. The parameter $\kappa_{ij} > 0$ determines the cost, in terms of productivity, of moving the *j*th input share α_{ij} away from its ideal share α_{ij}° . The last term captures the productivity penalty of deviating from an ideal labor share.

The distribution of the sectoral productivity shock ε_i in (1) is a key primitive of the model and an important input to firms' technique choices. We collect the sectoral shocks in the vector $\varepsilon = (\varepsilon_1, \ldots, \varepsilon_n)$, which we assume to be normally distributed, $\varepsilon \sim \mathcal{N}(\mu, \Sigma)$. The vector μ determines the expected level of sectoral productivities. The covariance matrix Σ , with typical element Σ_{ij} , determines both uncertainty about individual elements of ε , as well as their correlation across industries. We assume throughout that Σ is positive definite. The vector ε is the only source of uncertainty in this economy.¹⁰

In equilibrium, ε will have a direct impact on prices, and the moments (μ, Σ) will affect expectations about the price system. For instance, a sector with a high μ_i will have a low unit cost and therefore the price of good *i* will be low in expectation. Similarly, a high Σ_{ii} implies large productivity shocks and a volatile price. Since production techniques must be chosen before ε is realized, the beliefs (μ, Σ) affect the sourcing decisions of the firms. Returning to the example from the introduction, if carbon fiber prices are expected to increase or to be more volatile, a car manufacturer might switch to using steel instead for some components. If the change is large enough, the manufacturer may switch away from using carbon fiber altogether, in which case the link with carbon fiber suppliers would disappear from the production network.

Importantly, we impose the restriction that the representative firm in sector i can only adopt a single production technique α_i . Without this restriction, the firm would set up a continuum of individual plants, each with its own technique, to cover the whole set of available techniques \mathcal{A}_i . After the realization of the productivity shocks ε , the firm would only operate the plant that is best suited to the specific ε draw. All the other plants would remain idle. In reality, we think that fixed costs would prevent the firm from setting up all these plants. Information frictions might also impede the reallocation of sectoral demand to the best suited technique. To avoid burdening the

¹⁰This choice of distribution implies that the expectation $E[exp(\varepsilon)]$ of the vector of TFP shocks depends on the covariance matrix Σ . This assumption is common in the literature but implies that an increase in the variance Σ_{ii} of a sector *i* has a beneficial impact on its expected TFP. Through this channel the adverse effect of an increase in uncertainty is mitigated. One common way to correct for this effect is to remove half of the variance of ε from its mean. In Appendix I, we describe why such a correction is problematic in our model. We also discuss other potential corrections.

exposition of the model, we adopt this restriction in an ad hoc fashion here, but provide a possible microfoundation for it in Appendix A.

2.2 Household preferences

A risk-averse representative household supplies one unit of labor inelastically and chooses a consumption vector $C = (C_1, \ldots, C_n)$ to maximize

$$u\left(\left(\frac{C_1}{\beta_1}\right)^{\beta_1} \times \dots \times \left(\frac{C_n}{\beta_n}\right)^{\beta_n}\right),\tag{3}$$

where $\beta_i > 0$ for all i and $\sum_{i=1}^n \beta_i = 1$.¹¹ We refer to $Y = \prod_{i=1}^n (\beta_i^{-1}C_i)^{\beta_i}$ as aggregate consumption or, equivalently in this setting, GDP. The utility function u is CRRA with a coefficient of relative risk aversion $\rho \ge 1$.¹² The household makes consumption decisions after uncertainty is revealed and so in each state of the world it faces the budget constraint

$$\sum_{i=1}^{n} P_i C_i \le 1,\tag{4}$$

where P_i is the price of good *i* and where we use the wage as numeraire so that W = 1.

Firms are owned by the representative household and maximize expected profits discounted by the household's stochastic discount factor

$$\Lambda = u'(Y) \times 1/\overline{P},\tag{5}$$

where $\overline{P} = \prod_{i=1}^{n} P_i^{\beta_i}$ is the price index. The stochastic discount factor captures how much an extra unit of the numeraire contributes to the utility of the household in different states of the world.

From the optimization problem of the household it is straightforward to show that

$$y = -\beta' p, \tag{6}$$

where $y = \log Y$, $p = (\log (P_1), \ldots, \log (P_n))$ and $\beta = (\beta_1, \ldots, \beta_n)$.¹³ Log GDP is thus the negative of the sum of log prices weighted by the consumption shares β . Intuitively, as prices become lower

¹¹Several assumptions of the model can be relaxed at the cost of extra complications in the proofs. For instance, the model can handle 1) $\beta_i = 0$ for some goods, 2) Σ that is only positive semi-definite, and 3) TFP shifter functions $\{A_i\}_{i=1}^n$ that are only log concave.

¹²The case $0 < \rho < 1$ is straightforward to characterize but is somewhat unnatural since the household then seeks to increase the variance of log consumption. To see this, note that when log Y is normal, maximizing $E\left[Y^{1-\rho}\right]$ amounts to maximizing $E\left[\log Y\right] - \frac{1}{2}(\rho - 1) V\left[\log Y\right]$ such that $\rho \leq 1$ indicates whether the household likes uncertainty in log consumption or not. This is a consequence of the usual increase in the mean of a log-normal variable from an increase in the variance of the underlying normal variable. See Appendix I for a version of the model in which we correct for this term.

 $^{^{13}}$ See Appendix F.1 for a derivation of this equation and of the stochastic discount factor (5).

relative to wages, the household can purchase more goods and aggregate consumption increases. Equation (6) also implies that it is sufficient to derive the vector of prices to determine GDP.

2.3 Unit cost minimization

We solve the problem of a representative firm in two stages. In the first stage, the firm decides which production technique to use. Importantly, this choice is made before the random productivity vector ε is realized. In contrast, consumption, labor and intermediate inputs are chosen (and their respective markets clear) in the second stage, after the realization of ε . This timing captures that production techniques take time to adjust, as they might involve retooling a plant, teaching new processes to workers, negotiating contracts with new suppliers, etc. We begin by solving the second stage problem, by deriving the optimal input choice for a given production technique α_i . The resulting expressions are then used to solve the firm's first-stage problem of choosing α_i .

Under a given technique α_i , the cost minimization problem of a firm in sector *i* is

$$K_i(\alpha_i, P) = \min_{L_i, X_i} \left(L_i + \sum_{j=1}^n P_j X_{ij} \right), \text{ subject to } F(\alpha_i, L_i, X_i) \ge 1,$$
(7)

where $P = (P_1, \ldots, P_n)$ is the price vector, L_i is the labor input and $X_i = (X_{i1}, \ldots, X_{in})$ is the vector of intermediate inputs.

The solution to this problem implicitly defines the unit cost of production $K_i(\alpha_i, P)$, which plays an important role in our analysis. Since, for a given α_i , the firm operates a constant returns to scale technology, K_i does not depend on the scale of the firm and is only a function of the (relative) prices P. It is straightforward to show (and we do so in Appendix F.2) that with the production function (1) the unit cost function is

$$K_i(\alpha_i, P) = \frac{1}{e^{\varepsilon_i} A_i(\alpha_i)} \prod_{j=1}^n P_j^{\alpha_{ij}}.$$
(8)

Equation (8) is the standard unit cost for a Cobb-Douglas production function. It states that the cost of producing one unit of good *i* is equal to the geometric mean of the individual input prices (weighed by their respective shares) and adjusted for sectoral total factor productivity. The unit cost $K_i(\alpha_i, P)$ therefore rises when inputs become more expensive and declines when total factor productivity increases.

2.4 Technique choice

Given an expression for the unit cost of production, the first stage of the representative firm's problem is to pick a technique $\alpha_i \in \mathcal{A}_i$ to maximize expected profits, that is,

$$\alpha_i^* \in \arg \max_{\alpha_i \in \mathcal{A}_i} \mathbb{E} \left[\Lambda Q_i \left(P_i - K_i \left(\alpha_i, P \right) \right) \right].$$
(9)

Here, Q_i is the equilibrium demand for good *i*, and profits in different states of the world are weighed by the stochastic discount factor Λ of the household. The representative firm takes prices P, demand Q_i and the stochastic discount factor Λ as given and so the only term in (9) over which it has any control is the unit cost $K_i(\alpha_i, P)$. Problem (9) can therefore be written as

$$\alpha_i^* \in \arg\min_{\alpha_i \in \mathcal{A}_i} \mathbb{E}\left[\Lambda Q_i K_i\left(\alpha_i, P\right)\right].$$
(10)

The firm thus selects the technique $\alpha_i \in \mathcal{A}_i$ that minimizes the expected discounted value of the total cost of goods sold $Q_i K_i(\alpha_i, P)$, while taking into consideration that final consumption goods are valued differently across different states of the world, as captured by Λ . Because profits are discounted by Λ , firms effectively inherit the risk aversion of the representative household.

2.5 Equilibrium conditions

In equilibrium, competitive pressure from other firms in the same sector pushes prices to be equal to unit costs so that

$$P_i = K_i(\alpha_i, P) \text{ for all } i \in \{1, \dots, n\}.$$

$$(11)$$

For a given network $\alpha \in \mathcal{A}$, this equation, together with (8), allows us to fully characterize the price system as a function of the random productivity shocks ε .¹⁴

An equilibrium is defined by the optimality conditions of both the household and the firms holding simultaneously, together with the usual market clearing conditions.

Definition 1. An equilibrium is a choice of technique $\alpha^* = (\alpha_1^*, \ldots, \alpha_n^*)$ and a stochastic tuple $(P^*, C^*, L^*, X^*, Q^*, \Lambda^*)$ such that

- 1. (Optimal technique choice) For each $i \in \{1, ..., n\}$, the technique choice $\alpha_i^* \in \mathcal{A}_i$ solves (9) given prices P^* , demand Q_i^* and the stochastic discount factor Λ^* given by (5).
- 2. (Optimal input choice) For each $i \in \{1, ..., n\}$, factor demands per unit of output L_i^*/Q_i^* and X_i^*/Q_i^* are a solution to (7) given prices P^* and technique choice α_i^* .

¹⁴Even without imposing that production techniques are Cobb-Douglas, the system (11) yields a unique price vector P under standard assumptions. But the Cobb-Douglas structure implies that we can write the *distribution* of P in closed form, which allows us to characterize the technique choice problem in a tractable way.

- 3. (Consumer maximization) The consumption vector C^* maximizes (3) subject to (4) given prices P^* .
- 4. (Unit cost pricing) For each $i \in \{1, \ldots, n\}$, P_i^* solves (11) where $K_i(\alpha_i^*, P^*)$ is given by (8).
- 5. (Market clearing) For each $i \in \{1, \ldots, n\}$,

$$C_i^* + \sum_{j=1}^n X_{ji}^* = Q_i^* = F_i(\alpha_i^*, L_i^*, X_i^*), \text{ and } \sum_{i=1}^n L_i^* = 1.$$
(12)

Conditions 2 to 5 correspond to the standard competitive equilibrium conditions for an economy with a fixed production network. They imply that firms and the household optimize in a competitive environment and that all markets clear given equilibrium prices. Condition 1 emphasizes that the production techniques, and hence the production network represented by the matrix α^* , are equilibrium objects that depend on the primitives of the economy.

It is straightforward to extend the model along several dimensions without losing tractability. For instance, the model can accommodate disturbances that happen at the link level instead of at the sectoral level. To do so, we can simply think of a link between two producers as a fictitious "transport" sector that is also subject to shocks. It is also straightforward to extend the model to include multiple primary factors, like different types of labor, or wedges that create a gap between unit costs and prices. We work out this last extension in Appendix K.

On the other hand, certain ingredients are essential to keep the model tractable. Here the key challenge comes from the fact that technique choices affect equilibrium prices which in turn affect technique choices. The log-linearity implied by the Cobb-Douglas aggregators in (1) and (3) are needed to keep the equilibrium beliefs tractable. While this implies a unit elasticity of substitution in the production function (1), this elasticity only captures the response of intermediate inputs to realized prices conditional on a chosen production technique. Since firms' expectations affect their technique choice, the model is able to handle richer substitution patterns between expected prices and intermediate inputs, as we explore in more details in Section 5.

3 Equilibrium prices and GDP in a fixed-network economy

Before analyzing how the equilibrium production network α^* responds to changes in the environment, it is useful to first establish how prices and GDP depend on productivity under a fixed network. To this end, we first define two variables that will play a central role in our analysis. The first is the Leontief inverse $\mathcal{L}(\alpha) = (I - \alpha)^{-1}$, which can also be written as the geometric sum $\mathcal{L}(\alpha) = I + \alpha + \alpha^2 + \dots^{15}$ The element i, j of $\mathcal{L}(\alpha)$ captures the importance of sector j as an

¹⁵This last expression makes clear that \mathcal{L} is non-negative. Since $\sum_{j=1}^{n} \alpha_{ij} \leq \overline{\alpha}_i < 1$, $I - \alpha$ is strictly diagonally dominant and therefore invertible.

input in the production of good *i* by taking into account direct and indirect connections between them in the production network. It is also useful to define the Domar weight ω_i of sector *i* as the ratio of its sales to nominal GDP, i.e. $\omega_i = \frac{P_i Q_i}{P'C}$. As we show in Appendix E.2, this quantity is equal to $\omega_i(\alpha) = \beta' \mathcal{L}(\alpha) \mathbf{1}_i > 0$ in the model, where $\mathbf{1}_i$ is a column vector with a 1 as *i*th element and zeros elsewhere. Domar weights combine the preferences of the household with the Leontief inverse to provide an overall measure of the importance of a sector as a supplier in the aggregate economy. Domar weights are constant in a fixed-network economy but vary when firms are free to adjust their sourcing decisions.

With these definitions in hand, we establish a first result that links the vector of sectoral productivities with prices and GDP.

Lemma 1. For a fixed production network α ,

$$p(\alpha) = -\mathcal{L}(\alpha) \left(\varepsilon + a(\alpha)\right), \tag{13}$$

and

$$y(\alpha) = \omega(\alpha)'(\varepsilon + a(\alpha)), \qquad (14)$$

where $a(\alpha) = (\log A_i(\alpha_i), \dots, \log A_n(\alpha_n))$ and $\omega(\alpha) = (\omega_1(\alpha), \dots, \omega_n(\alpha)).$

Proof. All proofs are in Appendix E.

Lemma 1 describes how prices and GDP depend on 1) the vector of firm-level productivities and 2) the production network. We will describe both channels in turn.

First, consider the impact of the vector of log productivities $\varepsilon + a(\alpha)$. Since all the elements of $\omega(\alpha)$ and $\mathcal{L}(\alpha)$ are non-negative, an increase in productivity has a negative impact on log prices and a positive impact on log GDP. Intuitively, as firms become more productive, their unit costs decline, and competition forces them to sell at lower prices. From the perspective of GDP, higher productivity implies that the available labor can be transformed into more consumption goods. Second, the lemma makes clear that production techniques α matter for prices and GDP through two distinct channels. They have a direct impact on the productivity shifters $a(\alpha)$ because different techniques have different productivities. In addition, α affects prices and GDP through its impact on the Leontief inverse and the Domar weights. The matrix $\mathcal{L}(\alpha) = I + \alpha + \alpha^2 + \ldots$ in (13) implies that the price of good *i* depends not only on the productivity of producer *i* itself, but also on the productivity of its suppliers, and on the productivity of sectoral productivity on aggregate output depends on the sector's importance, as captured by its Domar weight.

Lemma 1 also shows that p and y are linear functions of the productivity vector ε and, as a result, inherit the normality of ε . It follows that the first and second moments of log GDP can be

written as

$$E[y(\alpha)] = \omega(\alpha)'(\mu + a(\alpha)) \text{ and } V[y(\alpha)] = \omega(\alpha)'\Sigma\omega(\alpha).$$
(15)

It is clear from these expressions that the production network α , through its impact on the Domar weights $\omega(\alpha)$, matters for the mean and the variance of log GDP. In addition, note that the covariance matrix Σ only affects expected log GDP through its influence on the structure of the network. It follows that whenever we discuss the response of expected log GDP to a change in uncertainty, the mechanism must operate through the endogenous reorganization of the network.

We conclude this section with a simple corollary, already known in the literature, that describes the impact of beliefs on the mean and the variance of log GDP. In what follows, we use partial derivatives to emphasize that the network α is kept fixed.

Corollary 1. For a fixed production network α , the following holds.

1. The impact of a change in expected TFP μ_i on expected log GDP E[y] is given by

$$\frac{\partial \operatorname{E}\left[y\right]}{\partial \mu_{i}} = \omega_{i}$$

2. The impact of a change in volatility Σ_{ij} on the variance of log GDP V [y] is given by¹⁶

$$\frac{\partial \operatorname{V}[y]}{\partial \Sigma_{ij}} = \begin{cases} \omega_i^2 & i = j, \\ 2\omega_i \omega_j & i \neq j. \end{cases}$$

The first part of the lemma demonstrates that for a fixed production network, Hulten's (1978) celebrated theorem also holds in expectational terms. That is, the change in expected log GDP following a change in the expected productivity of a sector i is equal to that sector's sales share ω_i . The second part of the lemma establishes a similar result for changes in volatility. We see that the impact of an increase in the uncertainty of a sector's TFP on the variance of log GDP is equal to the square of that sector's sales share. The corollary also describes how aggregate volatility responds to a change in the correlation between two sectors. In this case, the increase in V [y] is proportional to the product of the two industries' sales shares. Since Domar weights are always positive, an increase in correlation always leads to higher aggregate volatility. Intuitively, positively correlated shocks are unlikely to offset each other, and their expected aggregate impact is therefore larger.

Finally, Corollary 1 emphasizes that for a fixed network knowing the sales shares of every industry is sufficient to compute the impact of changes in μ and Σ on the moments of log GDP. In

¹⁶Throughout the paper, derivatives with respect to off-diagonal elements of Σ simultaneously change Σ_{ij} and Σ_{ji} to preserve the symmetry of Σ .

Section 5, we show that this is no longer true when firms can adjust their input shares in response to changes in the distribution of sectoral productivity. In fact, when the network is free to adjust, an increase in an element of μ can even lead to a decline in expected log GDP.

4 Equilibrium production network

In the full equilibrium the production network responds endogenously to changes in beliefs. To explore the mechanisms involved, we begin by characterizing how firms select a production technique in this environment. We then establish that an equilibrium exists under general conditions. We also show that there exists an efficient equilibrium and that its associated production network is characterized by a trade-off between the expected level and the volatility of GDP.

4.1 Technique choice

In the previous section, we described prices under a given network. Here, we use that information to characterize the problem of the representative firm in sector *i* that must choose a technique $\alpha_i \in \mathcal{A}_i$. It is convenient to work with the log of the stochastic discount factor $\lambda(\alpha^*) = \log \Lambda(\alpha^*)$, the log of the unit cost $k_i(\alpha_i, \alpha^*) = \log K_i(\alpha_i, P^*(\alpha^*))$ and the log of sectoral demand $q_i(\alpha^*) = \log Q_i(\alpha^*)$, where α^* denotes equilibrium network. The following lemma shows that these objects are normally distributed and describes how they influence the firm's problem.

Lemma 2. In equilibrium, $\lambda(\alpha^*)$, $k_i(\alpha_i, \alpha^*)$ and $q_i(\alpha^*)$ are normally distributed, and the technique choice of the representative firm in sector *i* solves

$$\alpha_{i}^{*} \in \arg\min_{\alpha_{i} \in \mathcal{A}_{i}} \mathbb{E}\left[k_{i}\left(\alpha_{i}, \alpha^{*}\right)\right] + \operatorname{Cov}\left[\lambda\left(\alpha^{*}\right), k_{i}\left(\alpha_{i}, \alpha^{*}\right)\right].$$
(16)

The terms on the right-hand side of (16) capture how beliefs and uncertainty affect the production network. The first term implies that the firm prefers to adopt techniques that provide, in expectation, a lower unit cost of production. Taking the expected value of the log of (8), we can write this term as

$$\mathbf{E}\left[k_{i}\left(\alpha_{i},\alpha^{*}\right)\right] = -\mu_{i} - a_{i}\left(\alpha_{i}\right) + \sum_{j=1}^{n} \alpha_{ij} \mathbf{E}\left[p_{j}\right]$$

so that, unsurprisingly, the firm prefers techniques that have high productivity a_i and that rely on inputs that are expected to be cheap.

The second term in (16) captures the importance of aggregate risk for the firm's decision. It implies that the firm prefers to have a low unit cost in states of the world in which the marginal utility of aggregate consumption is high. As a result, the coefficient of risk aversion ρ of the household indirectly determines how risk averse firms are. We can expand this term as $\text{Cov} [\lambda, k_i] = \text{Corr} [\lambda, k_i] \sqrt{\mathcal{V}[\lambda]} \sqrt{\mathcal{V}[k_i]}$, which implies that the firm tries to minimize the correlation of its unit cost with λ . Furthermore, since prices and GDP tend to move in opposite directions (see Lemma 1), Corr $[\lambda, k_i]$ is typically positive, and so firms seek to minimize the variance of their unit cost.¹⁷ This has several implications for their choice of suppliers. To see this, it is convenient to write $V[k_i]$ in terms of input prices as

$$V[k_i(\alpha_i, \alpha^*)] = \sum_{j=1}^n \alpha_{ij}^2 V[p_j] + \sum_{j \neq k} \alpha_{ij} \alpha_{ik} \operatorname{Cov}[p_j, p_k] + 2 \operatorname{Cov}\left[-\varepsilon_i, \sum_{j=1}^n \alpha_{ij} p_j\right] + \Sigma_{ii}.$$
 (17)

The variance of the unit cost can thus be decomposed into four channels, shown on the right-hand side of (17). The first term implies that the firm prefers inputs that have stable prices. The second term implies that the firm avoids techniques that rely on inputs with positively correlated prices and instead prefers to diversify its set of suppliers and adopt inputs whose variation in prices offset each other. The third term implies that the firm prefers inputs whose prices are positively correlated with its own productivity shocks. When the firm experiences a negative productivity shock, the prices of its inputs are then more likely to also be low, reducing the expected increase in its unit cost. Finally, the last term captures the fact that more volatile productivity contributes to more volatile unit costs.

We will explore in more details how beliefs affect the structure of the network in general equilibrium in the next section, but for now it is useful to highlight some of the key forces that affect the choice of technique of a firm by considering the following partial equilibrium example.

Example (Sourcing decisions in partial equilibrium). Consider again a car manufacturer (firm i) that must decide on the share of steel (good 1) and carbon fiber (good 2) to use in production. We show in Figure 1 how the solution α_i^* to problem (16) is affected by changes in the mean and the variance of p_2 . Panels (a) and (b) show that when good 2 is expected to be more expensive, firm i lowers α_{i2} and increases α_{i1} . A similar mechanism is at work when uncertainty about p_2 increases, as seen in panels (c) and (d). When V $[p_2]$ is large, the firm prefers to use a larger share of the relatively safer good 1. Notice that the share α_{i2} reaches zero when p_2 is expected to be sufficiently large or uncertain. In that case, firm i severs the link with the carbon fiber supplier and an input/output relationship disappears from the production network. In this example, we have assumed that steel and carbon fiber are substitutes, so that their input shares tend to move in opposite directions. We show in Section 5.2 that the model can also accommodate inputs that are complements.

¹⁷If the productivity shock of a sector *i* is strongly negatively correlated with that of the other sectors, it can be that $\operatorname{Corr}[\lambda, k_i] < 0$, in which case the representative firm in sector *i* seeks to be *more* volatile, to insure the household in states of low consumption.

Figure 1: Beliefs and input shares



Notes: Parameters: $\rho = 5$, A_i is as in (2) with $\kappa_i = (1/3, 1/3, 1/3)$, $a_i^0 = 0$, $\alpha_i^\circ = (1/3, 1/3)$. Input prices are $(p_1, p_2) \sim \mathcal{N}(\mathbf{E}[p], \mathbf{V}[p])$ with diagonal covariance matrix, $\mathbf{E}[p_1] = 0$ and $\mathbf{V}[p_1] = 0$.

4.2 Equilibrium existence and efficiency

The example above demonstrates how an individual firm's technique choice responds to changes in beliefs about input prices. However, prices are equilibrium objects that depend on the production network and, therefore, on the choices made by other firms. Here, we first show that there exists an equilibrium that satisfies the conditions in Definition 1. We then show that there exists a solution to a social planner's problem that coincides with a decentralized equilibrium. In this equilibrium, the production network strikes an optimal balance between maximizing the mean level of log GDP and minimizing its variance.

Existence of an equilibrium

Lemma 2 describes a self-map $\mathcal{K} : \mathcal{A} \to \mathcal{A}$ that can be used to define an equilibrium network α^* . At a fixed point of this mapping, we have that $\alpha_i^* = \mathcal{K}_i(\alpha^*)$ for all $i \in \mathcal{N}$, where $\mathcal{K}_i(\alpha^*)$ is the right-hand side of (16). Hence, such a fixed point describes an equilibrium network. Given an equilibrium network, it is then straightforward to compute prices from (13). From there, all other equilibrium quantities can be uniquely determined. The following proposition shows that \mathcal{K} has a fixed point and that an equilibrium exists.

Proposition 1. An equilibrium exists.

The proof uses that \mathcal{K} is a continuous mapping on the compact set \mathcal{A} . From Brouwer's fixed point theorem, we then know that there exists at least one element $\alpha^* \in \mathcal{A}$ such that $\alpha^* = \mathcal{K}(\alpha^*)$.

While Proposition 1 guarantees the existence of an equilibrium, it is silent about the number of such equilibria. However, the next subsection demonstrates that there always exists an efficient equilibrium, which provides a natural benchmark to study since any inefficient equilibrium would be the result of coordination failure among agents. While such coordination failures may exist in reality, they are not the focus of this paper.

Equilibrium efficiency

There is a representative household in the economy, and hence finding the set of Pareto efficient allocations amounts to solving the problem of a social planner that maximizes the utility function (3) of the household, subject to the resource constraints (12). The next proposition demonstrates that a solution to the planner's problem exists, and that it corresponds to an equilibrium.

Proposition 2. There exists an efficient equilibrium.

From here on, our analysis will focus on the efficient equilibrium.¹⁸ One key advantage of Proposition 2 is that it allows us to investigate the properties of the equilibrium by solving the problem of the social planner directly. This last point implies that we can characterize the equilibrium network as the outcome of a welfare maximization problem.

Corollary 2. The efficient equilibrium production network α^* solves

$$\mathcal{W} \equiv \max_{\alpha \in \mathcal{A}} \mathbb{E}\left[y\left(\alpha\right)\right] - \frac{1}{2}\left(\rho - 1\right) \mathbb{V}\left[y\left(\alpha\right)\right],\tag{18}$$

where \mathcal{W} is a measure of the welfare of the household, and y is log GDP as defined in (14).¹⁹

Corollary 2 follows directly from the fact that, by Proposition 2, the equilibrium network α^* must maximize the expected utility of the representative household. It is clear from the objective function (18) that the household prefers networks that strike a balance between maximizing expected log GDP E [$y(\alpha)$] and minimizing aggregate uncertainty V [$y(\alpha)$], with the relative risk aversion parameter ρ determining the importance of each term.²⁰

¹⁸While the mapping \mathcal{K} is in general not a contraction, iterating on that mapping turns out to be a convenient method for finding an equilibrium. We have not been able to prove that the equilibrium is unique, although we have also not been able to find an economy with multiple equilibria. We discuss conditions under which the solution to the planner's problem is generically unique in Appendix E.7. In particular, we establish a generic uniqueness result when $A_i(\alpha_i)$ takes the form (2), which we will adopt for our quantitative exercises.

¹⁹The variable \mathcal{W} defined in (18) is a convenient monotone transformation of the expected utility of the household, such that $E\left[Y^{1-\rho}\right](1-\rho)^{-1} = \exp\left((1-\rho)\mathcal{W}\right)(1-\rho)^{-1}$, and we adopt it as our main measure of welfare. If we denote the expected utility of the household by \mathbb{W} , we can write $(\mathbb{W}' - \mathbb{W}) / |\mathbb{W}| \approx (\rho - 1)(\mathcal{W}' - \mathcal{W})$ so that it is straightforward to convert changes in welfare between measures.

²⁰From (18), we see that when $\rho = 0$ the household actually likes uncertainty in log GDP. This is because of the properties of lognormal variables, as discussed in footnote 10. (18) also implies that uncertainty has no impact on the planner's decisions when $\rho = 1$. See Appendix I for more details.

5 Beliefs, the production network and aggregate outcomes

In this section, we explore how the mean μ and the covariance Σ of the productivity shocks affect the equilibrium structure of the production network and, through that mechanism, GDP and welfare.

5.1 Beliefs and Domar weights

In Corollary 1, we saw that the Domar weights are key objects to understand how changes in μ and Σ affect the expected level and the variance of GDP. In a fixed-network environment, these weights are fixed and do not respond to changes in beliefs. In contrast, when the network is endogenous, they are equilibrium objects that vary with μ and Σ . The next proposition describes the relationship between these quantities.

Proposition 3. The Domar weight ω_i is increasing in μ_i and decreasing in Σ_{ii} .

