

Endogenous Production Networks under Supply Chain Uncertainty*

Alexandr Kopytov

University of Rochester

Bineet Mishra

Cornell University

Kristoffer Nimark

Cornell University

Mathieu Taschereau-Dumouchel

Cornell University

November 11, 2023

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 tends to reduce 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 provide a basic calibration of the model using U.S. data to evaluate the importance of these mechanisms.

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

*We are grateful to the editor, three anonymous referees, Ivan Alfaro, Matteo Bizzarri, Hamed 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 get 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. Such shocks can propagate to the rest of the economy through input-output linkages, resulting in aggregate fluctuations. However, firms may mitigate this 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 macroeconomic aggregates.

Supply chain disruptions are one of the key challenges that business executives face and are responsible for substantial investments in risk-mitigation strategies (Ho et al., 2015). The COVID-19 pandemic provides a stark illustration of how uncertainty can disrupt supply relationships. Following the onset of the pandemic, many companies realized that their supply chains were more vulnerable than previously thought. A recent survey revealed that seventy percent of firm managers agreed that the pandemic pushed companies to favor higher supply-chain resiliency instead of simply purchasing from the lowest-cost supplier. Many also reported plans to diversify their supply chains across suppliers and geographies.¹

To study how supply chain uncertainty affects firms’ sourcing decisions and how these decisions affect the macroeconomy, we construct a model of endogenous network formation that builds on Acemoglu and Azar (2020). 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 using different *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 technique, a firm can marginally adjust the importance of a supplier or drop that supplier altogether. Consequently, these decisions, when aggregated, lead to changes in the production network along both the intensive and extensive margins.

After selecting production techniques, 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, and in contrast with Acemoglu and Azar (2020), firms’ beliefs about the distribution of sectoral productivities can influence their choice of production technique and, thus, the structure of the network. Firms compare profits across different states of the world using the repre-

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.

sentative household's stochastic discount factor and, as such, they inherit the household's attitude toward risk. Consequently, while a firm would generally prefer to purchase from a more productive supplier, it might decide otherwise if this supplier is also riskier. Such a supplier 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 consider this possibility and balance concerns between average productivity and stability when choosing their production techniques.

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

We prove that the unique equilibrium in this environment is efficient, so that the equilibrium production network can be understood as resulting from a social planner maximizing the utility of the representative household. That network 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. This trade-off implies a novel mechanism through which uncertainty can lower expected GDP. In the presence of uncertainty, firms prefer stable input prices and, as a result, move toward safer suppliers even though they might be less productive. Through this flight to safety process, less productive producers gain in importance, and aggregate productivity and GDP fall as a result. On the other hand, this supply chain reshuffling leads to a more resilient network that dampens the effect of shocks and reduces aggregate fluctuations.

We further show that in equilibrium the importance of a producer (measured by its sales share, or Domar weight) is greater when its productivity has a higher expected value or a lower variance. More broadly, the impact of beliefs on the economy depends crucially on substitution patterns that describe whether the Domar weights of two sectors tend to move together or in opposite directions after a change in the TFP process. These patterns depend on how technique choices affect productivity and on the covariance matrix of the TFP shocks. For instance, if sectors i and j are strongly positively correlated, the planner tends to make them move in opposite direction as to avoid too much risk exposure. In that case, an increase in the expected productivity of sector i is accompanied by a decline in the Domar weight of sector j .

Whether sectors are substitutes or complements also determines how the expected value and the variance of GDP respond to shifts in beliefs. We characterize conditions under which these changes are amplified or mitigated, compared to the fixed-network benchmark of Hulten's theorem (Hulten, 1978). We further show that when there is no uncertainty, Hulten's theorem applies in our setting even though the network is endogenous.

In some circumstances, the forces at work in the model can have counterintuitive implications for how the productivity process affects aggregate quantities. While an increase in expected pro-

ductivity or a decline in volatility always benefit 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 theorem is not a good guide to understanding changes in GDP, even as a first-order approximation. To understand why, imagine 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 riskier producers and, as a result, expected GDP might fall. We show that a similar mechanism also implies that an increase in the volatility of a sector’s productivity can lead to a decline in the variance of aggregate output.

We provide a basic calibration of the model using sectoral U.S. data. To isolate the impact of uncertainty, we compare our calibrated model to an alternative economy in which firms are unconcerned about risk when making sourcing decisions. Although this economy is similar to the baseline model during normal times, significant differences appear during high-volatility periods, such as the Great Recession. During that episode, firms responded to uncertainty by moving to safer but less productive suppliers. These decisions led to a meaningful reduction in the volatility of GDP, but the added stability came at the cost of an additional decline in expected GDP.

The model that we use for our quantitative analysis relies on some simplifying assumptions for tractability reasons. To verify the robustness of our findings, we provide additional empirical evidence that does not rely on the structure of the model. Taking advantage of rich firm-level U.S. data, we find that, as in the model, higher uncertainty leads to a decline in Domar weights, 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 from [Alfaro et al. \(2019\)](#) 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](#)). 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 can lead to lower aggregate output and measured TFP.²

There is a growing literature that studies how shocks propagate through production networks, in the spirit of early contributions by [Long and Plosser \(1983\)](#), [Dopor \(1999\)](#) and [Horvath \(2000\)](#).

²[Fernández-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 a long-lasting impact on economic aggregates.

[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.³ [Acemoglu et al. \(2017\)](#) and [Baqae and Farhi \(2019a\)](#) describe conditions under which production networks can generate fat-tailed aggregate output distributions. [Foerster et al. \(2011\)](#) and [Atalay \(2017\)](#) study the empirical contributions of sectoral shocks for aggregate fluctuations. [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.⁴

In most of this 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 choose production techniques in the presence of uncertainty, Hulten’s theorem is not a useful guide to how productivity affects expected GDP in our model.⁵ 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. [Oberfield \(2018\)](#) builds a model in which each firm selects a supplier to purchase from, and studies how changes in the environment affect the production network. Our model is closely related to [Acemoglu and Azar \(2020\)](#). As in that paper, we model endogenous network formation as a technique choice problem in which firms do not internalize the impact of their supply chain decisions on equilibrium objects. The key difference between the two models is in terms of timing. In [Acemoglu and Azar \(2020\)](#) firms know the realization of the shock when choosing their technique, while in our model that decision is made under uncertainty. As a result, in our setting uncertainty and beliefs influence the structure of the network and economic aggregates. [Taschereau-Dumouchel \(2020\)](#), [Acemoglu and Tahbaz-Salehi \(2020\)](#) and [Elliott et al. \(2022\)](#) study economies in which firms’ decisions to operate or not shape the production network. [Lim \(2018\)](#) and [Huneus \(2018\)](#) evaluate the importance of endogenous changes in the 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 network formation model using Indian microdata to study misallocation in input markets. [Bernard et al. \(2022\)](#) build a model of network formation to explain firm heterogeneity. [Grossman et al. \(2023\)](#) consider how policy can improve resiliency in a model with endogenous formation of supply links.⁶ A key distinguishing feature of our work is its focus on how uncertainty affects the structure of the production network and macroeconomic aggregates.

³Production networks are one mechanism through which granular fluctuations can emerge ([Gabaix, 2011](#)).

⁴Other studies have looked at the importance of production networks outside the business cycle literature. [Jones \(2011\)](#) investigates their importance to explain income differences between countries. [Barrot and Sauvagnat \(2016\)](#), [Boehm et al. \(2019\)](#) and [Carvalho et al. \(2021\)](#) study the propagation of shocks after natural disasters.

⁵[Baqae and Farhi \(2019a\)](#) investigate departures from Hulten’s theorem due to higher-order effects. Recent work that has studied production networks under distortions, where Hulten’s theorem generally does not hold, includes [Baqae \(2018\)](#), [Liu \(2019\)](#), [Baqae and Farhi \(2019b\)](#), [Bigio and La’O \(2020\)](#) and [Caliendo et al. \(2022\)](#).

⁶[Atalay et al. \(2011\)](#) show that a “preferential attachment” model can fit features of the U.S. firm-level production network. [Carvalho and Voigtländer \(2014\)](#) build a rule-based network formation model to study the diffusion of intermediate inputs. [Kopytov \(2023\)](#) studies financial interconnectedness and systemic risk under uncertainty.

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 characterize the equilibrium when the network is fixed. We then investigate the firms' technique choice problem in [Section 4](#) and consider the full equilibrium with a flexible network in [Section 5](#). In [Sections 6 and 7](#), we describe how the productivity process affects the production network, welfare and GDP. In [Section 8](#), we provide a basic calibration of the model. [Section 9](#) provides additional empirical evidence in support of the mechanisms. The last section concludes. All proofs are in [Appendix A](#) and [Supplemental Appendix D](#).

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 $i \in \{1, \dots, n\}$, 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 $A_i(\alpha_i)$. 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, $\alpha_i = (\alpha_{i1}, \dots, \alpha_{in})$, and to write the

corresponding production function as

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

where L_i is labor and $X_i = (X_{i1}, \dots, X_{in})$ 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.⁷

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 $\mathcal{A}_i = \{\alpha_i \in [0, 1]^n : \sum_{j=1}^n \alpha_{ij} \leq \bar{\alpha}_i\}$, where $0 < 1 - \bar{\alpha}_i < 1$ provides a lower bound on the share of labor in the production of good i . We denote by \mathcal{A} the Cartesian product $\mathcal{A}_1 \times \dots \times \mathcal{A}_n$, such that an element $\alpha \in \mathcal{A}$, which corresponds to a choice of inputs for each sector, fully characterizes the production network in this economy. 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.

The choice of technique influences the total factor productivity of sector i through $A_i(\alpha_i)$. This term is given by nature and represents how effective a combination of inputs is 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 uses aluminum, steel, car engines, etc. would be associated with a higher productivity. When deciding on its optimal production technique, firm i 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 unique vector of *ideal* input shares $\alpha_i^\circ \in \mathcal{A}_i$ that maximizes A_i and that represents the most productive way to combine intermediate inputs to produce good i . Without loss of generality, we normalize $A_i(\alpha_i^\circ) = 1$ for all i .⁸

Example. One example of a function $A_i(\alpha_i)$ that satisfies Assumption 1 is the quadratic form

$$\log A_i(\alpha_i) = \frac{1}{2} (\alpha_i - \alpha_i^\circ)^\top H_i (\alpha_i - \alpha_i^\circ), \quad (2)$$

⁷Namely, $[\zeta(\alpha_i)]^{-1} = \left(1 - \sum_{j=1}^n \alpha_{ij}\right)^{1-\sum_{j=1}^n \alpha_{ij}} \prod_{j=1}^n \alpha_{ij}^{\alpha_{ij}}$. 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.

⁸We further assume for all i that $\nabla \log A_i(\alpha_i^\circ) = 0$, where ∇ denotes the gradient. Since α_i° maximizes A_i , this assumption is potentially restrictive only if $\alpha_i^\circ \notin \text{int } \mathcal{A}_i$. All the results go through without it except that an extra term must be added to the approximation in Proposition 8.

where H_i is a negative-definite matrix that is also the Hessian of $\log A_i$. Throughout the paper, we will sometimes assume that A_i takes this form to more transparently describe the mechanisms at work. We also use this functional form in the quantitative section of the paper.

The distribution of sectoral productivity shock ε_i in (1) is a key primitive of the model and an important input into the firms' technique choice problem. We collect these shocks in the vector $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n)$, which we assume to be normally distributed, $\varepsilon \sim \mathcal{N}(\mu, \Sigma)$.⁹ The vector μ determines the expected level of sectoral productivities, while the covariance matrix Σ determines both uncertainty about individual elements of ε and their correlations across industries. We assume throughout that Σ is positive definite. The vector ε is the only source of uncertainty in this economy.

In equilibrium, ε will have a direct impact on prices, and its moments (μ, Σ) will affect expectations about the price system. For instance, a sector with a high μ_i will have a low expected 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 of good i . Since production techniques must be chosen before ε is realized, the beliefs (μ, Σ) will affect the sourcing decisions of the firms.

We impose the restriction that the representative firm in sector i can only adopt one 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 \mathcal{A}_i . After the realization of the productivity shocks, 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 exposition of the model, we adopt this restriction in an ad hoc fashion here, but provide a possible microfoundation for it in Supplemental Appendix E.

2.2 Household preferences

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

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

where $\beta_i > 0$ for all i and $\sum_{i=1}^n \beta_i = 1$. 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 $u(\cdot)$ is CRRA with a coefficient of relative risk aversion $\rho \geq 1$.¹⁰ The household makes consumption decisions after uncertainty is resolved

⁹This 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. A common way to correct for this effect is to remove half of the variance of ε from its mean. Supplemental Appendix H describes why such a correction is problematic in our model and discusses other potential corrections.

¹⁰The case $0 < \rho < 1$ is straightforward to characterize but is somewhat unnatural since the household then seeks to increase the variance of log consumption. Indeed, when $\log Y$ is normal, maximizing $E[(1 - \rho)^{-1} Y^{1-\rho}]$ amounts

and so in each state of the world it faces the budget constraint

$$\sum_{i=1}^n P_i C_i \leq 1, \quad (4)$$

where P_i is the price of good i , and the wage is used as the numeraire.

Firms are owned by the representative household and maximize expected profits discounted by the household's stochastic discount factor¹¹

$$\Lambda = u'(Y) / \bar{P}, \quad (5)$$

where $\bar{P} = \prod_{i=1}^n P_i^{\beta_i}$ 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

$$y = -\beta^\top p, \quad (6)$$

where $y = \log Y$, $p = (\log P_1, \dots, \log P_n)$ and $\beta = (\beta_1, \dots, \beta_n)$. Log GDP is thus the negative of the sum of log prices weighted by the consumption shares β . Intuitively, as prices decrease relative to wages, the household can purchase more goods, and aggregate consumption increases.

2.3 Unit cost minimization

We solve the problem of a given 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 the fact 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. Under a given technique α_i , the cost minimization problem of a firm in sector i is

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

The solution to this problem implicitly defines the unit cost of production $K_i(\alpha_i, P)$, which plays

to maximizing $E[\log Y] - \frac{1}{2}(\rho - 1)V[\log Y]$ 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 Supplemental Appendix H for a version of the model in which we correct for this term.

¹¹See Supplemental Appendix C.1 for the derivation of equations (5) and (6).

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 $P = (P_1, \dots, P_n)$. We show in Appendix C.2 that the production function (1) implies that

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

which is the standard Cobb-Douglas unit cost function. Equation (8) states that the cost of producing one unit of good i is equal to the geometric average of the individual input prices (weighted by their respective shares) adjusted for sectoral total factor productivity.

2.4 Technique choice

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

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

where Q_i is the equilibrium demand for good i , and where the profits in different states of the world are weighted by the household's stochastic discount factor Λ . The representative firm takes P , Q_i and Λ as given, and so the only term in (9) over which it has any control is the unit cost $K_i(\alpha_i, P)$. 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 .¹² Because profits are discounted by Λ , firms effectively inherit the risk attitude of the representative household.

2.5 Equilibrium conditions

In equilibrium, competitive pressure pushes prices to be equal to unit costs, so that

$$P_i = K_i(\alpha_i, P) \text{ for all } i \in \{1, \dots, n\}. \quad (10)$$

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 .¹³

An equilibrium is defined by the optimality conditions of both the household and the firms

¹²As usual, the presence of the stochastic discount factor in the firm problem comes from the implicit assumption that there are complete markets in this economy. Since agents can trade state- ε contingent claims, state prices reflect the marginal utility of the household in each state.

¹³Even without imposing that production techniques are Cobb-Douglas, the system (10) 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.

holding simultaneously, together with the usual market clearing conditions.

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

1. (Optimal technique choice) For each $i \in \{1, \dots, n\}$, 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 $i \in \{1, \dots, n\}$, factor demands per unit of output L_i^*/Q_i^* and X_i^*/Q_i^* are a solution to (7) given prices P^* and the chosen technique α_i^* .
3. (Consumer maximization) The consumption vector C^* maximizes (3) subject to (4) given prices P^* .
4. (Unit cost pricing) For each $i \in \{1, \dots, n\}$, P_i^* solves (10) where $K_i(\alpha_i^*, P^*)$ is given by (8).
5. (Market clearing) For each $i \in \{1, \dots, n\}$,

$$C_i^* + \sum_{j=1}^n X_{ji}^* = Q_i^* = F_i(\alpha_i^*, L_i^*, X_i^*), \text{ and } \sum_{i=1}^n L_i^* = 1. \quad (11)$$

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 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 or wedges between unit costs and prices. We work out this last extension in Supplemental Appendix K. In Supplemental Appendix L, we also consider additional sources of uncertainty in terms of 1) household preferences, 2) labor supply, and 3) distortions (e.g. due to government policies). We find that these sources of uncertainty either do not matter for the equilibrium network, matter only if they interact with the productivity shocks ε , or have a similar impact to the uncertainty about ε .

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 input shares, as we explore in more details in Section 6.

3 Equilibrium prices and GDP in a fixed-network economy

Before analyzing how the equilibrium production network responds to changes in the productivity process, it is useful to first establish how prices and GDP behave under a fixed network. To this end, we first define two objects 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 $\mathcal{L}(\alpha) = I + \alpha + \alpha^2 + \dots$. An element i, j of $\mathcal{L}(\alpha)$ captures the importance of sector j as an input in the production of good i by taking into account direct and indirect connections between the two sectors in the production network.

We also define the Domar weight ω_i of sector i as the ratio of its sales to nominal GDP, such that $\omega_i = \frac{P_i Q_i}{P^\top C}$. As we show in the proof of Corollary 1, the vector of Domar weights $\omega = (\omega_1, \dots, \omega_n)$ is equal to $\omega^\top = \beta^\top \mathcal{L}(\alpha) > 0$ in the model. 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. They are constant in a fixed-network economy but vary when firms are free to adjust sourcing decisions.