This proposition can be understood both from the perspective of individual producers as well as from the perspective of the social planner. Individual producers rely more on sectors whose prices are low and stable. As a result, these sectors are more important suppliers and their Domar weights are relatively high. From the planner's perspective, recall from (14) that the Domar weight of a sector captures its contribution to log GDP. Since the planner wants to increase and stabilize log GDP, it naturally increases the importance of more productive (larger μ_i) or less volatile (smaller Σ_{ii}) sectors in the production network. In Section 5.3 below we show that such an adjustment in the network is welfare-improving. But before doing so, we first discuss how changes in beliefs affect the precise structure of the equilibrium production network α .

5.2 Beliefs and the structure of the production network

While the Domar weights respond in an intuitive and unambiguous manner to changes in beliefs, the same is not true for the matrix α that describes the complete structure of the production network. In some cases an increase in the expected productivity of a sector *i* can even lead some of its customers to lower their usage of input *i*. In this section, we highlight that substitution patterns between input shares are key determinants of the impact of beliefs on the network. We then explore these patterns through two examples and introduce a useful approximation that allows us to analytically characterize them in terms of primitives. Finally, we use this approximation to characterize the impact of beliefs on the production network.

Input shares substitutabilities

The model can handle rich substitution patterns between input shares. To have a sense of what these patterns might involve, we can go back to our car manufacturer example. Suppose that the price of carbon fiber is expected to decrease (higher μ_{carbon}). The firm might respond by choosing a technique that is more intensive in carbon fiber and use that material instead of steel for key components. Steel and carbon fiber would then be substitutes, and their shares would tend to move in opposite directions. At the same time, the firm might purchase additional equipment that is needed to handle carbon fiber. This equipment and carbon fiber would be complements and their shares would tend to move together.

The model generates these substitution patterns through the shape of the TFP shifter functions $\{a_i\}_{i=1}^n$ and the set of available techniques \mathcal{A} . Consider as an example a function a_i such that $\frac{\partial^2 a_i(\alpha_i)}{\partial \alpha_i j \partial \alpha_{ik}} \geq 0$ for some $j \neq k$. This implies that an increase in α_{ij} increases the incentives to raise α_{ik} as well. It follows that these shares will tend to move in the same direction after a change in beliefs.²¹ If instead $\frac{\partial^2 a_i(\alpha_i)}{\partial \alpha_{ij} \partial \alpha_{ik}} < 0$, the shares α_{ij} and α_{ik} will tend to move in opposite direction. Substitution patterns can also arise through the set of admissible techniques \mathcal{A}_i . Suppose, for instance, that the minimum labor share constraint $\sum_{l=1}^n \alpha_{il} \leq \overline{\alpha}_i$ binds and that the expected price of good j falls. To increase the share of good j, a firm in sector i would have to lower its share of some other input, say k, to avoid violating the constraint. In this case the shares of j and k would behave as substitutes in the production of good i.

An example of factor share substitution patterns

To better illustrate these substitution patterns, consider the simple economy depicted in Figure 2. Sectors 1 to 3 produce their goods using only labor, while sector 4 can also use intermediate inputs produced by other sectors. We parametrize the TFP shifter function a_4 so that the factor shares of goods 1 and 2 are complements in the production of good 4, while the shares of goods 1 and 3 are substitutes. Specifically, we introduce two additional terms to the functional form (2) so that

$$a_4(\alpha_4) = a_4^{\circ} - \sum_{j=1}^4 \kappa_j \left(\alpha_{4j} - \alpha_{4j}^{\circ} \right)^2 - \psi_1 \left(\alpha_{41} - \alpha_{42} \right)^2 - \psi_2 \left(\alpha_{41} + \alpha_{43} - \left(\alpha_{41}^{\circ} + \alpha_{43}^{\circ} \right) \right)^2.$$
(19)

The third term in (19) creates the complementarity between goods 1 and 2, while the last term creates the substitutability between goods 1 and 3. It follows that goods 2 and 3 are also substitutes. The parameters $\psi_1 \ge 0$ and $\psi_2 \ge 0$ determine the strength of these substitution patterns.

Panel (a) shows the initial state of the economy. Sector 4 relies heavily on sectors 1 and 2 as suppliers (as shown by the thick arrows) and does not use good 3 in production. In panel (b), μ_1 declines slightly. As we can see, the input share of good 1 declines in response (thinner arrow). The complementarity between 1 and 2 then pushes firms in sector 4 to rely less on good 2, and the substitutability between 1 and 3 pushes them to start using good 3 as an input. In panel (c),

²¹We show in Appendix H that when all shares are complements in the production of all goods we can predict unambiguously how the production network responds to a change in beliefs.





Notes: Arrows represent the movement of goods: there is a solid blue arrow from j to i if $\alpha_{ij} > 0$. Dashed gray arrows indicate $\alpha_{ij} = 0$. The thickness of the arrow reflects the size of the input share. a is as in (2) with the additional terms (19). a_i for $i \in \{1, 2, 3\}$ are such that those sector only use labor. $a_i^0 = 0$ for all i. $\kappa_j = 1/10$ for all j, $\psi_1 = 1$ and $\psi_2 = 1/10$. $\alpha_{4j}^\circ = 1/3$ for all j. $\beta_i = 1/n$ for all i. $\Sigma = 0.3 \times I_{n \times n}$. $\mu_i = 1$ for all i in panel (a). Panel (b): same as panel (a) except $\mu_1 = 1/2$. Panel (c): same as panel (a) except $\mu_1 = -1/2$. The risk aversion of the household is $\rho = 5$.

 μ_1 declines some more. As a result, sector 4 stops using good 1 as an input completely and relies more heavily on good 3. The share α_{42} also decreases slightly.

Cascading flight to safety

When input shares are substitutes, a small change in the volatility of a single sector can push multiple producers to sequentially switch to safer suppliers. To give an example of that process, consider the economy depicted in Figure 3. Firms in sectors 4 to 7 can only use labor as an input, but firms in sectors 1 to 3 can each source inputs from two potential suppliers, indicated by the arrows. The model is parameterized such that the shares of these suppliers are substitutes.

Figure 3: Cascading impact of Σ_{44}



Notes: Arrows represent the movement of goods: there is a solid blue arrow from j to i if $\alpha_{ij} > 0$. Dashed gray arrows indicate $\alpha_{ij} = 0$. a is as in (2) with $a_i^0 = 0$ for all i, $\kappa_{ij} = 0$ if there is a potential link between two firms and infinity otherwise. $\alpha_{ij}^{\circ} = 0.5$ if there is a potential link, and 0 otherwise. $\mu = 0$ except for $\mu_4 = 0.1$. In the left figure, Σ is diagonal with $\Sigma_{ii} = 0.1$ for all i except $\Sigma_{44} = 1$. In the right figure $\Sigma_{44} = 0$. The risk aversion of the household is $\rho = 2$. $\beta_i = 1/n$ for all i.

When uncertainty about sector 4 is sufficiently high (left panel), sector 3 does not rely on good 4 for production. As uncertainty about sector 4 decreases (lower Σ_{44}), firms in sector 3, seeking a stable supply of goods, switch to using good 4 as an input. As a result, sector 3's price becomes less volatile which makes firms in sector 2 want to use good 3 in production instead of good 6. The same logic applies to firms in sector 1, which then also switch from good 5 to the good with the

now less volatile price provided by sector 2. A change in the uncertainty of a single sector can thus lead to a cascading movement to safety. Because uncertainty about each sector's input price is now lower, aggregate uncertainty also decreases.

A useful approximation

To further describe the substitution patterns that emerge in the model, it is useful to consider an economy in which the cost of deviating from the ideal shares α° is large. Concretely, let $a_i(\alpha_i) = -\bar{\kappa} \times \hat{a}_i(\alpha_i)$, where \hat{a}_i is such that $\hat{a}_i(\alpha_i^{\circ}) = 0$, and suppose that α_i° is in the interior of \mathcal{A}_i . We will characterize this economy when $\bar{\kappa} > 0$ is large. For that purpose, it will be convenient to work with the risk-adjusted transformation of the log price vector given by

$$\mathcal{R}^{\circ} \equiv \mathrm{E}\left[p^{\circ}\right] + \mathrm{cov}\left(p^{\circ}, \lambda^{\circ}\right)$$

where p° and λ° denote the equilibrium log price vector and the log stochastic discount factor under the ideal input shares α° . The vector \mathcal{R}° captures how (un)attractive each input is in risk-adjusted terms. A good *j* that is expensive on average (high $\mathbb{E}\left[p_{j}^{\circ}\right]$) or that is particularly risky, in the sense that its price is negatively correlated with consumption (high cov $\left(p_{j}^{\circ}, \lambda^{\circ}\right)$), has a high \mathcal{R}_{j}° and is overall less attractive as an input. We show in Appendix E.9 that we can write \mathcal{R}° in terms of the fundamentals of the economy as

$$\mathcal{R}^{\circ} = -\mathcal{L}^{\circ} \left[\mu - (\rho - 1) \Sigma \omega^{\circ} \right], \qquad (20)$$

where $\mathcal{L}^{\circ} = (I - \alpha^{\circ})^{-1}$ and $(\omega^{\circ})' = \beta' \mathcal{L}^{\circ}$ are the Leontief inverse matrix and the vector of Domar weights under the ideal shares.

We can then derive an approximate expression for the production network.

Proposition 4. For large $\bar{\kappa}$ the vector of input shares in sector *i* is approximately given by

$$\alpha_i \approx \alpha_i^{\circ} - \bar{\kappa}^{-1} \left(H_i^{\circ} \right)^{-1} \mathcal{R}^{\circ}, \tag{21}$$

where H_i° is the Hessian matrix of \hat{a}_i evaluated at the ideal shares α_i° .²²

Given (20) and since H_i° only depends on α_i° , (21) provides a closed-form characterization of the production network in terms of the primitives of the model.

Equation (21) shows that sectors with larger risk-adjusted prices \mathcal{R}° are overall less attractive as input providers and firms tend to rely less on them in production. How a change in \mathcal{R}° affects the input shares α_i depends on the inverse Hessian $(H_i^{\circ})^{-1}$, which captures the substitution patterns

²²We assume throughout that H_i° is invertible for all *i*.

embedded in the TFP shifters. If an element $(H_i^{\circ})_{jl}^{-1}$ is positive, we say that j and l are complements in the production of good i. If instead $(H_i^{\circ})_{il}^{-1}$ is negative, they are substitutes.

To better understand the role of $(H_i^{\circ})^{-1}$, it is convenient to focus on a single share α_{ij} and to write (21) as

$$\alpha_{ij} \approx \alpha_{ij}^{\circ} - \bar{\kappa}^{-1} (H_i^{\circ})_{jj}^{-1} \mathcal{R}_j^{\circ} - \bar{\kappa}^{-1} \sum_{l \neq j} (H_i^{\circ})_{jl}^{-1} \mathcal{R}_l^{\circ}.$$
(22)

Risk-adjusted prices \mathcal{R}° affect the share α_{ij} of input j in the production of good i through two channels. First, α_{ij} obviously depends on how attractive good j is. Since $(H_i^{\circ})_{jj}^{-1}$ is always positive, a lower \mathcal{R}_j° implies a higher α_{ij} through a *direct effect*, highlighted in (22). How large the direct effect is depends on the cost of deviating from the ideal share α_{ij}° , which is captured by $\bar{\kappa}^{-1} (H_i^{\circ})_{jj}^{-1} > 0.^{23}$

The share of good j used in sector i also depends on an *indirect effect* that operates through sectors other than j. Suppose for instance that the risk-adjusted price \mathcal{R}_l° of a good $l \neq j$ declines so that good l becomes more attractive. If the shares of good j and l are substitutes in the production of good i, that is if $(H_i^{\circ})_{jl}^{-1} < 0$, then α_{ij} declines as \mathcal{R}_l° falls. If instead they are complements, such that $(H_i^{\circ})_{jl}^{-1} > 0$, a more attractive input l makes input j more attractive as well, and α_{ij} increases.

The first-order approximation (21) together with (20) allows for a closed-form characterization of the impact of changes in (μ, Σ) on the equilibrium production network.

Corollary 3. For large $\bar{\kappa}$ the impact of an increase in μ_k on the network is approximately given by

$$\frac{d\alpha_{ij}}{d\mu_k} \approx \bar{\kappa}^{-1} (H_i^{\circ})_{jj}^{-1} \mathcal{L}_{jk}^{\circ} + \bar{\kappa}^{-1} \sum_{l \neq j} (H_i^{\circ})_{jl}^{-1} \mathcal{L}_{lk}^{\circ} , \qquad (23)$$

$$\underbrace{\underset{\text{direct effect of } k \text{ on } j}_{\text{through other suppliers } l \neq j}$$

and the impact of an increase in Σ_{km} on the network is approximately given by

$$\frac{d\alpha_{ij}}{d\Sigma_{km}} \approx \begin{cases} -\left(\rho-1\right)\omega_k^{\circ}\frac{d\alpha_{ij}}{d\mu_k} & k=m, \\ -\left(\rho-1\right)\left(\omega_m^{\circ}\frac{d\alpha_{ij}}{d\mu_k}+\omega_k^{\circ}\frac{d\alpha_{ij}}{d\mu_m}\right) & k\neq m, \end{cases}$$
(24)

where ω° is the vector of Domar weights at the ideal input shares.

If j relies on k as an input $(\mathcal{L}_{jk}^{\circ} > 0)$, an increase in μ_k lowers the risk-adjusted price \mathcal{R}_j and makes j more attractive. This direct effect pushes α_{ij} up. If another sector $l \neq j$ also relies on k $(\mathcal{L}_{lk}^{\circ} > 0)$, then an increase in μ_k makes l more attractive. This indirect channel can lead to either an increase or a decrease in α_{ij} , depending on whether j and l are complements or substitutes;

²³Since H_i° is invertible and that \hat{a}_i is strictly convex, H_i° is positive definite. It follows that its inverse is also positive definite and the diagonal elements $(H_i^{\circ})_{jj}^{-1}$ must be positive.

that is, whether $(H_i^{\circ})_{jl}^{-1}$ is positive or negative. If the indirect effect is strongly negative, it can dominate the direct effect and lead to an overall decline in α_{ij} .

Finally, (24) shows that the impact of a change in Σ is similar to that of a change in μ but with the opposite sign. In particular, the same substitution patterns apply. Notice also that a higher risk aversion ρ leads to a stronger impact of a change in Σ on the production network.

The examples of Figures 2 and 3 through the lens of the approximation

Proposition 4 and Corollary 3 can help us understand the example of Figure 2 in which the productivity of sector 1 declines. Direct computation of the inverse Hessian $(H_4^{\circ})^{-1}$ from (19) implies that $(H_4^{\circ})_{12}^{-1} > 0$, $(H_4^{\circ})_{13}^{-1} < 0$ and $(H_4^{\circ})_{23}^{-1} < 0$. In the production of good 4, the shares of goods 1 and 2 are thus complements, but the shares of goods 1 and 3 and goods 2 and 3 are substitutes. As we move from panel (a) to panels (b) and (c), the risk-adjusted price of sector 1, \mathcal{R}_1° , increases. Hence, the direct effect in (22) pushes α_{41} down, while the indirect effect pushes α_{42} down and α_{43} up.

We can also interpret the cascading downstream network adjustment in Figure 3 through the lens of the approximation. Consider for instance the response of sector 1 to the change in the uncertainty of sector 4, at the opposite end of the supply chain. Combining (23) and (24), we can write

$$\frac{d\alpha_{1j}}{d\Sigma_{44}} \approx -(\rho-1)\,\omega_4^{\circ}\left(\bar{\kappa}^{-1}\,(H_1^{\circ})_{jj}^{-1}\,\mathcal{L}_{j4}^{\circ} + \bar{\kappa}^{-1}\sum_{l\neq j}\,(H_1^{\circ})_{jl}^{-1}\,\mathcal{L}_{l4}^{\circ}\right).$$

The response of α_{12} thus depends on how much sector 2 relies on sector 4 as an input. This is captured by \mathcal{L}_{24}° , which is large in this example, because sector 2 is an (indirect) customer of good 4 under the ideal input shares. It follows that the decline in Σ_{44} has a strong positive direct effect on α_{12} . Instead, if we consider the response of α_{15} , this direct effect is absent because sector 5 uses only labor and never relies on sector 4 ($\mathcal{L}_{54}^{\circ} = 0$). Only the indirect effect remains, and since \mathcal{L}_{24}° is large, the relevant term here is $(H_1^{\circ})_{52}^{-1}$. This quantity is negative because 5 and 2 are substitutes in the production of good 1, so that a decline in Σ_{44} decreases α_{15} . The same logic applies to the responses of firms 2 and 3, and thus explains the cascading effect illustrated in Figure 3.²⁴

5.3 Implications for GDP and welfare

Above we analyzed how the production network responds to changes in beliefs about the productivity process. What the household ultimately cares about though is the level and variance of consumption, and we now turn to the implications of an endogenous production network for macroeconomic aggregates. Here we describe how GDP and welfare are affected by changes in the

²⁴The approximation of Proposition 4 and Corollary 3 assumes that the production network α is in the interior of \mathcal{A} , which is not the case in Figures 2 and 3, but (23) and (24) still capture the main forces that push the shares in response to changes in (μ, Σ) and are informative about the response of the network.

mean μ and the variance Σ of productivity when the network is endogenous. We also show that some changes to the productivity process can have surprising implications when the network itself responds to changes in the distribution of shocks.

Uncertainty and expected GDP

We begin with a general result that shows how GDP reacts to uncertainty in equilibrium.

Proposition 5. Expected log GDP E[y] reaches its maximum at $\Sigma = 0$.

Proposition 5 follows directly from Corollary 2. When there is no uncertainty ($\Sigma = 0$), the variance V [$y(\alpha)$] of log GDP is zero for all networks $\alpha \in \mathcal{A}$, so that the equilibrium network maximizes only E [y]. When, instead, the TFP vector is uncertain ($\Sigma \neq 0$), the equilibrium network also seeks to lower V [$y(\alpha)$], which necessarily lowers expected log GDP.

Proposition 5 establishes a novel mechanism through which uncertainty reduces expected log GDP. To understand why, consider the technique choice problem from the firm's perspective. When there is no uncertainty, firms do not worry about risk and move toward cheaper suppliers, which tend to also be the most productive ones, and toward techniques with higher TFP shifters. As a result, the aggregate economy is particularly productive, and E[y] is large. When some suppliers become risky, firms worry that their inputs might become expensive and, to prevent large fluctuations in their own unit cost, start purchasing from more stable but less productive suppliers. As a result, the aggregate economy becomes less productive on average and expected log GDP falls.

The endogenous response of the network is essential for the result of Proposition 5. Indeed, in our model uncertainty affects expected log GDP *only* through the endogenous response of the firms' sourcing decisions. As a result, the mechanism through which uncertainty lowers expected log GDP is only active when the production network is flexible. If instead the shares α were fixed, uncertainty would have no impact on E[y].

Beliefs and welfare

The next proposition characterizes the impact of μ and Σ on our measure of welfare \mathcal{W} , as defined in (18). Throughout this section, we again use partial differentiation to indicate that a derivative is taken while keeping the network α fixed.

Proposition 6. When the network α is free to adjust to changes in μ and Σ the following holds.

1. The impact of an increase in μ_i on expected welfare is given by

$$\frac{d\mathcal{W}}{d\mu_i} = \frac{\partial \operatorname{E}\left[y\right]}{\partial\mu_i} = \omega_i.$$
(25)

2. The impact of an increase in Σ_{ij} on expected welfare is given by

$$\frac{d\mathcal{W}}{d\Sigma_{ij}} = \begin{cases} -\frac{1}{2} \left(\rho - 1\right) \left(\frac{\partial \operatorname{E}[y]}{\partial \mu_i}\right)^2 = -\frac{1}{2} \left(\rho - 1\right) \omega_i^2 & i = j, \\ -\left(\rho - 1\right) \frac{\partial \operatorname{E}[y]}{\partial \mu_i} \frac{\partial \operatorname{E}[y]}{\partial \mu_j} = -\left(\rho - 1\right) \omega_i \omega_j & i \neq j. \end{cases}$$
(26)

This proposition follows directly from applying the envelope theorem to (18). Its first part provides a Hulten-like result. It states that the impact of an increase in μ_i on welfare is equal to its marginal impact on expected log GDP taking the network α as fixed. By Corollary 1, this quantity is also equal to the Domar weight ω_i of sector *i*. Since Domar weights are positive, it follows that an increase in μ_i always has a positive impact on welfare. The second part of the proposition provides a similar result for an increase in Σ_{ij} . In this case, the impact of the change is proportional to the product of the Domar weights ω_i and ω_j . Again, (26) implies that an increase in uncertainty must necessarily lower welfare when $\rho > 1$.

Amplification and dampening

One important consequence of the endogenous reorganization of the network is that changes in the process for ε that are beneficial to welfare are amplified while changes that are harmful are dampened. The following proposition establishes this result formally.

Proposition 7. Let $\alpha^*(\mu, \Sigma)$ be the equilibrium production network under (μ, Σ) and let $\mathcal{W}(\alpha, \mu, \Sigma)$ be the welfare of the household under the network α . The change in welfare after a change in beliefs from (μ, Σ) to (μ', Σ') satisfies the inequality

$$\underbrace{\mathcal{W}\left(\alpha^{*}\left(\mu',\Sigma'\right),\mu',\Sigma'\right)-\mathcal{W}\left(\alpha^{*}\left(\mu,\Sigma\right),\mu,\Sigma\right)}_{\text{Change in welfare under the flexible network}} \geq \underbrace{\mathcal{W}\left(\alpha^{*}\left(\mu,\Sigma\right),\mu',\Sigma'\right)-\mathcal{W}\left(\alpha^{*}\left(\mu,\Sigma\right),\mu,\Sigma\right)}_{\text{Change in welfare under the fixed network}}.$$
 (27)

Under a flexible network, the extra margin of adjustment thus allows firms to produce more efficiently and to better mitigate risk, which translates into higher welfare for the household. This operates through the Domar weights, as the importance for GDP of sectors facing positive shocks increase (see Proposition 3).

Beliefs and GDP

Proposition 6 shows that changes in the mean and the variance of ε have an intuitive impact on welfare when the network is endogenous, qualitatively working in the same direction as in the fixed network economy. The same is not true about the impact of such changes on the moments of log GDP. To see how the endogenous adjustment of the network can lead to counterintuitive implications for output, it is helpful to decompose the impact of a change in beliefs into its direct and indirect impacts. For instance, for a change in μ_i we can write

$$\frac{d \operatorname{E}[y]}{d\mu_{i}} = \underbrace{\frac{\partial \operatorname{E}[y]}{\partial\mu_{i}}}_{\text{direct impact with fixed network}} + \underbrace{\frac{\partial \operatorname{E}[y]}{\partial\alpha}\frac{d\alpha}{d\mu_{i}}}_{\text{network adjustment}}.$$
(28)

The first term on the right-hand side of (28) denotes the direct impact, keeping the network α fixed. The second term captures the impact of μ_i on the network α , and the impact of the change in the network on expected log GDP. When the network is fixed, this network adjustment term is zero and the full impact of the change in μ_i is simply equal to its direct impact. This is the situation that we explored in Corollary 1, which states that an increase in μ_i always has a positive impact on E [y]. But with an endogenous network the indirect effect can amplify, mitigate, or even completely overturn the direct impact, in which case an increase in μ_i can lower expected log GDP. When this happens, the Hulten-like result established in Corollary 1 ceases to work, even as a local approximation. A similar mechanism can also flip the impact of changes in uncertainty such that an increase in Σ_{ii} can lower the variance of log GDP. The following example illustrates how the endogenous adjustment of the network can lead to counterintuitive responses to changes in the productivity process.

Counterintuitive implications of changes in beliefs Consider the economy depicted in the left column of Figure 4, where sectors 4 and 5 use only labor to produce, while sectors 1 to 3 can also use goods 4 and 5 as intermediate inputs. Their TFP shifter functions are parametrized so that the shares of goods 4 and 5 are substitutes. Sector 4 is more productive and volatile than sector 5 ($\mu_4 > \mu_5$ and $\Sigma_{44} > \Sigma_{55}$).

Consider the impact of a positive shock to μ_5 . The solid blue lines in panels (a), (b) and (c) in the right column of Figure 4 illustrate the impact of this shock on E [y], V [y], and on welfare. Point O on the graphs represents the economy before the shock. As we can see, the initial increase in μ_5 has a negative impact on expected log GDP. To understand why, notice that for a small increase in μ_5 , sector 5 is still less productive (on average) than sector 4, but it now offers a better risk-reward trade-off given its lower variance. As a result, firms in sectors 1 to 3 increase their shares of good 5 and reduce their share of good 4. But since $\mu_4 > \mu_5$, this readjustment leads to a fall in expected log GDP for a small increase in μ_5 . At the same time, the variance of log GDP also declines because sector 5 is less volatile than sector 4. The implied changes in E [y] and V [y] thus have opposite impacts on welfare. By Proposition 6, the overall effect on welfare must be positive though, and this is indeed confirmed in panel (c). Naturally, as μ_5 keeps increasing expected log GDP eventually starts to increase as well.

To emphasize the importance of the endogenous network for this mechanism, we also show the effect of the same increase in μ_5 when the network is kept fixed (dashed red lines in the same

panels). From Corollary 1, the marginal impact of μ_5 on expected log GDP is equal to its Domar weight, and increasing μ_5 has a positive impact on E [y]. At the same time, the variance of log GDP is unaffected by changes in μ . While an increase in μ_5 is welfare-improving in this case, the effect is less pronounced than in the flexible network economy. Indeed, in the latter case the equilibrium network changes precisely to maximize the beneficial impact of the shock on welfare, as predicted by Proposition 7.

We can use the same economy to illustrate how an *increase* in an element of Σ can *lower* the variance of log GDP, and simultaneously lower welfare. Start from the economy in the left column of Figure 4 (point O) and suppose that the volatility of sector 4 goes up. In response, firms 1 to 3 start to purchase from sector 5 more actively. Because sector 5 is less volatile (recall that $\Sigma_{55} < \Sigma_{44}$ initially), the variance of log GDP declines (panel e). At the same time, expected log GDP goes down because sector 5 is also less productive on average than sector 4 (panel d). The combined effect on welfare is negative, as predicted by Proposition 6 (panel f). In this case, the reorganization of the network mitigates the adverse effect of the increase in volatility on welfare. Instead, when the network is fixed, an increase in Σ_{44} does not affect expected log GDP but leads to an increase in the variance of log GDP. As a result, welfare drops more substantially than when the network is endogenous, as predicted by Proposition 7.





Notes. Left column: Arrows represent the movement of goods: there is a solid blue arrow from j to i if $\alpha_{ij} > 0$. a is as in (2) with $a_i^0 = 0$ for all i, $\kappa_{ij} = 0$ if there is a potential link between two firms and very large otherwise. $\alpha_{ij}^\circ = 0.5$ if there is a potential link and zero otherwise. The labor share for firms 1, 2 and 3 is fixed at 0.5 and for firms 4 and 5 it is one (deviations are punished severely). The risk aversion of the household is $\rho = 2.5$. Household's utility weights are $\beta_1 = \beta_2 = \beta_3 = \frac{1}{3} - \epsilon$, $\beta_4 = \beta_5 = \frac{3}{2}\epsilon$, where ϵ is a very small positive number. $\mu = (0.1, 0.1, 0.1, -0.04)$, Σ is diagonal, with diag (Σ) = (0.2, 0.2, 0.2, 0.3, 0.05). Right column: In panels (a)-(c), μ_5 increase from -0.04 to 0.13. In panels (d)-(f), Σ_{44} increases from 0.3 to 0.4.

First-order approximation of GDP and welfare

Under the approximation of Proposition 4, we can derive closed-form equations for log GDP and welfare. For that purpose, it is useful to first approximate the equilibrium network and the Domar weights around their values at the ideal shares. That is, let

$$\alpha \approx \alpha^{\circ} + \Delta \alpha \text{ and } \omega \approx \omega^{\circ} + \Delta \omega, \tag{29}$$

where $\Delta \alpha$ is implicitly defined as $\Delta \alpha_i = -\bar{\kappa}^{-1} (H_i^\circ)^{-1} \mathcal{R}^\circ$ by (21) and where a first-order approximation of the Leontief inverse around α° implies $(\Delta \omega)' = \beta' \mathcal{L}^\circ \Delta \alpha \mathcal{L}^\circ$. Similarly, we can approximate the TFP shifter vector $a(\alpha)$ around α° as

$$[a(\alpha)]' \approx [a(\alpha^{\circ}) + \Delta a]' = -\bar{\kappa} \left[(\Delta \alpha_1)' H_1^{\circ} \Delta \alpha_1, \dots, (\Delta \alpha_i)' H_i^{\circ} \Delta \alpha_i, \dots, (\Delta \alpha_n)' H_n^{\circ} \Delta \alpha_n \right], \quad (30)$$

where the last equality follows because $a(\alpha^{\circ}) = 0$ and $\frac{\partial \hat{a}_i(\alpha_i^{\circ})}{\partial \alpha_{ij}} = 0$ for all *i* and *j*. Log GDP can then be approximated by (see Appendix B for the full derivations)

$$y \approx (\omega^{\circ})' \varepsilon + (\Delta \omega)' \varepsilon + (\omega^{\circ})' \Delta a.$$
(31)

The first term in (31) corresponds to log GDP at the ideal input shares α° , while the next two terms capture the impact of the deviation from α° . Specifically, the second term takes into account that the Domar weights deviate from ω° , and that this affects how sectoral shocks ε translate into GDP. The third term captures the fact that deviating from the ideal shares incurs a TFP loss, which also affects GDP. This last term is always negative.

Whether the second term is positive or negative depends on $\Delta \alpha$. Combining (21) and (20), we see that $\Delta \alpha$ depends on the beliefs (μ, Σ) and on the substitution patterns embedded in the TFP shifter functions $\{a_i\}_{i=1}^n$. Consider for instance an increase in μ_i for some sector *i*. In response other sectors increase their reliance on *i*, thus pushing ω_i up (Proposition 3). Through the second term in (31) this change magnifies the impact of *i* on GDP. At the same time, the deviation from the ideal shares leads to a decline in TFP, as captured by the third term in (31).

From (31), it is straightforward to derive closed-form approximations for E[y] and V[y], as well as for welfare \mathcal{W} , in terms of the primitives of the economy. We do so in Appendix B, where we also describe how these quantities vary in response to change in beliefs.²⁵

²⁵In Appendix J we also explore how changes in the beliefs (μ, Σ) affect the correlations between different sectors. Furthermore, we show in Appendix L that the planner's problem can be solved analytically in the special case in which the TFP shifter functions are $A_i(\alpha_i) = 1$ for all i and all $\alpha_i \in \mathcal{A}_i$. In this case, we can derive closed-form expressions for the Domar weights, GDP and welfare.

6 Endogenous network response in a calibrated model

To illustrate the quantitative implications of the mechanism, we calibrate the model to the United States economy. We rely on sectoral shares and productivity data constructed from the Bureau of Economic Analysis' input-output tables, and use a mix of direct estimation and indirect inference to choose values for the model parameters. We start by describing our data sources and calibration strategy. We then show that the calibrated model is able to replicate several important features of the data. Finally, we use the calibrated model to evaluate the role of beliefs in shaping the production network and to investigate how the changing structure of the network influences aggregate output and welfare.

6.1 Data

The Bureau of Economic Analysis (BEA) provides sectoral input-output tables that allow us to compute the intermediate input shares as well as the shares of final consumption accounted for by different sectors. We rely on the harmonized tables constructed by vom Lehn and Winberry (2021) that provide consistent annual data for n = 37 sectors over the period 1948-2020. Table 5 in Appendix C.1 provides the list of the sectors included in this data set.

From these data, we can compute the input shares α_{ijt} of each sector in each year t. The typical share α_{ij} in the data has an average of 0.0128 and a standard deviation over time of 0.0048, for a coefficient of variation of 0.37. We also use the input-output tables to compute sectoral total factor productivity, following the procedure in vom Lehn and Winberry (2021) closely. Specifically, sectoral TFP is measured as the Solow residual, i.e. the residual that remains after removing the contribution of input factors from a sector's gross output.²⁶

6.2 Calibration strategy

The three groups of parameters that we need to calibrate are 1) the household's preferences, i.e. the consumption shares β and the risk-aversion ρ , 2) the parameters κ and α° of the TFP shifter function (2), and 3) the processes for the exogenous sectoral productivity shocks, i.e. μ_t and Σ_t . Some of these parameters can be computed directly from the data. The other ones are estimated using a combination of indirect inference and standard time-series methods. Below, we describe the exact procedure used for each set of parameters.

²⁶We make three departures from vom Lehn and Winberry (2021) in constructing the TFP series. First, to be consistent with our model, we let the input shares α_{ijt} vary over time. Second, we do not smooth the resulting Solow residuals. Finally, we update the time series to include the years up to 2020.

Household preferences

Since the preference parameter β_i corresponds to the household's expenditure share of good i, we can pin down its value directly from the data by averaging the consumption share of good i over time.²⁷ The sectors with the largest consumption shares are "Real estate" (14%), "Retail trade" (12%) and "Health care" (11%).

The relative risk aversion parameter ρ determines to what extent firms are willing to trade off higher input prices for access to more stable suppliers. The literature uses a broad range of values for ρ and it is unclear a priori which one is best for our application. We therefore estimate ρ using a method of simulated moments (MSM) described below.

Endogenous productivity shifter

We adopt the functional form (2) for the productivity shifter $A_i(\alpha_{it})$. This functional form takes as inputs the ideal shares α_{ij}° , the actual shares α_{ijt} , the coefficients κ_{ij} and the constant a_i° . The ideal shares α_{ij}° are set to the time average of the input shares observed in the data.²⁸ We set the constant a_i° equal to the average TFP of sector *i*. The coefficients κ_{ij} , which determine how costly it is to deviate from the ideal shares in terms of productivity, are estimated using the MSM procedure described below. Without any restrictions the matrix κ would have $n \times (n+1) = 1406$ elements. To reduce the number of free parameters to estimate, we restrict κ to be of the form $\kappa = \kappa^i \kappa^j$ where κ^i is an $n \times 1$ column vector and κ^j is an $1 \times (n+1)$ row vector. The *k*th element of κ^i then scales the cost for producer *k* of changing the share of any of its inputs, and the *l*th element in κ^j scales the cost of changing the share of input *l* for any producer. We normalize the first element in κ^i to pin down the scale of κ^i and κ^j . The matrix κ then contains only 2n = 74free parameters to estimate.

Exogenous productivity process

The source of uncertainty in the model is the vector of productivity shocks $\varepsilon_t \sim \mathcal{N}(\mu_t, \Sigma_t)$. In the calibrated model, we allow μ_t and Σ_t to vary over time to account for changes in the stochastic process for ε_t over the sample period. To parameterize the evolution of μ_t and Σ_t , we first filter out the endogenous productivity shifter $A_i(\alpha_{it})$ and the normalization term $\zeta(\alpha_{it})$ from the measured sectoral TFP, $e^{\varepsilon_{it}}A_i(\alpha_{it})\zeta(\alpha_{it})$, implied by the production function (1). We then estimate the evolution of μ_t and Σ_t from the remaining component. To do so, we assume that ε_t follows a random walk with drift,

$$\varepsilon_t = \gamma + \varepsilon_{t-1} + u_t, \tag{32}$$

²⁷See Appendix G.3 for a version of the calibrated economy with time-varying β 's.

²⁸We experimented with an alternative calibration in which we include and estimate a *j*-specific shifter to α_{ij}° . The results are similar to our baseline calibration.

where γ is an $n \times 1$ vector of deterministic drifts and $u_t \sim \text{iid } \mathcal{N}(0, \Sigma_t)$ is a vector of shocks. We estimate γ by computing the average of the productivity growth rates $\Delta \varepsilon_t = \varepsilon_t - \varepsilon_{t-1}$ over time.

When making decisions in period t, firms know the past realizations of ε_t so that the conditional mean of ε_t is given by $\mu_t = \gamma + \varepsilon_{t-1}$. The covariance Σ_t of the innovation u_t is estimated using a rolling window that puts more weight on more recent observations to allow for time-varying uncertainty about sectoral productivity. Specifically, we estimate the covariance between sector *i* and *j* at time *t* by computing $\Sigma_{ijt} = \sum_{s=1}^{t-1} \phi^{t-s-1} u_{is} u_{js}$, where $0 < \phi < 1$ is a parameter that determines the relative weight of more recent observations. Its value is set to the sectoral average of the corresponding parameters of a GARCH(1,1) model estimated on each sector's productivity innovation u_{it} . In the calibrated economy, its value is $\phi = 0.47$. Note that this procedure implies that the time series for ε_t depends on the parameters of the TFP shifters. Therefore, the estimation of the stochastic process for sectoral productivity has to be done jointly with the estimation of κ .

Matching model and data moments

We use an indirect inference approach and estimate the parameters $\Theta \equiv \{\rho, \kappa\}$ by minimizing

$$\hat{\Theta} = \arg\min_{\Theta} \left(m\left(z \right) - m\left(\Theta \right) \right)' W\left(m\left(z \right) - m\left(\Theta \right) \right),$$

where m(z) is a vector of moments computed from the data, and $m(\Theta)$ is the vector of corresponding model-implied moments conditional on the parameters Θ . The moments that we target are the time series of the production shares α_{ijt} , normalized by their average in the data, and the demeaned time series of aggregate consumption growth, normalized by the average of its absolute value in the data. We target consumption since the stochastic discount factor of the household is central to the trade-off that firms face when choosing production techniques.²⁹

We match $n^2 \times T + T - 1$ moments with only 2n + 1 free parameters. The model is thus strongly over-identified. We use particle swarm optimization to find the global minimizer $\hat{\Theta}$ (Kennedy and Eberhart, 1995). The estimated coefficient of relative risk aversion $\hat{\rho}$ is 4.27, which is similar to values used or estimated in the macroeconomics literature. We provide additional details about the calibrated economy in Appendix C.2.

6.3 Model-implied moments and the data

We want our model to fit key features of the data that relate to 1) the structure of the production network, 2) how the network responds to changes in beliefs, and 3) how this response affects macroeconomic aggregates. As we have seen earlier, the Domar weights, and how they react to

²⁹To strike a balance between matching both the shares and consumption growth reasonably well, the weighting matrix W assigns a weight of $(n^2 \times T)^{-1}$ to the shares moments (recall that there are n^2 shares time series, each of length T) and a weight of $(T-1)^{-1}$ to the consumption growth moment (the length of the consumption growth time series is T-1).

changes in μ_t and Σ_t , play a central role for these mechanisms. In this section, we first describe the evolution of μ_t and Σ_t in the calibrated economy. We then report unconditional moments of the model-implied Domar weights and how they compare to the data. Finally, we look at the relationship between the Domar weights and the beliefs μ_t and Σ_t and verify that the correlations predicted by the mechanisms of the model are present in the data.

Time-variation in the aggregate productivity process

To illustrate the overall evolution of beliefs over our sample period, we compute two measures that capture the aggregate impact of changes in μ_t and Σ_t . The first measure is the Domar-weighted average growth in the conditional mean of productivity, defined as

$$\Delta \bar{\mu}_t = \sum_{j=1}^n \omega_{jt} \Delta \mu_{jt}.$$
(33)

We use the Domar weights ω_{jt} in this equation to properly reflect the importance of a sector for GDP, as implied by (14). The solid blue line in Figure 5 shows the evolution of $\Delta \bar{\mu}_t$ over the sample period. As expected, $\Delta \bar{\mu}_t$ tends to go below zero during NBER recessions and is positive during expansions.

To describe how aggregate uncertainty evolves in the calibrated economy, we also compute the within-period perceived standard deviation of log GDP. From (15), this can be written as

$$\sigma_{yt} = \sqrt{\mathcal{V}[y]} = \sqrt{\omega_t' \Sigma_t \omega_t}.$$
(34)

The red dashed line in Figure 5 represents the evolution of σ_{yt} over the sample period. While uncertainty is on average relatively low, especially during the Great Moderation era, spikes are clearly visible in the earlier years and, in particular, during the Great Recession of 2007-2009.³⁰

Unconditional Domar weights

Figure 6 shows the average Domar weight of each sector in the data (blue bars) and in the model (black line). The sectors with the highest Domar weights in the data are "Real estate", "Food and beverage", "Retail trade", "Finance and insurance" and "Health care". According to our theory (Proposition 6), changes in the expected level and variance of productivity in those sectors will have the largest effects on welfare.

The cross-sectional correlation between the average Domar weights in the model and in the data is 0.96, so that the calibrated model fits this important feature of the production network well. However, the average Domar weight in the model (0.032) is lower than its counterpart in the

 $^{{}^{30}\}sigma_{yt}$ pertains only to uncertainty about the stochastic part of TFP ε . As such, it does not capture overall economic uncertainty, which might also be affected by changes in employment, investment, monetary and fiscal policy, etc.



Figure 5: Domar-weighted TFP and uncertainty changes

Notes: Solid blue line: Domar-weighted average growth in the conditional mean of productivity, $\Delta \bar{\mu}_t = \sum_{j=1}^n \omega_{jt} \Delta \mu_{jt}$. Red dashed line: Domar-weighted conditional variance of productivity, $\sigma_{ut} = \sqrt{\omega'_{\lambda} \Sigma_t \omega_t}$. Shaded areas represent NBER recessions.

data (0.047). This is because the estimation also targets aggregate consumption growth. Given the observed variation in TFP, if the model were to match the Domar weights perfectly, consumption would be too volatile compared to the data. Under our calibration, the volatility of consumption growth in the model is 2.73%, close to its data target of 2.65% (row (6) of Table 1).³¹

Figure 6: Sectoral Domar weights in the data and the model



Notes: The Domar weights are computed for each sector in each year and then averaged over all time periods.

The model can account for about 40% of the observed average standard deviation of the Domar weights over time, as shown in row (2) of Table 1. Row (3) also reports that the coefficient of variation of the Domar weights in the model is 0.07 compared to 0.11 in the data. Once we take into account their relative scale, the model can thus account for a substantial portion of the

³¹Since there is no investment and that the only primary factor of production (labor) is in fixed supply, consumption and aggregate TFP are equal in the model. It follows that we cannot match the volatility of both quantities and the model somewhat overpredicts TFP volatility (see Table 1). Including an investment margin in the model, so that GDP no longer equals consumption, might improve the fit of the Domar weights while keeping consumption growth in the model as volatile as in the data.

variation in a key moment that characterizes the production network.³²

	Statistic	Data	Model
(1)	Average Domar weight $\bar{\omega}_j$	0.047	0.032
(2)	Standard deviation $\sigma(\omega_j)$	0.0050	0.0021
(3)	Coefficient of variation $\sigma(\omega_j)/\bar{\omega}_j$	0.107	0.067
(4)	$\operatorname{Corr}\left(\omega_{jt},\mu_{jt}\right)$	0.08	0.08
(5)	$\operatorname{Corr}\left(\omega_{jt}, \Sigma_{jjt}\right)$	-0.37	-0.31
(6)	Consumption growth volatility	2.65%	2.73%
(7)	TFP growth volatility	1.83%	2.73%

Table 1: Domar weights, consumption and TFP in the model and in the data

Notes: For each sector, we compute the time series of its Domar weight ω_{jt} , as well as its standard deviation $\sigma(\omega_j)$ and its mean $\bar{\omega}_j$. Rows (1) and (2) report cross-sectional averages of these statistics. Row (3) is the ratio of rows (2) and (1). Each period, we compute cross-sectional correlations of the Domar weights ω_{jt} with μ_{jt} and Σ_{jjt} (mean and variance of exogenous TFP ε_{jt}). Rows (4) and (5) report time-series averages of these correlations. Rows (6) and (7) compare consumption growth and TFP growth volatilies across the model and the data. The TFP data comes from Fernald (2014) and is not adjusted for capacity utilization.

Domar weights and beliefs

One of the key mechanisms of the model predicts that a decline in the expected productivity of a sector, or an increase in its variance, should lead firms to reduce the importance of that sector as an input provider, leading to a decline in its Domar weight. Proposition 3 makes this point formally for a single change in μ_i or Σ_{ii} . Of course, in the data multiple changes in μ_t and Σ_t occur at the same time, and it would be difficult to isolate the impact of a single change on the Domar weights. Instead, we look at simple cross-sector correlations between the Domar weights ω_{it} and the first (μ_{it}) and the second moments (Σ_{iit}) of sectoral TFPs, both in the data and in the model. These correlations provide a straightforward, albeit noisy, measure of the interrelations between ω_t , μ_t and Σ_t . As can be seen in rows (4) and (5) of Table 1, the predictions of the model are borne out in the data. The model can not only explain these correlations qualitatively but is also doing a good job quantitatively. The model is thus able to capture well the impact of beliefs on the structure of the production network. In Appendix C.4, we also show that the model is able to replicate features of the empirical correlations between sectoral sales.

6.4 The production network, welfare and output

To evaluate the importance of an endogenous production network for welfare and GDP, we compare the calibrated model to two sets of alternative economies. First, we compare our baseline

³²One reason why the Domar weights are less volatile in the model than in the data is that we assume that the $\{A_i\}_{i=1}^n$ functions are time invariant. In reality, technological changes might affect the shape of these functions which would translate into additional variation in the Domar weights.

model to an economy in which the network is kept completely fixed. This exercise therefore informs us about the overall impact of changes in the structure of the production network. We then investigate the role of uncertainty alone in shaping the production network. We do so by considering an economy in which production techniques are chosen as if $\Sigma_t = 0$, and a perfect foresight economy in which firms observe the realization of ε_t before making technique choices. These exercises allow us to isolate the impact that uncertainty has on the production network and, through that channel, on macroeconomic aggregates.

Fixed vs endogenous production networks

We start by comparing the baseline model to the alternative economy in which the production network cannot respond to changes in μ_t and Σ_t . Specifically, we fix the input share matrix α to its time average in the data and recompute all other equilibrium quantities. We find that the economy with a fixed network is on average 2.12% less productive than the economy with a flexible network (see Table 2). The intuition for this difference is straightforward. As some sectors of the economy become more productive, firms would like to take advantage of their cheaper inputs by relying more on them in production. With a flexible network this is possible, and the aggregate economy becomes more productive as a result. In contrast, when the network is fixed, firms are stuck with less productive suppliers and the economy produces less. Perhaps surprisingly, the fixed network economy is also slightly more stable than its flexible counterpart with a standard deviation of log GDP that is 0.13% smaller. To understand why, recall that the planner balances increasing the expected value of log GDP against lowering its variance. In the calibrated model, the planner is willing to suffer a slight increase in variance, compared to the fixed network alternative, for large gains in expected value. Overall, this comparison with a fixed network suggests that policy interventions that impede or slow down the reorganization of supply chains might have a sizable impact on welfare.

While the differences between the fixed and endogenous network economies are substantial, they are particularly large during volatile periods, when adjusting the network is most beneficial. In Appendix G.2, we show that allowing the network to adjust in response to shocks leads to large gains in expected GDP during the Great Recession.

The role of uncertainty

The differences between the flexible and fixed network economies come from variations in both μ_t and Σ_t . To isolate the role of uncertainty in shaping the production network and GDP, we consider two alternative economies. Our comparison below is focused on GDP (and its moments) and expected welfare, but Appendix C.3 also compares the input shares and the Domar weights in the baseline and the alternative economies.
	Baseline model compared to			
	Fixed network	as if $\Sigma_t = 0$	Known ε_t	
Expected log GDP, $E[y]$	+2.12%	-0.01%	+0.68%	
Expected st. dev. of log GDP, $\sqrt{\mathcal{V}[y]}$	+0.13%	-0.10%	-0.22%	
Expected welfare, \mathcal{W}	+2.11%	+0.01%	+0.71%	
Realized log GDP, y	+1.61%	+0.07%	-0.54%	

Table 2: Uncertainty, GDP and welfare in the post-war sample

Notes: Baseline variables minus their counterparts in the "fixed network", the "as if $\Sigma_t = 0$ ", the "known ε_t " alternatives.

The "as if $\Sigma_t = 0$ " economy We first consider an alternative economy in which firms choose their production technique as if $\Sigma_t = 0$. In practice we set $\Sigma = 0$ when solving the problem of the social planner (18) for the equilibrium network α^* . We then reintroduce uncertainty when computing the moments of GDP and welfare. Table 2 reports these moments. In line with the theory, the baseline economy is less productive and less volatile than the alternative. When $\Sigma_t \neq 0$, it is worthwhile from the firms' perspective to use suppliers that are less productive but safer, which translates into the observed differences in E[y] and V[y]. The differences are however fairly small. The reason for this is that for most of the sample period, uncertainty is relatively low, so that firms simply buy their inputs from the most productive suppliers without much concern for any risk involved.³³ We explore the sensitivity of our results to different values of ρ in Appendix G.1.

Since uncertainty changes over time, its impact on GDP and welfare is also time-varying. The panels in the left column of Figure 7 compare the baseline economy to the "as if $\Sigma_t = 0$ " alternative (denoted with tildes in the figure) over the sample period. The comparison is made in terms of expected log GDP (top panel), the standard deviation of log GDP (second panel), expected welfare (third panel), and realized log GDP (bottom panel). As we can see, the biggest differences occur during periods of high uncertainty such as the oil shocks of the 1970's, the dot-com bubble, and the Great Recession.

We know from Figure 5 that uncertainty was particularly high during the Great Recession and, accordingly, this is also when the differences between the two models are the largest. The top panel of Figure 7 shows that expected GDP in the baseline economy is about 0.25% lower in 2009 than in the alternative economy. Because of the large increase in uncertainty, firms adjust their production techniques toward safer but less productive suppliers to avoid potentially disastrous increases in costs. The result in terms of aggregate volatility is visible in the second panel, where we see that GDP is about 2.37% less volatile in 2009 in the baseline economy. While the implied reduction in expected log GDP in the baseline economy is sizable, the reduction in its variance is important enough to lead to an increase in welfare of about 0.36% when compared to the "as

 $^{^{33}}$ As in Lucas (1987), the utility cost of business cycle fluctuations is on average small in our model and the planner does not want to sacrifice much in terms of the level of GDP for a reduction in volatility.

if $\Sigma_t = 0$ " economy (third panel). Interestingly, realized log GDP, shown in the bottom panel, is substantially higher in the baseline economy than in the alternative. Essentially, firms were worried about particularly bad TFP draws and opted for safer suppliers, and then their fears were realized. The year 2009 saw bad TFP draws (as evident from Figure 5), and so the baseline economy fared about 2.67% better in terms of realized GDP compared to the alternative.

The perfect foresight ("known ε_t ") economy We also consider a second alternative economy in which firms can observe the random draw of the productivity vector ε_t prior to choosing their production techniques. One interpretation of this setup is that adopting a new technique is immediate, so that firms can wait to pick the best technique for a particular ε_t draw. Techniques and intermediate input choices are thus made simultaneously and conditional on observed prices.

In this alternative economy, beliefs (μ_t, Σ_t) , and in particular uncertainty, play no role in shaping the network and, from the planner's problem, the optimal network is simply the one that maximizes (realized) consumption. It follows that consumption (or GDP) is always larger than in the baseline model (bottom right panel in Figure 7). Unsurprisingly, the difference is particularly pronounced during episodes of high uncertainty, when knowing ε_t provides a larger advantage, and reaches a high of 3% during the Great Recession. On average, GDP is 0.54% larger than in the baseline economy suggesting a sizable impact of uncertainty on the economy (bottom row in Table 2).

The top three panels in the right column of Figure 7 show how the baseline and alternative economies differ in terms of expected log GDP, the standard deviation of log GDP, and (expected) welfare. Crucially, these measures are evaluated *before* ε is realized.³⁴ Welfare \mathcal{W} is always lower in the alternative economy because, by construction, \mathcal{W} is what the network in the baseline model maximizes. Furthermore, the optimal network in this economy does not seek to increase E[y] and reduce V[y]. As a result, E[y] is on average lower and V[y] is on average higher (right column in Table 2).

Overall, our findings suggest that allowing the production network to reorganize itself in response to changes in the productivity process can lead to large welfare gains. During volatile periods, uncertainty alone can play an important role in shaping the production network, with substantial consequences for expected and realized GDP, as well as for welfare. Our results therefore highlight the importance for firms of reorganizing supply lines during turbulent periods. It also suggests that policies, such as trade barriers, that would impede this reorganization might have significant side effects.

While the calibrated model suggests that the effect of uncertainty on welfare and output may be substantial, there are reasons to believe that the model estimates are conservative. In reality, many firms are owned by individuals whose consumption is likely to be more volatile than that of the representative household in the model. It is also likely that many entrepreneurs have a

³⁴Note that *realized* welfare in this economy is simply equal to realized log GDP.



Figure 7: The role of uncertainty in the postwar period

Notes: The differences between the series implied by the baseline model (without tildes) and the two alternatives (marked by tildes): the "as if $\Sigma_t = 0$ " alternative (left column) and the "known ε_t " alternative (right column). All economies are hit by the same shocks that are filtered out from the TFP data under our baseline model. All differences are expressed in percentage terms. Expected log GDP E[y] and expected standard deviation of log GDP V[y] are evaluated before ε_t is realized.

less diversified portfolio than the representative household and earn most of their income from a single firm. More uncertain consumption that covaries more strongly with firm profits would make entrepreneurs more sensitive to risk, hence making them more likely to take action to mitigate uncertainty about input prices. In addition, the model probably underestimates the risk associated with input prices. While the firms within each sector in our model are homogenous, in reality each sector is made up of heterogeneous firms that, in addition to sector-specific shocks, are also subject to firm-specific disturbances. To the extent that individual firms buy intermediate inputs from other individual firms, this will add to the input price uncertainty that they are facing.

Finally, it is important to note that for tractability the model abstracts from other forces that might affect the network and the Domar weights. For instance, technological changes might have an impact on the TFP shifter functions and push firms to adjust their technique choice. Similarly, the elasticity of substitution, given a technique, could differ from one, so that price changes would also affect Domar weights. Including these forces in the model may affect the quantitative response of the production network to changes in uncertainty.

7 Model-free evidence for the mechanisms

The model used in the previous section to study the impact of uncertainty on the production network and the macroeconomy, while tractable, imposes a lot of structure on the data. In this section, we present additional evidence in support of the main mechanisms of the paper that does not rely on this structure. Through firm-level regressions that closely follow Alfaro et al. (2019) we document that 1) higher uncertainty leads to a decline in Domar weight, and 2) network connections involving riskier suppliers are more likely to break down. We test these predictions at the firm level to take advantage of the abundance of data and of instrumental variables that are available at this level of aggregation. Appendix D describes the data and the instruments in detail.

7.1 Uncertainty and Domar weights

We first test the model's prediction that Domar weights decrease with supplier uncertainty. We use annual U.S. data from 1963 to 2016 provided by Compustat. Our main variables of interest are a firm's Domar weight, constructed by dividing its sales by nominal GDP, and a measure of its stock price volatility, which we use as a proxy for uncertainty.³⁵ We then regress the change in Domar weight on the change in stock price volatility. The results are presented in the first column of Table 3. In column (2), we follow Alfaro, Bloom, and Lin (2019) and address potential endogeneity concerns by instrumenting stock price volatility with industry-level exposure to ten aggregate sources of uncertainty shocks. In column (3), we use option prices to back out an implied measure of future volatility. In all cases, we find a negative and significant relationship between uncertainty and Domar weights. The effect is also economically large with a decline in Domar weight of about 18% following a doubling in firm-level volatility (roughly a 3.3 standard deviation volatility shock), according to the IV estimates. Overall, these results provide evidence that higher uncertainty leads to lower Domar weights, in line with the predictions of our theoretical model.

7.2 Uncertainty and link destruction

We conduct a similar exercise, this time at the firm-to-firm relationship level, to investigate whether higher supplier uncertainty is associated with a higher likelihood of link destruction. We proceed by combining the uncertainty data described above with data from 2003 to 2016 about firm-level supply relationships provided by Factset. We then regress a dummy variable that equals

 $^{^{35}}$ Ersahin et al. (2022) use textual analysis of earning conference calls to measure firm-level supply chain risk, and find that it is positively correlated with stock price volatility. They also find that firms respond to higher supply chain risks by switching to a wider range of less risky suppliers.

	Cha (1): OLS	nge in Domar weight (2): IV	(3): IV
$\Delta \text{Volatility}_{i,t-1}$	-0.058^{***} (0.004)	-0.137^{***} (0.034)	-0.218^{***} (0.073)
1st moment $10IV_{i,t-1}$	No	Yes	Yes
Type of volatility	Realized	Realized	Implied
Fixed effects	Yes	Yes	Yes
Observations	112,563	$27,\!380$	$17,\!151$
F-statistic		14.2	9.8

Table 3: Domar weights and uncertainty

Notes: Table presents OLS and 2SLS annual regression results of firm-level voltality. The dependent variable is the growth rate in Domar weight. Supplier Δ Volatility_{i,t-1} is the 1-year lagged change in firm-level volatility. Realized volatility is the 12-month standard deviation of daily stock returns from CRSP. Implied volatility is the 12-month average of daily (365-day horizon) implied volatility of at-the-money-forward call options from OptionMetrics. As in Alfaro et al. (2019), "we address endogeneity concerns on firm-level volatility by instrumenting with industry-level (3SIC) non-directional exposure to 10 aggregate sources of uncertainty shocks. These include the lagged exposure to annual changes in expected volatility of energy, currencies, and 10-year treasuries (as proxied by at-the-money forward-looking implied volatilities of oil, 7 widely traded currencies, and TYVIX) and economic policy uncertainty from Baker et al. (2016). [...] To tease out the impact of 2nd moment uncertainty shocks. These are labeled 1st moment 10IV_{t-1}." See Alfaro et al. (2019) for more details about the data and the construction of the instruments. All specifications include year×industry (2SIC) fixed effects. Standard errors (in parentheses) are clustered at the industry (3SIC) level. *F*-statistics are Kleibergen-Paap. *,**, **** indicate significance at the 10%, 5%, and 1% levels, respectively.

one in the last year of a relationship on the change in the supplier's stock price volatility. The results are presented in column (1) of Table 4. As in the last exercise, column (2) uses industrylevel sensitivity to aggregate shocks as instruments, and column (3) uses implied volatility from option prices as a measure of uncertainty. In all cases, we find a positive and statistically significant relationship between supplier volatility and the end of supply relationships, which is consistent with buyers moving away from riskier suppliers. The effect is also economically large with a doubling in volatility associated with a 12 percentage point increase in the likelihood that a relationship is destroyed, according to the IV estimates.