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

Lemma 1. *Under a given network α , the vector of log prices is given by*

$$p(\alpha) = -\mathcal{L}(\alpha)(\varepsilon + a(\alpha)), \quad (12)$$

and log GDP is given by

$$y(\alpha) = \omega(\alpha)^\top (\varepsilon + a(\alpha)), \quad (13)$$

where $a(\alpha) = (\log A_1(\alpha_1), \dots, \log A_n(\alpha_n))$.

Lemma 1 describes how prices and GDP depend on 1) the productivity vector $\varepsilon + a(\alpha)$ and 2) the production network α . 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 when the network is fixed. 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.

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 $\mathcal{L}(\alpha) = I + \alpha + \alpha^2 + \dots$ in (12) implies that the price of good i depends not only on i 's productivity, but also on the productivity of its suppliers, and on the productivity of their suppliers, and so on. These higher-order connections also matter for GDP and thus the impact 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 . The first and second moments of y can thus be written as

$$\mathbb{E}[y(\alpha)] = \omega(\alpha)^\top (\mu + a(\alpha)) \quad \text{and} \quad \mathbb{V}[y(\alpha)] = \omega(\alpha)^\top \Sigma \omega(\alpha). \quad (14)$$

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 under a fixed production network. 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 the moments of log GDP is given by*

$$\frac{\partial \mathbb{E}[y]}{\partial \mu_i} = \omega_i, \quad \text{and} \quad \frac{\partial \mathbb{V}[y]}{\partial \mu_i} = 0.$$

2. *The impact of a change in volatility Σ_{ij} on the moments of log GDP is given by¹⁴*

$$\frac{\partial \mathbb{E}[y]}{\partial \Sigma_{ij}} = 0, \quad \text{and} \quad \frac{\partial \mathbb{V}[y]}{\partial \Sigma_{ij}} = \omega_i \omega_j.$$

The first part of the corollary 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 corollary establishes a similar result for a change in Σ . It shows that the impact of an increase in the volatility of a sector's TFP on the variance of log GDP is equal to the square of that sector's sales share. This result also applies to a change in covariance, in which case the increase in $\mathbb{V}[y]$ is equal to the product of the two industries' sales shares. Since Domar weights are always positive, an increase in covariance 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, the Corollary shows that when the network is fixed, the covariance matrix Σ has no impact on $\mathbb{E}[y]$. It follows that whenever we discuss the response of expected log GDP

¹⁴Whenever we take derivatives with respect to off-diagonal elements of Σ , we simultaneously change Σ_{ij} and Σ_{ji} to preserve the symmetry of Σ , and divide the result by two.

to a change in uncertainty, the mechanism must operate through the endogenous reorganization of the network.

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 Section 7, 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 Firm decisions

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 $\lambda(\alpha^*) = \log \Lambda(\alpha^*)$ and the log of the unit cost $k_i(\alpha_i, \alpha^*) = \log K_i(\alpha_i, P^*(\alpha^*))$, where α^* denotes the *equilibrium* network. These quantities are normally distributed in equilibrium.

Using this notation, we can reorganize the problem of the firm (9) as¹⁵

$$\alpha_i^* \in \arg \min_{\alpha_i \in \mathcal{A}_i} \mathbb{E}[k_i(\alpha_i, \alpha^*)] + \text{Cov}[\lambda(\alpha^*), k_i(\alpha_i, \alpha^*)]. \quad (15)$$

The objective function in (15) captures how beliefs and uncertainty affect the production network. Its 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

$$\mathbb{E}[k_i(\alpha_i, \alpha^*)] = -\mu_i - a_i(\alpha_i) + \sum_{j=1}^n \alpha_{ij} \mathbb{E}[p_j],$$

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 (15) 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 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 $\text{Cov}[\lambda, k_i] = \text{Corr}[\lambda, k_i] \sqrt{V[\lambda]} \sqrt{V[k_i]}$, 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), $\text{Corr}[\lambda, k_i]$ is typically positive, and so firms seek to minimize the variance of their unit cost.¹⁶ This has several implications for their choice of suppliers. To see this, we can use (8)

¹⁵We show how to derive this equation as part of the proof of Lemma 2.

¹⁶If i 's productivity shock is strongly negatively correlated with that of the other sectors, it can be that

to write

$$V[k_i(\alpha_i, \alpha^*)] = \sum_{j=1}^n \alpha_{ij}^2 V[p_j] + \sum_{j \neq k} \alpha_{ij} \alpha_{ik} \text{Cov}[p_j, p_k] + 2 \text{Cov}\left[-\varepsilon_i, \sum_{j=1}^n \alpha_{ij} p_j\right] + \Sigma_{ii}. \quad (16)$$

The variance of the unit cost can thus be decomposed into four channels. 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 shock, the price of its inputs are then more likely to be low, reducing the expected increase in its unit cost. Finally, the last term captures the fact that a more volatile productivity ε_i contributes to a more volatile unit cost.

Risk-adjusted prices

At an equilibrium network α^* , we can simplify the technique choice problem of the firm by introducing a risk-adjusted version of sectoral prices.

Lemma 2. *In equilibrium, the technique choice problem of the representative firm in sector i is*

$$\alpha_i^* \in \arg \max_{\alpha_i \in \mathcal{A}_i} a_i(\alpha_i) - \sum_{j=1}^n \alpha_{ij} \mathcal{R}_j(\alpha^*), \quad (17)$$

where

$$\mathcal{R}(\alpha^*) = E[p(\alpha^*)] + \text{Cov}[p(\alpha^*), \lambda(\alpha^*)] \quad (18)$$

is the vector of equilibrium risk-adjusted prices, and where

$$E[p(\alpha^*)] = -\mathcal{L}(\alpha^*)(\mu + a(\alpha^*)) \text{ and } \text{Cov}[p(\alpha^*), \lambda(\alpha^*)] = (\rho - 1) \mathcal{L}(\alpha^*) \Sigma [\mathcal{L}(\alpha^*)]^\top \beta.$$

This lemma shows that all the equilibrium information needed for the firm's problem is contained in the vector of risk-adjusted prices \mathcal{R} . This quantity provides an overall measure of the desirability of an input that depends on its expected price and on how its price covaries with the stochastic discount factor. This latter term implies that goods that are cheap when aggregate consumption is low are particularly attractive as inputs.

Lemma 2 implies that the TFP shifter a_i plays a crucial role in determining how a change in risk-adjusted prices affects firm i 's chosen input shares. To see this, we can take the first-order

$\text{Corr}[\lambda, k_i] < 0$, in which case i seeks to be *more* volatile to insure the household in states of low consumption.

condition for an interior solution of problem (17) and use the implicit function theorem to write

$$\frac{\partial \alpha_{ij}}{\partial \mathcal{R}_k} = [H_i^{-1}(\alpha_i)]_{jk}, \quad (19)$$

where H_i^{-1} is the inverse of the Hessian matrix of a_i and where $[\cdot]_{jk}$ denotes its element j, k . This equation implies that if a good k becomes marginally more expensive or more risky (higher \mathcal{R}_k), firm i responds by changing its share α_{ik} of good k by $[H_i^{-1}(\alpha_i)]_{kk}$. Since a_i is strictly concave by Assumption 1, the diagonal elements of H_i^{-1} are negative, and so a higher \mathcal{R}_k always leads to a decline in α_{ik} . The size of that decline depends on the curvature of a_i .¹⁷

Whether the increase in \mathcal{R}_k leads to a decline or an increase in the share of other inputs $j \neq k$ depends on whether the shares of j and k are complements or substitutes in the production of good i . If $[H_i^{-1}]_{jk} > 0$ we say that they are *substitutes*, and in that case a higher risk-adjusted price \mathcal{R}_k leads to an increase in α_{ij} . As the firm decreases α_{ik} , the incentives embedded in a_i to increase α_{ij} get stronger, and the firm substitutes α_{ij} for α_{ik} . In contrast, if $[H_i^{-1}]_{jk} < 0$ we say that the shares of j and k are *complements*, and an increase in \mathcal{R}_k leads to a decline in α_{ij} . One sufficient condition for a Hessian matrix H_i to feature complementarities for all sectors is $[H_i]_{jk} \geq 0$ for all $j \neq k$.¹⁸

This notion of substitution and complementarity embedded in H_i^{-1} applies ex-ante, before uncertainty is realized, and when firms can adjust their input shares. It is not to be confused with the usual elasticity of substitution between goods, which would be computed ex-post, once the shares are fixed, and which equals one in our setup given the Cobb-Douglas nature of production.

While (19) is only valid at an interior solution of the firm's problem, the forces that it captures are also at work when some of the constraints embedded in $\alpha_i \in \mathcal{A}_i$ bind. But these constraints can also increase the degree of substitution between input shares. Suppose, for instance, that the minimum labor share constraint $\sum_{l=1}^n \alpha_{il} \leq \bar{\alpha}_i$ binds, and that the risk-adjusted 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 .

Example: Substitutability and complementarity in partial equilibrium

To show how the substitution patterns embedded in a_i affect technique choices, we can revisit the car manufacturer example from the introduction. Suppose that this manufacturer primarily uses steel (input 1) to produce cars, and that it relies on equipment (input 2) such as milling machines and lathes to transform raw steel into usable components. As before, the manufacturer

¹⁷The inverse Hessian frequently appears when doing comparative statics using the implicit function theorem. See for instance Chapter 17.G in Mas-Colell et al. (1995). See Supplemental Appendix C.3 for the derivation of (19).

¹⁸In this case $-H_i$ is an M-matrix and therefore inverse-positive. Intuitively, $[H_i]_{jk} \geq 0$ implies that a higher α_{ij} increases the TFP benefit of raising α_{ik} .

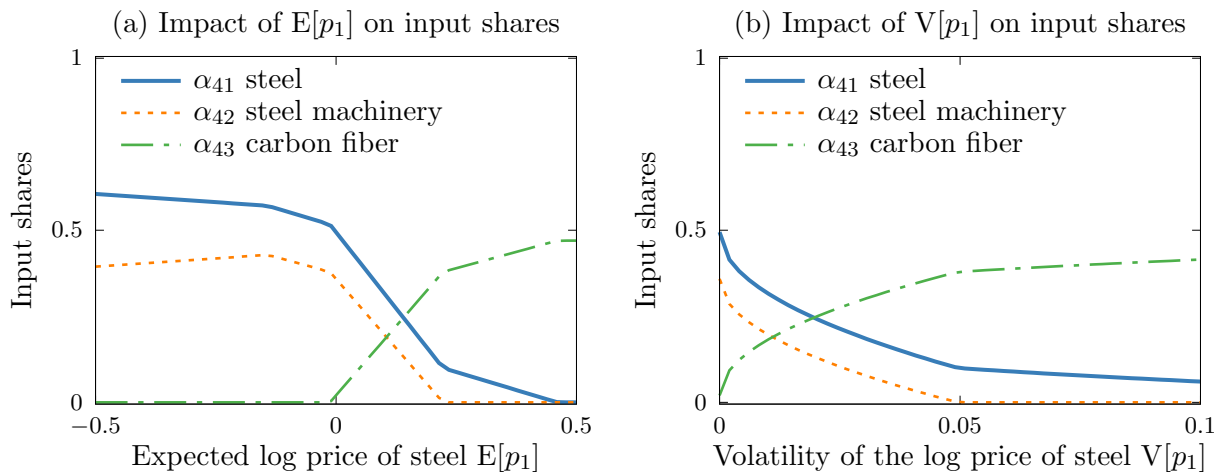
also has the option to purchase carbon fiber (input 3) to replace steel components if needed. It would be natural to endow this manufacturer (sector $i = 4$) with a TFP shifter function of the form

$$a_4(\alpha_4) = - \sum_{j=1}^4 \kappa_j (\alpha_{4j} - \alpha_{4j}^\circ)^2 - \psi_1 (\alpha_{41} - \alpha_{42})^2 - \psi_2 [(\alpha_{41} + \alpha_{43}) - (\alpha_{41}^\circ + \alpha_{43}^\circ)]^2, \quad (20)$$

where $\kappa_j > 0$, $\psi_1 > 0$ and $\psi_2 > 0$. From the second term, we see that any increase in the share α_{41} of steel would incentivize the firm to increase the share α_{42} of steel machinery as well. Inputs 1 and 2 are therefore complements in the production of cars. In contrast, the third term implies that any increase in the share α_{41} of steel would make it optimal to reduce the share α_{43} of carbon fiber, and so the shares of inputs 1 and 3 are substitutes. These patterns can be confirmed by computing the inverse Hessian of a_4 directly and inspecting the off-diagonal terms. The parameters $\psi_1 > 0$ and $\psi_2 > 0$ determine the strength of these substitution-complementarity patterns.

Figure 1 shows what happens to the production technique chosen by this car manufacturer if the risk-adjusted price of steel increases. In panel (a) the increase in \mathcal{R}_1 comes from a higher expected price $E[p_1]$, while in panel (b) the price of steel becomes more volatile (higher $V[p_1]$). Naturally, when the risk-adjusted price of steel rises, the manufacturer relies less on steel in production, and α_{41} falls. Since steel machinery is only useful when steel is used in production, the share α_{42} falls as well. If the increase in \mathcal{R}_1 is large enough, the manufacturer severs the link with its steel and steel machinery suppliers completely so that both $\alpha_{41} = \alpha_{42} = 0$. At the same time, as steel becomes more expensive in risk-adjusted terms, the firm finds a carbon fiber supplier and progressively increases the share α_{43} .

Figure 1: Impact of rising the risk-adjusted price of steel



Notes: a_4 as in (20) with $\psi_1 = \psi_2 = 1$, $\alpha_{4j}^\circ = 1/3$, $\kappa_j = 1/10$ for $j \neq 4$, and $\kappa_4 = \infty$ and $\alpha_{44}^\circ = 0$. $\sqrt{V[\lambda]} \text{Corr}[p_j, \lambda] = 1$, $E[p_2] = -0.05$, $E[p_3] = 0.05$, and $V[p_2] = V[p_3] = 0.1$. Panel (a): $V[p_1] = 0$. Panel (b): $E[p_1] = 0.05$.

5 Equilibrium existence, uniqueness and efficiency

In the previous section we characterized how an individual firm's technique choice depends on risk-adjusted prices. However, prices are equilibrium objects that depend on the production network and, therefore, on the choices made by other firms. In this section, we consider the full equilibrium mapping and show that there exists a unique equilibrium and that it is efficient. To prove these results, we rely on the problem of the social planner, and on the fact that the set of equilibria coincides with the set of efficient allocations.

5.1 The efficient allocation

There is a representative household in the economy, and so finding the set of Pareto efficient allocations amounts to solving the problem of a social planner that maximizes the utility function (3) subject to the resource constraints (11). The following lemma characterizes production networks that solve that problem.

Lemma 3. *An efficient production network α^* solves*

$$\mathcal{W} := \max_{\alpha \in \mathcal{A}} W(\alpha, \mu, \Sigma),$$

where \mathcal{W} is a measure of the welfare of the household, and where

$$W(\alpha, \mu, \Sigma) := \mathbb{E}[y(\alpha)] - \frac{1}{2}(\rho - 1) \mathbb{V}[y(\alpha)], \quad (21)$$

is welfare under a given network α .¹⁹

Lemma 3 follows directly from the fact that an efficient network must maximize the expected utility of the representative household. It further shows that the household favors networks associated with high expected log GDP $\mathbb{E}[y(\alpha)]$ and low aggregate uncertainty $\mathbb{V}[y(\alpha)]$. The risk aversion parameter ρ determines the relative importance of these two terms.

Recasting household welfare in terms of Domar weights

Since Domar weights play a crucial role in determining the expected value and the variance of GDP, it will be useful to recast the problem of the social planner in the space of ω . Using (14), we can write the objective function (21) as

$$\omega^\top \mu + \omega^\top a(\alpha) - \frac{1}{2}(\rho - 1) \omega^\top \Sigma \omega. \quad (22)$$

¹⁹ \mathcal{W} is a convenient monotone transformation of the expected utility of the household, such that $\mathbb{E}[Y^{1-\rho}](1-\rho)^{-1} = \exp((1-\rho)\mathcal{W})(1-\rho)^{-1}$, and we adopt it as our 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.

The only term in this expression that does not depend exclusively on ω is $\omega^\top a(\alpha)$, which corresponds to the contribution of the TFP shifter functions (a_1, \dots, a_n) to aggregate TFP. We want to write this object in terms of ω alone. For that purpose, notice that several networks α are consistent with a given Domar weight vector ω , but that not all of them are equivalent in terms of welfare. Indeed, to achieve a given ω the planner will only select the network α that maximizes welfare, which amounts to maximizing $\omega^\top a(\alpha)$.

Formally, consider the optimization problem

$$\bar{a}(\omega) := \max_{\alpha \in \mathcal{A}} \omega^\top a(\alpha), \quad (23)$$

subject to the definition of the Domar weights given by $\omega^\top = \beta^\top \mathcal{L}(\alpha)$. We refer to the value function \bar{a} as the *aggregate TFP shifter* function. It provides the maximum value of TFP $\omega^\top a(\alpha)$ that can be achieved under the constraint that the Domar weights must be equal to some given vector ω . We denote by $\alpha(\omega)$ the solution to (23). Since both $\bar{a}(\omega)$ and $\alpha(\omega)$ depend exclusively on the TFP shifter functions (a_1, \dots, a_n) and on the preference vector β , these two functions will be invariant, for a given ω , to the changes in beliefs (μ, Σ) that we consider in the next sections.