	Dummy for last year of supply relationship			
	(1): OLS	(2): IV	(3): IV	
Δ Volatility _{t-1} of supplier	0.026**	0.097***	0.144^{**}	
	(0.012)	(0.035)	(0.063)	
1st moment $10IV_{t-1}$ of supplier	No	Yes	Yes	
Type of volatility	Realized	Realized	Implied	
Fixed effects	Yes	Yes	Yes	
Observations	$35,\!629$	$35,\!620$	$26,\!195$	
<i>F</i> -statistic	_	22.9	10.39	

Table 4: Link destruction and supplier volatility

Notes: Table presents OLS and 2SLS annual regression results of firm-level voltality. The dependent variable is a dummy variable that equals one in the last year of a supply relatonship and zero otherwise. We limit the sample to relationships that have lasted at least five years. The IV estimates remain significant when relationships of other lengths are considered. Supplier Δ Volatility_{t-1} is the 1-year lagged change in supplier-level volatility. Realized volatility is the 12-month standard deviation of daily stock returns from CRSP. Implied volatility is the 12-month average of daily (365-day horizon) implied volatility of at-the-money-forward call options from OptionMetrics. As in Alfaro et al. (2019), "we address endogeneity concerns include the lagged exposure to annual changes in expected volatility of energy, currencies, and 10-year treasuries (as proxied by at-the-money forward-looking implied volatilities of oil, 7 widely traded currencies, and TYVIX) and economic policy uncertainty from Baker et al. (2016). [...] To tease out the impact of 2nd moment uncertainty shocks from 1st moment aggregate shocks we also include as controls the lagged literctional industry 3SIC exposure to changes in the price of each of the 10 aggregate instruments (i.e., 1st moment return shocks). These are labeled 1st moment $10IV_{t-1}$." See Alfaro et al. (2019) for more details about the data and the construction of the instruments. All specifications include year×customer×supplier industry (2SIC) levels. *F*-statistics are Kleibergen-Paap. *,***, indicate significance at the 10,5,%, and 1% levels. *F*, statistics are Kleibergen-Paap. *,***, indicate significance at the 10,5%, and 1% levels, *F*, statistics are Kleibergen-Paap. *,***, indicate significance at the 10,5%, and 1% levels, *F*, statistics are Kleibergen-Paap. *,***

8 Conclusion

We construct a model in which agents' beliefs about productivity affect the structure of the production network and, through that channel, other macroeconomic aggregates such as output and welfare. We prove that there exists an equilibrium that is efficient and that this equilibrium is characterized by a trade-off between the expected level and the volatility of GDP. We also prove that uncertainty, through its effect on the network, unambiguously lowers expected log GDP. When calibrated to the United States economy, the model predicts that the impact of uncertainty on the network can have a sizable effect on GDP and welfare, especially during periods of high uncertainty such as the Great Recession.

The model is tractable and can serve as a framework to study other related questions. For instance, with some adjustments, our closed economy model could be extended to study uncertainty about international supply chains. Such a model could inform recent policy discussions about onshoring by spelling out both the benefits and the costs of reallocating production to locations with lower geopolitical risk. It would also be natural to extend the analysis in this paper to a model calibrated to firm-level data and to allow firms to enter and exit. However, such an extension would be more involved, as it would necessitate moving away from the perfect competition framework proposed here. Finally, we believe that in reality dynamic considerations might play an important role when firms are deciding to create relationships with suppliers, and so a dynamic version of our model can be a worthwhile extension.

References

- ACEMOGLU, D. AND P. D. AZAR (2020): "Endogenous Production Networks," *Econometrica*, 88, 33–82.
- 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 (2017): "Microeconomic Origins of Macroeconomic Tail Risks," *American Economic Review*, 107, 54–108.
- ACEMOGLU, D. AND A. TAHBAZ-SALEHI (2020): "Firms, Failures, and Fluctuations: The Macroeconomics of Supply Chain Disruptions," Working paper.
- ALESSANDRIA, G., S. YAR KHAN, A. KHEDERKARIAN, C. MIX, AND K. J. RUHL (2022): "The Aggregate Effects of Global and Local Supply Chain Disruptions: 2020-2022," Working paper.
- ALFARO, I., N. BLOOM, AND X. LIN (2019): "The Finance Uncertainty Multiplier," Working paper.
- ATALAY, E. (2017): "How Important Are Sectoral Shocks?" American Economic Journal: Macroeconomics, 9, 254–280.
- ATALAY, E., A. HORTACSU, J. ROBERTS, AND C. SYVERSON (2011): "Network Structure of Production," *Proceedings of the National Academy of Sciences*, 108, 5199–5202.
- BAKER, S. R., N. BLOOM, AND S. J. DAVIS (2016): "Measuring Economic Policy Uncertainty," *Quarterly Journal of Economics*, 131, 1593–1636.
- BAQAEE, D. R. (2018): "Cascading Failures in Production Networks," *Econometrica*, 86, 1819–1838.
- BAQAEE, D. R. AND E. FARHI (2019a): "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem," *Econometrica*, 87, 1155–1203.
- (2019b): "Productivity and Misallocation in General Equilibrium," *Quarterly Journal of Economics*, 135, 105–163.
- BARROT, J.-N. AND J. SAUVAGNAT (2016): "Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks," *Quarterly Journal of Economics*, 131, 1543–1592.
- BERNARD, A. B., E. DHYNE, G. MAGERMAN, K. MANOVA, AND A. MOXNES (2022): "The Origins of Firm Heterogeneity: A Production Network Approach," *Journal of Political Economy*, 130, 000–000.
- BIGIO, S. AND J. LA'O (2020): "Distortions in Production Networks," Quarterly Journal of Economics, 135, 2187–2253.
- BLOOM, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77, 623–685. ——— (2014): "Fluctuations in Uncertainty," *Journal of Economic Perspectives*, 28, 153–76.
- BLOOM, N., M. FLOETOTTO, N. JAIMOVICH, I. SAPORTA-EKSTEN, AND S. J. TERRY (2018): "Really Uncertain Business Cycles," *Econometrica*, 86, 1031–1065.

- BOEHM, C. E., A. FLAAEN, AND N. PANDALAI-NAYAR (2019): "Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tohoku Earthquake," *Review of Economics* and Statistics, 101, 60–75.
- BOEHM, J. AND E. OBERFIELD (2020): "Misallocation in the Market for Inputs: Enforcement and the Organization of Production," *Quarterly Journal of Economics*, 135, 2007–2058.
- CALIENDO, L., F. PARRO, AND A. TSYVINSKI (2022): "Distortions and the Structure of the World Economy," *American Economic Journal: Macroeconomics*, 14, 274–308.
- CARVALHO, V. AND X. GABAIX (2013): "The Great Diversification and its Undoing," *American Economic Review*, 103, 1697–1727.
- CARVALHO, V. AND N. VOIGTLÄNDER (2014): "Input Diffusion and the Evolution of Production Networks," Working paper.
- CARVALHO, V. M., M. NIREI, Y. U. SAITO, AND A. TAHBAZ-SALEHI (2021): "Supply Chain Disruptions: Evidence from the Great East Japan Earthquake," *Quarterly Journal of Economics*, 136, 1255–1321.
- Cox, G. (2020): "Almost Sure Uniqueness of a Global Minimum without Convexity," Annals of Statistics, 48, 584–606.
- DAVID, J. M., L. SCHMID, AND D. ZEKE (2022): "Risk-Adjusted Capital Allocation and Misallocation," *Journal of Financial Economics*, 145, 684–705.
- DHYNE, E., A. K. KIKKAWA, M. MOGSTAD, AND F. TINTELNOT (2021): "Trade and Domestic Production Networks," *Review of Economic Studies*, 88, 643–668.
- DUPOR, B. (1999): "Aggregation and Irrelevance in Multi-Sector Models," *Journal of Monetary Economics*, 43, 391–409.
- ELLIOTT, M., B. GOLUB, AND M. V. LEDUC (2022): "Supply Network Formation and Fragility," American Economic Review, 112, 2701–47.
- ERSAHIN, N., M. GIANNETTI, AND R. HUANG (2022): "Supply Chain Risk: Changes in Supplier Composition and Vertical Integration," Working paper.
- FAJGELBAUM, P. D., E. SCHAAL, AND M. TASCHEREAU-DUMOUCHEL (2017): "Uncertainty Traps," *Quarterly Journal of Economics*, 132, 1641–1692.
- FERNALD, J. (2014): "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," Working paper.
- FERNÁNDEZ-VILLAVERDE, J., P. GUERRÓN-QUINTANA, J. F. RUBIO-RAMÍREZ, AND M. URIBE (2011): "Risk Matters: The Real Effects of Volatility Shocks," *American Economic Review*, 101, 2530–2561.
- FOERSTER, A. T., P.-D. G. SARTE, AND M. W. WATSON (2011): "Sectoral versus Aggregate Shocks: A Structural Factor Analysis of Industrial Production," *Journal of Political Economy*, 119, 1–38.

- GABAIX, X. (2011): "The Granular Origins of Aggregate Fluctuations," *Econometrica*, 79, 733–772.
- HO, W., T. ZHENG, H. YILDIZ, AND S. TALLURI (2015): "Supply Chain Risk Management: A Literature Review," International Journal of Production Research, 53, 5031–5069.
- HORVATH, M. (2000): "Sectoral Shocks and Aggregate Fluctuations," Journal of Monetary Economics, 45, 69–106.
- HULTEN, C. R. (1978): "Growth Accounting with Intermediate Inputs," *Review of Economic Studies*, 45, 511–518.
- HUNEEUS, F. (2018): "Production Network Dynamics and the Propagation of Shocks," Working paper.
- JONES, C. I. (2011): "Intermediate Goods and Weak Links in the Theory of Economic Development," American Economic Journal: Macroeconomics, 3, 1–28.
- JURADO, K., S. C. LUDVIGSON, AND S. NG (2015): "Measuring Uncertainty," American Economic Review, 105, 1177–1216.
- KENNEDY, J. AND R. EBERHART (1995): "Particle Swarm Optimization," in *Proceedings of ICNN'95-international conference on neural networks*, IEEE, vol. 4, 1942–1948.
- KOPYTOV, A. (2018): "Booms, Busts, and Common Risk Exposures," Working paper.
- LIM, K. (2018): "Endogenous Production Networks and the Business Cycle," Working paper.
- LIU, E. (2019): "Industrial Policies in Production Networks," Quarterly Journal of Economics, 134, 1883–1948.
- LONG, J. B. AND C. I. PLOSSER (1983): "Real Business Cycles," *Journal of Political Economy*, 91, 39–69.
- LUCAS, R. E. (1987): Models of Business Cycles, vol. 26, Basil Blackwell Oxford.
- NIEUWERBURGH, S. V. AND L. VELDKAMP (2006): "Learning Asymmetries in Real Business Cycles," *Journal of Monetary Economics*, 53, 753–772.
- OBERFIELD, E. (2018): "A Theory of Input-Output Architecture," *Econometrica*, 86, 559–589.
- TASCHEREAU-DUMOUCHEL, M. (2020): "Cascades and Fluctuations in an Economy with an Endogenous Production Network," Working paper.
- VOM LEHN, C. AND T. WINBERRY (2021): "The Investment Network, Sectoral Comovement, and the Changing U.S. Business Cycle," *Quarterly Journal of Economics*.
- 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.
- ZURICH INSURANCE GROUP (2015): "Potential Business Impact of 'Loss of the Main Supplier' for Small and Medium Enterprises (SMEs) in 2015," Global survey report.

Online Appendix

A Microfoundation for the "one technique" restriction

In the main text, we made the ad hoc assumption that each sector can only adopt one production technique. Without this restriction, a large number of production techniques might be adopted and, after the shock ε is realized, only the technique that is best suited to this specific realization of ε would produce. In practice, we believe that several frictions might prevent this type of behavior. For instance, information frictions might make it impossible to redirect demand to the plant with the best technique after the shock is realized. Alternatively, engineers might be needed to explore how to set up a production technique, and there might be economies of scales pushing firms to adopt the same technique to save on engineering costs.

In this appendix, we propose one possible microfoundation for the "one technique" restriction. This microfoundation relies on decentralized trade for goods and on an information friction that prevents buyers from targeting specific producers. To describe this microfoundation, we first go over the economic agents in this environment. As in the main text, we still assume that there are nsectors/goods, but we are now explicit about the firms that operate within a sector. Specifically, in each sector $i \in \{1, \ldots, n\}$ there is a continuum of firms indexed by $l \in [0, 1]$. Each firm l operates a plant that can produce using a single production technique $\alpha_i^l \in \mathcal{A}_i$. We assume that physical restrictions, such as available factory space, prevent a plant for adopting multiple techniques. Different firms/plants in the same sector are however free to adopt different techniques.

Transactions between buyers (the household and intermediate firms) and sellers are conducted through shoppers. These shoppers are sent out by the buyers to meet sellers and negotiate terms of trade. Each shopper is atomistic, can purchase a measure one of goods and is matched with a seller at random. It follows that if in the market for good *i* there is a total demand of Q_i , each producer *l* is matched with a mass $Q_i dl$ of shoppers. Importantly, we assume that shoppers do not observe anything about the producers before they meet, and so cannot direct their search in any way.

If a shopper from firm m in sector j (or from the household) meets producer l in sector i, they agree to trade at a price \tilde{P}_{il}^{jm} through a protocol described below.³⁶ From these prices we can compute the effective price paid by a firm in j (or by the household) for goods i. Since m sends a continuum of shoppers to all producers in sector i, the effective price it pays is equal to the average price

$$\tilde{P}_i^{jm} = \int_0^1 \tilde{P}_{il}^{jm} dl.$$

³⁶For notational convenience, let j = 0 denote the household and assume that there is a unit mass of "sub-households" indexed by $m \in [0, 1]$.

Individual prices \tilde{P}_{il}^{jm} are set by splitting the joint surplus of the match through Nash bargaining. Specifically, if we denote the marginal benefit to the buyer of acquiring good *i* as B_i^{jm} , then the transaction price is such that the surplus of the seller is equal to a fraction $0 \le \varsigma \le 1$ of the total surplus. That is to say,

$$\tilde{P}_{il}^{jm} - K_i \left(\alpha_i^l, \left\{ \tilde{P}_k^{il} \right\}_{k \in \{1, \dots, n\}} \right) = \varsigma \left(B_i^{jm} - K_i \left(\alpha_i^l, \left\{ \tilde{P}_k^{il} \right\}_{k \in \{1, \dots, n\}} \right) \right), \tag{35}$$

where $K_i\left(\alpha_i^l, \left\{\tilde{P}_k^{il}\right\}_{k \in \{1,...,n\}}\right)$ is the unit cost of producer l in sector i under a chosen technique α_i^l .

From this last equation, we can write the technique choice problem of firm l. Since techniques are chosen before uncertainty is realized, we must average (35) across all states of the world, taking into account the stochastic discount factor of the household and the varying demand (mass of shoppers) for the good. It follows that firm l in sector i picks a production technique α_i^l to maximize

$$\mathbf{E}\left[\Lambda\sum_{j=0}^{n}Q_{ji}dl\int_{0}^{1}\varsigma\left(B_{i}^{jm}-K_{i}\left(\alpha_{i}^{l},\left\{\tilde{P}_{k}^{il}\right\}_{k\in\{1,\dots,n\}}\right)\right)dm\right],$$

where $Q_{ji}dl$ denotes the demand from sector j for goods produced by firm l in sector i. Since α_i^l only affects this expression through K_i , this maximization problem is equivalent to minimizing

$$\mathbf{E}\left[\Lambda Q_i K_i \left(\alpha_i^l, \left\{\tilde{P}_k^{il}\right\}_{k \in \{1,\dots,n\}}\right)\right],\tag{36}$$

where $Q_i = \sum_{j=0}^{n} Q_{ji}$ is total demand for sector *i*. Notice that this technique choice problem would be the same as the one described by (9) in the main text if the vector of input prices did not depend on the specific buying firm *l* and if all prices were equal to unit costs. To get that outcome, we now take the limit $\varsigma \to 0$, so that the bargaining power of the sellers goes to zero. In that case, (35) implies that

$$\tilde{P}_{il}^{jm} = K_i \left(\alpha_i^l, \left\{ \tilde{P}_k^{il} \right\}_{k \in \{1, \dots, n\}} \right),$$

and so \tilde{P}_{il}^{jm} does not depend on the identity of the buyer, i.e. on j or m. It follows that effective demand $\tilde{P}_{k}^{il} = \int_{0}^{1} \tilde{P}_{ks}^{il} ds$ does not depend on i and l, and we therefore can write $\tilde{P}_{k}^{il} = P_{k}$. Finally, this implies that the cost minimization problem (36) does not depend on the specific identity l of the firm. Given that the TFP shifter function is log concave, all firms in sector i therefore make the same technique choice α_{i} , have the same unit cost $K_{i}(\alpha_{i}, P)$ where $P = (P_{1}, \ldots, P_{n})$, and that all prices are equal to unit cost, as in the model in the main text.

B First order approximation of the equilibrium network

In Section 5, we described how beliefs affect the structure of the production network when the cost of deviating from the ideal shares is large. In this appendix, we are explicit about these derivations. Recall that for this exercise we assume that the TFP shifter $a_i(\alpha_i)$ can be written as $a_i(\alpha_i) = -\bar{\kappa} \times \hat{a}_i(\alpha_i)$ with $\hat{a}_i(\alpha_i^\circ) = 0$ and $\frac{\partial \hat{a}_i(\alpha_i^\circ)}{\partial \alpha_{ij}} = 0$. Throughout, we also assume that the Hessian H_i° of \hat{a}_i evaluated at α_i° , i.e. the matrix with typical element $[H_i^\circ]_{jk} = \frac{\partial^2 \hat{a}_i(\alpha_i^\circ)}{\partial \alpha_{ij}\partial \alpha_{ik}}$, is positive definite for all i.

GDP and welfare when $\bar{\kappa}$ is large

We now derive approximate equations for GDP and welfare when the cost of deviating from the ideal shares is large. These equations are described in the main text in Section 5.3. Recall from (14) that GDP is the inner product of the Domar weight and sectoral TFP vectors. Using the standard formula for the derivative of a matrix inverse, we can approximate the vector of Domar weights around α° as

$$\omega' = \beta' \mathcal{L} \approx (\omega^{\circ})' + \underbrace{\beta' \mathcal{L}^{\circ} (\Delta \alpha) \mathcal{L}^{\circ}}_{(\Delta \omega)'}, \tag{37}$$

where $\omega^{\circ} = \beta' \mathcal{L}(\alpha^{\circ})$ and $\Delta \alpha$ is the change in the production network as defined in (50). In (37) we implicitly define $\Delta \omega$ for future reference.

Multiplying (37) with an approximation of the sectoral TFP vector, we can write log GDP as

$$y \approx (\omega^{\circ} + \Delta \omega)' \times \left(\varepsilon - \bar{\kappa} \begin{bmatrix} (\Delta \alpha_1)' H_1^{\circ} \Delta \alpha_1 \\ \vdots \\ (\Delta \alpha_i)' H_i^{\circ} \Delta \alpha_i \\ \vdots \\ (\Delta \alpha_n)' H_n^{\circ} \Delta \alpha_n \end{bmatrix} \right),$$
(38)

where we have used the fact that $a(\alpha^{\circ}) = 0$ and $\frac{\partial \hat{a}_i(\alpha_i^{\circ})}{\partial \alpha_{ij}} = 0$ for all *i* and *j*. Keeping only linear terms in $\bar{\kappa}^{-1}$ (from (50), $\Delta \alpha$ is of order $\bar{\kappa}^{-1}$, which means that Δa is also of order $\bar{\kappa}^{-1}$) we find

$$y = (\omega^{\circ})' \varepsilon + (\Delta \omega)' \varepsilon + (\omega^{\circ})' \Delta a, \qquad (31)$$

which is the equation displayed in the main text. It is then straightforward to compute the expected

value and the variance of log GDP, to a first order, as

$$\mathbf{E}[y] \approx (\omega^{\circ})' \mu + (\Delta \omega)' \mu + (\omega^{\circ})' \Delta a,$$
$$\mathbf{V}[y] \approx (\omega^{\circ})' \Sigma \omega^{\circ} + 2 (\Delta \omega)' \Sigma \omega^{\circ}.$$

Using (18), we can then compute welfare as

$$\mathcal{W} \approx \underbrace{(\omega^{\circ})' \mu + (\Delta \omega)' \mu + (\omega^{\circ})' \Delta a}_{\mathrm{E}[y]} - \frac{1}{2} \left(\rho - 1\right) \underbrace{\left((\omega^{\circ})' \Sigma \omega^{\circ} + 2 \left(\Delta \omega\right)' \Sigma \omega^{\circ}\right)}_{\mathrm{V}[y]}.$$

These equations provide closed-form characterizations of the key aggregate quantities in the economy.

It is then straightforward to compute how log GDP responds to changes in beliefs. The derivatives of E[y] are given by

$$\begin{split} &\frac{d \operatorname{E}\left[y\right]}{d\mu_{i}} \approx \omega_{i}^{\circ} + \Delta \omega_{i} + \left(\frac{d\Delta\omega}{d\mu_{i}}\right)' \mu + (\omega^{\circ})' \frac{d\Delta a}{d\mu_{i}}, \\ &\frac{d \operatorname{E}\left[y\right]}{d\Sigma_{ii}} \approx \left(\frac{d\Delta\omega}{d\Sigma_{ii}}\right)' \mu + (\omega^{\circ})' \frac{d\Delta a}{d\Sigma_{ii}}, \end{split}$$

where

$$\frac{d\Delta\omega}{d\mu_i} \approx \beta' \mathcal{L}^{\circ} \left(\frac{d\Delta\alpha}{d\mu_i}\right) \mathcal{L}^{\circ},$$
$$\frac{d\Delta\omega}{d\Sigma_{ii}} \approx \beta' \mathcal{L}^{\circ} \left(\frac{d\Delta\alpha}{d\Sigma_{ii}}\right) \mathcal{L}^{\circ},$$
$$\frac{d\Delta a_j}{d\mu_i} \approx 2 \left(\frac{d\Delta\alpha_j}{d\mu_i}\right)' H_j^{\circ} \Delta\alpha_j,$$
$$\frac{d\Delta a_j}{d\Sigma_{ii}} \approx 2 \left(\frac{d\Delta\alpha_j}{d\Sigma_{ii}}\right)' H_j^{\circ} \Delta\alpha_j,$$

and where $\frac{d\Delta\alpha}{d\mu_i}$ and $\frac{d\Delta\alpha}{d\Sigma_{ii}}$ are given by (23) and (24). Similarly, the derivatives of V [y] are

$$\frac{d \operatorname{V}[y]}{d\mu_{i}} \approx 2 \left(\frac{d\Delta\omega}{d\mu_{i}}\right)' \Sigma\omega^{\circ},$$
$$\frac{d \operatorname{V}[y]}{d\Sigma_{ii}} \approx 2 \left(\frac{d\Delta\omega}{d\Sigma_{ii}}\right)' \Sigma\omega^{\circ} + (\omega_{i}^{\circ})^{2} + 2\Delta\omega_{i}\omega_{i}^{\circ},$$

and the derivatives of \mathcal{W} can be computed by combining the underlying derivatives of $\mathbf{E}[y]$ and $\mathbf{V}[y]$.

C Additional results related to the calibrated economy

C.1 List of the sectors in our analysis

Table 5 provides the list of the sectors in the vom Lehn and Winberry (2021) dataset and in the calibrated economy

Mining	Utilities
Construction	Wood products
Nonmetallic minerals	Primary metals
Fabricated metals	Machinery
Computer and electronic manufacturing	Electrical equipment manufacturing
Motor vehicles manufacturing	Other transportation equipment
Furniture and related manufacturing	Misc. manufacturing
Food and beverage manufacturing	Textile manufacturing
Apparel manufacturing	Paper manufacturing
Printing products manufacturing	Petroleum and coal manufacturing
Chemical manufacturing	Plastics manufacturing
Wholesale trade	Retail trade
Transportation and warehousing	Information
Finance and insurance	Real estate and rental services
Professional and technical services	Management of companies and enterprises
Administrative and waste management services	Educational services
Health care and social assistance	Arts and entertainment services
Accommodation	Food services
Other services	

Table 5: The 37 sectors used in our analysis

Notes: Sectors are classified according to the NAICS-based BEA codes. See vom Lehn and Winberry (2021) for details of the data construction.

C.2 Details of the calibrated economy

Cost of deviating from the ideal input shares

The overall mean of the elements of the calibrated cost matrix $\hat{\kappa}$ is 194 with a standard deviation of 447. The average and the standard deviation of the elements of the estimated vector $\hat{\kappa}^i$ are 14.11 and 17.56, respectively. The analogous statistics for $\hat{\kappa}^j$ are 13.73 and 17.18. To interpret these numbers, it helps to transform them into what they imply for productivity. If we increase one input share from its ideal value by one standard deviation in one sector, the average TFP loss for that sector is 0.06%.

To better understand the structure of the $\hat{\kappa}$ matrix, Figure 8 shows for each sector the elements of the vectors $\hat{\kappa}^i$ and $\hat{\kappa}^j$. As we can see, the amount of variation across sectors is quite substantial. The sectors with the highest $\hat{\kappa}^i$ are "Misc. manufacturing" and "Machinery", indicating that it is particularly costly for these sectors to deviate from their ideal input shares. The sectors with the highest $\hat{\kappa}^{j}$ are "Petroleum", "Furniture" and "Health care", implying that all firms tend to find it costly to adjust their input share of these sectors.



Figure 8: The calibrated costs of deviating from the ideal input shares

Sectoral total factor productivity

The estimated drift vector $\hat{\gamma}$ features substantial variation across sectors, indicating sizable dispersion in the trajectory of sectoral TFP. "Computer and electronic manufacturing" has the highest average annual growth in the sample, with ε_{it} growing 5.6% faster than the average sector. At the other end of the spectrum, productivity in "Food services" shrank by -2.9% per year relative to the average sector.

Similarly, the estimated covariance matrix $\hat{\Sigma}_t$ suggests that there is also substantial dispersion in uncertainty across sectors. The most volatile productivity is found in "Electrical equipment" with an average $\sqrt{\hat{\Sigma}_{iit}}$ of 38.0%, and the least volatile sector is "Real estate" with an average $\sqrt{\hat{\Sigma}_{iit}}$ of 1.8%. There is also a lot of variation across sectors in how much volatility changes over time. The standard deviation of $\sqrt{\hat{\Sigma}_{iit}}$ is largest for the "Electrical equipment" sector at 25.6% and smallest for "Real estate" at 1.1%.

C.3 Input shares and Domar weights

In this appendix, we compare the behavior of the input shares and the Domar weights in the baseline economy and in the two alternative economies introduced in Section 6: the one with $\Sigma_t = 0$ and the economy in which production techniques are chosen after observing ε_t . As we can see from Table 6 all versions of the model perform almost identically in terms of average shares and Domar weights. The standard deviations differ however across models. Specifically, the baseline model, in which firms care about uncertainty, features standard deviations that 4% to 8% lower than in the alternatives, depending on the precise comparison.

	Statistic	Version of the model				
		Data	Baseline	$\Sigma_t = 0$	Known ε_t	
(1)	Average Domar weight $\bar{\omega}_i$	0.047	0.032	0.032	0.032	
(2)	Standard deviation $\sigma(\omega_j)$	0.0050	0.0021	0.0022	0.0023	
(3)	Coefficient of variation $\sigma(\omega_j)/\bar{\omega}_j$	0.107	0.066	0.070	0.070	
(4)	Average share, $\bar{\alpha}_{ij}$	0.013	0.008	0.008	0.008	
(5)	Standard derivation $\sigma(\alpha_{ij})$	0.0048	0.0023	0.0024	0.0024	
(6)	Coefficient of variation $\sigma(\alpha_{ij})/\bar{\alpha}_{ij}$	0.37	0.30	0.31	0.31	

Table 6: Domar weights and input shares in the model and in the data

Notes: For each sector, we compute the time series of its Domar weight ω_{jt} , as well as their their mean $\bar{\omega}_j$ and standard deviation $\sigma(\omega_j)$. Rows (1) and (2) report the cross-sectional average of these statistics. Row (3) is the ratio of rows (2) and (1). For each pair of sectors, we compute the time series of the input share α_{ijt} , as well as their their mean $\bar{\alpha}_{ij}$ and standard deviation $\sigma(\alpha_{ij})$. Rows (4) and (5) report the cross-sectional average of these statistics. Row (6) is the ratio of rows (5) and (4). The "Baseline" model is the model in which risk-averse firms choose techniques before TFP shocks ε are realized. The " $\Sigma_t = 0$ " model is the model in which firms choose techniques after the TFP shocks ε are realized.

The response of the network to uncertainty differs particularly strongly across models during periods of high uncertainty. To show this, we compute changes in sectoral Domar weights $\Delta \omega_{it} = \omega_{it} - \omega_{i,t-1}$ in the baseline model and in the two alternatives. As usual, we denote changes in sectoral Domar weights in the alternative models by tildes, i.e. $\Delta \tilde{\omega}_{it}$. We then compute the cross-sectional average of the absolute differences between $\Delta \omega_{it}$ and $\Delta \tilde{\omega}_{it}$, and normalize it by the cross-sectional average of standard deviations of $|\Delta \omega_{it}|$. This measure captures the difference between models in how Domar weights change over time.

Figure 9 shows the resulting graphs. In the top panel, the " $\Sigma_t = 0$ " model is used as alternative. In the bottom panel, the "known ε_t " model is used as alternative. In the top panel the differences are particularly pronounced during high-uncertainty episodes, when risk-averse firms actively switch to safer production inputs. In the bottom panel, the Domar weights deviate from the baseline model much more. This is because the production network adapts to the specific ε_t draw, and so the differences are visible even in relatively tranquil times.

C.4 Correlations between sectors

In this section, we report the correlation between sectoral output in the calibrated model and in the data. We focus on growth rates to accommodate different trends in the data and in the model. The results are presented in Table 7. In the first column, we see that sectoral sales are positively correlated in both the model and in the data data, although the model correlation is a bit weaker. We also report these correlations during periods of low and high TFP growth and uncertainty growth, as measured by (33) and (34). We see that in the data these correlations are lower in good times, when TFP growth is robust and uncertainty growth is low. The model is able to replicate that ranking. Intuitively, in bad times consumption is low and so the household is particularly





Notes: Panel (a): difference between the series implied by the baseline model (without tildes) and the "as if $\Sigma_t = 0$ " alternative (with tildes); Panel (b): difference between the baseline model (without tildes) and the alternative in which firms choose techniques after TFP shocks ε are realized (with tildes). Both economies are hit by the same shocks that are filtered out from the TFP data under our baseline model. Both series are normalized by the cross-sectional average of the standard deviations of growths in sectoral Domar weights.

worried about bad shocks. To avoid them, firms rely more on the most stable producers. As firms are mostly purchasing from the same sectors, sectoral outputs become more correlated.

	All years	TFP growth		Uncertainty growth	
		Low	High	Low	High
Model	0.18	0.22	0.13	0.16	0.20
Data	0.36	0.37	0.34	0.32	0.38

Table 7: Correlations in sectoral sales growth

Notes: Correlation of sectoral sales growth in the model and in the data. TFP growth and Uncertainty growth are measured as in Figure 5. We use high/low to refer to years with TFP growth or uncertainty growth above/below corresponding median levels.

D More details about the regressions of Section 7

In this section, we provide more details about the regressions presented in Tables 3 and 4. The firm-level production network data comes from the Factset Revere database and covers the period from 2003 to 2016. We limit the sample to relationships that have lasted at least five years. The IV estimates remain significant when relationships of other lengths are considered. The firm-level uncertainty data comes from Alfaro et al. (2019) and was downloaded from Nicholas Bloom's website at https://nbloom.people.stanford.edu. We thank the authors for sharing their data.

Alfaro et al. (2019) describes how the data is constructed in detail, and we only include here a summary of how the instruments are computed. The instruments are created by first computing the industry-level sensitivity to each aggregate shock c, where c is either the price of oil, one of seven exchange rates, the yield on 10-year U.S. Treasury Notes and the economic policy uncertainty index of Baker et al. (2016). As Alfaro et al. (2019) explain, "for firm i in industry j, sensitivity^c = β_j^c is estimated as follows

$$r_{i,t}^{riskadj} = \alpha_j + \sum_c \beta_j^c \cdot r_t^c + \epsilon_{i,t},$$

where $r_{i,t}^{riskadj}$ is the daily risk-adjusted return of firm i, r_t^c is the change in the price of commodity c, and α_j is industry j's intercept. [...] Estimating the main coefficients of interest, β_j^c , at the SIC 3-digit level (instead of at the firm-level) reduces the role of idiosyncratic noise in firm-level returns, increasing the precision of the estimates. [...] We allow these industry-level sensitivities to be time-varying by estimating them using 10-year rolling windows of daily data." The instruments $z_{i,t-1}^c$ are then computed as follows:

$$z_{i,t-1}^c = \left|\beta_{j,t-1}^c\right| \cdot \Delta \sigma_{t-1}^c$$

where $\Delta \sigma_{t-1}^c$ denotes the volatility of the aggregate variable c. As a result, instruments vary on the 3-digit SIC industry-by-year level. As in Alfaro et al. (2019), we also include in the IV regressions the first moments associated with each aggregate series c ("1st moment 10IV_{*i*,*t*-1}" in Tables 3 and 4) to isolate the impact of changes in their second moment alone. Note that we control for year×customer×supplier industry (2-digit SIC) fixed effects in Tables 3 and 4. Therefore, instruments and control variables used in columns (2) and (3) exhibit nontrivial variation within fixed-effects bins. At the same time, such rich fixed effects allow us to compare how a given customer firm in a given year reacts to different volatility shocks hitting its suppliers within the same 2-digit SIC industry.

E Proofs

This section contains the proofs of the formal results from the main text.

E.1 Proof of Lemma 1

Lemma 1. For a fixed production network α ,

$$p(\alpha) = -\mathcal{L}(\alpha) \left(\varepsilon + a(\alpha)\right), \tag{13}$$

and

$$y(\alpha) = \omega(\alpha)'(\varepsilon + a(\alpha)), \qquad (14)$$

where $a(\alpha) = (\log A_i(\alpha_i), \dots, \log A_n(\alpha_n))$ and $\omega(\alpha) = (\omega_1(\alpha), \dots, \omega_n(\alpha)).$

Proof. Combining the unit cost equation (8) with the equilibrium condition (11) and taking the log we find that, for all i,

$$p_i = -\varepsilon_i - a_i \left(\alpha_i\right) + \sum_{j=1}^n \alpha_{ij} p_j, \tag{39}$$

where $a_i(\alpha_i) = \log(A_i(\alpha_i))$. This is a system of linear equations whose solution is (13). The log price vector is also normally distributed since it is a linear transformation of normal random variable. Combining with (6) yields (14).

E.2 Proof of Corollary 1

Corollary 1. For a fixed production network α , the following holds.

1. The impact of a change in expected TFP μ_i on expected log GDP E[y] is given by

$$\frac{\partial \operatorname{E}\left[y\right]}{\partial \mu_{i}} = \omega_{i}.$$

2. The impact of a change in volatility Σ_{ij} on the variance of log GDP V [y] is given by³⁷

$$\frac{\partial \mathbf{V}[y]}{\partial \Sigma_{ij}} = \begin{cases} \omega_i^2 & i = j, \\ 2\omega_i \omega_j & i \neq j. \end{cases}$$

Proof. (15) implies that $\frac{\partial E[y(\alpha)]}{\partial \mu_i} = \beta' \mathcal{L}(\alpha) \mathbf{1}_i$. Since P'C = WL = 1 by the household's budget constraint, we need to show that $\beta' \mathcal{L}(\alpha) \mathbf{1}_i = P_i Q_i$ to complete the proof of the first result. From (53), we know that $P_i C_i = \beta_i$. Using Shepard's Lemma together with the marginal pricing equation (11), we can find the firm's factor demands equations

$$P_j X_{ij} = \alpha_{ij} P_i Q_i$$
$$L_i = \left(1 - \sum_{j=1}^n \alpha_{ij}\right) P_i Q_i.$$
(40)

³⁷Throughout the paper, derivatives with respect to off-diagonal elements of Σ simultaneously change Σ_{ij} and Σ_{ji} to preserve the symmetry of Σ .

Using these results, we can write the market clearing condition (12) as

$$P_i Q_i = \beta_i + \sum_{j=1}^n \alpha_{ji} P_j Q_j.$$

Solving the linear system implies

$$\beta' \mathcal{L}(\alpha) \, \mathbf{1}_i = P_i Q_i,\tag{41}$$

which proves the first part of the proposition (recall that P'C = 1 by the household's budget constraint). For the second part of the result, differentiating (15) with respect to Σ_{ij} and holding Σ symmetric yields

$$\frac{\partial \operatorname{V}\left[y\left(\alpha\right)\right]}{\partial \Sigma_{ij}} = \begin{cases} \beta' \mathcal{L}\left(\alpha\right) \mathbf{1}_{i} \mathbf{1}_{i}' \mathcal{L}\left(\alpha\right)' \beta & i = j, \\ \beta' \mathcal{L}\left(\alpha\right) \left[\mathbf{1}_{i} \mathbf{1}_{j}' + \mathbf{1}_{j} \mathbf{1}_{i}'\right] \mathcal{L}\left(\alpha\right)' \beta & i \neq j \end{cases} = \begin{cases} \omega_{i}^{2} & i = j, \\ 2\omega_{i}\omega_{j} & i \neq j, \end{cases}$$

which is the result.

E.3 Proof of Lemma 2

Lemma 2. In equilibrium, $\lambda(\alpha^*)$, $k_i(\alpha_i, \alpha^*)$ and $q_i(\alpha^*)$ are normally distributed, and the technique choice of the representative firm in sector *i* solves

$$\alpha_{i}^{*} \in \arg\min_{\alpha_{i} \in \mathcal{A}_{i}} \mathbb{E}\left[k_{i}\left(\alpha_{i}, \alpha^{*}\right)\right] + \operatorname{Cov}\left[\lambda\left(\alpha^{*}\right), k_{i}\left(\alpha_{i}, \alpha^{*}\right)\right].$$
(16)

Proof. We first consider the stochastic discount factor. (55) in Appendix F.1 shows that aggregate consumption can be written as a function of prices. Combining that equation with (5) we can write $\lambda = \log(\Lambda)$ as

$$\lambda(\alpha^*) = -(1-\rho)\sum_{i=1}^n \beta_i p_i(\alpha^*).$$
(42)

Taking the log of (8) yields

$$k_i(\alpha_i, \alpha^*) = -(\varepsilon_i + a(\alpha_i)) + \sum_{j=1}^n \alpha_{ij} p_j(\alpha^*).$$
(43)

Both $\lambda(\alpha^*)$ and $k_i(\alpha_i, \alpha^*)$ are normally distributed since they are linear combinations of ε and the log price vector, which is normally distributed by Lemma 1.

Turning to the firm problem (10), we can write

$$\alpha_{i}^{*} \in \arg\min_{\alpha_{i} \in \mathcal{A}_{i}} \mathbb{E}\left[\Lambda \frac{\beta' \mathcal{L}(\alpha^{*}) \mathbf{1}_{i}}{P_{i}} K_{i}(\alpha_{i}, P)\right]$$

where we have used (41). We can drop $\beta' \mathcal{L}(\alpha^*) \mathbf{1}_i \geq 0$ since it is a deterministic scalar that does not depend on α_i . Rewriting this equation in terms of log quantities yields

$$\alpha_{i}^{*} \in \arg\min_{\alpha_{i} \in \mathcal{A}} \mathbb{E}\left[\exp\left[\lambda\left(\alpha^{*}\right) - p_{i}\left(\alpha^{*}\right) + k_{i}\left(\alpha_{i},\alpha^{*}\right)\right]\right],$$

where we emphasize that λ and p_i depend only on the equilibrium technique choice α^* . All the terms in the exponential are normally distributed. We can therefore use the expression for the expected value of a lognormal distribution and write

$$\alpha_{i}^{*} \in \arg\min_{\alpha_{i} \in \mathcal{A}} \exp\left\{ \mathbb{E}\left[\lambda\left(\alpha^{*}\right) - p_{i}\left(\alpha^{*}\right) + k_{i}\left(\alpha_{i},\alpha^{*}\right)\right] + \frac{1}{2} \operatorname{V}\left[\lambda\left(\alpha^{*}\right) - p_{i}\left(\alpha^{*}\right) + k_{i}\left(\alpha_{i},\alpha^{*}\right)\right] \right\}.$$

Taking away the exponentiation, as it is a monotone transformation, and $\mathbb{E}\left[\lambda\left(\alpha^{*}\right)-q_{i}\left(\alpha^{*}\right)\right]$ since it does not affect the minimization yields

$$\alpha_i^* \in \arg\min_{\alpha_i \in \mathcal{A}} \mathbb{E}\left[k_i\left(\alpha_i, \alpha^*\right)\right] + \frac{1}{2} \operatorname{V}\left[\lambda\left(\alpha^*\right) - p_i\left(\alpha^*\right) + k_i\left(\alpha_i, \alpha^*\right)\right].$$
(44)

This expression can be written as

$$\alpha_{i}^{*} \in \arg\min_{\alpha_{i} \in \mathcal{A}} \mathbb{E}\left[k_{i}\left(\alpha_{i}, \alpha^{*}\right)\right] + \frac{1}{2} \operatorname{V}\left[\lambda\left(\alpha^{*}\right)\right] + \frac{1}{2} \operatorname{V}\left[k_{i}\left(\alpha_{i}, \alpha^{*}\right) - p_{i}\left(\alpha^{*}\right)\right] + \operatorname{Cov}\left[\lambda\left(\alpha^{*}\right), k_{i}\left(\alpha_{i}, \alpha^{*}\right)\right] + \operatorname{Cov}\left[\lambda\left(\alpha^{*}\right), -p_{i}\left(\alpha^{*}\right)\right],$$

where we can drop V $[\lambda(\alpha^*)]$ and 2 Cov $[\lambda(\alpha^*), -p_i(\alpha^*)]$ as they do not affect the maximization. Finally, we can expand V $[k_i(\alpha_i, \alpha^*) - p_i(\alpha^*)]$ to get

$$\alpha_{i}^{*} \in \arg\min_{\alpha_{i} \in \mathcal{A}} \mathbb{E}\left[k_{i}\left(\alpha_{i}, \alpha^{*}\right)\right] + \frac{1}{2} \mathbb{E}\left[\left(k_{i}\left(\alpha_{i}, \alpha^{*}\right) - p_{i}\left(\alpha^{*}\right) - \mathbb{E}\left[k_{i}\left(\alpha_{i}, \alpha^{*}\right) - p_{i}\left(\alpha^{*}\right)\right]\right)^{2}\right] + \operatorname{Cov}\left[\lambda\left(\alpha^{*}\right), k_{i}\left(\alpha_{i}, \alpha^{*}\right)\right].$$

Taking the first-order condition with respect to α_{ik} , we find

$$\frac{1}{2} \mathbf{E} \left[2 \left(k_i \left(\alpha_i, \alpha^* \right) - p_i \left(\alpha^* \right) - \mathbf{E} \left[k_i \left(\alpha_i, \alpha^* \right) - p_i \left(\alpha^* \right) \right] \right) \left(\frac{dk_i \left(\alpha_i, \alpha^* \right)}{d\alpha_{ik}} - \mathbf{E} \left[\frac{dk_i \left(\alpha_i, \alpha^* \right)}{d\alpha_{ik}} \right] \right) \right] \\ + \mathbf{E} \left[\frac{dk_i \left(\alpha_i, \alpha^* \right)}{d\alpha_{ik}} \right] + \mathbf{Cov} \left[\lambda \left(\alpha^* \right), \frac{dk_i \left(\alpha_i, \alpha^* \right)}{d\alpha_{ik}} \right] + \eta_i + \nu_{ik} = 0,$$

where η_i is the Lagrange multiplier on $\sum_{j=1}^n \alpha_{ij} \leq \overline{\alpha}_i$ and ν_{ik} is the multiplier on $\alpha_{ik} \geq 0$. At an equilibrium, $\alpha = \alpha^*$ and $k_i (\alpha_i^*, \alpha^*) = p_i (\alpha^*)$, and so

$$\mathbf{E}\left[\frac{dk_{i}\left(\alpha_{i}^{*},\alpha^{*}\right)}{d\alpha_{ik}}\right] + \mathbf{Cov}\left[\lambda\left(\alpha^{*}\right),\frac{dk_{i}\left(\alpha_{i}^{*},\alpha^{*}\right)}{d\alpha_{ik}}\right] + \eta_{i} + \nu_{ik} = 0,$$

describes the equilibrium choice of firm i. Notice that this equilibrium first-order condition can also come from the problem

$$\alpha_{i}^{*} \in \arg\min_{\alpha_{i} \in \mathcal{A}} \mathbb{E}\left[k_{i}\left(\alpha_{i}, \alpha^{*}\right)\right] + \operatorname{Cov}\left[\lambda\left(\alpha^{*}\right), k_{i}\left(\alpha_{i}, \alpha^{*}\right)\right],$$

which completes the proof.

E.4 Proof of Proposition 1

Proposition 1. An equilibrium exists.

Proof. We proceed in three steps. First, we show that there is a unique technique α_i that solves the problem of the firm, i.e. \mathcal{K}_i is a function. Second, we show that that function is continuous. Finally, we use a fixed-point theorem to show the existence of an equilibrium.

Step 1. We show that the right-hand side of (44) is a strictly concave function. First, note that from (43) we can write

$$E[k_i(\alpha_i, \alpha^*)] = E\left[-(\varepsilon_i + a(\alpha_i)) + \sum_{j=1}^n \alpha_{ij} p_j(\alpha^*)\right]$$
$$= -a(\alpha_i) + E\left[-\varepsilon_i - \alpha'_i \mathcal{L}(\alpha^*)(\varepsilon + a(\alpha^*))\right]$$

which is strictly convex in α_i since $a(\alpha_i) = \log A_i(\alpha_i)$ is strictly concave by Assumption 1.

Similarly, combining (42) and (43) we can write

$$\frac{1}{2} \operatorname{V} \left[\lambda\left(\alpha^{*}\right) - p_{i}\left(\alpha^{*}\right) + k_{i}\left(\alpha_{i},\alpha^{*}\right)\right] = \frac{1}{2} \operatorname{V} \left[-\left(\varepsilon_{i} + a\left(\alpha_{i}\right)\right) - p_{i}\left(\alpha^{*}\right) + \sum_{j=1}^{n}\left(\alpha_{ij} - \left(1 - \rho\right)\beta_{j}\right)p_{j}\left(\alpha^{*}\right)\right].$$

We can remove the term $a(\alpha_i)$ from the variance as it is not stochastic. Combining with the equilibrium price equation (13), we get

$$\frac{1}{2} \operatorname{V} \left[\lambda \left(\alpha^* \right) - p_i \left(\alpha^* \right) + k_i \left(\alpha_i, \alpha^* \right) \right] = \frac{1}{2} \operatorname{V} \left[-\varepsilon_i + 1'_i \mathcal{L} \left(\alpha^* \right) \left(\varepsilon + a \left(\alpha^* \right) \right) - \left(\alpha_i - \left(1 - \rho \right) \beta \right)' \mathcal{L} \left(\alpha^* \right) \left(\varepsilon + a \left(\alpha^* \right) \right) \right] \\ = \frac{1}{2} \operatorname{V} \left[-\varepsilon_i - \left(\alpha_i - 1_i - \left(1 - \rho \right) \beta \right)' \mathcal{L} \left(\alpha^* \right) \left(\varepsilon + a \left(\alpha^* \right) \right) \right]$$

where 1_i is a column vector with a 1 in element *i* and zeros elsewhere. Once again we can drop the term in $a(\alpha^*)$ as it is non stochastic. Define the row vector *B* as

$$B(\alpha_i, \alpha^*) = -(\alpha_i - 1_i - (1 - \rho)\beta)' \mathcal{L}(\alpha^*) - 1'_i,$$

where $\beta = (\beta_1, \ldots, \beta_n)$ is a column vector. Then

$$V\left[\lambda\left(\alpha^{*}\right)-p_{i}\left(\alpha^{*}\right)+k_{i}\left(\alpha_{i},\alpha^{*}\right)\right]=B\left(\alpha_{i},\alpha^{*}\right)\Sigma B\left(\alpha_{i},\alpha^{*}\right)',$$

where Σ is the covariance matrix of ε . The right-hand side will have a term that is linear in α_i , and that therefore does not affect the concavity of the expression, and the quadratic term

$$\alpha_i' \mathcal{L}(\alpha^*) \Sigma \mathcal{L}(\alpha^*)' \alpha_i.$$

The matrix $\mathcal{L}(\alpha^*) \Sigma \mathcal{L}(\alpha^*)'$ is a covariance matrix and hence positive semi-definite. The expression $V[\lambda(\alpha^*) - p_i(\alpha^*) + k_i(\alpha_i, \alpha^*)]$ is therefore convex in α_i . Since the sum of a strictly convex function and a convex function is strictly convex, the expression

$$\mathbf{E}\left[k_{i}\left(\alpha_{i},\alpha^{*}\right)\right] + \frac{1}{2}\mathbf{V}\left[\lambda\left(\alpha^{*}\right) - p_{i}\left(\alpha^{*}\right) + k_{i}\left(\alpha_{i},\alpha^{*}\right)\right]$$

is strictly convex in the vector α_i .

To complete this first step, note that the set of techniques \mathcal{A} is convex. Since the problem of the firm involves the minimization of strictly convex function on a convex set it has a unique minimizer. The mapping $\mathcal{K}_i(\alpha^*)$ is therefore a function for every i and every $\alpha^* \in \mathcal{A}$.

Step 2. We now show that \mathcal{K}_i is continuous. To simplify the notation, define

$$g_i(\alpha, \alpha^*) = \mathbf{E} \left[k_i(\alpha, \alpha^*) \right] + \frac{1}{2} \mathbf{V} \left[\lambda(\alpha^*) - p_i(\alpha^*) + k_i(\alpha, \alpha^*) \right].$$

where we have temporarily removed the subscript i on the column vector α_i to avoid cluttering the notation.

We will first show that $g_i(\alpha, \alpha^*)$ is continuous. From (42) and (43), λ and k are continuous functions of α and linear functions of $p(\alpha^*)$. It therefore suffices to show that $p(\alpha^*)$ is continuous. From (13), we see that $p(\alpha^*)$ is continuous since $\mathcal{L}(\alpha^*)$, as a matrix inverse, in continuous and $a(\alpha^*)$ is continuous by Assumption 1. So $g_i(\alpha, \alpha^*)$ is continuous.

We now turn to the proof of the continuity of \mathcal{K}_i . We have already shown that g is strictly convex in α so there is a unique minimizer $\mathcal{K}_i(\alpha^*) = \arg \min_{\alpha} g_i(\alpha, \alpha^*)$. Take a sequence $\alpha_k^* \to \alpha_\star^*$ and let $\alpha_k = \mathcal{K}_i(\alpha_k^*)$ and $\alpha_\star = \mathcal{K}_i(\alpha_\star^*)$. Choose any subsequence $I \subset \mathbb{N}$, then α_k has an accumulation point α'_k since \mathcal{A} is compact. Since $g(\alpha_k, \alpha_k^*) \leq g(\alpha, \alpha_k^*)$ for all $\alpha \in \mathcal{A}$ and $k \in I$ we have, by continuity of g, that $g_i(\alpha'_k, \alpha^*) \leq g_i(\alpha, \alpha^*)$ for all $\alpha \in \mathcal{A}$ and since the minimizer is unique it must be that $\alpha'_k = \alpha_\star$. As a result, $\alpha_k \to \alpha_\star$ and κ_i is continuous.

Step 3. We have shown that the mapping $\mathcal{K}_i(\alpha^*)$ is continuous for all i = 1, ..., n. Define the mapping $\mathcal{K}(\alpha^*) = (\mathcal{K}_1(\alpha^*), ..., \mathcal{K}_n(\alpha^*))$. Then $\mathcal{K}(\alpha^*)$ is a continuous mapping from \mathcal{A} (a compact and convex set) to itself. Therefore, by Brouwer's fixed-point theorem \mathcal{K} has a fixed point and an equilibrium exists.

E.5 Proof of Proposition 2

Proposition 2. There exists an efficient equilibrium.

Proof. Since we only have one agent in the economy, any Pareto efficient allocation must maximize the utility of the representative household. Under a given network and a given productivity shock ε the first welfare theorem applies, and the equilibrium is efficient. The consumption chosen by the planner is therefore given by (14). Taking a step back, the efficient production network must therefore solve

$$\max_{\alpha \in \mathcal{A}} \operatorname{E} \left[u\left(Y\right) \right] = \max_{\alpha \in \mathcal{A}} \frac{1}{1-\rho} \operatorname{E} \left[\exp\left(\left(1-\rho\right)\log Y\right) \right]$$
$$= \max_{\alpha \in \mathcal{A}} \frac{1}{1-\rho} \exp\left(\left(1-\rho\right)\operatorname{E} \left[\log Y\right] + \frac{1}{2}\left(1-\rho\right)^2 \operatorname{V} \left[\log Y\right] \right)$$
$$= \max_{\alpha \in \mathcal{A}} \operatorname{E} \left[\log Y\right] - \frac{1}{2}\left(\rho-1\right) \operatorname{V} \left[\log Y\right]$$
(45)

where we have used the fact that $\log Y$ is normally distributed. Note that this problem involves maximizing a continuous function over a compact set so that a solution in \mathcal{A} exists by the extreme value theorem. The rest of the proof compares the first-order conditions of the planner and of the firms in the equilibrium.

First-order conditions of the planner. Using (15), we can write (45) as

$$\max_{\alpha \in \mathcal{A}} \beta' \mathcal{L}(\alpha) \left(\mu + a(\alpha)\right) + \frac{1}{2} \left(1 - \rho\right) \beta' \mathcal{L}(\alpha) \Sigma \mathcal{L}(\alpha)' \beta.$$

The first-order conditions are

$$0 = \beta' \left(\frac{\partial}{\partial \alpha_{ij}} \mathcal{L}(\alpha) \right) \left(\mu + a(\alpha) + \frac{1}{2} (1 - \rho) \Sigma \mathcal{L}(\alpha)' \beta \right) + \beta' \mathcal{L}(\alpha) \left(\frac{\partial}{\partial \alpha_{ij}} a(\alpha) + \frac{1}{2} (1 - \rho) \Sigma \left(\frac{\partial}{\partial \alpha_{ij}} \mathcal{L}(\alpha) \right)' \beta \right) + \underline{\mu}_{ij} - \gamma_i$$

where $\underline{\mu}_{ij}$ are the Lagrange multipliers on the constraints on $\alpha_{ij} \ge 0$ and γ_i is the Lagrange multiplier on the constraint $\sum_j \alpha_{ij} \le \overline{\alpha}_i$.³⁸ Now,

$$\frac{\partial}{\partial \alpha_{ij}} \mathcal{L}(\alpha) = \frac{\partial}{\partial \alpha_{ij}} \left(I - \alpha\right)^{-1} = -\left(I - \alpha\right)^{-1} \left[\frac{\partial}{\partial \alpha_{ij}} \left(I - \alpha\right)\right] \left(I - \alpha\right)^{-1}$$
(46)

$$= (I - \alpha)^{-1} [O_{ij}] (I - \alpha)^{-1} = \mathcal{L} (\alpha) O_{ij} \mathcal{L} (\alpha)$$
(47)

³⁸Note that $\sum_{j} \alpha_{ij} \leq \overline{\alpha}_i < 1$ implies that $\alpha_{ij} < 1$, so we do not need to explicitly consider a constraint $\alpha_{ij} \leq 1$.

where $O_{ij} = 1_i 1'_j$ is a matrix full of zero except for a one at element (i, j). Plugging back in and grouping terms yields

$$0 = \beta' \mathcal{L}(\alpha) \, \mathbf{1}_{i} \mathbf{1}_{j}' \mathcal{L}(\alpha) \, [\mu + a(\alpha)] + \beta' \mathcal{L}(\alpha) \, \mathbf{1}_{i} \frac{\partial}{\partial \alpha_{ij}} a_{i}(\alpha) + (1 - \rho) \, \beta' \mathcal{L}(\alpha) \, \mathbf{1}_{i} \mathbf{1}_{j}' \mathcal{L}(\alpha) \, \Sigma \mathcal{L}(\alpha)' \, \beta + \underline{\mu}_{ij} - \gamma_{i}$$

Since $\beta' \mathcal{L}(\alpha) \mathbf{1}_i$ is a strictly positive scalar we can divide the whole equation by it to find

$$0 = 1'_{j}\mathcal{L}(\alpha)\left[\mu + a(\alpha)\right] + \frac{\partial}{\partial\alpha_{ij}}a_{i}(\alpha) + (1-\rho)1'_{j}\mathcal{L}(\alpha)\Sigma\mathcal{L}(\alpha)'\beta$$

$$+ \left(\beta'\mathcal{L}(\alpha)1_{i}\right)^{-1}\left(\underline{\mu}_{ij} - \gamma_{i}\right).$$
(48)

Firms first-order conditions. We can repeat similar steps for the equilibrium. Combining (44) with (13), (42) and (43), we find that firm *i*'s problem can be written as

$$\alpha_i^* = \arg\min_{\alpha_i \in \mathcal{A}_i} -a(\alpha_i) - \alpha_i' \mathcal{L}(\alpha^*) (\mu + a(\alpha^*)) + \frac{1}{2} \left((\alpha_i - 1_i - (1 - \rho)\beta)' \mathcal{L}(\alpha^*) + 1_i' \right) \Sigma \left((\alpha_i - 1_i - (1 - \rho)\beta)' \mathcal{L}(\alpha^*) + 1_i' \right)'.$$

Differentiating with respect to α_{ij} we can write the first-order conditions as

$$0 = -\frac{\partial a(\alpha_i)}{\partial \alpha_{ij}} - 1'_j \mathcal{L}(\alpha^*)(\mu + a(\alpha^*)) + \frac{1}{2}(1'_j \mathcal{L}(\alpha^*)) \Sigma((\alpha_i - 1_i - (1 - \rho)\beta)' \mathcal{L}(\alpha^*) + 1'_i)' + \frac{1}{2}((\alpha_i - 1_i - (1 - \rho)\beta)' \mathcal{L}(\alpha^*) + 1'_i) \Sigma \mathcal{L}(\alpha^*)' 1_j + \underline{\mu}^e_{ij} - \gamma^e_i$$

or

_

$$0 = -\frac{\partial a(\alpha_i)}{\partial \alpha_{ij}} - 1'_j \mathcal{L}(\alpha^*) (\mu + a(\alpha^*)) + ((\alpha_i - 1_i - (1 - \rho)\beta)' \mathcal{L}(\alpha^*) + 1'_i) \Sigma \mathcal{L}(\alpha^*)' 1_j + \underline{\mu}^e_{ij} - \gamma^e_i,$$

where the Lagrange multipliers have a superscript e to indicate the equilibrium. In equilibrium $\alpha = \alpha^*$ and so

$$-\frac{\partial a\left(\alpha_{i}^{*}\right)}{\partial \alpha_{ij}}-1_{j}^{\prime}\mathcal{L}\left(\alpha^{*}\right)\left(\mu+a\left(\alpha^{*}\right)\right)+\left(\left(\alpha_{i}^{*}-1_{i}-\left(1-\rho\right)\beta\right)^{\prime}\mathcal{L}\left(\alpha^{*}\right)+1_{i}^{\prime}\right)\Sigma\mathcal{L}\left(\alpha^{*}\right)^{\prime}1_{j}+\underline{\mu}_{ij}^{e}-\gamma_{i}^{e}=0.$$

Finally, we can show that $(1_i - \alpha_i^*)' \mathcal{L}(\alpha^*) - 1_i' = 0$ by right-multiplying both sides by $(\mathcal{L}(\alpha^*))^{-1}$. As a result, the first-order conditions become

$$-\frac{\partial a\left(\alpha_{i}^{*}\right)}{\partial \alpha_{ij}}-1_{j}^{\prime}\mathcal{L}\left(\alpha^{*}\right)\left(\mu+a\left(\alpha^{*}\right)\right)-\left(1-\rho\right)\beta^{\prime}\mathcal{L}\left(\alpha^{*}\right)\Sigma\mathcal{L}\left(\alpha^{*}\right)^{\prime}1_{j}+\underline{\mu}_{ij}^{e}-\gamma_{i}^{e}=0.$$

Notice that these are the same first-order conditions (up to a normalization of the Lagrange multipliers) as the planner's (equation 48). The complementary slackness conditions are also the same in both problems. As a result, any equilibrium allocation also satisfies the planner's first-order conditions and vice versa. Since the planner's first-order conditions are necessary to characterize the efficient allocation, it follows that there is an equilibrium that coincides with the efficient allocation. Note however that the planner's problem is in general not concave and that the first-order conditions are therefore not sufficient. \Box

E.6 Proof of Corollary 2

Corollary 2. The efficient equilibrium production network α^* solves

$$\mathcal{W} \equiv \max_{\alpha \in \mathcal{A}} \operatorname{E}\left[y\left(\alpha\right)\right] - \frac{1}{2}\left(\rho - 1\right) \operatorname{V}\left[y\left(\alpha\right)\right],\tag{18}$$

where W is a measure of the welfare of the household, and y is log GDP as defined in (14).