Example. We can solve explicitly for $\bar{a}(\omega)$ and $\alpha(\omega)$ under the quadratic TFP shifter function specified in (2). At an interior solution $\alpha \in \text{int } \mathcal{A}$, the optimal production network $\alpha(\omega)$ that solves (23) for a given vector of Domar weights ω is²⁰

$$\alpha_i(\omega) - \alpha_i^\circ = H_i^{-1} \left(\sum_{j=1}^n \omega_j H_j^{-1} \right)^{-1} \left(\omega - \beta - \sum_{j=1}^n \omega_j \alpha_j^\circ \right), \quad (24)$$

for all i , and the associated value function \bar{a} is

$$\bar{a}(\omega) = \frac{1}{2} \sum_{i=1}^n \omega_i (\alpha_i(\omega) - \alpha_i^\circ)^\top H_i (\alpha_i(\omega) - \alpha_i^\circ). \quad (25)$$

From (24), it is straightforward to show that the gradients ∇a_i of the TFP shifter functions are all equal to each other such that $\nabla a_i = \nabla a_j$ for all i, j .²¹ It follows that at an interior solution, input shares must be such that the marginal TFP benefit $[\nabla a_i]_k$ of increasing α_{ik} is equal across all sectors i .

We can use $\bar{a}(\omega)$ to recast the planner's problem in the space of Domar weights.

²⁰See Supplemental Appendix C.4 for the full derivation.

²¹It is clear from (24) that $H_i(\alpha_i - \alpha_i^\circ) = H_j(\alpha_j - \alpha_j^\circ)$ for all i, j . Furthermore, since $a_i(\alpha)$ is a quadratic function given by (2), we have that $\nabla a_i = H_i(\alpha_i - \alpha_i^\circ)$, where $\nabla a_i = \frac{da_i}{d\alpha_i}$ denotes the gradient vector of a_i .

Corollary 2. *The efficient Domar weight vector ω^* solves*

$$\mathcal{W} = \max_{\omega \in \mathcal{O}} \underbrace{\omega^\top \mu + \bar{a}(\omega)}_{\mathbb{E}[y]} - \frac{1}{2}(\rho - 1) \underbrace{\omega^\top \Sigma \omega}_{\mathbb{V}[y]}, \quad (26)$$

where $\mathcal{O} = \{\omega \in \mathbb{R}_+^n : \omega \geq \beta \text{ and } 1 \geq \omega^\top (\mathbf{1} - \bar{\alpha})\}$ and $\bar{a}(\omega)$ is given by (23).

The set \mathcal{O} contains the vectors ω that are feasible given the restriction that the corresponding network $\alpha(\omega)$ must belong to \mathcal{A} . The first inequality in its definition follows from $\alpha_{ij} \geq 0$ for all i, j . The second inequality, where $\mathbf{1}$ denotes the $n \times 1$ all-one column vector, follows from $\sum_j \alpha_{ij} \leq \bar{\alpha}_i$ for all i .

One key advantage of the optimization problem (26) over (21) is that its choice variable is a vector instead of a matrix. This makes the comparative static results presented in the next section simpler and more transparent. In addition, the recast objective function (26) has attractive properties, as the following lemma shows.

Lemma 4. *The objective function of the planner's problem (26) is strictly concave. Furthermore, there is a unique vector of Domar weights ω^* that solves that problem, and there is a unique production network $\alpha(\omega^*)$ associated with that solution.*

This lemma shows that there is a unique efficient network in this economy. It also implies that first-order conditions are sufficient to characterize that network, such that we can easily solve for it using standard numerical methods.

5.2 Fundamental properties of the equilibrium

Having characterized the problem of the social planner, we can go back to the equilibrium and establish some of its basic properties. The following proposition follows from the fact that there are no frictions or externalities in the environment and that all markets are competitive.

Proposition 1. *There exists a unique equilibrium, and it is efficient.*

The proof of this proposition establishes that the set of equilibria coincides with the set of efficient allocations. Since by Lemma 4 there is a unique efficient allocation, it follows that there is also a unique equilibrium.

Proposition 1 implies that we can investigate the properties of the equilibrium by solving the problem of the social planner directly. This will prove useful when characterizing how the equilibrium network and aggregate quantities respond to changes in the productivity process.

6 Beliefs and the production network

In this section, we characterize how beliefs (μ, Σ) affect the equilibrium production network. We begin with a general result that describes how a change in a sector's risk or expected TFP impacts its own Domar weight. We then provide an expression that characterizes how the full vector of Domar weights responds to a marginal change in (μ, Σ) . Finally, we investigate how beliefs affect the structure of the underlying production network α . As we only consider the equilibrium network from now on, we lighten the notation by dropping the superscript $*$ when referring to equilibrium variables.

6.1 Domar weights

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

Proposition 2. *The Domar weight ω_i of sector i is (weakly) increasing in μ_i and (weakly) decreasing in Σ_{ii} .*

This proposition can be understood from both the perspective of an individual producer and 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 larger. From the planner's perspective, recall from (13) that the Domar weight of a sector captures its contribution to log GDP. Since the planner wants to increase and stabilize GDP, it naturally increases the importance of more productive (larger μ_i) and less volatile (smaller Σ_{ii}) sectors in the production network.

Risk-adjusted productivity shocks

Proposition 2 describes how the Domar weight of a sector responds to a change in its own TFP process, and it holds generally. At an interior equilibrium, we can also characterize how any change in beliefs affects the full vector ω . For that purpose, we introduce a risk-adjusted version of the productivity vector ε defined as

$$\mathcal{E} = \underbrace{\mu}_{E[\varepsilon]} - \underbrace{(\rho - 1) \Sigma \omega}_{\text{Cov}[\varepsilon, \lambda]}. \quad (27)$$

The vector \mathcal{E} captures how higher exposure to the productivity process ε affects the representative household's utility. It depends on how productive each sector i is in expectation, and on how its

ε_i covaries with the stochastic discount factor λ . If we denote by $\mathbf{1}_i$ the column vector with a 1 as i th element and zeros elsewhere, we can write

$$\frac{\partial \mathcal{E}}{\partial \mu_i} = \mathbf{1}_i, \quad (28)$$

such that an increase in μ_i makes sector i more attractive. It however leaves the risk-adjusted TFP of other sectors unchanged. Similarly, for a change in Σ_{ij} , we can compute

$$\frac{\partial \mathcal{E}}{\partial \Sigma_{ij}} = -\frac{1}{2} (\rho - 1) (\omega_j \mathbf{1}_i + \omega_i \mathbf{1}_j), \quad (29)$$

such that an increase in variance Σ_{ii} , by adding aggregate risk to the economy, decreases the risk-adjusted TFP of sector i . The intensity of that effect depends on the risk aversion of the household ρ and, through ω_i , on the importance of i as a supplier. Similarly, an increase in covariance Σ_{ij} , $i \neq j$, decreases the risk-adjusted TFP of both sectors i and j . Again, this effect is stronger when the household is more risk averse. In what follows, we refer to a change that increases \mathcal{E} as *beneficial*, and to a change that decreases \mathcal{E} as *adverse*.

Using the definition of \mathcal{E} , we can write the first-order condition of the planner's problem (26) at an interior solution as

$$\nabla \bar{a}(\omega) + \mathcal{E} = 0, \quad (30)$$

where $\nabla \bar{a}$ is the gradient of the aggregate TFP shifter function \bar{a} . This first-order condition shows that the planner balances the benefit of a sector in terms of risk-adjusted TFP against its impact on the aggregate TFP shifter.

Response of the Domar weight vector to changes in beliefs

The first-order condition (30) allows us to characterize how the entire vector of Domar weights responds to any change in the productivity process in a unified way. Applying the implicit function theorem to (30) yields the following result.

Proposition 3. *Let γ denote either the mean μ_i or an element of the covariance matrix Σ_{ij} . If $\omega \in \text{int } \mathcal{O}$, then the response of the equilibrium Domar weights to a change in γ is given by*

$$\frac{d\omega}{d\gamma} = \underbrace{-\mathcal{H}^{-1}}_{\text{propagation}} \times \underbrace{\frac{\partial \mathcal{E}}{\partial \gamma}}_{\text{impulse}}, \quad (31)$$

where the $n \times n$ negative definite matrix \mathcal{H} is given by

$$\mathcal{H} = \nabla^2 \bar{a} + \frac{d\mathcal{E}}{d\omega}, \quad (32)$$

and where the matrix $\nabla^2 \bar{a}$ is the Hessian of the aggregate TFP shifter function \bar{a} , and $\frac{d\mathcal{E}}{d\omega} = -\frac{d\text{Cov}[\varepsilon, \lambda]}{d\omega} = -(\rho - 1)\Sigma$ is the Jacobian matrix of the risk-adjusted TFP vector \mathcal{E} .²²

The response of the Domar weights to a change in beliefs, as given by (31), can be decomposed into an *impulse* component and a *propagation* component. The impulse captures the direct impact of the change on risk-adjusted TFP. It is simply given by the *partial* derivative of \mathcal{E} with respect to the moment of interest (see (28) and (29) above). This impulse is then propagated through \mathcal{H}^{-1} to capture its full equilibrium effect on the Domar weights.

Just as H_i^{-1} captured *local* substitution patterns between inputs in the problem of firm i , \mathcal{H}^{-1} captures *global*, economy-wide substitution patterns between sectors. If $\mathcal{H}_{ij}^{-1} < 0$, we say that i and j are global complements. If instead $\mathcal{H}_{ij}^{-1} > 0$, we say that i and j are global substitutes.

The following corollary justifies this terminology by showing that the sign of \mathcal{H}_{ij}^{-1} is sufficient to characterize how Domar weights respond to a change in the productivity process.

Corollary 3. *If $\omega \in \text{int } \mathcal{O}$, then the following holds.*

1. *An increase in the expected value μ_i or a decline in the variance Σ_{ii} leads to an increase in ω_j if i and j are global complements, and to a decline in ω_j if i and j are global substitutes.*
2. *An increase in the covariance Σ_{ij} , $i \neq j$, leads to a decline in ω_k if k is global complement with i and j , and to an increase in ω_k if k is global substitute with i and j .*

This corollary shows that if sectors are global complements they tend to move together after a change in beliefs. If they are substitutes instead, they tend to move in opposite directions. Indeed, by Proposition 2, a beneficial change to a sector i leads to an increase in its Domar weight ω_i . This direct effect then contributes to further adjustments of the Domar weights through \mathcal{H}^{-1} . Corollary 3 shows that for sectors that are complements with i , this indirect effect leads to an increase in their Domar weights. When they are substitutes, ω_j declines instead.

It is clear from (32) that global substitution patterns are determined by the shape of the TFP shifter functions (a_1, \dots, a_n) through $\nabla^2 \bar{a}(\omega)$, and by the household's risk perception through $-(\rho - 1)\Sigma$. We will explore these two channels in turn.

Σ and global substitution patterns

The following lemma describes how an increase in covariance Σ_{ij} between any two sectors affects the degree of global substitution between them.

²²This proposition focuses on an interior equilibrium, such that $\omega \in \text{int } \mathcal{O}$, but this restriction can be satisfied even if some shares α_{ij} are equal to zero. Indeed, for $\omega_j \geq \beta_j$ to bind, it must be that $\alpha_{ij} = 0$ for all i . Furthermore, Proposition 3 can be extended to include some binding constraints. When $\omega_i \geq \beta_i$ binds, ω_i is not affected by a marginal change in beliefs. We can therefore exclude these constrained Domar weights from the application of the implicit function theorem. It follows that a version of (31) holds for unconstrained Domar weights, as we show in Supplemental Appendix F.

Lemma 5. *An increase in the covariance Σ_{ij} induces stronger global substitution between i and j , in the sense that $\frac{\partial \mathcal{H}_{ij}^{-1}}{\partial \Sigma_{ij}} > 0$.*

Intuitively, if the correlation between ε_i and ε_j becomes larger, the planner has stronger incentives to lower ω_j after an increase in ω_i in order to reduce aggregate risk. From (32), we see that the strength of that diversification mechanism depends on the household's risk aversion through ρ .

$\nabla^2 \bar{a}$ and global substitution patterns

The curvature of the aggregate TFP shifter function \bar{a} , as captured by its Hessian $\nabla^2 \bar{a}$, also contributes to global substitution patterns. Intuitively, if a higher ω_i raises the marginal TFP benefit of increasing ω_j , sectors i and j tend to move together, which pushes these sectors to be global complements. Clearly, the local TFP shifter functions (a_1, \dots, a_n) play a key role in shaping \bar{a} such that the local substitution patterns matter for the global ones. The next lemma establishes sufficient conditions under which local complementarities translate into global complementarities.

Lemma 6. *Suppose that all input shares are (weak) local complements in the production of all goods, that is $[H_i^{-1}]_{kl} \leq 0$ for all i and all $k \neq l$. If $\alpha \in \text{int } \mathcal{A}$, there exists a scalar $\bar{\Sigma} > 0$ such that if $\|\Sigma\| \leq \bar{\Sigma}$, all sectors are global complements, that is $\mathcal{H}_{ij}^{-1} < 0$ for all $i \neq j$.*

This result shows that if all input shares are local complements, then sectors are also global complements, if the covariance matrix Σ is small enough. This last condition ensures that the substitution forces from diversification that are described in Lemma 5 do not dominate the complementarities coming from the TFP shifters (a_1, \dots, a_n) .

Lemma 6 also shows that sectors are global complements even if the local TFP shifters are neutral in the sense that $[H_i^{-1}]_{kl} = 0$ for all i and all $k \neq l$. This suggests that the equilibrium forces of the model, on their own, create global complementarities between sectors. To understand why, suppose that a sector i becomes more attractive, for instance due to an increase in μ_i . Any other sector j that relies either directly or indirectly on i ($\mathcal{L}_{ji} > 0$) would benefit from that change, and also become more attractive. By itself, this triggers an increase in Domar weights throughout the network and a shift away from labor. Through this mechanism, the model generates global complementarities between sectors, even under TFP shifter functions that do not feature local complementarities.

We also consider how local substitution can lead to global substitution. To do so, it is convenient to parametrize H_i to be able to tractably adjust the strength of local substitution. For that purpose, let

$$H_i^{-1} = \begin{bmatrix} -1 & \frac{s}{n-1} & \cdots & \frac{s}{n-1} \\ \frac{s}{n-1} & -1 & & \vdots \\ \vdots & & \ddots & \frac{s}{n-1} \\ \frac{s}{n-1} & \cdots & \frac{s}{n-1} & -1 \end{bmatrix}, \quad (33)$$

where we impose $-(n-1) < s < 1$ to guarantee that H_i^{-1} is negative definite. When $s < 0$ all input shares are complements in the production of good i , and when $s > 0$ they are substitutes. The next lemma describes sufficient conditions under which local substitution imply global substitution.

Lemma 7. *Suppose that all the TFP shifter functions (a_1, \dots, a_n) take the form (2), with $\alpha_i^\circ = \alpha_j^\circ$ for all i, j , and that H_i^{-1} is of the form (33) for all i . If $\alpha \in \text{int } \mathcal{A}$, there exists a scalar $\bar{\Sigma} > 0$ and a threshold $0 < \bar{s} < 1$ such that if $\|\Sigma\| \leq \bar{\Sigma}$ and $s > \bar{s}$, then all sectors are global substitutes, that is $\mathcal{H}_{ij}^{-1} > 0$ for all $i \neq j$.*

This result shows that global substitution emerges if local substitution forces are sufficiently strong (s sufficiently close to 1) to overcome the natural forces of the model that push for complementarity between sectors. Again, this result requires that $\|\Sigma\| \leq \bar{\Sigma}$ to limit the complementarity forces that could arise, for instance, from two sectors that are strongly negatively correlated.

An approximate equation for the equilibrium Domar weights

Propositions 2 and 3 describe how the equilibrium Domar weights respond to a marginal change in beliefs but they are silent about which sectors will have large or small Domar weights in equilibrium. Given the structure of the TFP shifter function \bar{a} , solving the planner's problem (26) for ω must in general be done using numerical methods. We can however derive approximate equations for ω using a Taylor expansion of $\nabla \bar{a}$. The ideal shares α° , as they lead to the highest values of the TFP shifters (a_1, \dots, a_n) , provide a natural point around which to do this approximation. Denote by $\omega^\circ = [\mathcal{L}(\alpha^\circ)]^\top \beta$ the vector of Domar weights associated with α° . Then, if the equilibrium network ω is close to ω° , we can write

$$\nabla \bar{a}(\omega) \approx \nabla \bar{a}(\omega^\circ) + \nabla^2 \bar{a}(\omega^\circ)(\omega - \omega^\circ). \quad (34)$$

This approximation is accurate if, for instance, the cost of deviating from the ideal shares embedded in the local TFP shifters is large. We work out that case formally in Supplemental Appendix I.

With this approximation, the first-order condition (30) becomes linear in ω , and we can solve for the equilibrium Domar weights.

Lemma 8. *If $\omega \in \text{int } \mathcal{O}$, the equilibrium Domar weights are approximately given by*

$$\omega = \omega^\circ - [\mathcal{H}^\circ]^{-1} \mathcal{E}^\circ + O(\|\omega - \omega^\circ\|^2), \quad (35)$$

where the superscript \circ indicates that \mathcal{H} and \mathcal{E} are evaluated at ω° .

This proposition provides an approximate expression for the equilibrium Domar weights in terms of the global substitution patterns embedded in $[\mathcal{H}^\circ]^{-1}$ and the expected attractiveness of all sectors, as captured by the risk-adjusted productivity \mathcal{E}° . Suppose that a sector i is endowed with a

productivity process ε_i that is high in expectation or that has a low covariance with the stochastic discount factor. In this case, \mathcal{E}_i° is large and, since the diagonal elements of $[\mathcal{H}^\circ]^{-1}$ are negative, ω_i tends to be larger than ω_i° .²³ In addition, the large \mathcal{E}_i° contributes to increasing the Domar weights of all sectors that are global complements with i , and to decreasing the Domar weights of sectors that are global substitutes with it.