Proof. This is an intermediate result that was proven at (45) in the proof of Proposition 2.

E.7 Generic uniqueness of the efficient equilibrium

Consider the planner's objective function from (18): $\mathcal{W}(\alpha; z) = \beta' \mathcal{L}(\alpha) (\mu + a(\alpha, z)) + \frac{1}{2} (1 - \rho) \beta' \mathcal{L}(\alpha) \Sigma \mathcal{L}(\alpha)' \beta$, where z is a vector of parameters, which includes μ, Σ, β, ρ and any additional parameters of the $a(\alpha, z)$ function. Define a space \mathcal{Z} on the set of parameters z. We endow this space with an absolutely continuous probability measure \mathbb{P} . We will call the solution to that problem generically unique if the set \mathcal{Z}^* for which \mathcal{W} has multiple maximizers is almost surely empty, i.e. $\mathbb{P}(z \in \mathcal{Z}^*) = 0$.

Our proof strategy relies on Lemma 1 from Cox (2020).

Proposition A1. Suppose that $A_i(\alpha_i)$ takes the form (2) and all elements of the κ matrix are positive.³⁹ Then the Pareto efficient equilibrium is generically unique.

Proof. Lemma 1 of Cox (2020) requires that three properties be satisfied.

- 1. The set $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$, where \mathcal{A}_i is the set of feasible production technique for firm *i*, must be a disjoint union of finitely or countably many second-countable Hausdorff manifolds, possibly with boundary or corner. This assumption is satisfied in our case since \mathcal{A} is a manifold in \mathbb{R}^{n^2} of dimension $n^2 n$.
- 2. We need $\mathcal{W}(\alpha, z)$ to be differentiable with respect to z and the derivative to be continuous with respect to α and z. This is satisfied in our case given the form (2).

 $^{^{39}}$ This is the functional form we use for the quantitative analyses. In our calibrated model, all elements of the κ matrix are positive.

3. It must be that for all $\alpha_1, \alpha_2 \in \mathcal{A}$ such that $\alpha_1 \neq \alpha_2$ we have $\frac{d\mathcal{W}(\alpha_1, z)}{dz} \neq \frac{d\mathcal{W}(\alpha_2, z)}{dz}$, where the derivative here indicates the gradient. We prove this by contraposition. For that purpose, take $\alpha_1, \alpha_2 \in \mathcal{A}$ such that $\frac{d\mathcal{W}(\alpha_1, z)}{dz} = \frac{d\mathcal{W}(\alpha_2, z)}{dz}$. We are going to show that it implies that $\alpha_1 = \alpha_2$. From Proposition 6, it must be that $\frac{d\mathcal{W}(\alpha_1, z)}{d\mu_i} = \omega_i(\alpha_1, z) = \omega_i(\alpha_2, z) = \frac{d\mathcal{W}(\alpha_2, z)}{d\mu_i}$. Since this is true for all *i*, it follows that the vector of Domar weights must be the same, that it $\omega(\alpha_1, z) = \omega(\alpha_2, z) > 0$. Next, differentiate $\mathcal{W}(\alpha, z)$ with respect to α_{il}° to write

$$\frac{d\mathcal{W}(\alpha, z)}{d\alpha_{il}^{\circ}} = 2\omega_i \left[\kappa_{il} \left(\alpha_{il} - \alpha_{il}^{\circ} \right) + \kappa_{i0} \left(\sum_{j=1}^n \alpha_{ij} - \sum_{j=1}^n \alpha_{ij}^{\circ} \right) \right].$$

Suppose by contradiction that $\alpha_1 \neq \alpha_2$. Then there exists a pair i, l such that $(\alpha_{il})_1 \neq (\alpha_{il})_2$. Without loss of generality, suppose that $(\alpha_{il})_1 > (\alpha_{il})_2$. Then it must be that $\sum_{j=1}^n (\alpha_{ij})_1 < \sum_{j=1}^n (\alpha_{ij})_2$ for $\frac{dW(\alpha_{1,z})}{d\alpha_{il}^\circ} = \frac{dW(\alpha_{2,z})}{d\alpha_{il}^\circ}$ to hold. Therefore, there exists l' such that $(\alpha_{il'})_1 < (\alpha_{il'})_2$. But then it must be $\frac{dW(\alpha_{1,z})}{d\alpha_{il'}^\circ} < \frac{dW(\alpha_{2,z})}{d\alpha_{il'}^\circ}$. Therefore, we have a contradiction and $\alpha_1 = \alpha_2$.

We have shown that the three properties required by Lemma 1 of Cox (2020) are satisfied. It follows that $\mathbb{P}(z \in \mathbb{Z}^*) = 0$ and the planner's solution is generically unique. As a result, there is a generically unique efficient equilibrium.

E.8 Proof of Proposition 3

Proposition 3. The Domar weight ω_i is increasing in μ_i and decreasing in Σ_{ii} .

Proof. Fix the initial mean and variance-covariance matrix at μ^0 and Σ^0 , and denote the optimal network by $\alpha^*(\mu^0, \Sigma^0)$. Now, consider an increase in μ_i from μ_i^0 to μ_i^1 (holding other elements of μ^0 and Σ^0 fixed). The welfare changes from $\mathcal{W}\left(\alpha^*\left(\mu_i^0, \mu_{-i}^0, \Sigma^0\right); \mu_i^0, \mu_{-i}^0, \Sigma^0\right)$ to $\mathcal{W}\left(\alpha^*\left(\mu_i^1, \mu_{-i}^0, \Sigma^0\right); \mu_i^1, \mu_{-i}^0, \Sigma^0\right)$, which, by Proposition 6, can be written as

$$\mathcal{W}\left(\alpha^{*}\left(\mu_{i}^{1},\mu_{-i}^{0},\Sigma^{0}\right);\mu_{i}^{1},\mu_{-i}^{0},\Sigma^{0}\right)=\mathcal{W}\left(\alpha^{*}\left(\mu_{i}^{0},\mu_{-i}^{0},\Sigma^{0}\right);\mu_{i}^{0},\mu_{-i}^{0},\Sigma^{0}\right)+\int_{\mu_{i}^{0}}^{\mu_{i}^{1}}\omega_{i}\left(\mu_{i},\mu_{-i}^{0},\Sigma^{0}\right)d\mu_{i}.$$

Now suppose instead that the network is fixed at its original value $\alpha^* (\mu_i^0, \mu_{-i}^0, \Sigma^0)$. From Equations (15), a change in μ_i affects welfare only through its impact on expected GDP. By Lemma 1, the change in welfare can be written as

$$\mathcal{W}\left(\alpha^{*}\left(\mu_{i}^{0},\mu_{-i}^{0},\Sigma^{0}\right);\mu_{i}^{1},\mu_{-i}^{0},\Sigma^{0}\right) = \mathcal{W}\left(\alpha^{*}\left(\mu_{i}^{0},\mu_{-i}^{0},\Sigma^{0}\right);\mu_{i}^{0},\mu_{-i}^{0},\Sigma^{0}\right) + \omega_{i}\left(\mu_{i}^{0},\mu_{-i}^{0},\Sigma^{0}\right)\left(\mu_{i}^{1}-\mu_{i}^{0}\right)$$

Since the initial network $\alpha^* \left(\mu_i^0, \mu_{-i}^0, \Sigma^0 \right)$ is attainable at $\left(\mu_i^1, \mu_{-i}^0, \Sigma^0 \right)$, it must be that $\mathcal{W} \left(\alpha^* \left(\mu_i^1, \mu_{-i}^0, \Sigma^0 \right); \mu_i^1, \mu_{-i}^0, \Sigma^0 \right) \geq \mathcal{W} \left(\alpha^* \left(\mu_i^0, \mu_{-i}^0, \Sigma^0 \right); \mu_i^1, \mu_{-i}^0, \Sigma^0 \right)$. Because μ_i^1 can be picked

arbitrary close to μ_i^0 , it must therefore be that $\omega_i\left(\mu_i^1, \mu_{-i}^0, \Sigma^0\right) \ge \omega_i\left(\mu_i^0, \mu_{-i}^0, \Sigma^0\right)$, or $\frac{d\omega_i}{d\mu_i} \ge 0$.

For the second part of the proposition, recall that $\frac{dW}{d\Sigma_{ii}} = (1-\rho)\omega_i^2$ by Proposition 6. Using analogous steps, we then can establish the second part of this proposition.

E.9 Proof of Proposition 4

Proposition 4. For large $\bar{\kappa}$ the vector of input shares in sector *i* is approximately given by

$$\alpha_i \approx \alpha_i^{\circ} - \bar{\kappa}^{-1} \left(H_i^{\circ} \right)^{-1} \mathcal{R}^{\circ}, \tag{21}$$

where H_i° is the Hessian matrix of \hat{a}_i evaluated at the ideal shares α_i° .

Proof. An interior solution to the planner's problem must satisfy the first-order conditions (48), which in our setup can be written as

$$-\frac{\partial \hat{a}_{i}\left(\alpha_{i}\right)}{\partial \alpha_{ij}}+1_{j}^{\prime}\mathcal{L}\left(\alpha\right)\left[\bar{\kappa}^{-1}\mu-\hat{a}\left(\alpha\right)\right]-\bar{\kappa}^{-1}\left(\rho-1\right)1_{j}^{\prime}\mathcal{L}\left(\alpha\right)\Sigma\mathcal{L}\left(\alpha\right)^{\prime}\beta=0,\tag{49}$$

for all *i* and *j*. These conditions implicitly define the efficient network as a function of $\bar{\kappa}$, and so we write $\alpha = \alpha$ ($\bar{\kappa}$). We characterize this mapping when $\bar{\kappa}$ is large or, equivalently, when $\bar{\kappa}^{-1}$ is small. Given our assumptions on \hat{a}_i , it must be that in the limit as $\bar{\kappa} \to \infty$ we have $\alpha_i = \alpha_i^{\circ}$. By Taylor's theorem, we can write

$$\begin{pmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{pmatrix} = \begin{pmatrix} \alpha_1^{\circ} \\ \vdots \\ \alpha_n^{\circ} \end{pmatrix} + \begin{pmatrix} \frac{d\alpha_1}{d\bar{\kappa}^{-1}} \\ \vdots \\ \frac{d\alpha_n}{d\bar{\kappa}^{-1}} \end{pmatrix}_{\bar{\kappa}^{-1}=0} \times \bar{\kappa}^{-1} + o\left(\bar{\kappa}^{-1}\right).$$

To compute the first-order term, we use the implicit function theorem together with (49) to write

$$\begin{pmatrix} \frac{d\alpha_1}{d\bar{\kappa}^{-1}} \\ \vdots \\ \frac{d\alpha_n}{d\bar{\kappa}^{-1}} \end{pmatrix}_{\bar{\kappa}^{-1}=0} = (H^{\circ})^{-1} \times \begin{pmatrix} \mathcal{L}^{\circ}\mu - (\rho - 1) \mathcal{L}^{\circ}\Sigma\omega^{\circ} \\ \vdots \\ \mathcal{L}^{\circ}\mu - (\rho - 1) \mathcal{L}^{\circ}\Sigma\omega^{\circ} \end{pmatrix},$$

where $(H^{\circ})^{-1}$ is a block-diagonal matrix with the inverse Hessians $\{(H_i^{\circ})^{-1}\}_{i=1}^n$ on the main diagonal. Therefore, the linearized equilibrium shares are

$$\alpha_i \approx \alpha_i^{\circ} + \underbrace{\bar{\kappa}^{-1} \left(H_i^{\circ} \right)^{-1} \left[\mathcal{L}^{\circ} \mu - (\rho - 1) \mathcal{L}^{\circ} \Sigma \omega^{\circ} \right]}_{\Delta \alpha_i}, \tag{50}$$

where we implicitly define the linear term $\Delta \alpha_i$ for future reference.

We can write (50) in terms of prices. Since $a_i(\alpha_i^\circ) = 0$ for all *i*, it follows from (13) that the

equilibrium expected log price vector evaluated at the ideal shares is given by $E[p(\alpha^{\circ})] = -\mathcal{L}^{\circ}\mu$. Similarly, from (42) we can write

$$\operatorname{cov}\left(p\left(\alpha^{\circ}\right),\lambda\left(\alpha^{\circ}\right)\right) = \operatorname{cov}\left(p\left(\alpha^{\circ}\right),\left(\rho-1\right)\beta'p\left(\alpha^{\circ}\right)\right)$$
$$= \left(\rho-1\right)\mathcal{L}^{\circ}\Sigma\left(\mathcal{L}^{\circ}\right)'\beta = \left(\rho-1\right)\mathcal{L}^{\circ}\Sigma\omega^{\circ},$$

where $\lambda(\alpha^{\circ})$ is the log of the stochastic discount factor under the ideal input shares. It follows that we can write (50) as

$$\alpha_i \approx \alpha_i^{\circ} + \bar{\kappa}^{-1} \left(H_i^{\circ} \right)^{-1} \underbrace{\left[- \operatorname{E} \left[p^{\circ} \right] - \operatorname{cov} \left(p^{\circ}, \lambda^{\circ} \right) \right]}_{-\mathcal{R}^{\circ}}, \tag{21}$$

which is the equation that we use in the main text.

E.10 Proof of Corollary 3

Corollary 3. For large $\bar{\kappa}$ the impact of an increase in μ_k on the network is approximately given by

and the impact of an increase in Σ_{km} on the network is approximately given by

$$\frac{\partial \alpha_{ij}}{\partial \Sigma_{km}} \approx \begin{cases} -\left(\rho - 1\right) \omega_k^{\circ} \frac{d\alpha_{ij}}{d\mu_k} & k = m, \\ -\left(\rho - 1\right) \left(\omega_m^{\circ} \frac{d\alpha_{ij}}{d\mu_k} + \omega_k^{\circ} \frac{d\alpha_{ij}}{d\mu_m}\right) & k \neq m, \end{cases}$$
(24)

where ω° is the vector of Domar weights at the ideal input shares.

Proof. The proof follows by directly differentiating (50).

E.11 Proof of Proposition 5

Proposition 5. Expected log GDP E[y] reaches its maximum at $\Sigma = 0$.

Proof. The proof follows from Lemma 2. Without uncertainty $(\Sigma = 0)$, the term $V[c(\alpha)]$ is 0 for all α , and so α is set to maximize $E[c(\alpha)]$. When uncertainty is introduced, the objective function also depends on $V[c(\alpha)]$ and so E[c] is no longer maximized.

E.12 Proof of Proposition 6

Proposition 6. When the network α is free to adjust to changes in μ and Σ the following holds.

1. The impact of an increase in μ_i on expected welfare is given by

$$\frac{d\mathcal{W}}{d\mu_i} = \frac{\partial \operatorname{E}\left[y\right]}{\partial \mu_i} = \omega_i.$$
(25)

2. The impact of an increase in Σ_{ij} on expected welfare is given by

$$\frac{d\mathcal{W}}{d\Sigma_{ij}} = \begin{cases} -\frac{1}{2} \left(\rho - 1\right) \left(\frac{\partial \operatorname{E}[y]}{\partial \mu_i}\right)^2 = -\frac{1}{2} \left(\rho - 1\right) \omega_i^2 & i = j, \\ -\left(\rho - 1\right) \frac{\partial \operatorname{E}[y]}{\partial \mu_i} \frac{\partial \operatorname{E}[y]}{\partial \mu_j} = -\left(\rho - 1\right) \omega_i \omega_j & i \neq j. \end{cases}$$
(26)

Proof. Recall from Lemma 2 that the equilibrium α^* solves the welfare-maximization problem

$$\mathcal{W}(\mu, \Sigma) = \max_{\alpha \in \mathcal{A}} \left\{ E\left[y\left(\alpha\right)\right] - \frac{1}{2}\left(\rho - 1\right) V\left[y\left(\alpha\right)\right] \right\}.$$

Since that the objective function and the constraints are continuously differentiable functions of α , we can apply the envelope theorem, such that

$$\frac{d\mathcal{W}}{d\mu_{i}} = \frac{\partial \operatorname{E}\left[y\right]}{\partial \mu_{i}} = \beta' \mathcal{L}\left(\alpha\right) \mathbf{1}_{i} = \omega_{i},$$

and

$$\frac{d\mathcal{W}}{d\Sigma_{ij}} = -\frac{1}{2} \left(\rho - 1\right) \frac{\partial \operatorname{V}\left[y\left(\alpha\right)\right]}{\partial \Sigma_{ij}} = \left(1 - \rho\right) \beta' \mathcal{L}\left(\alpha\right) \left(1_{i} 1_{j}'\right) \mathcal{L}\left(\alpha\right)' \beta = \left(1 - \rho\right) \omega_{i} \omega_{j},$$

where we used the expressions for the expectation and the variance of output given by (15).

E.13 Proof of Proposition 7

Proposition 7. Let $\alpha^*(\mu, \Sigma)$ be the equilibrium production network under (μ, Σ) and let $\mathcal{W}(\alpha, \mu, \Sigma)$ be the welfare of the household under the network α . The change in welfare after a change in beliefs from (μ, Σ) to (μ', Σ') satisfies the inequality

$$\underbrace{\mathcal{W}\left(\alpha^{*}\left(\mu',\Sigma'\right),\mu',\Sigma'\right)-\mathcal{W}\left(\alpha^{*}\left(\mu,\Sigma\right),\mu,\Sigma\right)}_{\text{Change in welfare under the flexible network}} \geq \underbrace{\mathcal{W}\left(\alpha^{*}\left(\mu,\Sigma\right),\mu',\Sigma'\right)-\mathcal{W}\left(\alpha^{*}\left(\mu,\Sigma\right),\mu,\Sigma\right)}_{\text{Change in welfare under the fixed network}}.$$
 (27)

Proof. By definition, the change in welfare under the flexible network is

$$\mathcal{W}\left(\alpha^{*}\left(\mu',\Sigma'\right),\mu',\Sigma'\right)-\mathcal{W}\left(\alpha^{*}\left(\mu,\Sigma\right),\mu,\Sigma\right).$$

By Proposition 2, $\alpha^*(\mu', \Sigma')$ maximizes welfare under (μ, Σ) so that

$$\mathcal{W}\left(\alpha^{*}\left(\mu',\Sigma'\right),\mu',\Sigma'\right) \geq \mathcal{W}\left(\alpha^{*}\left(\mu,\Sigma\right),\mu',\Sigma'\right).$$

Combining the two expression gives the result.

Supplemental appendix (Not for publication)

F Additional derivations

This appendix contains additional derivations that are used in the main text.

F.1 Derivation of the stochastic discount factor

The Lagrange multiplier on the budget constraint of the household captures the value of an extra unit of the numeraire and serves as stochastic discount factor for firms to compare profits across states of the world. The following lemma shows how to derive the expression in the main text.

Lemma 3. The Lagrange multiplier on the budget constraint of the household (4) is

$$\Lambda = \frac{u'\left(Y\right)}{\overline{P}},$$

where $Y = \prod_{i=1}^{n} (\beta_i^{-1} C_i)^{\beta_i}$ and $\overline{P} = \prod_{i=1}^{n} P_i^{\beta_i}$.

Proof. The household makes decisions after the realization of the state of the world ε . The statespecific maximization problem has a concave objective function and a convex constraint set so that first-order conditions are sufficient to characterize optimal decisions. The Lagrangian is

$$u\left(\left(\frac{C_1}{\beta_1}\right)^{\beta_1} \times \dots \times \left(\frac{C_n}{\beta_n}\right)^{\beta_n}\right) - \Lambda\left(\sum_{i=1}^n P_i C_i - 1\right)$$

and the first-order condition with respect to C_i is

$$\beta_i u'(Y) Y = \Lambda P_i C_i. \tag{51}$$

Summing over i on both sides and using the binding budget constraint yields

$$u'(Y)Y = \Lambda, \tag{52}$$

which, together with (51), implies that

$$P_i C_i = \beta_i. \tag{53}$$

We can also plug back the first-order condition in $Y = \prod_{i=1}^{n} (\beta_i^{-1} C_i)^{\beta_i}$ to find

$$Y = \prod_{i=1}^{n} \left(\beta_i^{-1} C_i\right)^{\beta_i} = \prod_{i=1}^{n} \left(\beta_i^{-1} \frac{\beta_i u'\left(Y\right) Y}{\Lambda P_i}\right)^{\beta_i}$$
$$\Lambda = u'\left(Y\right) \prod_{i=1}^{n} P_i^{-\beta_i}$$
(54)

which, combined with (52), yields

$$Y = \prod_{i=1}^{n} P_i^{-\beta_i}.$$
 (55)

This last equation implicitly defines a price index $\overline{P} = \prod_{i=1}^{n} P_i^{\beta_i}$ such that $\overline{P}Y = 1$. Combining that last equation with (52) yields the result.

F.2 Derivation of the unit cost function

The cost minimization problem of the firm is

$$K_{i}(\alpha_{i}, P) = \min_{L_{i}, X_{i}} \left(L_{i} + \sum_{j=1}^{n} P_{j} X_{ij} \right)$$

subject to $F(\alpha_{i}, L_{i}, X_{i}) \ge 1$,

where F is given by (1). The first-order conditions are

$$L_{i} = \theta \left(1 - \sum_{j=1}^{n} \alpha_{ij} \right) F(\alpha_{i}, L_{i}, X_{i}),$$
$$P_{j}X_{ij} = \theta \alpha_{ij}F(\alpha_{i}, L_{i}, X_{i}),$$

where θ is the Lagrange multiplier. Plugging these expressions back into the objective function, we see that $K_i(\alpha_i, P) = \theta$ since $F(\alpha_i, L_i, X_i) = 1$ at the optimum. Now, plugging the first-order conditions in the production function we find

$$1 = e^{\varepsilon_i} A_i(\alpha_i) \theta \prod_{j=1}^n P_j^{-\alpha_{ij}},$$

which is the result.

G Additional results related to the calibrated economy

G.1 Sensitivity to the risk aversion parameter ρ

In this section, we investigate the sensitivity of our results to the value of the risk aversion parameter ρ . To do so, we solve the model for different values of ρ without recalibrating the matrix κ . We then compare this economy to the alternative with $\Sigma = 0$. Not surprisingly, we find that ignoring uncertainty is costlier for higher values of ρ (Table 8).

	Comparison with $\Sigma = 0$ model		
	$\rho = 2$	$\rho = 4.27$	$\rho = 10$
		(baseline)	
Expected log GDP $E[y(\alpha)]$	+0.001%	+0.008%	+0.033%
Std. dev. of log GDP $\sqrt{V[y(\alpha)]}$	+0.038%	+0.105%	+0.208%
Welfare \mathcal{W}	-0.001%	-0.010%	-0.057%

Table 8: Uncertainty, GDP and welfare: the Role of risk aversion

Notes: Baseline economy variables minus their counterparts in the $\Sigma = 0$ alternative for different values of risk aversion ρ .

The economy also responds to the spike in uncertainty during the Great Recession much more for $\rho = 10$ (Figure 10). Specifically, if $\rho = 10$, the network adjusts such that the standard deviation of GDP is almost 4.2% lower in 2009 relative to the risk-neutral alternative (second panel, yellow crossed line). Although this adjustment is associated with a sizable decline in expected GDP (-0.8%; first panel), welfare raises substantially (1.8%; third panel). This is because the representative household enjoys a larger utility gain from a reduction in uncertainty under a higher risk aversion parameter.

G.2 Great Recession: Flexible vs fixed network

In this section, we explore the role of network flexibility during the Great Recession—the period in which the economy was hit by large adverse shocks (see Figure 5). Specifically, we fix the network α at its 2006 pre-recession level and then hit the economy with the same shocks as in the baseline economy with endogenous network. Figure 11 shows how the baseline economy compares to the fixed-network alternative (denoted with tildes in the figure) over the years 2006 to 2012. We find that expected GDP (top panel) is higher under the flexible network. This is because firms are able to respond to changes in TFP and move away from sectors that are expected to perform badly. When doing so, firms become exposed to more productive but also more volatile suppliers, which results in an increase in GDP volatility (second panel). However, the first effect dominates, and welfare is quite substantially higher when the network is allowed to adjust (third panel). Interestingly, the economy with a flexible network does substantially worse in terms of realized GDP (bottom panel) during the Great Recession years. As evident from the two top panels, firms optimally choose to



Figure 10: The role of uncertainty during the Great Recession: Role of risk aversion

Notes: The differences between the series implied by the models featuring various degrees of risk aversion (without tildes) and the "as if $\Sigma_t = 0$ " alternative (with tildes). All differences are expressed in percentage terms.

be exposed to more productive but riskier suppliers. During the Great Recession, some of those risks were realized, pushing realized GDP down for the baseline case.

G.3 Time-varying consumption shares

In this appendix, we consider a version of the calibrated economy in which we let the β preference vector change over time to match the observed consumption shares in the data. Figure 12 shows the difference between that economy and the $\Sigma = 0$ alternative in which uncertainty has no impact on the production network. As we can see, this figure is quite similar to Figure 7 (left column) in the main text, suggesting that allowing β to change over time does not have a large effect on the impact of uncertainty on the network.



Figure 11: The role of network flexibility during the Great Recession

Notes: The differences between the series implied by the full model (without tildes) and the model in which the network is fixed at its 2006 level (with tildes). All differences are expressed in percentage terms.


Figure 12: The role of uncertainty in the postwar period with time-varying β

Notes: The differences between the series implied by the baseline model with time-varying β (without tildes) and the "as if $\Sigma_t = 0$ " alternative (with tildes). Both economies are hit by the same shocks that are filtered out from the TFP data under our baseline model. All differences are expressed in percentage terms.

H Network response when shares are complements

In this appendix, we show that when input shares are complement, we can unambiguously sign the impact of changes in beliefs on the production network. We consider economies in which the functions (a_1, \ldots, a_n) satisfy the following property.

Assumption 2 (Weak Complementarity). For all *i*, a_i satisfies $\frac{\partial^2 a_i(\alpha_i)}{\partial \alpha_{ij} \partial \alpha_{ik}} \ge 0$ for all $j \neq k$.