6.2 The production network

In the previous section, we described how a change in beliefs affects the vector of Domar weights. While Domar weights are key objects that influence aggregate outcomes, they do not provide a complete description of the underlying production network. In this section, we extend our analysis and characterize how beliefs affect the individual links in the equilibrium network α .

Proposition 4. *If $\alpha \in \text{int } \mathcal{A}$, there exists a scalar $\bar{\Sigma} > 0$ such that if $\|\Sigma\| \leq \bar{\Sigma}$ the following holds.*

1. *(Complementarity) Suppose that input shares are local complements in the production of good i , that is $[H_i^{-1}]_{kl} < 0$ for all $k \neq l$. Then a beneficial change to k ($\partial \mathcal{E}_k / \partial \gamma > 0$) increases α_{ij} for all j .*
2. *(Substitution) Suppose that the conditions of Lemma 7 about the TFP shifters (a_1, \dots, a_n) hold. Then there exists a threshold $0 < \bar{s} < 1$ such that if $s > \bar{s}$, a beneficial change to k ($\partial \mathcal{E}_k / \partial \gamma > 0$) decreases α_{ij} for all i and all $j \neq k$, and increases α_{ik} for all i .*

Point 1 shows that if all inputs are local complements in the production of good i , all shares α_{ij} tend to move together. After a beneficial change to a given sector k , firms in sector i increase their reliance on k which, through complementarity, leads to an increase in i 's reliance on other sectors as well. If instead local substitution forces are sufficiently strong (point 2), a beneficial change to the productivity process of firm k still leads to a higher reliance on sector k , but in this case the forces embedded in H_i push for a decline in other shares. The proof of Proposition 4 also provides an explicit expression for the derivative $d\alpha_{ij}/d\gamma$ in terms of the gradient of $\alpha(\omega)$ and of $d\omega/d\gamma$.

An approximate equation for the equilibrium production network

As for the Domar weights, one must in general use numerical methods to find the equilibrium network α . We can, however, derive an approximation for the equilibrium production network when the cost of deviating from the ideal shares α° is large. Specifically, let $a_i(\alpha_i) = \bar{\kappa} \times \hat{a}_i(\alpha_i)$ and suppose that $\alpha_i^\circ \in \text{int } \mathcal{A}_i$. The parameter $\bar{\kappa} > 0$ captures how costly it is for the firms to deviate from α° in terms of TFP loss. When $\bar{\kappa}$ is large, we can use perturbation theory to derive an approximate equation for α (Judd and Guu, 2001; Schmitt-Grohé and Uribe, 2004).

²³Recall that \mathcal{H}^{-1} is negative definite by Proposition 3.

Lemma 9. *If $\alpha \in \text{int } \mathcal{A}$, the equilibrium input shares in sector i are approximately given by*

$$\alpha_i = \alpha_i^\circ + \bar{\kappa}^{-1} \left(\hat{H}_i^\circ \right)^{-1} \mathcal{R}^\circ + O(\kappa^{-2}), \quad (36)$$

where \hat{H}_i° is the Hessian of \hat{a}_i at α_i° , and where the vector of risk-adjusted prices at α° is given by


$$\mathcal{R}^\circ = -\mathcal{L}^\circ \mu + (\rho - 1) \mathcal{L}^\circ \Sigma \omega^\circ.$$

Recall from (19) that H_i^{-1} describes how a marginal change in \mathcal{R} affects α_i in the problem of firm i . The approximation (36) captures the same forces. It shows that the deviation of α_i from α_i° depends, approximately, on the vector of risk-adjusted prices \mathcal{R} evaluated at the ideal shares α° . Intuitively, when it is costly for firms to deviate from α° , we can evaluate the equilibrium prices *as if* firms chose α° and use these prices to compute the sourcing decisions of the firm. By Lemma 9, these decisions provide a first-order approximation of the true equilibrium network.²⁴

Example: cascading link destruction

To illustrate what type of network adjustments the model can generate, we consider an example in which a small change in the volatility of a single sector can push multiple producers to sequentially switch to safer suppliers, creating a cascade of adjustments. Consider the economy depicted in Figure 2. As indicated by the arrows, firms in sectors 1 to 3 can source inputs from two potential suppliers. The model is parametrized such that the shares of these suppliers are local substitutes. Firms in sectors 4 to 7, in contrast, can only use labor in production.

Figure 2: Cascading impact of a change in 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 (69) in the appendix with $\kappa_{ij} = 0$ 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, $\Sigma = 0$. In the right figure $\Sigma = 0$ except $\Sigma_{44} = 1$. The risk aversion of the household is $\rho = 2$. $\beta_i = 1/n$ for all i .

When uncertainty about sector 4 is sufficiently low ($\Sigma_{44} \rightarrow 0$; left panel), sectors 1 to 3 rely, directly or indirectly, on sector 4 as a supplier. As Σ_{44} increases (right panel), firms in sector 3,

²⁴See Supplemental Appendix I for more details and for second-order approximations for key model quantities.

seeking a more stable supply of goods, switch to using good 7 as an input instead. But this change implies a higher risk-adjusted price for sector 3, which makes firms in sector 2 want to use good 6 in production instead of good 2. The same logic then applies to firms in sector 1. A change in the uncertainty of a single sector can thus lead to a cascading movement to safety that affects far-away sectors.

We can interpret this cascading network adjustment through the lens of Lemma 9. Differentiating the expression with respect to Σ_{44} yields

$$\frac{d\alpha_{ij}}{d\Sigma_{44}} = \bar{\kappa}^{-1} (\rho - 1) \omega_4^\circ \left(\underbrace{\left[\left(\hat{H}_i^\circ \right)^{-1} \right]_{jj} \mathcal{L}_{j4}^\circ}_{\text{direct effect of } \Sigma_{44} \text{ on } j} + \underbrace{\sum_{l \neq j} \left[\left(\hat{H}_i^\circ \right)^{-1} \right]_{jl} \mathcal{L}_{l4}^\circ}_{\text{indirect effect of } \Sigma_{44} \text{ through other suppliers } l \neq j} \right) + O(\bar{\kappa}^{-2}). \quad (37)$$

Equation (37) states that if a firm j relies on sector 4 as an input (either immediate or distant, such that $\mathcal{L}_{j4}^\circ > 0$), an increase in Σ_{44} makes j less attractive. This direct effect pushes α_{ij} down (recall that $[H_i^{-1}]_{jj} < 0$ by the concavity of a_i). There is also an indirect effect that operates through the second term in (37). If another sector $l \neq j$ also relies on 4 ($\mathcal{L}_{l4}^\circ > 0$), then an increase in Σ_{44} makes l less attractive as well. This indirect channel can lead to either a decrease or an increase in α_{ij} , depending on whether j and l are complements or substitutes in the production of i ; that is, whether $[(H_i^\circ)^{-1}]_{jl}$ is negative or positive.

Sector 1, for instance, has two potential suppliers, sectors 2 and 5, with associated shares α_{12} and α_{15} . The direct effect of an increase in uncertainty Σ_{44} on α_{12} is strongly negative since sector 2 relies heavily on 4 (large \mathcal{L}_{24}°). The indirect effect through sector 5 is however zero since sector 5 does not rely on sector 4 in production ($\mathcal{L}_{54}^\circ = 0$). Furthermore, the contribution through the indirect effect of all other sectors is also zero since sector 1 never uses them in production and hence $[(H_1^\circ)^{-1}]_{2l} = 0$ for $l \neq 2$ and $l \neq 5$. It follows that (37) predicts a decline in α_{12} , and this is indeed what we see in Figure 2.

Instead, if we consider the response of α_{15} , the direct effect is absent because sector 5 does not rely on sector 4 ($\mathcal{L}_{54}^\circ = 0$). Since sector 2 is 1's only other possible connection, only the indirect effect through that sector remains. The relevant term here is $[(H_1^\circ)^{-1}]_{52} \mathcal{L}_{24}^\circ$, which is positive because $\mathcal{L}_{24}^\circ > 0$, and the shares of goods 5 and 2 are substitutes in the production of good 1, $[(H_1^\circ)^{-1}]_{52} > 0$. Therefore, an increase in Σ_{44} leads to a larger α_{15} . The same logic applies to the responses of firms 2 and 3, thus explaining the cascading effect illustrated in Figure 2.²⁵

²⁵Lemma 9 assumes that $\alpha \in \text{int } \mathcal{A}$, which is not the case in Figure 2, but it still captures the main forces that push the shares in response to changes in (μ, Σ) and is therefore informative about the response of the network.

7 Implications for GDP and welfare

Above, we analyzed how the production network responds to changes in beliefs (μ, Σ) , but what ultimately matters for welfare is the level and the variance of GDP. In this section, we describe how these objects are affected by changes in (μ, Σ) when the network is endogenous.

7.1 Beliefs and welfare

The next result compares how beliefs affect our measure of welfare, defined in (21), under a flexible and a fixed network.

Proposition 5. *Let γ denote either the mean μ_i or an element of the covariance matrix Σ_{ij} . Under an endogenous network, welfare responds to a marginal change in γ as if the network were fixed at its equilibrium value α^* , that is*

$$\frac{d\mathcal{W}(\mu, \Sigma)}{d\gamma} = \frac{\partial W(\alpha^*, \mu, \Sigma)}{\partial \gamma}.$$

This proposition is a direct consequence of the envelope theorem: Since the equilibrium network is welfare-maximizing, any marginal movement around that network must have no impact on welfare. It follows that as beliefs change, their impact on the production network does not affect welfare at the margin.

While this proposition shows that the flexibility of the network plays no role for the response of welfare to a *marginal* change in beliefs, this is generally not true for non-infinitesimal changes. In that case, shifts in (μ, Σ) that are beneficial to welfare are amplified, compared to the fixed-network benchmark, while changes that are harmful are dampened (see Proposition 2). Indeed, if we denote by $\alpha^*(\mu, \Sigma)$ the equilibrium production network under (μ, Σ) and by $W(\alpha, \mu, \Sigma)$ welfare under a network α , we can write that the difference in welfare after a change in beliefs from (μ, Σ) to (μ', Σ') satisfies the inequality

$$\underbrace{\mathcal{W}(\mu', \Sigma') - \mathcal{W}(\mu, \Sigma)}_{\text{Change in welfare under a flexible network}} \geq \underbrace{W(\alpha^*(\mu, \Sigma), \mu', \Sigma') - W(\alpha^*(\mu, \Sigma), \mu, \Sigma)}_{\text{Change in welfare under a fixed network}}. \quad (38)$$

This result follows directly from the fact that a flexible network provides an extra margin of adjustment to the planner and thus cannot leave the household worse off than under a fixed network.²⁶

We can also use Proposition 5 to show that the impact of a change in (μ, Σ) on \mathcal{W} is completely determined by the equilibrium Domar weights and the coefficient of relative risk aversion ρ .

²⁶We provide a proof of this result in Supplemental Appendix C.5.

Corollary 4. *The impact of an increase in μ_i on welfare is given by*

$$\frac{d\mathcal{W}}{d\mu_i} = \omega_i, \quad (39)$$

and the impact of an increase in Σ_{ij} on welfare is given by

$$\frac{d\mathcal{W}}{d\Sigma_{ij}} = -\frac{1}{2}(\rho - 1)\omega_i\omega_j. \quad (40)$$

This proposition follows directly from Corollary 1 and Proposition 5. Its first part provides a Hulten-like result for welfare in an endogenous network economy: Equation (39) states that the impact of an increase in μ_i on welfare is equal to the Domar weight ω_i of the affected sector. Since Domar weights are positive, increasing μ_i always has a positive impact on welfare. The second part of the proposition provides a similar result for an increase in uncertainty or covariance. In this case, the impact of the change is proportional to the product of the relevant Domar weights, and an increase in Σ_{ij} lowers welfare when $\rho > 1$. Intuitively, with a higher Σ_{ii} , the economy features more uncertainty which the household dislikes. Similarly, when sectoral shocks are more positively correlated, they offset each other less, such that the volatility of consumption increases and welfare falls.

7.2 Beliefs and GDP

Under an endogenous network, changes in beliefs also affect GDP through their impact on the production network. In this section, we analyze this link explicitly, starting with a general result that describes how GDP reacts to the presence of uncertainty.

Proposition 6. *The presence of uncertainty lowers expected log GDP, in the sense that $E[y]$ is largest when $\Sigma = 0$.*

This proposition follows directly from Lemma 3. Without uncertainty ($\Sigma = 0$), the variance $V[y]$ of log GDP is zero for all networks $\alpha \in \mathcal{A}$. The social planner then maximizes $E[y]$ only. When, instead, the productivity vector ε is uncertain ($\Sigma \neq 0$), the planner also seeks to lower $V[y]$ which necessarily lowers expected log GDP in equilibrium.

Proposition 6 establishes a novel mechanism through which uncertainty reduces expected log GDP. To understand that mechanism, 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 be the most productive ones, and toward techniques with higher TFP. As a result, the aggregate economy is maximally productive, and $E[y]$ is large. When some suppliers become risky, customers worry about a possible increase in input costs and 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 6. Indeed, in our model uncertainty affects expected log GDP *only* through the endogenous response of the network. If the shares α were fixed, uncertainty would have no impact on $E[y]$.

Response of GDP to a marginal change in beliefs

The previous proposition states that expected GDP is maximized in the absence of any uncertainty, but we can also consider the impact of a marginal change in beliefs on the moments of GDP. To do so, we first provide a result that connects the responses of $E[y]$ and $V[y]$ under an endogenous network to their counterparts under a fixed network.

Corollary 5. *Let γ denote either the mean μ_i or an element of the covariance matrix Σ_{ij} . The equilibrium response to a change in beliefs γ must satisfy*

$$\underbrace{\frac{dE[y]}{d\gamma} - \frac{\partial E[y]}{\partial \gamma}}_{\text{Excess response of } E[y]} = \frac{1}{2}(\rho - 1) \underbrace{\left(\frac{dV[y]}{d\gamma} - \frac{\partial V[y]}{\partial \gamma} \right)}_{\text{Excess response of } V[y]}. \quad (41)$$

The left-hand side of (41) is the response of $E[y]$ to the change in γ in the flexible-network economy (full derivatives) in excess of its fixed-economy response (partial derivatives). The right-hand side involves the same quantity for $V[y]$. Corollary 5 is a direct consequence of Proposition 5. Since the response of welfare to a marginal change in beliefs must be the same under a flexible and a fixed network, a larger increase in $E[y]$ under a flexible network must come at the cost of a larger increase in the variance $V[y]$. This fundamental tension between $E[y]$ and $V[y]$ comes from the fact that the equilibrium network was efficient before the change in the productivity process and already optimally traded off increasing $E[y]$ against reducing $V[y]$.

We now turn to a key result, which describes how GDP responds to marginal changes in beliefs.

Proposition 7. *If $\omega \in \text{int } \mathcal{O}$, the following holds.*

1. *The impact of an increase in μ_i on log GDP is given by*

$$\frac{dE[y]}{d\mu_i} = \underbrace{\omega_i}_{\text{Fixed network}} - (\rho - 1)\omega^\top \Sigma \mathcal{H}^{-1} \frac{\partial \mathcal{E}}{\partial \mu_i}, \quad \text{and} \quad \frac{dV[y]}{d\mu_i} = \underbrace{0}_{\text{Fixed network}} - 2\omega^\top \Sigma \mathcal{H}^{-1} \frac{\partial \mathcal{E}}{\partial \mu_i}.$$

2. *The impact of an increase in Σ_{ij} on log GDP is given by*

$$\frac{dE[y]}{d\Sigma_{ij}} = \underbrace{0}_{\text{Fixed network}} - (\rho - 1)\omega^\top \Sigma \mathcal{H}^{-1} \frac{\partial \mathcal{E}}{\partial \Sigma_{ij}}, \quad \text{and} \quad \frac{dV[y]}{d\Sigma_{ij}} = \underbrace{\omega_i \omega_j}_{\text{Fixed network}} - 2\omega^\top \Sigma \mathcal{H}^{-1} \frac{\partial \mathcal{E}}{\partial \Sigma_{ij}}.$$

The first part of Proposition 7 describes how log GDP responds to an increase in μ_i . On impact, sector i becomes more productive, which has a *direct* effect of ω_i on $E[y]$. This is the standard

Hulten's theorem effect that occurs when the network is kept fixed (Corollary 1). When the network is flexible, a reorganization also occurs to take advantage of the new μ . Corollary 5 implies that this excess response of $E[y]$ can be computed from the excess response of $V[y]$, such that

$$\frac{dE[y]}{d\mu_i} - \frac{\partial E[y]}{\partial \mu_i} \propto \underbrace{\frac{dV[y]}{d\omega^\top}}_{\frac{dV[y]}{d\omega^\top}} \times \underbrace{\left(-\mathcal{H}^{-1} \frac{\partial \mathcal{E}}{\partial \mu_i}\right)}_{\frac{d\omega}{d\mu_i}} - 0,$$

where we used (14) and Proposition 3 to compute $dV[y]/d\mu_i$, and where $\partial V[y]/\partial \mu_i = 0$ by Corollary 1. It follows that the response of the moments of log GDP to a change in μ_i depends on how that change affects the Domar weights $(-\mathcal{H}^{-1} \partial \mathcal{E} / \partial \mu_i)$ and on how that movement in Domar weights influences the variance of log GDP ($2\omega^\top \Sigma$).

A similar reasoning applies for changes in Σ_{ij} (point 2 of the proposition). On impact, a higher Σ_{ij} leads to an increase in $V[y]$ by the fixed-network term $\omega_i \omega_j$, and the ensuing reorganization of the network can amplify or dampen that direct effect. If $V[y]$ increases by more than $\omega_i \omega_j$, welfare maximization implies that $E[y]$ must also increase, as the result shows.

The role of risk and of the global substitution patterns

For a given equilibrium, one can compute the expressions in Proposition 7 to fully characterize how GDP would respond to a change in beliefs. This response, in turn, depends on the risk structure Σ of the economy and on the global substitution patterns embedded in \mathcal{H}^{-1} . We now explore these two channels more thoroughly.

We can readily characterize the impact of beliefs when there is no uncertainty.

Corollary 6. *Without uncertainty ($\Sigma = 0$) the moments of GDP respond to changes in beliefs as if the network were fixed, such that*

$$\frac{dE[y]}{d\mu_i} = \frac{\partial E[y]}{\partial \mu_i} = \omega_i, \quad \text{and} \quad \frac{dV[y]}{d\Sigma_{ij}} = \frac{\partial V[y]}{\partial \Sigma_{ij}} = \omega_i \omega_j.$$

When $\Sigma = 0$, the Domar weights are sufficient to characterize the behavior of GDP, even though the production network is flexible and can respond to changes in beliefs. It follows that uncertainty is essential for the economy to depart from Hulten's theorem. Intuitively, without uncertainty, the network maximizes expected log GDP, such that at an interior equilibrium $dE[y]/d\alpha = 0$. It follows that even if the network responds to a marginal change in beliefs, this reorganization has no impact on $E[y]$. Corollary 6 shows that this logic also applies when the equilibrium is not interior.

In contrast, when there is uncertainty, whether a change in beliefs amplifies or dampens the fixed-network effect depends crucially on the global substitution patterns embedded in \mathcal{H}^{-1} . The next result describes what happens when sectors are global complements.

Corollary 7. *Suppose that $\omega \in \text{int } \mathcal{O}$. There exists a threshold $\bar{\Sigma} < 0$ such that if $\Sigma_{kl} > \bar{\Sigma}$ for all k, l , then the following holds.*

1. *If all sectors are global complements with sector i , that is $\mathcal{H}_{ik}^{-1} < 0$ for $k \neq i$, then*

$$\frac{dE[y]}{d\mu_i} > \omega_i, \quad \text{and} \quad \frac{dV[y]}{d\mu_i} > 0.$$

2. *If all sectors are global complements with sectors i and j , that is $\mathcal{H}_{ik}^{-1} < 0$ and $\mathcal{H}_{jk}^{-1} < 0$ for $k \neq i, j$, then*

$$\frac{dE[y]}{d\Sigma_{ij}} < 0, \quad \text{and} \quad \frac{dV[y]}{d\Sigma_{ij}} < \omega_i \omega_j.$$

The first part of the corollary shows that under global complementarities expected log GDP responds to expected TFP by more than when the network is fixed. Effectively, the network is reorganized to amplify the positive impact of the change in beliefs on $E[y]$. Intuitively, after the increase in μ_i the Domar weight of sector i increases (Proposition 2). Because of the global complementarities, this causes all the other Domar weights to rise as well (Corollary 3). As long as the covariances Σ_{ij} are not too negative, this simultaneous increase in Domar weights pushes the variance of log GDP up. From Proposition 7 it then follows that $E[y]$ increases by more than ω_i . A similar mechanism explains the impact of a change in Σ_{ii} and Σ_{ij} on the moments of GDP, but in this case the economy responds by less than predicted by Hulten's theorem.

We can also explore how GDP responds to changes in beliefs under global substitutabilities.

Corollary 8. *Suppose that $\omega \in \text{int } \mathcal{O}$. Then there exist thresholds $\underline{\Sigma} > 0$ and $\bar{\Sigma} > 0$ such that,*

1. *If all sectors are global substitutes with sector i , that is $\mathcal{H}_{ik}^{-1} > 0$ for $k \neq i$, and sector i is not too risky while other sectors are sufficiently risky in the sense that $\Sigma_{ji} < \underline{\Sigma}$ for all j and $\Sigma_{jk} > \bar{\Sigma}$ for all $j, k \neq i$, then*

$$\frac{dE[y]}{d\mu_i} < \omega_i, \quad \text{and} \quad \frac{dV[y]}{d\mu_i} < 0.$$

2. *If all sectors are global substitutes with sectors i and j , that is $\mathcal{H}_{ik}^{-1} > 0$ and $\mathcal{H}_{jk}^{-1} > 0$ for $k \neq i, j$, and sectors i and j are not too risky while other sectors are sufficiently risky in the sense that $\Sigma_{li} < \underline{\Sigma}$ and $\Sigma_{lj} < \underline{\Sigma}$ for all l , and $\Sigma_{lk} > \bar{\Sigma}$ for all $l, k \neq i$ and $l, k \neq j$, then*

$$\frac{dE[y]}{d\Sigma_{ij}} > 0, \quad \text{and} \quad \frac{dV[y]}{d\Sigma_{ij}} > \omega_i \omega_j.$$

After an increase in μ_i , the Domar weight of sector i increases (Proposition 2) which pushes $V[y]$ up, but if Σ_{ii} is small, this increase in $V[y]$ is also small. Because other sectors are global substitutes

with i , the increase in ω_i leads to a decline in all the other Domar weights. If the variances of those sectors are large relative to Σ_{ii} , this decline in Domar weights leads to a substantial decrease in $V[y]$. By the logic of Proposition 7, this implies that $E[y]$ must increase by less than its fixed-network term ω_i . Through a similar mechanism, an increase in Σ_{ii} leads to an increase in $V[y]$ that is larger than under a fixed network. In this case, $E[y]$ *increases* in response to the higher Σ_{ii} , such that uncertainty can be beneficial to expected log GDP at the margin.²⁷

Counterintuitive implications of changes in beliefs

Corollaries 7 and 8 establish sufficient conditions under which the response of GDP to beliefs can be larger or smaller than predicted by Hulten’s theorem in the fixed-network economy. But the endogenous adjustment of the network can also have more extreme consequences: In some cases, an increase in μ can lead to a *decline* in $E[y]$ and an increase in Σ can lead to a *decline* in $V[y]$. To understand why, consider a producer with (on average) low but stable productivity. The high price of its good 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 expected GDP might fall as a result. A similar mechanism implies that an increase in the volatility of a sector’s productivity can lead to a decline in $V[y]$. In what follows, we provide an example that explicitly illustrates how such counterintuitive effects may arise.

In the economy depicted in Figure 3, sectors 4 and 5 use only labor to produce, while sectors 1 to 3 can also use goods 4 and 5 as inputs. The local TFP shifter functions are such that for $i \in \{1, 2, 3\}$ the shares of goods 4 and 5 are either local substitutes with $[H_i^{-1}]_{45} > 0$ in panels (a) to (c), or local complements with $[H_i^{-1}]_{45} < 0$ in panels (d) to (f). Sector 4 is more productive and volatile than sector 5 ($\mu_4 > \mu_5$ and $\Sigma_{44} > \Sigma_{55}$).

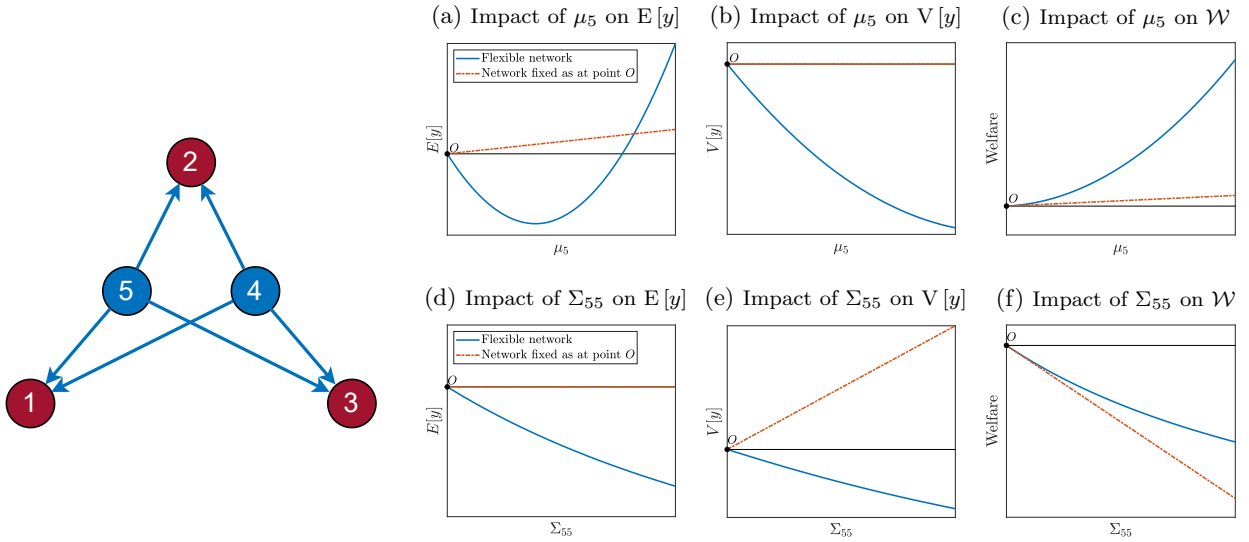
Consider the impact of a positive shock to μ_5 when inputs 4 and 5 are substitutes. The solid blue lines in panels (a) to (c) illustrate the impact of this change, and point O represents the economy before the change. As we can see, the initial *increase* in μ_5 has a *negative* impact on expected log GDP. To understand why, notice that for a small increase in μ_5 , sector 5 is still less productive (in expectation) than sector 4, but it now offers a better risk-reward trade-off. As a result, sectors 1 to 3 increase their shares of good 5 and, since 4 and 5 are substitutes, reduce their shares of good 4. But since $\mu_4 > \mu_5$ this readjustment leads to a fall in $E[y]$ for a small increase in μ_5 . At the same time, $V[y]$ also declines because sector 5 is less volatile than sector 4, in line with Proposition 7. The implied changes in $E[y]$ and $V[y]$ thus have opposite impacts on welfare. By Corollary 4, the overall effect on welfare must be positive though, and this is indeed confirmed in panel (c). Naturally, as μ_5 keeps increasing, $E[y]$ eventually starts to increase as well.

²⁷This does not contradict Proposition 6 as Corollary 8 only applies at the margin when $\Sigma_{jk} > \bar{\Sigma} > 0$ for all $j, k \neq i$. Eliminating uncertainty altogether would still lead to an increase in $E[y]$.

To emphasize the role of the endogenous network for this mechanism, Figure 3 also shows the effect of the same increase in μ_5 when the network is kept fixed (dashed red lines). 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, $V[y]$ 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 adjusts precisely to maximize the beneficial impact of the change in beliefs on welfare, as implied by (38).

We can use a small variation of this economy to illustrate how an *increase* in an element of Σ can *lower* the variance of log GDP, and simultaneously lower welfare. Start again from the economy in the left column of Figure 3 (point O) but suppose that inputs 4 and 5 are complements in the production of goods 1 to 3. Consider an increase in the volatility of sector 5. In response, sectors 1 to 3 start to rely less on sector 5. But since inputs 4 and 5 are complements, sectors 1 to 3 also reduce their shares of input 4, thus increasing the overall share of labor which is a safe input. As a result, the variance of log GDP declines (panel e). Expected log GDP also goes down by Proposition 7 (panel d). The combined effect on welfare is negative, as predicted by Corollary 4 (panel f). In this case, the reorganization of the network mitigates the adverse effect of the increase in volatility on welfare. Instead, if the network is fixed, an increase in Σ_{55} does not affect expected log GDP but leads to an increase in the variance of log GDP. As a result, welfare drops substantially more than under an endogenous network, as implied by (38).

Figure 3: The non-monotone impact of beliefs on GDP



Notes. There is an arrow from j to i if $\alpha_{ij} > 0$. Household: $\rho = 2.5$ and $\beta_1 = \beta_2 = \beta_3 = \frac{1}{3} - \epsilon$, $\beta_4 = \beta_5 = \frac{2}{3}\epsilon$, where $\epsilon > 0$ is very small. $\mu = (0.1, 0.1, 0.1, 0.1, -0.08)$, Σ is diagonal, with $\text{diag}(\Sigma) = (0.2, 0.2, 0.2, 0.2, 0.02)$. a is as in (2) with $\alpha_{14}^o = \alpha_{15}^o = \alpha_{24}^o = \alpha_{25}^o = \alpha_{34}^o = \alpha_{35}^o = 0.25$; all other α_{ij}^o are zero. $H_4 = H_5$ are matrices with -50 on the diagonal. $H_1 = H_2 = H_3$ with $[H_1]_{11} = [H_1]_{22} = [H_1]_{33} = -50$, $[H_1]_{44} = [H_1]_{55} = -2$. In panels (a)-(c), μ_5 goes from -0.08 to 0.1 ; 4 and 5 are substitutes, $[H_1]_{45} = -1.9$. In panels (d)-(f), Σ_{55} goes from 0.02 to 0.2 ; 4 and 5 are complements, $[H_1]_{45} = 1.9$.

8 A basic calibration of the model

The analysis above highlights the economic forces that determine how the production network, GDP and welfare respond to changes in the productivity process. Clearly, the model is too stylized to capture all the fluctuations in the production network observed in reality, and other mechanisms, not present in our model, may also be important in practice. With that caveat in mind, we present in this section results from a basic calibration of the model to the United States economy to get a sense of the quantitative potential of our main mechanisms.

Below, we first describe how the model is parameterized and briefly go over which features of the US economy the model matches well, and in what dimensions it falls short. Finally, we explore how beliefs shape the production network and investigate how the changing structure of the network influences aggregate output and welfare in our stylized model. We keep the analysis succinct but provide more details in Appendix B.

8.1 Parametrization

The Bureau of Economic Analysis (BEA) provides U.S. sectoral input-output tables for $n = 37$ sectors at an annual frequency from 1948 to 2020. From these data, we compute the input shares α_{ijt} of each sector in each year t , the average consumption expenditure share of each sector β_i , and sectoral TFP measured as the Solow residual.

To calibrate the model, we need to make explicit assumptions about the process for TFP. For the endogenous productivity shifter $A_i(\alpha_{it})$ we adopt a particular version of form (2) which includes a diagonal component for H_i are a penalty for deviating from an ideal labor share (see (69) in the appendix). We set the ideal shares $(\alpha_1^o, \dots, \alpha_n^o)$ equal to the time average of the input shares observed in the data. The exogenous sectoral productivity process ε_t is assumed to follow a random walk with drift,

$$\varepsilon_t = \gamma + \varepsilon_{t-1} + u_t, \quad (42)$$

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 further assume that firms know γ and ε_{t-1} at time t , so that the conditional mean and the covariance of beliefs are given by $\mu_t = \gamma + \varepsilon_{t-1}$ and Σ_t . Importantly, we allow uncertainty Σ_t to vary over time and estimate it from TFP data using a rolling window that puts more weight on more recent observations.

We use a simple moment-matching strategy to pin down the 1) relative risk aversion parameter ρ of the household, 2) the TFP shifter functions H_i and 3) the time-varying beliefs (μ_t, Σ_t) . We describe this procedure in Appendix B.

The calibrated coefficient of relative risk aversion $\hat{\rho}$ is 4.3, which is similar to values used or estimated in the macroeconomics literature. Our procedure also provides time-series for the vector

μ_t and the matrix Σ_t , and we aggregate these variables across sectors to obtain economy-wide measures of the expected value $\bar{\mu}_t$ and the variance $\bar{\Sigma}_t$ of aggregate TFP. As we might expect, these measures are cyclical, with $\bar{\mu}_t$ falling and $\bar{\Sigma}_t$ rising during recessions. Overall, our measure of aggregate uncertainty $\bar{\Sigma}_t$ has been relatively stable since 1980, with occasional sharp spikes, most notably during the Great Recession of 2007–2009 (see Figure 5 in Appendix B.3).

We next assess how well the calibrated model fits key moments in the data. As we have seen above, the Domar weights, and how they react to changes in μ_t and Σ_t , are central for the mechanisms of the model. The model is able to roughly replicate features of the empirical Domar weights, with a cross-sectional correlation between the time-averaged Domar weights in the model and in the data of 0.96. However, the average Domar weight in the model (0.03) is lower than its data counterpart (0.05).²⁸ Overall, the model can account for about 40% of the over-time standard deviation of Domar weights, which indicates that other mechanisms, such as technological progress that might expand the set of available techniques, might be at work in reality.

The mechanisms of the model predict that a decline in the expected productivity of a sector μ_i , or an increase in its variance Σ_{ii} , should push firms to reduce the importance of that sector as an input provider, leading to a decline in its Domar weight. Reassuringly, these correlations are visible in the data, where $\text{Corr}(\omega_{jt}, \mu_{jt}) = 0.1$, and $\text{Corr}(\omega_{jt}, \Sigma_{jjt}) = -0.4$. The calibrated model is also able to roughly match these correlations, and the corresponding numbers are 0.1 and -0.3 .

8.2 The production network, welfare and output

To evaluate the quantitative potential 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 model to an economy in which the network is kept completely fixed at its sample average. 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 1) an economy in which production techniques are chosen *as if* $\Sigma_t = 0$,²⁹ and 2) a perfect-foresight economy in which firms observe the realization of ε_t before making technique choices (the “known ε_t ” economy).³⁰ In both cases, uncertainty is irrelevant for decisions, and so these exercises allow us to isolate the impact of uncertainty on the production network and, through that channel, on macroeconomic aggregates.

We find that expected log GDP in the “fixed network” economy is 2.1% lower than in our baseline calibration with a flexible network. Intuitively, as some sectors become more productive

²⁸We explain in Appendix B that this discrepancy can be explained by our choice to target consumption growth instead of GDP growth in the estimation.