Assumption 2 defines a weak complementarity property between the shares that a producer allocates to its suppliers. It states that as a firm increases the share of one input, the marginal benefit of increasing the share of the other inputs weakly increases as well. In the context of the functional form (2), weak complementarity is satisfied if $\kappa_{i0} \leq 0.40$

The following lemma shows that the impact of μ and Σ on the equilibrium network is straightforward when Assumption 2 holds.

Lemma 4. Let $\alpha^* \in int(\mathcal{A})$ be the equilibrium network and suppose that Assumption 2 holds. There exists a scalar $\overline{\Sigma} > 0$ such that if $|\Sigma_{ij}| < \overline{\Sigma}$ for all i, j, there is a neighborhood around α^* in which

- (i) an increase in μ_j leads to an increase in the shares α_{kl}^* for all k, l;
- (ii) an increase in Σ_{jj} leads to a decline in the shares α_{kl}^* for all k, l;
- (iii) an increase in Σ_{ij} leads to a decline in the shares α_{kl}^* for all k, l.

Part (i) of this lemma shows that when μ_j increases there is a widespread increase in input shares throughout the economy. To understand this result, it is useful to decompose the impact of the change into three channels: 1) the direct impact, 2) the indirect impact, and 3) the complementarity effect. First, the increase in μ_j makes good *j* cheaper in expectation which pushes all of *j*'s direct customers to increase their share of *j* in production. Second, all of *j*'s customers now benefit from cheaper input prices, which makes their own goods cheaper through competition, and so other firms are also increasing their share of these goods into production (indirect effect). Finally, these increases in shares from the direct and indirect effects push firms to adopt techniques with higher input shares because of the complementarities implied by Assumption 2. Taking these effects together, all shares α in the economy increase, and so the entire production structure moves away from labor.

Parts (ii) and (iii) of Lemma 4 provide similar results for increases in uncertainty and in correlations. As discussed in Section 4.1, firms prefer suppliers with stable and uncorrelated prices. As a result, the additional risk introduced by a higher Σ_{jj} pushes firm j's direct and indirect customers to reduce their exposure to j. Similarly, an increase in the covariance Σ_{ij} pushes firms to

⁴⁰Note that $\kappa_{i0} < 0$ does not necessarily break the concavity requirement of Assumption 1.

avoid inputs i and j. The complementarity effect is also at work, and so firms overall move toward production techniques that are more labor intensive.⁴¹

Proof of Lemma 4

Proof. Point (i). Away from the constraints, the first-order conditions of the planner are

$$F_{jk} := \frac{\partial a_j}{\partial \alpha_{jk}} + 1'_k \mathcal{L} \left(\mu + \alpha\right) + (1 - \rho) \beta' \mathcal{L} \Sigma \mathcal{L}' 1_k = 0.$$

We use the implicit function theorem to investigate how α changes with μ_i . First, differentiate F_{jk} with respect to μ_i to find

$$\frac{\partial F_{jk}}{\partial \mu_i} = \mathbf{1}'_k \mathcal{L} \mathbf{1}_i = \mathcal{L}_{ki}.$$

It follows that

$$\frac{\partial F}{\partial \mu_i} = \begin{pmatrix} \left(\frac{\partial F_{1.}}{\partial \mu_i}\right)' \\ \left(\frac{\partial F_{2.}}{\partial \mu_i}\right)' \\ \dots \\ \left(\frac{\partial F_{n.}}{\partial \mu_i}\right)' \end{pmatrix} = \mathbf{1}_{n \times 1} \otimes (\mathcal{L}\mathbf{1}_i),$$

where $1_{n \times 1}$ is an $n \times 1$ column vector of ones, $\frac{\partial F}{\partial \mu_i}$ is an $n^2 \times 1$ column vector which consists of the *n* column vectors $\left(\frac{\partial F_{j.}}{\partial \mu_i}\right)'$ with elements $\left(\frac{\partial F_{jk}}{\partial \mu_i}\right)_{k=1,\dots,n}^{r}$. Next, differentiate F_{jk} with respect to α_{lm} to get

$$\begin{aligned} \frac{\partial F_{jk}}{\partial \alpha_{lm}} &= \frac{\partial^2 a_j}{\partial \alpha_{jk} \partial \alpha_{lm}} + 1'_k \mathcal{L} \mathbf{1}_l \frac{\partial a_l}{\partial \alpha_{lm}} + 1'_k \mathcal{L} \left(\mathbf{1}_l \mathbf{1}'_m \right) \mathcal{L} \left(\mu + a \right) \\ &+ \left(1 - \rho \right) \mathbf{1}'_k \mathcal{L} \Sigma \left(\beta' \mathcal{L} \mathbf{1}_l \mathbf{1}'_m \mathcal{L} \right)' + \left(1 - \rho \right) \beta' \mathcal{L} \Sigma \left(\mathbf{1}'_k \mathcal{L} \mathbf{1}_l \mathbf{1}'_m \mathcal{L} \right)' \\ &= \frac{\partial^2 a_j}{\partial \alpha_{jk} \partial \alpha_{lm}} + \left(1 - \rho \right) \mathbf{1}'_k \mathcal{L} \Sigma \left(\beta' \mathcal{L} \mathbf{1}_l \mathbf{1}'_m \mathcal{L} \right)' + \mathcal{L}_{kl} \left[\underbrace{\frac{\partial a_l}{\partial \alpha_{lm}} + \mathbf{1}'_m \mathcal{L} \left(\mu + a \right) + \left(1 - \rho \right) \beta' \mathcal{L} \Sigma \mathcal{L}' \mathbf{1}_m}_{=F_{lm} = 0} \right] \end{aligned}$$

where we use the first-order condition to set the last term to 0.

⁴¹The assumption that $\alpha^* \in int(\mathcal{A})$ in Lemma 4 is needed to avoid potential substitution patterns between firms. For instance, if $\sum_{k=1}^{n} \alpha_{ik} = \overline{\alpha}_i$ for a given firm *i*, an increase in μ_j might lead to a decline in some $\alpha_{ik}, k \neq j$, to accommodate an increase in α_{ij} . The restriction on Σ is needed to prevent a strong uncertainty feedback. For example, if all firms increase their reliance on sector j (e.g., due to an increase in μ_j or a reduction in Σ_{jj}), the economy's exposure to j's risk may become so large that it would be optimal to reduce α_{kj} for some k. This does not happen when Σ is sufficiently small.

Now, denote by A the $n^2 \times n^2$ block-diagonal matrix with the n blocks A_1, A_2, \ldots, A_n along the main diagonal such that $(A_j)_{kl} = \left(\frac{\partial^2 a_j}{\partial \alpha_{jk} \partial \alpha_{jl}}\right)_{k,l=1,\ldots,n}$. Denote by D the $n \times n^2$ matrix $(1-\rho)\left[\left(\beta^T \mathcal{L}\right) \otimes (\mathcal{L}\Sigma\mathcal{L}')\right]$, and by B the $n^2 \times n^2$ matrix that consists of n copies of D, i.e. $B = 1_{n \times 1} \otimes D$. Then, by the implicit function theorem, we have

$$\begin{pmatrix} \left(\frac{\partial \alpha_{1.}}{\partial \mu_{i}}\right)'\\ \left(\frac{\partial \alpha_{2.}}{\partial \mu_{i}}\right)'\\ \dots\\ \left(\frac{\partial \alpha_{n.}}{\partial \mu_{i}}\right)' \end{pmatrix} = -(A+B)^{-1}\frac{\partial F}{\partial \mu_{i}}.$$
(56)

We will now show that when $\Sigma = 0$ (and so B = 0), all the elements on the left-hand side of (56) are positive. Since the right-hand side of (56) is continuous in the elements of Σ , the left-hand side will remain positive for small Σ .

We first establish that the elements of $-A^{-1}$ are positive. Since a_i is strictly concave by Assumption 1, A_i is strictly negative definite for all *i*. As, in addition, Weak Complementarity (Assumption 2) holds, $-A_i$ is a (non-singular) M-matrix and so its inverse $-A_i^{-1}$ is nonnegative. The diagonal elements of $-A_i^{-1}$ are also strictly positive. To see this, note that since A_i is Hermitian, so is A_i^{-1} , and we know from the Rayleigh quotient that

$$\lambda_{\min}\left(A_{i}^{-1}\right) \leq \frac{x'A_{i}^{-1}x}{x'x} \leq \lambda_{\max}\left(A_{i}^{-1}\right),$$

where $\lambda_{\min}(A_i^{-1})$ and $\lambda_{\max}(A_i^{-1})$ are the smallest and largest eigenvalues of A_i^{-1} , respectively, and where x is any nonzero vector. By setting $x = 1_t$, the tth basis vector we get $\lambda_{\min}(A_i^{-1}) \leq (A_i^{-1})_{tt} \leq \lambda_{\max}(A_i^{-1})$. Since the eigenvalues of A_i are strictly negative by Assumption 1, we know that $\lambda_{\min}(A_i^{-1}) = 1/\lambda_{\max}(A_i)$ and $\lambda_{\max}(A_i^{-1}) = 1/\lambda_{\min}(A_i)$. We therefore have that $(\lambda_{\max}(A_i))^{-1} \leq (A_i^{-1})_{tt} \leq (\lambda_{\min}(A_i))^{-1}$, and so all diagonal elements of A_i^{-1} are strictly negative and bounded away from zero by some number $0 > \overline{A} \geq [A_i^{-1}]_{tt}$, and so the diagonal elements of $-A_i^{-1}$ are positive.

Now, due to the block-diagonal structure of A, it is true that $-A^{-1}$ is a matrix with all positive diagonal elements and nonnegative off-diagonal elements. Notice also that all elements of $\frac{\partial F}{\partial \mu_i}$ are elements of the Leontief inverse matrix $\mathcal{L} = I + \alpha + \alpha^2 + \ldots$ and are positive since $\alpha_i \in int(\mathcal{A}_i)$ for all i.

Now, in the case of no uncertainty, $\Sigma = 0$, B = 0 and the right-hand side of 56 must be strictly positive and so is the vector of $\frac{\partial \alpha_{kl}}{\partial \mu_i}$. In this case, both parts of the Lemma hold. If there is uncertainty $(\Sigma = 0)$, the result still holds if all the elements of Σ are sufficiently close to zero. Indeed, $-(A+B)^{-1}$ is continuous in Σ and, thus there exists $\overline{\Sigma} > 0$ such that if $|\Sigma_{ij}| < \overline{\Sigma}$ for all $i, j \in \mathcal{N}^2$ then elements of $-(A+B)^{-1} \frac{\partial F}{\partial \mu_i}$ have the same signs as the corresponding elements of $-A^{-1}\frac{\partial F}{\partial \mu_i}$.⁴²

Point (ii). The proof is analogous to that of point (i). We differentiate the first order conditions with respect to a diagonal element of Σ

$$\frac{\partial F}{\partial \Sigma_{ii}} = (1-\rho) \left[\mathbb{1}_{n \times 1} \otimes \left(\left(\beta' \mathcal{L} \mathbb{1}_i \right) \left(\mathcal{L} \iota_i \right) \right) \right] = (1-\rho) \left(\beta' \mathcal{L} \mathbb{1}_i \right) \frac{\partial F}{\partial \mu_i}.$$

Since $\rho > 1$ and $\omega_i = \beta' \mathcal{L} 1_i > 0$ for all $i \in \mathcal{N}$, the result follows from the same steps as in point (i)

Point (iii). The proof is analogous to that of point (i). We differentiate the first order conditions with respect to an off-diagonal element of Σ . To preserve the symmetry of Σ , we simultaneously change Σ_{ij} and Σ_{ji} to find

$$\frac{\partial F}{\partial \Sigma_{ij}} = (1-\rho) \left[\mathbb{1}_{n \times 1} \otimes \left(\left(\beta' \mathcal{L} \mathbb{1}_i \right) \left(\mathcal{L} \iota_j \right) + \left(\beta' \mathcal{L} \mathbb{1}_j \right) \left(\mathcal{L} \iota_i \right) \right) \right] = (1-\rho) \left[\left(\beta' \mathcal{L} \mathbb{1}_i \right) \frac{\partial F}{\partial \mu_j} + \left(\beta' \mathcal{L} \mathbb{1}_j \right) \frac{\partial F}{\partial \mu_i} \right].$$

Since $\rho > 1$ and $\omega_i = \beta' \mathcal{L} 1_i > 0$ for all $i \in \mathcal{N}$, the result follows from the same steps as in point (i).

I Alternative specifications for the distribution of ε

The parametrization of the shock process ε that we use in the model is common in the uncertainty literature (see for instance, Bloom et al., 2018), but has the implication that a change in the covariance matrix Σ has a direct impact on expected GDP E [Y], and so can affect decisions even when the household is risk neutral ($\rho = 0$). This happens because the mean of a log-normal variable like GDP is an increasing function of the variance of the underlying normal distribution. A common approach used by many papers is to undo this effect by removing half of the variance from the mean of the normal distribution. Unfortunately, such a change is problematic in our setup.

In this appendix, we first describe that in our setting there is no parametrization of ε such that 1) ε is normally distributed, 2) Σ does not affect decisions when $\rho = 0$, and 3) the distribution of ε does not depend on endogenous objects. We then consider a version of the model in which the distribution of ε is such that changes in Σ do not affect any decision when $\rho = 0$. This specification is however conceptually problematic as the distribution of ε depends on endogenous equilibrium objects. Finally, we consider a specification in which we adjust the mean of ε so that changes in Σ have no effect on $E[e^{\varepsilon}]$. In that case, the expectation of firm-level TFP shocks is unaffected by Σ ; however, the expectation of macroeconomic aggregates, e.g. E[Y], still depend on Σ .

⁴²Note that $(A+B)^{-1}$ exists for small Σ because the eigenvalues of A are strictly negative (and so det $(A) \neq 0$ and A is invertible) and that the determinant of A+B is a continuous function of Σ . Note also that as Σ moves away from 0 the optimal matrix α changes and so do A and \mathcal{L} . But these changes are continuous so the strict inequality $-(A+B)^{-1}\frac{\partial F}{\partial \mu_i} > 0$ is preserved for small enough Σ .

I.1 How to parametrize ε so that a risk-neutral household does not respond to uncertainty

In this subsection we describe how ε must be parametrized so that a risk-neutral household $(\rho = 0)$ does not change its behavior in response to changes in uncertainty Σ . For that purpose, it is useful to go back to central equations of the model that hold whenever ε is normally distributed. Hulten's theorem implies that for any given network α , log GDP is given by $y = \omega(\alpha)'(\varepsilon + a(\alpha))$, where $\omega(\alpha)$ is the vector of Domar weights. Together with CRRA preferences, this implies that the social planner's problem can be written as

$$\mathcal{W} \equiv \max_{\alpha \in \mathcal{A}} \operatorname{E} \left[y\left(\alpha\right) \right] - \frac{1}{2} \left(\rho - 1\right) \operatorname{V} \left[y\left(\alpha\right) \right]$$
$$= \max_{\alpha \in \mathcal{A}} \omega\left(\alpha\right)' \left(\operatorname{E} \left[\varepsilon\right] + a\left(\alpha\right) \right) - \frac{1}{2} \left(\rho - 1\right) \omega\left(\alpha\right)' \operatorname{V} \left[\varepsilon\right] \omega\left(\alpha\right)$$

In the benchmark model we have $\varepsilon \sim \mathcal{N}(\mu, \Sigma)$ and clearly Σ matters for the planner's decisions when $\rho = 0$. Suppose instead that $\varepsilon \sim \mathcal{N}\left(\mu - \frac{1}{2}B, \Sigma\right)$ where *B* is some quantity that can depend on Σ and that would make α^* invariant to Σ when $\rho = 0$. Plugging in the planner's problem, we find

$$\mathcal{W} = \max_{\alpha \in \mathcal{A}} \omega \left(\alpha\right)' \left(\mu - \frac{1}{2}B + a\left(\alpha\right)\right) - \frac{1}{2}\left(\rho - 1\right)\omega\left(\alpha\right)'\Sigma\omega\left(\alpha\right)$$
$$= \max_{\alpha \in \mathcal{A}} \omega \left(\alpha\right)'\left(\mu + a\left(\alpha\right)\right) - \frac{1}{2}\rho\omega\left(\alpha\right)'\Sigma\omega\left(\alpha\right) + \frac{1}{2}\omega\left(\alpha\right)'\left(\Sigma\omega\left(\alpha\right) - B\right)$$

For Σ to have no influence when $\rho = 0$ we therefore need the last term to be zero, which requires $B = \Sigma \omega(\alpha)$. In other words, this requires that the distribution of firm-level TFP shocks itself depends on endogenous equilibrium objects, namely the Domar weights $\omega(\alpha)$. This is problematic for at least two important reasons. First, we cannot think of a good reason why the distribution of productivity shocks that affect one industry would depend on the production technique chosen by another industry. Why that dependence would operate through Domar weights is also unclear. Second, the parametrization $\varepsilon \sim \mathcal{N} \left(\mu - \frac{1}{2}\Sigma\omega(\alpha), \Sigma\right)$ potentially introduces an externality in the economy: when deciding on its input shares α_i , firm *i* is modifying the TFP process of all other firms in the economy. This would create a gap between the efficient and the equilibrium allocations.

I.2 A model in which risk considerations are absent when $\rho = 0$

Here, we propose a distribution for ε such that 1) changes in Σ do not affect decisions when the household is risk neutral ($\rho = 0$), and 2) the equilibrium coincides with the solution to the planner's problem. Note that simply setting $B = \Sigma \omega(\alpha)$ does not accomplish this because of the externalities mentioned above. Specifically, we assume that

$$\varepsilon \sim \mathcal{N}\left(\mu - g\left(\alpha, \alpha^*, \Sigma\right), \Sigma\right),$$
(57)

where

$$g(\alpha, \alpha^*, \Sigma) = \frac{1}{2} \Sigma \mathcal{L}(\alpha^*)' \beta + \frac{1}{2} (\alpha - \alpha^*)' \mathcal{L}(\alpha^*) \Sigma \mathcal{L}(\alpha^*)' \beta.$$
(58)

The term α^* in this expression is the equilibrium network, so that in equilibrium we have $g(\alpha^*, \alpha^*, \Sigma) = \frac{1}{2} \Sigma \mathcal{L}(\alpha^*)^T \beta$. When making decisions, the representative firm in sector *i* chooses $\alpha_i = (\alpha_{i1}, \ldots, \alpha_{in})$ but takes α^* as given.

A few comments are in order. First, this specification implies that the distribution of shocks depends on endogenous equilibrium objects. This is clearly conceptually problematic, but it is, as we have discussed above, required for the result. We are not arguing that this specification is desirable or plausible. Our goal here is to explore the conditions under which decisions are unaffected by Σ under risk neutrality. Second, instead of assuming that g shifts the mean of ε , we could equivalently include it in the TFP shifter A. In that case, A would depend on equilibrium objects, unlike in the baseline specification. Third, the specification (57)–(58) differs from the one discussed above, $\varepsilon \sim \mathcal{N} \left(\mu - \frac{1}{2}\Sigma\omega(\alpha), \Sigma\right)$, which made the planner's problem unaffected by Σ under $\rho = 0$. Notice that both specification coincide *in equilibrium* but extra terms are required in (57)–(58) to ensure that the decentralized equilibrium allocation is efficient.

Once production techniques have been chosen and a specific realization of ε has been drawn, the distribution of ε has no impact on the economy. Therefore, Lemmas 1 and 2 also hold under this alternative specification. Furthermore, the existence proof of Proposition 1 can be straightforwardly extended to this new setting. The next proposition adapts Proposition 2 and Corollary 2 from the main text to the specification (57)–(58).

Proposition 8. Let the distribution of ε be as in (57)–(58). There exists an efficient equilibrium, and the efficient equilibrium network α^* solves

$$\mathcal{W} \equiv \max_{\alpha \in \mathcal{A}} \underbrace{\beta' \mathcal{L}(\alpha) \left(\mu + a(\alpha)\right)}_{\log E[Y(\alpha)]} - \frac{1}{2} \rho \underbrace{\beta' \mathcal{L}(\alpha) \Sigma \mathcal{L}(\alpha)' \beta}_{V[y(\alpha)]}.$$
(59)

Proof. First-order conditions of the planner. We first consider the problem of the social planner. In that case, the second term in (58) disappears and by (14) we therefore have

$$y(\alpha) \sim \mathcal{N}\left(\beta'\mathcal{L}(\alpha)\left(\mu - \frac{1}{2}\Sigma\mathcal{L}(\alpha)'\beta + a(\alpha)\right), \beta'\mathcal{L}(\alpha)\Sigma\mathcal{L}(\alpha)'\beta\right).$$
(60)

Using the usual expression for the expected value of a lognormal variable, we can write expected GDP as

$$E[Y(\alpha)] = \exp\left(\beta'\mathcal{L}(\alpha)\left(\mu + a(\alpha)\right)\right).$$

Notice that this expression does not depend on Σ . Combining the planner's problem given by (18) with (60), we can derive (59). Following the same steps as in the proof of Proposition 2, we obtain the following first-order conditions,

$$0 = 1'_{j} \mathcal{L} (\alpha) [\mu + a (\alpha)] + \frac{\partial}{\partial \alpha_{ij}} a_{i} (\alpha) - \rho 1'_{j} \mathcal{L} (\alpha) \Sigma \mathcal{L} (\alpha)' \beta$$

$$+ (\beta' \mathcal{L} (\alpha) 1_{i})^{-1} (\underline{\mu}_{ij} - \gamma_{i}),$$
(61)

where $\underline{\mu}_{ij}$ are the Lagrange multipliers on the constraints $\alpha_{ij} \ge 0$ and γ_i is the Lagrange multiplier on the constraint $\sum_j \alpha_{ij} \le \bar{\alpha}_i$.

First-order conditions of a firm in equilibrium. We can repeat similar steps for the equilibrium. Combining (44) with (13), (42) and (43), we find that the problem of the representative firm in sector i can be written as

$$\alpha_{i}^{*} = \arg\min_{\alpha_{i}\in\mathcal{A}_{i}}\frac{1}{2}\left(\alpha_{i}-\alpha_{i}^{*}\right)'\mathcal{L}\left(\alpha^{*}\right)\Sigma\mathcal{L}\left(\alpha^{*}\right)'\beta - a\left(\alpha_{i}\right) - \alpha_{i}'\mathcal{L}\left(\alpha^{*}\right)\left(\mu - \frac{1}{2}\Sigma\mathcal{L}\left(\alpha^{*}\right)'\beta + a\left(\alpha^{*}\right)\right) + \frac{1}{2}\left(\left(\alpha_{i}-1_{i}-\left(1-\rho\right)\beta\right)'\mathcal{L}\left(\alpha^{*}\right) + 1_{i}'\right)\Sigma\left(\left(\alpha_{i}-1_{i}-\left(1-\rho\right)\beta\right)'\mathcal{L}\left(\alpha^{*}\right) + 1_{i}'\right)'.$$

Differentiating with respect to α_{ij} we can write the first-order conditions as

$$0 = 1'_{j}\mathcal{L}(\alpha^{*})\Sigma\mathcal{L}(\alpha^{*})'\beta - \frac{\partial a(\alpha_{i})}{\partial \alpha_{ij}} - 1'_{j}\mathcal{L}(\alpha^{*})(\mu + a(\alpha^{*})) + (1'_{j}\mathcal{L}(\alpha^{*}))\Sigma((\alpha_{i} - 1_{i} - (1 - \rho)\beta)'\mathcal{L}(\alpha^{*}) + 1'_{i})' - \mu^{e}_{ij} + \gamma^{e}_{i},$$

where $\underline{\mu}_{ij}^e$ are the Lagrange multipliers on the constraints $\alpha_{ij} \ge 0$ and γ_i^e is the Lagrange multiplier on the constraint $\sum_j \alpha_{ij} \le \bar{\alpha}_i$. In equilibrium $\alpha = \alpha^*$ and so the above expression simplifies to

$$\frac{\partial a\left(\alpha_{i}^{*}\right)}{\partial \alpha_{ij}} + 1_{j}^{\prime} \mathcal{L}\left(\alpha^{*}\right)\left(\mu + a\left(\alpha^{*}\right)\right) - \rho \beta^{\prime} \mathcal{L}\left(\alpha^{*}\right) \Sigma \mathcal{L}\left(\alpha^{*}\right)^{\prime} 1_{j} + \underline{\mu}_{ij}^{e} - \gamma_{i}^{e} = 0,$$

where we used the fact that $(1_i - \alpha_i^*)' \mathcal{L}(\alpha^*) - 1_i' = 0$. These are the same first-order conditions (up to a normalization of the Lagrange multipliers) as the planner's (equation 61). The complementary slackness conditions are also the same in both problems. As a result, any equilibrium allocation satisfies the planner's first-order conditions and vice versa, and there exists an efficient equilibrium.

Notice that the objective function (59) is identical to (18) in the main text once we replace E[y]and V[y] by their expressions from (15), except that the term multiplying $V[y(\alpha)]$ is ρ instead of $\rho - 1$. It follows that, unlike in the benchmark model, the decisions of a risk-neutral household ($\rho = 0$) are unaffected by uncertainty Σ about the TFP shock ε , which was our goal with this exercise.

Changes to analytical results

We have already described how some of the formal analytical results of Section 2 were not affected by the change in specification. This is also true for almost all the results in Section 5. The only exception is Proposition 6 where $(\rho - 1)$ should replaced by ρ . It follows that this model behaves in a very similar way to the one in the main body of the paper.

I.3 Making the expectation of firm-level TFP shocks independent of Σ

One specification for ε which is used in the literature is $\varepsilon \sim \mathcal{N} \left(\mu - \frac{1}{2} \operatorname{diag}(\Sigma), \Sigma\right)$. This adjustment implies that the expected value of firm-level TFP $\operatorname{E} \left[\exp\left(\varepsilon_{i}\right)\right] = \exp\left(\mu_{i} - \frac{1}{2}\Sigma_{ii} + \frac{1}{2}\Sigma_{ii}\right) = \exp\left(\mu_{i}\right)$ does not depend on Σ . Changes to Σ are therefore closer to pure changes in uncertainty. But, as follows from the discussion above, Σ still matters for decisions even when the household is risk neutral.

Changes to analytical results

Almost all our analytical results are unaffected by this change in specification. This is the case for the results that pertain to the existence and the efficiency of an equilibrium, and the results that consider a change in μ . Of the results that consider changes in Σ , Proposition 3 and Lemma 4 describe how an increase in Σ leads to a reduction in Domar weights and in shares α in the baseline model. In the new specification, an increase in Σ also leads to a decline in $E[\varepsilon]$, which also pushes for smaller Domar weights and shares.⁴³ As both effects go in the same direction, the statements of Proposition 3 and Lemma 4 remain the same with the new specification. Proposition 5 also remains valid under the new specification. In this case, the proof goes through since any $\Sigma \neq 0$ has an adverse effect on E[y]. The only result that is meaningfully affected by the change is point 2 of Proposition 6 which describes how a change in Σ affects welfare W. Using the envelope theorem, it is straightforward to see that this statement becomes

$$\frac{d\mathcal{W}}{d\Sigma_{ij}} = \begin{cases} -\frac{1}{2}\omega_i - \frac{1}{2}\left(\rho - 1\right)\omega_i^2 & i = j, \\ -\left(\rho - 1\right)\omega_i\omega_j & i \neq j, \end{cases}$$

which implies that uncertainty has a larger negative impact on welfare in this alternative specification.

⁴³Here, it is important that only the diagonal of Σ is included in the correction to the mean of ε . These terms are positive and so any increase in elements of Σ has a detrimental impact on $E[\varepsilon]$.

Changes to quantitative results

We also investigate the implications of this change in specification for our quantitative model. To do so, we consider an alternative economy, denoted with tildes, in which

$$\tilde{\varepsilon} \sim \mathcal{N}\left(\tilde{\mu} - \frac{1}{2} \operatorname{diag}\left(\tilde{\Sigma}\right), \tilde{\Sigma}\right).$$

If we were to calibrate this economy, we would find that

$$\tilde{\mu}_t - \frac{1}{2} \text{diag}\left(\tilde{\Sigma}_t\right) = \mu_t,$$
$$\tilde{\Sigma}_t = \Sigma_t,$$

where μ_t and Σ_t are the mean and covariance of ε_t in our baseline calibration. That is because the calibration matches the vector of sectoral TFP perfectly. If we remove $\frac{1}{2}\Sigma_{ii}$ from the expectation of ε_i , the estimation would increase the expectation $\tilde{\mu}_t$ to compensate and match the data.

We reproduce the exercise in the left column of Figure 7 in this setup. This amounts to comparing the economy described above with an alternative in which the production network is chosen as if $\tilde{\Sigma}_t = 0$. The results are presented in Figure 13. Overall, we find that uncertainty has a larger impact on the economy in this setting than in the baseline model of Section 6. As in Figure 7, the variance of log GDP is smaller in the baseline model. Expected log GDP is however quite different, with E[y] larger in the baseline than in the alternative model. This is because the network in the alternative economy is not well adapted to the TFP process. In the alternative model, firms choose production techniques as if $E[\tilde{\varepsilon}] = \tilde{\mu}$, when in reality $E[\tilde{\varepsilon}] = \tilde{\mu} - \frac{1}{2} \text{diag}\tilde{\Sigma}$. This implies that firms ignore the fact that risky suppliers (i.e. those with high Σ_{ii}) are also less productive on average, which results in a decline in expected log GDP relative to the baseline model (first panel of Figure 13). Given that the alternative model performs worse than the baseline both in terms of E[y] and V[y], the welfare losses in the alternative model are substantial (third panel).