²⁹Specifically, we set $\Sigma = 0$ when solving the problem of the social planner (21) for the equilibrium network α^* . We then reintroduce uncertainty when computing the moments of GDP and welfare.

³⁰One interpretation 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.

over time, the goods that they produce become cheaper, and firms would like to rely more on them. With a flexible network this is possible, and the aggregate economy becomes more productive as a result. The difference in welfare between the two models is about 2.1% as well.

When we isolate the role of uncertainty, however, these numbers become smaller. In line with the theory, the baseline economy is on average less productive and less volatile than under the “as if $\Sigma_t = 0$ ” alternative but the numbers are small, on the order of 0.01% for $E[y]$ and 0.10% for $V[y]$. This suggests that, for most of the sample period, uncertainty is sufficiently low that firms simply buy their inputs from the most productive suppliers without much concern for any risk involved.³¹

The differences between our calibrated economy and the “no uncertainty” alternatives are however larger during high-uncertainty episodes like the Great Recession.³² The top row of Figure 4 shows that expected log GDP in the baseline economy is about 0.25% lower in 2009 than in the alternative “as if $\Sigma_t = 0$ ” economy. Because of the large increase in uncertainty, firms adjust their production techniques toward safer but less productive suppliers to avoid potentially large increases in costs. The result in terms of aggregate volatility is visible in the top-right panel, where we see that log GDP is about 2.4% less volatile in 2009 in the baseline economy. Interestingly, realized log GDP, shown in the left-bottom panel, is substantially higher in the baseline economy than in the “as if $\Sigma_t = 0$ ” alternative. Essentially, firms took out an insurance against particularly bad TFP draws and opted for safer suppliers. When these fears were realized, this insurance policy paid off so that the baseline economy fared about 2.7% better in terms of realized log GDP compared to the alternative.

The right-bottom panel provides the same information for the “known ε_t ” alternative. In this case, 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 realized consumption (or GDP) is always larger than in the baseline model. 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.³³

Overall, our findings suggest that, while uncertainty might have a limited impact on the economy on average, it may play a larger role in shaping the production network during high-uncertainty periods, with consequences for expected and realized GDP, as well as for welfare. Given the stylized nature of the model, these findings should be interpreted with caution. The model abstracts from other forces that might affect the production network, such as changes in demand and technological

³¹As in Lucas (1987), the utility cost of business cycles 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 its volatility. We provide the same moments for the “known ε_t ” economy in Supplemental Appendix B.4.


³²The differences between our calibrated and fixed-network economies are also particularly large during volatile periods, when adjusting the network is most beneficial. In Supplemental Appendix J.3, we show that allowing the network to adjust leads to large gains in expected GDP during the Great Recession.

³³Since ε_t is known in this exercise, $E[y] = \mathcal{W} = y$ and $V[y] = 0$ and, so we do not report these moments in Figure 4. Alternatively, one can compute $E[y]$, $V[y]$ and \mathcal{W} before ε_t is known but still assuming that the production network is chosen optimally for the given realized draw of ε_t . We report these moments in Appendix B.4.


progress that would expand the set of production techniques. Similarly, the production function might not be Cobb-Douglas in reality, in which case changes in prices would affect Domar weights. We also made the implicit assumption that it takes one year (the frequency of our data) for firms to change production techniques. While this assumption might be reasonable for some sectors, it is likely that the time it takes to retool a factory varies significantly by industry, or even depending on what the new and the old techniques are.³⁴ While we believe that the mechanisms that we explore in this paper would still be present in a richer model, more work would be needed to fully assess their importance.

Figure 4: The role of uncertainty in the postwar period

First row: “as if $\Sigma_t = 0$ ” as the alternative




(a) Difference in expected log GDP [%]




(b) Difference in expected st. dev. of log GDP [%]

Left column: “as if $\Sigma_t = 0$ ” as alternative



(c) Difference in realized log GDP [%]

Right column: “known ε_t ” as alternative



(d) Difference in realized log GDP [%]

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 (panels (a) to (c)) and the “known ε_t ” alternative (panel (d)). 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 $\sqrt{V[y]}$ are evaluated before ε_t is realized.

9 Model-free evidence for the mechanisms

The model proposed in this paper relies on simplifying assumptions for tractability. In this section, we present additional evidence in support of the main mechanisms of the model that does not rely on this structure. Through firm-level regressions that closely follow [Alfaro, Bloom, and Lin \(2019\)](#) we document that 1) higher uncertainty about a firm leads to a decline in its 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. Supplemental Appendix [G](#) describes the

³⁴In the car industry, General Motors took about one year to retool a factory for electric vehicle production ([Lutz, 2021](#)), but it took Ford eight weeks to switch from using steel to aluminum for the body of the F-150 ([Dean, 2015](#)).

data and the instruments in detail.

9.1 Uncertainty and Domar weights

We first test the model’s prediction that Domar weights decrease with 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 1. In column (2), we follow [Alfaro et al. \(2019\)](#) 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.

Table 1: Domar weights and uncertainty

	Change in Domar weight		
	(1): OLS	(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

Notes: Table presents OLS and 2SLS annual regression results of firm-level volatility. The dependent variable is the growth rate in Domar weight. Supplier $\Delta\text{Volatility}_{i,t-1}$ 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 from 1st moment aggregate shocks we also include as controls the lagged directional 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 $_{i,t-1}$.” See [Alfaro et al. \(2019\)](#) for more details about the data and the construction of the instruments. All specifications include year \times 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.

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.

9.2 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 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 2. As in the last exercise, column (2) uses industry-level 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.

Table 2: Link destruction and supplier volatility

	Dummy for last year of supply relationship		
	(1): OLS	(2): IV	(3): IV
$\Delta\text{Volatility}_{t-1}$ of supplier	0.026** (0.012)	0.097*** (0.035)	0.144** (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
F -statistic	—	22.9	10.39

Notes: Table presents OLS and 2SLS annual regression results of firm-level volatility. The dependent variable is a dummy variable that equals one in the last year of a supply relationship 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 $\Delta\text{Volatility}_{t-1}$ 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 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 from 1st moment aggregate shocks we also include as controls the lagged directional 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 \times customer \times supplier industry (2SIC) fixed effects. Standard errors (in parentheses) are two-way clustered at the customer and the supplier industry (3SIC) levels. F -statistics are Kleibergen-Paap. *, **, *** indicate significance at the 10%, 5%, and 1% levels, respectively.

10 Conclusion

We construct a model in which agents’ beliefs about productivity affect the structure of the production network and, through that channel, macroeconomic aggregates such as output and welfare. We prove that the unique equilibrium is efficient, and that it is characterized by a trade-off between the expected level and the volatility of GDP. We also prove that the presence of

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 potentially have a sizable effect on GDP and welfare during periods of high uncertainty such as the Great Recession.

The model is tractable and can serve as a framework to study various related questions. For instance, with adjustments, our closed economy model could be adapted 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 our analysis 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 could 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 Tōhoku 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.
- DAVID, J. M., L. SCHMID, AND D. ZEKE (2022): “Risk-Adjusted Capital Allocation and Misallocation,” *Journal of Financial Economics*, 145, 684–705.
- DEAN, J. (2015): “How Ford’s largest truck factory was completely overhauled in 8 weeks,” *Popular Mechanics*.
- 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-VILLAYERDE, 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.
- GROSSMAN, G. M., E. HELPMAN, AND A. SABAL (2023): “Resilience in Vertical Supply Chains,” Tech. rep., National Bureau of Economic Research.
- 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.
- JUDD, K. L. AND S.-M. GUU (2001): “Asymptotic methods for asset market equilibrium analysis,” *Economic theory*, 18, 127–157.
- 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. (2023): “Booms, Busts, and Common Risk Exposures,” *Journal of Finance*, 78, 3299–3341.
- 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.
- LUTZ, H. (2021): “Retooling plants for evs will save GM Billions,” *Automotive News*.
- MAS-COLELL, A., M. D. WHINSTON, AND J. R. GREEN (1995): *Microeconomic Theory*, Oxford University Press.
- 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.
- SCHMITT-GROHÉ, S. AND M. URIBE (2004): “Solving dynamic general equilibrium models using a second-order approximation to the policy function,” *Journal of economic dynamics and control*, 28, 755–775.
- TASCHEREAU-DUMOUCHEL, M. (2020): “Cascades and Fluctuations in an Economy with an Endogenous Production Network,” Working paper.
- VOM LEHN, C. AND T. WINBERRY (2022): “The investment network, sectoral comovement, and the changing US business cycle,” *The Quarterly Journal of Economics*, 137, 387–433.
- 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 Proofs

A.1 Proof of Lemma 1

Lemma 1. *Under a given network α , the vector of log prices is given by*

$$p(\alpha) = -\mathcal{L}(\alpha)(\varepsilon + a(\alpha)), \quad (12)$$

and log GDP is given by

$$y(\alpha) = \omega(\alpha)^\top (\varepsilon + a(\alpha)), \quad (13)$$

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

Proof. Combining the unit cost equation (8) with the equilibrium condition (10) and taking the log we find that, for all i ,

$$p_i = -\varepsilon_i - a_i(\alpha_i) + \sum_{j=1}^n \alpha_{ij} p_j, \quad (43)$$

where $a_i(\alpha_i) = \log(A_i(\alpha_i))$. This is a system of linear equations whose solution is (12). The log price vector is also normally distributed since it is a linear transformation of normal random variable. Combining with (6) yields (13). \square

A.2 Proof of Corollary 1

The proof of Corollary 1 is Supplemental Appendix D.

A.3 Proof of Lemma 2

Lemma 2. *In equilibrium, the technique choice problem of the representative firm in sector i is*

$$\alpha_i^* \in \arg \max_{\alpha_i \in \mathcal{A}_i} a_i(\alpha_i) - \sum_{j=1}^n \alpha_{ij} \mathcal{R}_j(\alpha^*), \quad (17)$$

where

$$\mathcal{R}(\alpha^*) = \mathbb{E}[p(\alpha^*)] + \text{Cov}[p(\alpha^*), \lambda(\alpha^*)] \quad (18)$$

is the vector of equilibrium risk-adjusted prices, and where

$$\mathbb{E}[p(\alpha^*)] = -\mathcal{L}(\alpha^*)(\mu + a(\alpha^*)) \text{ and } \text{Cov}[p(\alpha^*), \lambda(\alpha^*)] = (\rho - 1) \mathcal{L}(\alpha^*) \Sigma [\mathcal{L}(\alpha^*)]^\top \beta.$$

Proof. We first consider the stochastic discount factor. Equation (76) in Appendix C.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

$$\lambda(\alpha^*) = -(1 - \rho) \sum_{i=1}^n \beta_i p_i(\alpha^*). \quad (44)$$

Taking the log of (8) yields

$$k_i(\alpha_i, \alpha^*) = -(\varepsilon_i + a(\alpha_i)) + \sum_{j=1}^n \alpha_{ij} p_j(\alpha^*). \quad (45)$$

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 9, we can write

$$\alpha_i^* \in \arg \min_{\alpha_i \in \mathcal{A}_i} \mathbb{E} \left[\Lambda \frac{\beta^\top \mathcal{L}(\alpha^*) \mathbf{1}_i}{P_i} K_i(\alpha_i, P) \right],$$

where we have used (78). We can drop $\beta^\top \mathcal{L}(\alpha^*) \mathbf{1}_i > 0$ since it is a deterministic scalar that does not depend on α_i . Rewriting this equation in terms of log quantities yields

$$\alpha_i^* \in \arg \min_{\alpha_i \in \mathcal{A}} \mathbb{E} \exp[\lambda(\alpha^*) - p_i(\alpha^*) + k_i(\alpha_i, \alpha^*)],$$

where we emphasize that λ and p_i depend only on the equilibrium technique choice α^* . The terms $\lambda(\alpha^*)$, $p_i(\alpha^*)$, and $k_i(\alpha_i, \alpha^*)$ are normally distributed. We can therefore use the expression for the expected value of a lognormal distribution and write

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

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

$$\alpha_i^* \in \arg \min_{\alpha_i \in \mathcal{A}} \mathbb{E}[k_i(\alpha_i, \alpha^*)] + \frac{1}{2} \mathbb{V}[\lambda(\alpha^*) - p_i(\alpha^*) + k_i(\alpha_i, \alpha^*)]. \quad (46)$$

This expression can be written as

$$\begin{aligned} \alpha_i^* \in \arg \min_{\alpha_i \in \mathcal{A}} & \mathbb{E}[k_i(\alpha_i, \alpha^*)] + \frac{1}{2} \mathbb{V}[\lambda(\alpha^*)] + \frac{1}{2} \mathbb{V}[k_i(\alpha_i, \alpha^*) - p_i(\alpha^*)] \\ & + \text{Cov}[\lambda(\alpha^*), k_i(\alpha_i, \alpha^*)] + \text{Cov}[\lambda(\alpha^*), -p_i(\alpha^*)], \end{aligned}$$

where we can drop $\mathbb{V}[\lambda(\alpha^*)]$ and $\text{Cov}[\lambda(\alpha^*), -p_i(\alpha^*)]$ as they do not affect the minimization.

Finally, we can expand $V[k_i(\alpha_i, \alpha^*) - p_i(\alpha^*)]$ to get

$$\begin{aligned} \alpha_i^* \in \arg \min_{\alpha_i \in \mathcal{A}} E[k_i(\alpha_i, \alpha^*)] + \frac{1}{2} E \left[(k_i(\alpha_i, \alpha^*) - p_i(\alpha^*) - E[k_i(\alpha_i, \alpha^*) - p_i(\alpha^*)])^2 \right] \\ + \text{Cov}[\lambda(\alpha^*), k_i(\alpha_i, \alpha^*)]. \end{aligned}$$

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

$$\begin{aligned} \frac{1}{2} E \left[2(k_i(\alpha_i, \alpha^*) - p_i(\alpha^*) - E[k_i(\alpha_i, \alpha^*) - p_i(\alpha^*)]) \left(\frac{dk_i(\alpha_i, \alpha^*)}{d\alpha_{ik}} - E \left[\frac{dk_i(\alpha_i, \alpha^*)}{d\alpha_{ik}} \right] \right) \right] \\ + E \left[\frac{dk_i(\alpha_i, \alpha^*)}{d\alpha_{ik}} \right] + \text{Cov} \left[\lambda(\alpha^*), \frac{dk_i(\alpha_i, \alpha^*)}{d\alpha_{ik}} \right] + \gamma_i - \chi_{ik} = 0, \end{aligned}$$

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

$$E \left[\frac{dk_i(\alpha_i^*, \alpha^*)}{d\alpha_{ik}} \right] + \text{Cov} \left[\lambda(\alpha^*), \frac{dk_i(\alpha_i^*, \alpha^*)}{d\alpha_{ik}} \right] + \gamma_i^* - \chi_{ik}^* = 0$$

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

$$\alpha_i^* \in \arg \min_{\alpha_i \in \mathcal{A}} E[k_i(\alpha_i, \alpha^*)] + \text{Cov}[\lambda(\alpha^*), k_i(\alpha_i, \alpha^*)].$$

Finally, note that

$$\begin{aligned} \arg \min_{\alpha_i \in \mathcal{A}} E[k_i(\alpha_i, \alpha^*)] + \text{Cov}[\lambda(\alpha^*), k_i(\alpha_i, \alpha^*)] &= \arg \min_{\alpha_i \in \mathcal{A}} -\mu_i - a_i(\alpha_i) + \sum_{j=1}^n \alpha_{ij} E[p_j] \\ &\quad + \text{Cov} \left[\lambda(\alpha^*), -\varepsilon_i - a_i(\alpha_i) + \sum_{j=1}^n \alpha_{ij} p_j \right] \\ &= \arg \min_{\alpha_i \in \mathcal{A}} -a_i(\alpha_i) + \sum_{j=1}^n \alpha_{ij} E[p_j] \\ &\quad + \text{Cov}[\lambda(\alpha^*), -\varepsilon_i] + \sum_{j=1}^n \alpha_{ij} \text{Cov}[\lambda(\alpha^*), p_j] \\ &= \arg \min_{\alpha_i \in \mathcal{A}} -a_i(\alpha_i) + \sum_{j=1}^n \alpha_{ij} \mathcal{R}_j(\alpha^*), \end{aligned}$$

which completes the proof. □

A.4 Proof of Lemma 3

Lemma 3. *An efficient production network α^* solves*

$$\mathcal{W} \equiv \max_{\alpha \in \mathcal{A}} W(\alpha, \mu, \Sigma),$$

where \mathcal{W} is a measure of the welfare of the household, and where

$$W(\alpha, \mu, \Sigma) \equiv \mathbb{E}[y(\alpha)] - \frac{1}{2}(\rho - 1) \mathbb{V}[y(\alpha)], \quad (21)$$

is welfare under a given network α .

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 (13). It follows that the efficient production network must solve

$$\begin{aligned} \max_{\alpha \in \mathcal{A}} \mathbb{E}[u(Y)] &= \max_{\alpha \in \mathcal{A}} \frac{1}{1 - \rho} \mathbb{E}[\exp((1 - \rho) \log Y)] \\ &= \max_{\alpha \in \mathcal{A}} \frac{1}{1 - \rho} \exp\left((1 - \rho) \mathbb{E}[\log Y] + \frac{1}{2}(1 - \rho)^2 \mathbb{V}[\log Y]\right) \\ &= \max_{\alpha \in \mathcal{A}} \mathbb{E}[\log Y] - \frac{1}{2}(\rho - 1) \mathbb{V}[\log Y] \end{aligned} \quad (47)$$

where we have used the fact that $\log Y$ is normally distributed. \square

A.5 Proof of Corollary 2

Corollary 2. *The efficient Domar weight vector ω^* solves*

$$\mathcal{W} = \max_{\omega \in \mathcal{O}} \underbrace{\omega^\top \mu + \bar{a}(\omega)}_{\mathbb{E}[y]} - \frac{1}{2}(\rho - 1) \underbrace{\omega^\top \Sigma \omega}_{\mathbb{V}[y]}, \quad (26)$$

where $\mathcal{O} = \{\omega \in \mathbb{R}_+^n : \omega \geq \beta \text{ and } 1 \geq \omega^\top (\mathbf{1} - \bar{\alpha})\}$ and $\bar{a}(\omega)$ is given by (23).

Proof. Using (14) and the definition of Domar weights, the original planning problem (21) can be written

$$\mathcal{W} = \max_{\alpha \in \mathcal{A}, \omega} \omega^\top \mu + \omega^\top a(\alpha) - \frac{1}{2}(\rho - 1) \omega^\top \Sigma \omega,$$

subject to $\beta^\top (I - \alpha)^{-1} = \omega^\top$, which is equivalent to

$$\max_{\omega \in \mathcal{O}} \omega^\top \mu + \left[\max_{\alpha \in \mathcal{A}} \omega^\top a(\alpha) \right] - \frac{1}{2}(\rho - 1) \omega^\top \Sigma \omega,$$

where the inner problem is subject to $\beta^\top (I - \alpha)^{-1} = \omega^\top$ and where we can limit the feasible set of the outside problem to \mathcal{O} since for $\omega \notin \mathcal{O}$ the inner constraints could never be satisfied. This last problem is the same as (26) because, by (23), $\bar{a}(\omega) = \max_{\alpha \in \mathcal{A}} \omega^\top a(\alpha)$ subject to $\beta^\top (I - \alpha)^{-1} = \omega^\top$. \square

A.6 Proof of Lemma 4

Lemma 4. *The objective function of the planner's problem (26) is strictly concave. Furthermore, there is a unique vector of Domar weights ω^* that solves that problem, and there is a unique production network $\alpha(\omega^*)$ associated with that solution.*

Proof. We first show that the value function $\bar{a}(\omega)$ defined by (23) is strictly concave. Consider the maximization problem

$$\bar{a}(\omega) = \max_{\alpha \in \mathcal{A}} \omega^\top a(\alpha),$$

subject to $\beta^\top (I - \alpha)^{-1} = \omega^\top$. Since $I - \alpha$ is always invertible for $\alpha \in \mathcal{A}$ we can rewrite this constraint as the affine relationship

$$\alpha^\top \omega = \omega - \beta. \quad (48)$$

Take two feasible points ω^0 and ω^1 , and let $\alpha^0 \in \mathcal{A}$ and $\alpha^1 \in \mathcal{A}$ be their respective maximizers. Consider the convex combination α^t defined component-by-component as

$$\alpha_i^t = \frac{\omega_i^0}{\omega_i^0 + \omega_i^1} \alpha_i^0 + \frac{\omega_i^1}{\omega_i^0 + \omega_i^1} \alpha_i^1.$$

We will show that α^t is a feasible point for $\omega^t = \frac{\omega^0 + \omega^1}{2}$. First notice that $\alpha^t \geq 0$ and that

$$\sum_j \alpha_{ij}^t = \frac{\omega_i^0}{\omega_i^0 + \omega_i^1} \sum_j \alpha_{ij}^0 + \frac{\omega_i^1}{\omega_i^0 + \omega_i^1} \sum_j \alpha_{ij}^1 \leq \bar{\alpha}_i,$$

so that $\alpha^t \in \mathcal{A}$. Next, since (48) holds for (α^0, ω^0) and (α^1, ω^1) , we can write

$$\sum_j \omega_j^0 \alpha_{ji}^0 = \omega_i^0 - \beta_i, \text{ and } \sum_j \omega_j^1 \alpha_{ji}^1 = \omega_i^1 - \beta_i.$$

Summing these two equations up, we get

$$\begin{aligned}\sum_j (\omega_j^0 \alpha_{ji}^0 + \omega_j^1 \alpha_{ji}^1) &= \omega_i^0 + \omega_i^1 - 2\beta_i, \\ \sum_j \frac{\omega_j^0 + \omega_j^1}{2} \left(\frac{\omega_j^0}{\omega_j^0 + \omega_j^1} \alpha_{ji}^0 + \frac{\omega_j^1}{\omega_j^0 + \omega_j^1} \alpha_{ji}^1 \right) &= \frac{\omega_i^0 + \omega_i^1}{2} - \beta_i, \\ \sum_j \omega_j^t \alpha_{ji}^t &= \omega_i^t - \beta_i,\end{aligned}$$

which implies that (48) holds for (α^t, ω^t) . Therefore, α^t is a feasible point for ω^t .

Consider the value function at ω^t :

$$\bar{a}(\omega^t) = \bar{a}\left(\frac{\omega^0 + \omega^1}{2}\right) \geq \sum_i \omega_i^t a_i(\alpha_i^t) = \sum_i \frac{\omega_i^0 + \omega_i^1}{2} a_i\left(\frac{\omega_i^0}{\omega_i^0 + \omega_i^1} \alpha_i^0 + \frac{\omega_i^1}{\omega_i^0 + \omega_i^1} \alpha_i^1\right),$$

where the inequality follows since α^t might not be a maximizer for ω^t . From the strict concavity of a_i , we find

$$\bar{a}(\omega^t) > \sum_i \frac{\omega_i^0 + \omega_i^1}{2} \left(\frac{\omega_i^0}{\omega_i^0 + \omega_i^1} a_i(\alpha_i^0) + \frac{\omega_i^1}{\omega_i^0 + \omega_i^1} a_i(\alpha_i^1) \right) = \frac{1}{2} \bar{a}(\omega^0) + \frac{1}{2} \bar{a}(\omega^1).$$

This holds for any feasible ω^0 and ω^1 , and so \bar{a} is midpoint strictly concave. By the Theorem of Maximum \bar{a} is also continuous, and so \bar{a} is therefore strictly concave. It follows that the objective function (26) is also strictly concave, which proves the first part of the statement.

Since the objective (26) is strictly concave, the feasible set is convex, there is a unique maximizer so there is a unique solution ω^* to the planner's problem. Now, notice that the objective function (23) is strictly concave since a_i is strictly concave for all i . The feasible set (the intersection of (48) and \mathcal{A}) is convex so there is once again a unique maximizer. It follows that for each ω there is a unique α that solves (23), and there is therefore a unique α^* associated with ω^* . \square

A.7 Proof of Proposition 1

Proposition 1. *There exists a unique equilibrium, and it is efficient.*

Proof. For a given production network α and a given draw of the random TFP vector ε , the economy is standard, and the equilibrium is unique. The first welfare theorem also applies and so the allocation is efficient. We therefore only need to focus on the choice of network under uncertainty. An equilibrium network $\alpha^* \in \mathcal{A}$ is fully characterized by a solution to (17) and where

$\mathcal{R}(\alpha^*)$ is given by (18) which can be written in terms of primitives as

$$\mathcal{R}(\alpha^*) = \underbrace{-\mathcal{L}(\alpha^*)(\mu + a(\alpha^*))}_{\mathbb{E}[p]} + \underbrace{(\rho - 1)\mathcal{L}(\alpha^*)\Sigma[\mathcal{L}(\alpha^*)]^\top \beta}_{\text{Cov}(p, \lambda)}.$$

Since the objective function is strictly concave and the constraint set is defined by affine functions, it follows that $\alpha^* \in \mathcal{A}$ is an equilibrium network if there exists Lagrange multipliers $\chi_{ij}^e \geq 0$ and $\gamma_i^e \geq 0$ such that 1) the first-order conditions of the firms

$$\frac{\partial a_i}{\partial \alpha_i}(\alpha^*) + \mathcal{L}(\alpha^*)(\mu + a(\alpha^*)) - (\rho - 1)\mathcal{L}(\alpha^*)\Sigma\mathcal{L}(\alpha^*)^\top \beta + \chi_i^e - \gamma_i^e \mathbf{1} = 0, \quad (49)$$

evaluated at α^* are satisfied, and 2) the complementary slackness conditions

$$-\chi_{ij}^e \alpha_{ij}^* = 0, \quad (50)$$

$$\gamma_i^e \left(\sum_{j=1}^n \alpha_{ij}^* - \bar{\alpha}_i \right) = 0, \quad (51)$$

are satisfied for all i, j .

Next, consider the social planner's problem given by (26) and subject to the constraints in Corollary 2. By Lemma 4, the objective function is strictly concave and the constraint set is defined by affine function. It follows that an allocation $\alpha \in \mathcal{A}$ is efficient if there exist nonnegative Lagrange multipliers $\hat{\chi}$ and $\hat{\gamma}$ such that 1) the first-order conditions

$$\mu + \nabla \bar{a} - (\rho - 1)\Sigma\omega + \hat{\chi} - \hat{\gamma}(\mathbf{1} - \bar{\alpha}) = 0, \quad (52)$$

where $\omega^\top = \beta^\top \mathcal{L}(\alpha)$ and where $\nabla \bar{a}$ is the derivative of the aggregate TFP shifter (23), are satisfied and the 2) complementary slackness conditions

$$-\hat{\chi}_i(\omega_i - \beta_i) = 0, \quad (53)$$

$$\hat{\gamma}(\omega^\top(\mathbf{1} - \bar{\alpha}) - 1) = 0, \quad (54)$$

are satisfied for all i . To derive $\nabla \bar{a}$, we can use the problem (23). The objective function of this problem is strictly concave (see proof of Lemma 4) and the constraint set is convex.³⁶ It follows that the unique maximizer is characterized by the first-order condition

$$\omega_i \frac{\partial a_i}{\partial \alpha_{ij}} - \zeta_j \omega_i + \tilde{\chi}_{ij} - \tilde{\gamma}_i = 0 \Leftrightarrow \zeta_j = \frac{\partial a_i}{\partial \alpha_{ij}} + \tilde{\chi}_{ij} - \tilde{\gamma}_i, \quad (55)$$

³⁶Recall that the constraint set is given by $\alpha \in \mathcal{A}$ and an affine function (48).

and the complementary slackness conditions

$$\alpha^\top \omega - \omega + \beta = 0, \quad (56)$$

$$-\tilde{\chi}_{ij} \alpha_{ij} = 0, \quad (57)$$

$$\tilde{\gamma}_i \left(\sum_{j=1}^n \alpha_{ij} - \bar{\alpha}_i \right) = 0, \quad (58)$$

for all i, j and where $\tilde{\chi}_{ij} = \frac{\tilde{\chi}_{ij}}{\omega_i}$ and $\tilde{\gamma}_i = \frac{\tilde{\gamma}_i}{\omega_i}$. Applying the envelope theorem to (23), we obtain

$$\nabla \bar{a} = a(\alpha) + (I - \alpha) \zeta = a(\alpha) + (I - \alpha) \left(\frac{\partial a_i}{\partial \alpha_i} + \tilde{\chi}_i - \tilde{\gamma} \right),$$

where we use (55) to express ζ . Plugging this expression in (52), we get

$$\begin{aligned} \mu + a(\alpha) + (I - \alpha) \left(\frac{\partial a_i}{\partial \alpha_i} + \tilde{\chi}_i - \tilde{\gamma} \right) - (\rho - 1) \Sigma \omega + \hat{\chi} - \hat{\gamma} (\mathbf{1} - \bar{\alpha}) &= 0 \Leftrightarrow \\ \frac{\partial a_i}{\partial \alpha_i} + \mathcal{L}(\alpha) (\mu + a(\alpha)) - (\rho - 1) \mathcal{L}(\alpha) \Sigma \mathcal{L}(\alpha)^\top \beta + \tilde{\chi}_i + \mathcal{L}(\alpha) \hat{\chi} - (\hat{\gamma} \mathcal{L}(\alpha) (\mathbf{1} - \bar{\alpha}) + \tilde{\gamma}) &= 0, \end{aligned} \quad (59)$$

where the second line follows from the first by left-multiplying by $\mathcal{L}(\alpha) = (I - \alpha)^{-1}$.

Now, we will show that the equilibrium and efficiency conditions coincide. Suppose that we have a solution to the planner's problem $(\alpha^p, \omega^p, \tilde{\chi}, \hat{\chi}, \tilde{\gamma}, \hat{\gamma}, \zeta)$. Consider the candidate equilibrium $(\alpha^e, \omega^e, \chi^e, \gamma^e)$ where $\alpha^e = \alpha^p$, $\omega^e = \omega^p$, $\chi_i^e = \tilde{\chi}_i + \mathcal{L}(\alpha^p) \hat{\chi}$ for all i , and $\gamma^e = \hat{\gamma} \mathcal{L}(\alpha^p) (\mathbf{1} - \bar{\alpha}) + \tilde{\gamma}$. First, note that since $\tilde{\chi}, \hat{\chi}, \tilde{\gamma}, \hat{\gamma}$ are nonnegative, so are χ^e, γ^e . Next, the candidate equilibrium satisfies the first-order condition (49). The first complementary slackness condition (50) is also satisfied. Indeed, suppose that $\alpha_{ij}^p > 0$, which implies that $\omega_i^p > \beta_i$, then $\tilde{\chi}_{ij} = 0$ and $\hat{\chi}_j = 0$, such that $\chi_{ij}^e = 0$, and the condition is satisfied. If instead the constraint $\alpha_{ij}^p \geq 0$ binds such that $\tilde{\chi}_{ij} > 0$, we have $\chi_{ij}^e > 0$. Furthermore, from the first-order condition (49) $\alpha_{ij}^e = 0$, so the condition is satisfied. For the second complementary slackness condition (51), if $\sum_{j=1}^n \alpha_{ij}^p < \bar{\alpha}_i$ for some i then $\tilde{\gamma}_i = 0$ and $\hat{\gamma} = 0$. It follows that $\gamma^e = 0$ and the condition is satisfied. If instead the constraint binds such that $\sum_{j=1}^n \alpha_{ij}^p = \bar{\alpha}_i$ for some i , then $\sum_{j=1}^n \alpha_{ij}^e = \bar{\alpha}_i$, and the second complementary slackness condition (51) is satisfied. It follows that any efficient allocation can be decentralized as an equilibrium allocation. Since we know that an efficient allocation exists (it is the outcome of an optimization problem on a compact set), this proves that an efficient equilibrium exists.

Suppose instead that we have an equilibrium $(\alpha^e, \omega^e, \chi^e, \gamma^e)$ where $\omega^e = \mathcal{L}(\alpha^e)^\top \beta$, and consider the candidate efficient allocation $(\alpha^p, \omega^p, \tilde{\chi}, \hat{\chi}, \tilde{\gamma}, \hat{\gamma}, \zeta)$, where $\alpha^p = \alpha^e$, $\omega^p = \omega^e$, $\tilde{\chi} = \chi^e$, $\hat{\chi} = 0$, $\tilde{\gamma} = \gamma^e$, $\hat{\gamma} = 0$, $\zeta = -\mathcal{L}(\alpha^e) (\mu + a(\alpha^e)) + (\rho - 1) \mathcal{L}(\alpha^e) \Sigma \mathcal{L}(\alpha^e)^\top \beta$. Note that the first-order conditions (59) of the planner are satisfied. Next, notice that the complementary slackness conditions (53)-

(54) are always satisfied. Finally, the first-order conditions (55) and the complementary slackness conditions (56)-(58) are also satisfied. For the condition (56), note that since $\omega^e = \mathcal{L}(\alpha^e)^\top \beta$ we have $(\alpha^p)^\top \omega^p - \omega^p + \beta = 0$ and so the condition is satisfied. For the condition (57), if $\alpha_{ij}^e > 0$ then $\chi_{ij}^e = \tilde{\chi}_{ij} = 0$ and the condition is satisfied. If instead the constraint $\alpha_{ij}^e \geq 0$ binds such that $\alpha_{ij}^e = 0$, then $\alpha_{ij}^p = 0$ and the condition is also satisfied. Finally, for the condition (58), if $\sum_{j=1}^n \alpha_{ij}^e < \bar{\alpha}_i$ then $\gamma_i^e = 0$ and so $\tilde{\gamma} = 0$, so that the condition is satisfied. If instead $\sum_{j=1}^n \alpha_{ij}^p = \bar{\alpha}_i$, it must be that $\sum_{j=1}^n \alpha_{ij}^e = \bar{\alpha}_i$, and so (58) is satisfied. We have therefore shown that any equilibrium corresponds to an efficient allocation. By Lemma 4, the objective function of the planner is strictly concave and its constraint set is convex. It follows that there is a unique efficient allocation and therefore a unique equilibrium. \square

A.8 Proof of Proposition 2

Proposition 2. *The Domar weight ω_i of sector i is (weakly) increasing in μ_i and (weakly) decreasing in Σ_{ii} .*

Proof. Note that by the maximum theorem applied to (26), ω_i is a continuous function of μ and Σ . We consider the comparative statics with respect to μ_i first. We proceed by contradiction. Suppose that ω_i is not an increasing function of μ_i . Then, by continuity of ω_i as a function of μ_i , there exists a point (μ^0, Σ^0) and an interval (μ_i^0, μ_i^1) such that $\omega_i(\mu_i, \mu_{-i}^0, \Sigma^0) < \omega_i(\mu_i^0, \mu_{-i}^0, \Sigma^0)$ for any $\mu_i \in (\mu_i^0, \mu_i^1)$. Denote the optimal network at (μ, Σ) by $\alpha^*(\mu, \Sigma)$. 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). From Corollary 4, we can write the change in welfare as

$$\mathcal{W}(\mu_i^1, \mu_{-i}^0, \Sigma^0) = \mathcal{W}(\mu_i^0, \mu_{-i}^0, \Sigma^0) + \int_{\mu_i^0}^{\mu_i^1} \omega_i(\mu_i, \mu_{-i}^0, \Sigma^0) d\mu_i.$$