Figure 13: The role of uncertainty when $\varepsilon \sim \mathcal{N}\left(\mu - \frac{1}{2}\operatorname{diag}(\Sigma), \Sigma\right)$

Notes: The differences between the series implied by the baseline model with $\varepsilon \sim \mathcal{N}\left(\mu - \frac{1}{2}\operatorname{diag}(\Sigma), \Sigma\right)$ (without tildes) and the "as if $\tilde{\Sigma}_t = 0$ " alternative (with tildes). Both economies are hit by the same shocks that are filtered out from the TFP data under our baseline model. All differences are expressed in percentage terms.

J Sectoral correlations

In this appendix, we explore the model's implications for sectoral correlations. By combining (13) and (41), we can write the covariance matrix of log sectoral output q as

$$V[q] = \mathcal{L}(\alpha) \Sigma \left(\mathcal{L}(\alpha)\right)', \qquad (62)$$

or, if we focus on a specific element,

$$[V[q]]_{ij} = \sum_{k=1}^{n} \underbrace{[\mathcal{L}(\alpha)]_{ik} [\mathcal{L}(\alpha)]_{jk} \Sigma_{kk}}_{\text{loading on same supplier}} + \sum_{k \neq l} \underbrace{[\mathcal{L}(\alpha)]_{ik} [\mathcal{L}(\alpha)]_{jl} \Sigma_{kl}}_{\text{loading on correlated suppliers}} .$$
(63)

This expression is intuitive. It highlights that the correlation between any two sectors i and j depends on how much they rely on common suppliers, and how much they rely on correlated suppliers. The first summation in (63) captures the first channel. If firms i and j both rely on firm k as a supplier, so that $[\mathcal{L}(\alpha)]_{ik} > 0$ and $[\mathcal{L}(\alpha)]_{jk} > 0$, then a shock to k jointly affects the production of i and j, which contributes to a positive correlation between the two sectors. Notice that this first term is always positive. The second summation in (63) captures the dependence of i and j on correlated suppliers. Suppose that sector i relies on k as an input and sector j relies on l as an input. Any covariance Σ_{kl} between k and l will then translate into a covariance between i and j. Unlike the first channel, the contribution of this second term to the covariance can be negative if $\Sigma_{kl} < 0$.

Correlations under weak complementarity

Changes in beliefs (μ, Σ) , through their impact on the network α , modify the correlation patterns highlighted in (63). Consider first a change in μ_i . Differentiating (62) yields

$$\frac{d \operatorname{V}[q]}{d\mu_{i}} = \mathcal{L}(\alpha) \frac{d\alpha}{d\mu_{i}} \operatorname{V}[q] + \operatorname{V}[q] \left(\frac{d\alpha}{d\mu_{i}}\right)' (\mathcal{L}(\alpha))'.$$
(64)

It follows that the impact on μ_i on the covariance matrix of q depends crucially on how the production network α responds to the change. As we have discussed in Section 5.2, this response depends in general on the specific structure of the economy. There is however one special case in which we know how α reacts. Under our weak complementarity assumption (Assumption 2), an increase in μ_i leads to an increase in α by Lemma 4. This in turn results in an increase in all the sectoral output correlations.

Lemma 5. Suppose that the assumptions of Lemma 4 hold and that $\Sigma_{ij} \ge 0$ for all *i* and *j*. Then an increase in μ_k leads to an increase in $[V[q]]_{ij}$ for all *i* and *j*.

Proof. By Lemma 4, such a change in μ_k leads to an increase in α . Since $\Sigma_{ij} \ge 0$ for all i and j, the right-hand side of 64 is positive.

Similarly, for a change in Σ_{ij} , differentiating (62) yields

$$\frac{d \operatorname{V}[q]}{d\Sigma_{ij}} = \mathcal{L}(\alpha) \frac{d\alpha}{d\Sigma_{ij}} \operatorname{V}[q] + \operatorname{V}[q] \left(\frac{d\alpha}{d\Sigma_{ij}}\right)' \left(\mathcal{L}(\alpha)\right)' + \mathcal{L}(\alpha) \operatorname{1}_{i} \operatorname{1}'_{j} \left(\mathcal{L}(\alpha)\right)'.$$
(65)

This expression is similar to (64) except that we must now take into account the direct impact of the change in Σ on the covariance matrix (last term in the equation). In particular, under the conditions of Proposition 5, an increase in Σ_{ij} implies a reduction in input shares α which pushes V[q] down (the first two terms in the right-hand side of (65) are negative). The direct effect, however, works in the opposite direction.

Approximate impact of beliefs on sectoral correlations

When Assumption 2 does not hold, we can rely on an approximation to characterize sectoral correlations. To simplify the notation, denote $V[q] \equiv \Sigma^q$ and note that by (62) we can write

$$\Sigma_{kl}^q = \mathbf{1}_k' \mathcal{L} \Sigma \mathcal{L}' \mathbf{1}_l$$

Using the same approximation as in Appendix B, we can write

$$\begin{split} \Sigma_{kl}^{q} &\approx \Sigma_{kl}^{\circ q} + \bar{\kappa}^{-1} \left[\mathbf{1}_{k}^{\prime} \left(\sum_{i,j} \frac{\partial \mathcal{L}}{\partial \alpha_{ij}} \frac{\partial \alpha_{ij}}{\partial \bar{\kappa}^{-1}} \right) \Sigma \mathcal{L}^{\prime} \mathbf{1}_{l} + \mathbf{1}_{k}^{\prime} \mathcal{L} \Sigma \left(\sum_{i,j} \frac{\partial \mathcal{L}}{\partial \alpha_{ij}} \frac{\partial \alpha_{ij}}{\partial \bar{\kappa}^{-1}} \right)^{\prime} \mathbf{1}_{l} \right]_{\bar{\kappa}^{-1} = 0} \\ &= \Sigma_{kl}^{\circ q} + \bar{\kappa}^{-1} \left[\mathbf{1}_{k}^{\prime} \sum_{i,j} \left[\frac{\partial \alpha_{ij}}{\partial \bar{\kappa}^{-1}} \right]_{\bar{\kappa}^{-1} = 0} \mathcal{L}^{\circ} \mathbf{1}_{i} \mathbf{1}_{j}^{\prime} \mathcal{L}^{\circ} \Sigma \left(\mathcal{L}^{\circ} \right)^{\prime} \mathbf{1}_{l} + \mathbf{1}_{k}^{\prime} \mathcal{L}^{\circ} \Sigma \left(\sum_{i,j} \mathcal{L}^{\circ} \mathbf{1}_{i} \mathbf{1}_{j}^{\prime} \mathcal{L}^{\circ} \frac{\partial \alpha_{ij}}{\partial \bar{\kappa}^{-1}} \right)^{\prime} \mathbf{1}_{l} \right] \\ &= \Sigma_{kl}^{\circ q} + \bar{\kappa}^{-1} \sum_{i,j} \left[\frac{\partial \alpha_{ij}}{\partial \bar{\kappa}^{-1}} \right]_{\bar{\kappa}^{-1} = 0} \left[\mathcal{L}_{ki}^{\circ} \Sigma_{jl}^{\circ p} + \mathcal{L}_{li}^{\circ} \Sigma_{kj}^{\circ p} \right] \\ &= \Sigma_{kl}^{\circ q} + \bar{\kappa}^{-1} \sum_{i,j} \mathbf{1}_{j}^{\prime} \left(\mathbf{H}_{i}^{\circ} \right)^{-1} \left[\mathcal{L}^{\circ} \mu - \left(\rho - 1 \right) \mathcal{L}^{\circ} \Sigma \left(\mathcal{L}^{\circ} \right)^{\prime} \beta \right] \left[\mathcal{L}_{ki}^{\circ} \Sigma_{jl}^{\circ q} + \mathcal{L}_{li}^{\circ} \Sigma_{kj}^{\circ q} \right], \end{split}$$

where we used (50) and where $\Sigma^{\circ q} = \mathcal{L}(\alpha^{\circ}) \Sigma (\mathcal{L}(\alpha^{\circ}))'$ is V[q] evaluated at the ideal shares. It follows that

$$\frac{d\Sigma_{kl}^{q}}{d\mu_{m}} \approx \bar{\kappa}^{-1} \sum_{i,j} 1_{j}^{\prime} (H_{i}^{\circ})^{-1} \mathcal{L}^{\circ} 1_{m} \left[\mathcal{L}_{ki}^{\circ} \Sigma_{jl}^{\circ q} + \mathcal{L}_{li}^{\circ} \Sigma_{kj}^{\circ q} \right]
= \bar{\kappa}^{-1} \sum_{i,j} \left[\sum_{s} (H_{i}^{\circ})_{js}^{-1} \mathcal{L}_{sm}^{\circ} \right] \left[\mathcal{L}_{ki}^{\circ} \Sigma_{jl}^{\circ q} + \mathcal{L}_{li}^{\circ} \Sigma_{kj}^{\circ q} \right].$$
(66)

The overall effect of μ_m on Σ_{kl}^q can be decomposed in two terms. The first term reflects how the input share α_{ij} respond to an increase in μ_m . The second term reflects how this adjustment in α_{ij} affect the covariance matrix Σ_{kl}^q . Specifically, the term $\sum_s (H_i^\circ)_{js}^{-1} \mathcal{L}_{sm}^\circ$ captures the response of α_{ij} following an increase in μ_m . In a nutshell, we look at all sectors s that depend on m directly on indirectly $(\mathcal{L}_{sm}^\circ > 0)$. Firms in these sectors become more attractive after the increase in μ_m . If j is complement with s in the production of good i, that is $(H_i^\circ)_{js}^{-1} > 0$, then that pushes for a higher α_{ij} . The second term, $\mathcal{L}_{ki}^\circ \Sigma_{jl}^{\circ q} + \Sigma_{kj}^\circ \mathcal{L}_{li}^\circ$, captures the exposure of k and l to that change. Suppose that α_{ij} increases after the increase in μ_m . If k and l rely on i $(\mathcal{L}_{ki}^\circ > 0$ and $\mathcal{L}_{li}^\circ > 0)$ then their covariance Σ_{kl}^q will increase if both also covary with j $(\Sigma_{jl}^{\circ q} > 0$ and $\Sigma_{kj}^\circ > 0)$. We repeat that process for all sectors i and j to get the full impact of the increase in μ_m on Σ_{kl}^q .

We can also look at the derivatives with respect to Σ ,

$$\frac{d\Sigma_{kl}^{q}}{d\Sigma_{mm}} = \mathcal{L}_{km}^{\circ}\mathcal{L}_{lm}^{\circ} - \bar{\kappa}^{-1} \left(\rho - 1\right) \sum_{i,j} \left[\sum_{s} \left(H_{i}^{\circ} \right)_{js}^{-1} \mathcal{L}_{sm}^{\circ} \right] \omega_{m}^{\circ} \left[\mathcal{L}_{ki}^{\circ} \Sigma_{jl}^{\circ q} + \mathcal{L}_{li}^{\circ} \Sigma_{kj}^{\circ q} \right]
+ \bar{\kappa}^{-1} \sum_{i,j} 1_{j}^{\prime} \left(H_{i}^{\circ} \right)^{-1} \left[\mathcal{L}^{\circ} \mu - \left(\rho - 1\right) \mathcal{L}^{\circ} \Sigma \left(\mathcal{L}^{\circ} \right)^{\prime} \beta \right] \left[\mathcal{L}_{ki}^{\circ} \mathcal{L}_{jm}^{\circ} \mathcal{L}_{lm}^{\circ} + \mathcal{L}_{km}^{\circ} \mathcal{L}_{jm}^{\circ} \mathcal{L}_{li}^{\circ} \right],$$

where we have used the fact that $\frac{d\Sigma_{kl}^{\circ q}}{d\Sigma_{mm}} = \mathcal{L}_{km}^{\circ} \mathcal{L}_{lm}^{\circ}$. The first term, $\mathcal{L}_{km}^{\circ} \mathcal{L}_{lm}^{\circ} \geq 0$, captures the direct exposure of k and l to m. That is, if both k and l rely on m ($\mathcal{L}_{km}^{\circ} > 0$ and $\mathcal{L}_{lm}^{\circ} > 0$), and m becomes more volatile, the covariance between k and l increases. Note that this term always dominates if changes in shares are small, that is, if $\bar{\kappa}$ is large enough. The last two terms appear because equilibrium input shares deviate from ideal values. To understand the second term, notice that $\left[\sum_{s} (H_{i}^{\circ})_{js}^{-1} \mathcal{L}_{sm}^{\circ}\right] \omega_{m}^{\circ}$ captures the response of α_{ij} to an increase in Σ_{mm} . If k and l are connected to i ($\mathcal{L}_{ki}^{\circ} > 0$ and $\mathcal{L}_{li}^{\circ} > 0$) then their covariance Σ_{kl}^{q} will increase if both also covary with j ($\Sigma_{jl}^{\circ q} > 0$ and $\Sigma_{kj}^{\circ q} > 0$). Finally, the last term captures the direct change in the covariance between j and l and between k and j. These changes in covariances are scaled by the deviation of the input share α_{ij} from its ideal value, i.e. by $\Delta \alpha_{ij}$ given by (50).

K Wedges and inefficient allocation

In this appendix, we consider a version of the competitive economy of Section 2 with wedges. To do so, we modify our setup as in Acemoglu and Azar (2020). Specifically, we assume that firms in industry *i* sell their goods at a markup $\tau_i \geq 0$ over their unit cost. A fraction $\zeta_i \in [0, 1]$ of the revenue from the distortions is rebated to the representative household. The remaining $1 - \zeta_i$ share is pure waste. We assume that τ_i and ζ_i are exogenous and do not depend on ε . Below, we first describe the decentralized equilibrium with wedges and then show that there exists a distorted "planner" whose decisions coincide with the distorted equilibrium. Finally, we characterize how distortions affect equilibrium outcomes.

K.1 A distorted equilibrium

Several parts of the model are not affected by the wedges. In particular, the objective function of the household remains unchanged. Its budget constraint must however be adjusted to take into account the profits generated by the wedges. It becomes

$$\sum_{i=1}^{n} P_i C_i \le 1 + T\left(\alpha\right),$$

where $T(\alpha) = \sum_{i=1}^{n} \zeta_i \frac{\tau_i}{1+\tau_i} P_i Q_i$ is the rebate due to distortions. As we show below, T depends on α but not on ε , which justifies the notation $T(\alpha)$. In the absence of distortions ($\tau_i = 0$ for all i) or if distortions are pure waste ($\zeta_i = 0$ for all i), T = 0. The additional term in the budget constraint implies a different stochastic discount factor Λ . From the first-order conditions of the household we can write

$$\lambda = (\rho - 1) \beta' p - \rho \log (1 + T), \qquad (67)$$

where $\lambda = \log \Lambda$ is the log of the stochastic discount factor.

On the side of the firm, the cost minimization problem (7) is unaffected by the wedges, and so the unit cost K_i conditional on a technique α_i and a price vector P can still be written as (8). Similarly, the technique choice problem of the firm, conditional on prices, is unaffected and is still defined by (9). We will see however that in equilibrium the wedges affect the technique choices of the firms through their impact on prices.

Equilibrium conditions

We now turn to the market clearing conditions and the pricing equations. Those are affected by the wedges. In particular, the pricing equation (11) becomes

$$P_i = (1 + \tau_i) K_i(\alpha, P), \qquad (68)$$

such that prices are set at a markup over unit cost. We can combine this equation with (8) to write an expression for log prices as a function of the network α ,

$$p = -\mathcal{L}(\alpha) \left(\varepsilon + a(\alpha) - \log(1+\tau)\right), \tag{69}$$

where $\log(1 + \tau)$ is a column vector with typical element $\log(1 + \tau_i)$. As we can see, wedges τ affect prices as productivity shifters.

The market clearing condition (12) for good *i* must also be adjusted for the potential loss in resources. It becomes

$$Q_i\left(1 - (1 - \zeta_i)\frac{\tau_i}{1 + \tau_i}\right) = C_i + \sum_j X_{ji}.$$
(70)

We can use these equations to find an expression for the rebate to the household T. Combining

the first-order conditions of the firms with (68) and (70), we get

$$T(\alpha) = \begin{pmatrix} \zeta_{1}\tau_{1} \\ \zeta_{2}\tau_{2} \\ \cdots \\ \zeta_{n}\tau_{n} \end{pmatrix}' \begin{bmatrix} \begin{pmatrix} 1+\zeta_{1}\tau_{1} & 0 & \cdots & 0 \\ 0 & 1+\zeta_{2}\tau_{2} & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & \cdots & 1+\zeta_{n}\tau_{n} \end{pmatrix} - \beta \begin{pmatrix} \zeta_{1}\tau_{1} \\ \zeta_{2}\tau_{2} \\ \cdots \\ \zeta_{n}\tau_{n} \end{pmatrix}' - \alpha' \end{bmatrix}^{-1} \beta, \quad (71)$$

and so we can fully characterize the stochastic discount factor (67) for a given network α . Note also that T = 0 whenever $\zeta = 0$ or $\tau = 0$.

Using the expression for T together with the price vector p, we can write log GDP as

$$y = \beta' \mathcal{L}(\alpha) \left(\varepsilon + a(\alpha) - \log(1+\tau)\right) + \log\left(1 + T(\alpha, \zeta, \tau)\right).$$
(72)

We see that the wedges have two different impacts on GDP. First, the distortions effectively lead to a decline in productivity through $\log (1 + \tau)$. At the same time, the part of these distortions that is rebated to the household has a positive impact on GDP through T.

Finally, in the following lemma we characterize the comparative statics of the rebate amount.

Lemma 6. Holding everything else equal, $T(\alpha)$ increases in α_{ij} for all i, j.

Proof. Denote $\chi_i = \zeta_i \tau_i$. Rewrite (71) as

$$T(\alpha, \chi) = \chi' \left[I + \operatorname{diag}(\chi) - \beta \chi' - \alpha' \right]^{-1} \beta,$$

where diag (χ) is a diagonal matrix with the vector χ on its main diagonal. Note that $I + \text{diag}(\chi) - \chi'\beta - \alpha$ is a diagonally dominant matrix with a positive main diagonal. Therefore, it is an M-matrix whose inverse has all nonnegative inputs. Therefore, $[I + \text{diag}(\chi) - \beta\chi' - \alpha']^{-1}$ is a nonnegative matrix as well. Differentiating $T(\alpha, \chi)$ with respect to α_{ij} yields

$$\frac{\partial T}{\partial \alpha_{ij}} = \chi' \left[I + \operatorname{diag}\left(\chi\right) - \beta \chi' - \alpha' \right]^{-1} \left(\mathbf{1}_j \mathbf{1}_i^T \right) \left[I + \operatorname{diag}\left(\chi\right) - \beta \chi' - \alpha' \right]^{-1} \beta.$$

Since all vectors and matrices in the right-hand side of the above expression are nonnegative, $\frac{\partial T}{\partial \alpha_{ij}} \geq 0.$

Equilibrium network

We now have all the ingredients to write down the technique choice problem of the firms. Following steps analogous to those of the proof of Proposition 8, we find that the equilibrium network α^* must satisfy

$$-\frac{\partial a\left(\alpha_{i}^{*}\right)}{\partial \alpha_{ij}}-1_{j}^{\prime}\mathcal{L}\left(\alpha^{*}\right)\left(\mu+a\left(\alpha^{*}\right)-\log\left(1+\tau\right)\right)-\left(1-\rho\right)\beta^{\prime}\mathcal{L}\left(\alpha^{*}\right)\Sigma\mathcal{L}\left(\alpha^{*}\right)^{\prime}1_{j}+\underline{\mu}_{ij}^{e}-\gamma_{i}^{e}=0, \quad (73)$$

where as before $\underline{\mu}_{ij}^e$ are the Lagrange multipliers on the constraints $\alpha_{ij} \geq 0$ and γ_i^e is the Lagrange multiplier on the constraint $\sum_j \alpha_{ij} \leq \bar{\alpha}_i$. Note that the fraction rebated ζ does not show up in (73) because, first, firms do not internalize how their decisions affect the household's stochastic discount factor and, second, the rebate amount T is non-stochastic and thus does not affect the covariance of firms' unit costs with the stochastic discount factor (67). However, firms' individual decisions are distorted by τ since prices are above unit costs. For instance, if the wedge τ_i in industry i is positive, firms perceive industry i as effectively less productive on average. This in turn affects their production technique choice.

K.2 A distorted planner's problem

In the main text, we exploit the fact that the equilibrium allocation can be written as the outcome of the planner's optimization problem to derive several results. Here, because of the distortions, it is no longer true that the equilibrium coincides with the planner's allocation. We can however derive the problem of a *distorted* fictitious planner whose preferred allocation coincides with the distorted equilibrium. We can then take advantage of that optimization problem to characterize the distorted equilibrium.

We define this distorted planner's problem as

$$\mathcal{W}^{d} \equiv \max_{\alpha \in \mathcal{A}} \mathbb{E}\left[y^{d}\left(\alpha\right)\right] - \frac{1}{2}\left(\rho - 1\right) \mathbb{V}\left[y^{d}\left(\alpha\right)\right],\tag{74}$$

where

$$y^{d}(\alpha) = \beta' \mathcal{L}(\alpha) \left(\varepsilon + a(\alpha) - \log(1+\tau)\right)$$

is what log GDP *would be* if nothing was rebated to the household. Taking the first-order conditions of this distorted planner's problem yields (73). Therefore, we can analyze this problem instead of the decentralized equilibrium.

K.3 Characterizing the distorted equilibrium

We can use the distorted planner's problem to characterize the distorted equilibrium. The next proposition describes how the Domar weights are affected by beliefs (μ, Σ) and the wedges τ .

Proposition 9. The Domar weight ω_i of sector *i* is increasing in μ_i , decreasing in Σ_{ii} and decreasing in τ_i .

Proof. The proof is similar to that of Proposition 3. We can write

$$\mathcal{W}^{d} = \max_{\alpha \in \mathcal{A}} \beta' \mathcal{L}(\alpha) \left(\mu + a(\alpha) - \log(1 + \tau)\right) - \frac{1}{2} \left(\rho - 1\right) \beta' \mathcal{L}(\alpha) \Sigma \left(\mathcal{L}(\alpha)\right)' \beta,$$

so that, by the envelope theorem,

$$\frac{d\mathcal{W}^d}{d\mu} = \omega, \ \frac{d\mathcal{W}^d}{d\Sigma} = -\left(\rho - 1\right)\omega\omega', \text{ and } \frac{d\mathcal{W}^d}{d\log\left(1 + \tau\right)} = -\omega$$

With these derivatives in hand, we can follow analogous steps to those in the proof of Proposition 3 to find that an increase in μ_i or a decline in Σ_{ii} leads to a higher ω_i . Finally, an increase in τ_i is equivalent to a decline in μ_i , and so the last part of the proposition follows.

Next, we investigate how changes in beliefs and wedges affect welfare, which can be written as

$$\mathcal{W}\left(\alpha^{*}\right) = \mathcal{W}^{d}\left(\alpha^{*}\right) + \log\left(1 + T\left(\alpha^{*}\right)\right),$$

where α^* solves (74).

Proposition 10. In the distorted equilibrium, the following holds.

1. The impact of an increase in μ_i on welfare is given by

$$\frac{d\mathcal{W}}{d\mu_{i}} = \omega_{i} + \frac{1}{1 + T\left(\alpha^{*}\right)} \frac{\partial T}{\partial \alpha} \frac{d\alpha^{*}}{d\mu_{i}}.$$

2. The impact of an increase in Σ_{ii} on welfare is given by

$$\frac{d\mathcal{W}}{d\Sigma_{ii}} = -\frac{1}{2} \left(\rho - 1\right) \omega_i^2 + \frac{1}{1 + T\left(\alpha^*\right)} \frac{\partial T}{\partial \alpha} \frac{d\alpha^*}{d\Sigma_{ii}}.$$

3. If $\zeta = 0$, then the results of Proposition 6 hold in the distorted equilibrium.

Proof. Taking derivatives of \mathcal{W} yields

$$\frac{d\mathcal{W}}{d\mu_{i}} = \frac{d\mathcal{W}^{d}\left(\alpha^{*}\right)}{d\mu_{i}} + \frac{1}{1+T\left(\alpha^{*}\right)}\frac{\partial T}{\partial\alpha}\frac{d\alpha^{*}}{d\mu_{i}} = \omega_{i} + \frac{1}{1+T\left(\alpha^{*}\right)}\frac{\partial T}{\partial\alpha}\frac{d\alpha^{*}}{d\mu_{i}}$$

Similar steps yield the expression for the derivative with respect to Σ_{ii} . Note that T = 0 when $\zeta = 0$ and so $\frac{\partial T}{\partial \alpha} = 0$ in this case. We then get the expressions of Proposition 6.

We see from these equations that changes in μ and Σ have two effects on welfare. There is a term that is the same as in the efficient allocation, as described by Proposition 6. But there is also a second term that reflects how changes in the production network lead to a higher or smaller rebate to the household. For example, under Assumption 2, an increase in μ_i leads to an increase in all shares, i.e. all elements of $\frac{d\alpha^*}{d\mu_i}$ are positive. This in turn increases the rebated amount T by Lemma (6) and, as a result, $\frac{dW}{d\mu_i} > \omega_i$.

L Special case with $A_i(\alpha_i) \equiv 1$

The problem of the social planner must in general be solved numerically. There is however one special case that can be solved analytically and that sheds light on the forces at work in the economy. If the TFP shifter functions $A_i(\alpha_i)$ are identically equal to 1 for all i,⁴⁴ we can write the planner's problem as

$$\mathcal{W} \equiv \max_{\alpha \in \mathcal{A}} \omega(\alpha)' \mu - \frac{1}{2} \left(\rho - 1\right) \omega(\alpha)' \Sigma \omega(\alpha), \qquad (75)$$

where the Domar weights are defined $\omega(\alpha)' = \beta'(I-\alpha)^{-1}$. Notice that the objective function only depends on Domar weights ω . We can then recast the constraint set $\alpha \in \mathcal{A}$ to be in terms of ω . This leads to the restrictions

$$\omega \ge \beta, \tag{76}$$
$$\omega' \left(1 - \bar{\alpha}\right) \le 1.$$

By combining these equations with (75), we have an optimization problem that is cast only in terms of ω . Since its objective function is strictly concave and the constraints are linear the first-order conditions are necessary and sufficient to characterize its solution. The first-order condition for an interior solution yields

$$\omega = \frac{1}{\rho - 1} \Sigma^{-1} \mu, \tag{77}$$

which makes explicit the importance of beliefs μ and Σ for the production network. To describe the forces at work in this equation, we will focus on the case $\mu > 0.45$ If we focus on a particular sector *i*, we can write

$$\omega_i = \frac{1}{\rho - 1} \sum_{j=1}^n \left[\Sigma^{-1} \right]_{ij} \mu_j,$$

where $[\Sigma^{-1}]_{ij}$ is the the typical element of Σ^{-1} . Clearly, ω_i depends on μ_i , and the diagonal term $[\Sigma^{-1}]_{ii}$ determines the strength of that dependence. Remember that we can interpret the inverse $1/[\Sigma^{-1}]_{ii}$ of the *i*th diagonal element of Σ^{-1} as the variance of the component of ε_i that is uncorrelated with the other elements of ε . If this variance increases, industry *i* becomes more risky and its Domar weight ω_i shrinks ceteris paribus.

The Domar weight ω_i also depends on the other sectors in the economy. Here the dependence

⁴⁴Assumption 1 is not satisfied in that case but the existence and efficiency results go through as long as Σ is positive definite.

⁴⁵If some elements of μ are negative, the constraint (76) tends to bind and we need to include Lagrange multipliers in (77).

operates through the off-diagonal terms in Σ^{-1} . These terms are informative about the partial correlations ρ_{ij} , such that

$$\rho_{ij} = -\frac{\left[\Sigma^{-1}\right]_{ij}}{\sqrt{\left[\Sigma^{-1}\right]_{ii}}\sqrt{\left[\Sigma^{-1}\right]_{jj}}}.$$

It follows that an increase in partial correlation ρ_{ij} between sector *i* and *j* leads to a decline in ω_i . Intuitively, the planner wants to diversify the risk embedded in ε and therefore prefers to avoid correlated sectors.

We can also use (77) together with (14) and (18) to write log GDP and welfare as

$$y = \frac{1}{\rho - 1} \mu' \Sigma^{-1} \varepsilon$$
 and $\mathcal{W} = \frac{1}{2} \frac{1}{\rho - 1} \mu' \Sigma^{-1} \mu$.

Furthermore, when the shocks are uncorrelated, we can write

$$E[y] = \frac{1}{\rho - 1} \sum_{i=1}^{n} \frac{\mu_i^2}{\Sigma_{ii}}.$$

This equation shows that an increase in uncertainty leads to a decline in expected log GDP, which is a more general version of Proposition 5 for this particular setup.