Suppose instead that the network is fixed at its original value $\alpha^*(\mu_i^0, \mu_{-i}^0, \Sigma^0)$. Equations (14) imply that under a fixed network the change in μ_i affects welfare only through its impact on expected log GDP. By Corollary 1, the change in welfare can thus be written as

$$W(\alpha^*(\mu_i^0, \mu_{-i}^0, \Sigma^0); \mu_i^1, \mu_{-i}^0, \Sigma^0) = W(\alpha^*(\mu_i^0, \mu_{-i}^0, \Sigma^0); \mu_i^0, \mu_{-i}^0, \Sigma^0) + \omega_i(\mu_i^0, \mu_{-i}^0, \Sigma^0) (\mu_i^1 - \mu_i^0).$$

But since the initial network $\alpha^*(\mu_i^0, \mu_{-i}^0, \Sigma^0)$ is feasible at $(\mu_i^1, \mu_{-i}^0, \Sigma^0)$, welfare maximization implies that $\mathcal{W}(\mu_i^1, \mu_{-i}^0, \Sigma^0) = W(\alpha^*(\mu_i^1, \mu_{-i}^0, \Sigma^0); \mu_i^1, \mu_{-i}^0, \Sigma^0) \geq W(\alpha^*(\mu_i^0, \mu_{-i}^0, \Sigma^0); \mu_i^1, \mu_{-i}^0, \Sigma^0)$, and so

$$\int_{\mu_i^0}^{\mu_i^1} \omega_i(\mu_i, \mu_{-i}^0, \Sigma^0) d\mu_i \geq \omega_i(\mu_i^0, \mu_{-i}^0, \Sigma^0) (\mu_i^1 - \mu_i^0). \quad (60)$$

Since we have assumed by contradiction that $\omega_i(\mu_i, \mu_{-i}^0, \Sigma^0) < \omega_i(\mu_i^0, \mu_{-i}^0, \Sigma^0)$ for all $\mu_i \in (\mu_i^0, \mu_i^1)$, it follows that

$$\int_{\mu_i^0}^{\mu_i^1} \omega_i(\mu_i, \mu_{-i}^0, \Sigma^0) d\mu_i < \int_{\mu_i^0}^{\mu_i^1} \omega_i(\mu_i^0, \mu_{-i}^0, \Sigma^0) d\mu_i = \omega_i(\mu_i^0, \mu_{-i}^0, \Sigma^0) (\mu_i^1 - \mu_i^0),$$

which contradicts (60). Therefore, ω_i is an increasing function of μ_i .

For the second part of the proposition, recall that $\frac{d\mathcal{W}}{d\Sigma_{ii}} = -\frac{1}{2}(\rho - 1)\omega_i^2$ by Corollary 4. Using analogous steps, we can then establish the second part of this proposition. \square

A.9 Proof of Proposition 3

Proposition 3. *Let γ denote either the mean μ_i or an element of the covariance matrix Σ_{ij} . If $\omega \in \text{int } \mathcal{O}$, then the response of the equilibrium Domar weights to a change in γ is given by*

$$\frac{d\omega}{d\gamma} = \underbrace{-\mathcal{H}^{-1}}_{\text{propagation}} \times \underbrace{\frac{\partial \mathcal{E}}{\partial \gamma}}_{\text{impulse}}, \quad (31)$$

where the $n \times n$ negative definite matrix \mathcal{H} is given by

$$\mathcal{H} = \nabla^2 \bar{a} + \frac{d\mathcal{E}}{d\omega}, \quad (32)$$

and where the matrix $\nabla^2 \bar{a}$ is the Hessian of the aggregate TFP shifter function \bar{a} , and $\frac{d\mathcal{E}}{d\omega} = \frac{d\text{Cov}[\varepsilon, \lambda]}{d\omega} = -(\rho - 1)\Sigma$ is the Jacobian matrix of the risk-adjusted TFP vector \mathcal{E} .

Proof. At an interior solution, the first-order conditions of (26) are

$$F(\omega, \mu, \Sigma) := \mu + \nabla \bar{a} - (\rho - 1)\Sigma\omega = 0,$$

where $\nabla \bar{a}$ is the gradient of \bar{a} . Differentiating with respect to ω , we find that

$$\frac{dF}{d\omega} = \nabla^2 \bar{a} - (\rho - 1)\Sigma,$$

where $\nabla^2 \bar{a}$ is the Hessian matrix of \bar{a} . From the implicit function theorem, it follows that

$$\frac{d\omega}{d\gamma} = -[\nabla^2 \bar{a} - (\rho - 1)\Sigma]^{-1} \frac{\partial F}{\partial \gamma}.$$

If $\gamma = \mu_i$, we have

$$\frac{\partial F}{\partial \gamma} = \frac{\partial F}{\partial \mu_i} = \mathbf{1}_i = \frac{\partial \mathcal{E}}{\partial \mu_i},$$

where $\mathbf{1}_i$ is a column vector of zeros except for 1 at element i . If $\gamma = \Sigma_{ij}$, we have

$$\frac{\partial F}{\partial \gamma} = \frac{\partial F}{\partial \Sigma_{ij}} = -\frac{1}{2}(\rho - 1)(\omega_j \mathbf{1}_i + \omega_i \mathbf{1}_j) = \frac{\partial \mathcal{E}}{\partial \Sigma_{ij}},$$

where, if $i \neq j$, we differentiate with respect to Σ_{ij} and Σ_{ji} simultaneously to preserve the symmetry of the covariance matrix and divide by two to preserve the scale. Finally, in the proof of Lemma 4, we show that $\nabla^2 \bar{a}$ is negative definite. It follows from (32) that \mathcal{H} and its inverse are also negative definite. \square

Proposition 3 can be extended to handle the case in which some of the constraints $\omega_i \geq \beta_i$ bind with strictly positive Lagrange multipliers. We show how this can be done in Supplemental Appendix F.

A.10 Proofs of Corollary 3, Lemmas 5, 6 and 7, and of Proposition 4

These proofs are in Supplemental Appendix D.

A.11 Proof of Lemma 8

Lemma 8. *If $\omega \in \mathcal{O}$, the equilibrium Domar weights are approximately given by*

$$\omega = \omega^\circ - [\mathcal{H}^\circ]^{-1} \mathcal{E}^\circ + O\left(\|\omega - \omega^\circ\|^2\right), \quad (35)$$

where the superscript \circ indicates that \mathcal{H} and \mathcal{E} are evaluated at ω° .

Proof. At an interior solution, the first-order conditions of (26) are

$$\mu + \nabla \bar{a}(\omega) - (\rho - 1) \Sigma \omega = 0. \quad (61)$$

The first-order Taylor expansion of $\nabla \bar{a}(\omega)$ around ω° is

$$\nabla \bar{a}(\omega) = \nabla \bar{a}(\omega^\circ) + \nabla^2 \bar{a}(\omega^\circ) (\omega - \omega^\circ) + O\left(\|\omega - \omega^\circ\|^2\right).$$

Plugging it into (61), we get

$$\omega - \omega^\circ = -[\nabla^2 \bar{a}(\omega^\circ) - (\rho - 1) \Sigma]^{-1} [\mu - (\rho - 1) \Sigma \omega^\circ + \nabla \bar{a}(\omega^\circ)] + O\left(\|\omega - \omega^\circ\|^2\right).$$

From (32), we can write $\mathcal{H}^\circ = \nabla^2 \bar{a}(\omega^\circ) - (\rho - 1) \Sigma$. From (27), $\mathcal{E}^\circ = \mu - (\rho - 1) \Sigma \omega^\circ$. Therefore,

$$\omega - \omega^\circ = -[\mathcal{H}^\circ]^{-1} [\mathcal{E}^\circ + \nabla \bar{a}(\omega^\circ)] + O\left(\|\omega - \omega^\circ\|^2\right).$$

Next, by the envelope theorem applied to (23) we find

$$\nabla \bar{a}(\omega^\circ) = a(\alpha^\circ) + (I - \alpha^\circ) \zeta = 0, \quad (62)$$

where ζ is the vector of Lagrange multipliers associated with the constraint $\alpha^\top \omega^\circ = \omega^\circ - \beta$. To find these multipliers, recall from (55) that the first-order conditions of the problem (23) are

$$\zeta_i = \frac{\partial a_i}{\partial \alpha_{ij}} + \tilde{\chi}_{ij} - \tilde{\gamma}_i. \quad (55)$$

Now recall that $\frac{\partial a_i}{\partial \alpha_{ij}}(\alpha_i^\circ) = 0$ for all i, j by the definition of α_i° . It follows that if set $\alpha_i = \alpha_i^\circ$ for all i and all the Lagrange multipliers $\tilde{\chi}_{ij}, \tilde{\gamma}_i$ equal to zero, then the first-order conditions and complementary slackness conditions are satisfied, and by definition of ω° , the constraint $(\alpha^\circ)^\top \omega^\circ = \omega^\circ - \beta$ is also satisfied. This is therefore the (unique) solution to that optimization problem. It follows from (62) that $\nabla \bar{a}(\omega^\circ) = a(\alpha^\circ) = 0$ where the last equality comes from our normalization. \square

A.12 Proof of Lemma 9

The proof of Lemma 9 is a special case of Proposition 9 in Supplemental Appendix I.

A.13 Proof of Proposition 5

Proposition 5. *Let γ denote either the mean μ_i or an element of the covariance matrix Σ_{ij} . Under an endogenous network, welfare responds to a marginal change in γ as if the network were fixed at its equilibrium value α^* , that is*

$$\frac{d\mathcal{W}(\mu, \Sigma)}{d\gamma} = \frac{\partial W(\alpha^*, \mu, \Sigma)}{\partial \gamma}.$$

Proof. Recall from Lemma 3 that the equilibrium α^* solves the welfare-maximization problem

$$\mathcal{W}(\mu, \Sigma) = \max_{\alpha \in \mathcal{A}} W(\alpha, \mu, \Sigma), \quad (63)$$

where

$$W(\alpha, \mu, \Sigma) = \mathbb{E}[y(\alpha)] - \frac{1}{2}(\rho - 1) \mathbb{V}[y(\alpha)], \quad (64)$$

is welfare under a given network α and beliefs (μ, Σ) . Both $\mathbb{E}[y]$ and $\mathbb{V}[y]$ depend on beliefs through (14). Since the objective function (63) and its associated constraints are continuously differentiable functions of α , and since the constraint $\alpha \in \mathcal{A}$ does not depend on beliefs, the envelope theorem immediately implies that

$$\frac{d\mathcal{W}(\mu, \Sigma)}{d\gamma} = \frac{\partial W(\alpha^*, \mu, \Sigma)}{\partial \gamma},$$

where the right-hand side is the change in welfare keeping the network constant at α^* . \square

A.14 Proof of Corollary 4

Corollary 4. *The impact of an increase in μ_i on welfare is given by*

$$\frac{d\mathcal{W}}{d\mu_i} = \omega_i, \quad (39)$$

and the impact of an increase in Σ_{ij} on welfare is given by

$$\frac{d\mathcal{W}}{d\Sigma_{ij}} = -\frac{1}{2}(\rho - 1)\omega_i\omega_j. \quad (40)$$

Proof. Combining (64) with Corollary 14, it is immediate to show that

$$\frac{\partial W(\alpha^*, \mu, \Sigma)}{\partial \mu_i} = \omega_i,$$

and

$$\frac{dW(\alpha^*, \mu, \Sigma)}{d\Sigma_{ij}} = -\frac{1}{2}(\rho - 1)\omega_i\omega_j.$$

Putting these expressions together with Proposition 5 yields the result. \square

A.15 Proof of Proposition 6

Proposition 6. *The presence of uncertainty lowers expected log GDP, in the sense that $E[y]$ is largest when $\Sigma = 0$.*

Proof. The proof follows from Corollary 3. From (14), define

$$\mathcal{Y}(\alpha, \mu, \Sigma) = E[y(\alpha)] = \omega(\alpha)^\top (\mu + a(\alpha)), \quad (65)$$

and

$$\mathcal{V}(\alpha, \mu, \Sigma) = V[y(\alpha)] = \omega(\alpha)^\top \Sigma \omega(\alpha), \quad (66)$$

as the expected value and the variance of log GDP under the network α and the beliefs (μ, Σ) . Let $\alpha^*(\mu, \Sigma)$ denote an optimal network (a solution to (21)) under the beliefs (μ, Σ) .

Fix μ . We first establish that $\alpha^*(\mu, 0)$ maximizes $\mathcal{Y}(\alpha, \mu, 0)$. To see this, note that (66) implies that $\mathcal{V}(\alpha, \mu, 0) = 0$ for all pairs (α, μ) . The problem (21) of the social planner with $\Sigma = 0$ can therefore be written as

$$\max_{\alpha \in \mathcal{A}} \mathcal{Y}(\alpha, \mu, 0) - \frac{1}{2}(\rho - 1)\mathcal{V}(\alpha, \mu, 0) = \max_{\alpha \in \mathcal{A}} \mathcal{Y}(\alpha, \mu, 0) = \mathcal{Y}(\alpha^*(\mu, 0), \mu, 0),$$

where the second equality comes from the definition of $\alpha^*(\mu, 0)$.

Next, notice that

$$\mathcal{Y}(\alpha^*(\mu, 0), \mu, 0) \geq \mathcal{Y}(\alpha^*(\mu, \Sigma), \mu, 0) = \mathcal{Y}(\alpha^*(\mu, \Sigma), \mu, \Sigma), \quad (67)$$

where the inequality comes from the fact that $\alpha^*(\mu, 0)$ maximizes $\mathcal{Y}(\alpha, \mu, 0)$, and the equality comes from the fact that $\mathcal{Y}(\alpha, \mu, \Sigma)$, given by (65), does not explicitly depend on Σ . Since (67) holds for any Σ , it follows that expected log GDP $\mathcal{Y}(\alpha^*(\mu, \Sigma), \mu, \Sigma)$ is maximized at $\Sigma = 0$, which is the desired result. \square

A.16 Proof of Corollary 5

The proof of Corollary 5 is in Supplemental Appendix I.

A.17 Proof of Proposition 7

Proposition 7. *If $\omega \in \text{int } \mathcal{O}$, the following holds.*

1. *The impact of an increase in μ_i on log GDP is given by*

$$\frac{dE[y]}{d\mu_i} = \underbrace{\omega_i}_{\text{Fixed network}} - (\rho - 1) \omega^\top \Sigma \mathcal{H}^{-1} \frac{\partial \mathcal{E}}{\partial \mu_i}, \quad \text{and} \quad \frac{dV[y]}{d\mu_i} = \underbrace{0}_{\text{Fixed network}} - 2\omega^\top \Sigma \mathcal{H}^{-1} \frac{\partial \mathcal{E}}{\partial \mu_i}.$$

2. *The impact of an increase in Σ_{ij} on log GDP is given by*

$$\frac{dE[y]}{d\Sigma_{ij}} = \underbrace{0}_{\text{Fixed network}} - (\rho - 1) \omega^\top \Sigma \mathcal{H}^{-1} \frac{\partial \mathcal{E}}{\partial \Sigma_{ij}}, \quad \text{and} \quad \frac{dV[y]}{d\Sigma_{ij}} = \underbrace{\omega_i \omega_j}_{\text{Fixed network}} - 2\omega^\top \Sigma \mathcal{H}^{-1} \frac{\partial \mathcal{E}}{\partial \Sigma_{ij}}.$$

Proof. Differentiating (14) with respect to μ_i yields

$$\frac{dV[y]}{d\mu_i} = 2\omega^\top \Sigma \frac{d\omega}{d\mu_i},$$

which, together with (31), yields

$$\frac{dV[y]}{d\mu_i} = -2\omega^\top \Sigma \mathcal{H}^{-1} \frac{\partial \mathcal{E}}{\partial \mu_i}. \quad (68)$$

Next, from (39) we find

$$\frac{dW}{d\mu_i} = \omega_i = \frac{dE[y]}{d\mu_i} - \frac{1}{2}(\rho - 1) \frac{dV[y]}{d\mu_i},$$

which we can combine with the previous equation to get

$$\frac{dE[y]}{d\mu_i} = \omega_i - (\rho - 1) \omega^\top \Sigma \mathcal{H}^{-1} \frac{\partial \mathcal{E}}{\partial \mu_i}.$$

Similarly, differentiating (14) with respect to Σ_{ij} yields

$$\frac{dV[y]}{d\Sigma_{ij}} = \omega_i \omega_j + 2\omega^\top \Sigma \frac{d\omega}{d\Sigma_{ij}} = \omega_i \omega_j - 2\omega^\top \Sigma \mathcal{H}^{-1} \frac{\partial \mathcal{E}}{\partial \Sigma_{ij}}.$$

From (40), we can write

$$\frac{d\mathcal{W}}{d\Sigma_{ij}} = -\frac{1}{2}(\rho - 1)\omega_i \omega_j = \frac{dE[y]}{d\Sigma_{ij}} - \frac{1}{2}(\rho - 1)\frac{dV[y(\alpha)]}{d\Sigma_{ij}},$$

which we can combine with the previous equation to find

$$\frac{dE[y]}{d\Sigma_{ij}} = -(\rho - 1)\omega^\top \Sigma \mathcal{H}^{-1} \frac{\partial \mathcal{E}}{\partial \Sigma_{ij}}.$$

□

A.18 Proof of Corollary 6

Corollary 6. *Without uncertainty ($\Sigma = 0$) the moments of GDP respond to changes in beliefs as if the network were fixed, such that*

$$\frac{dE[y]}{d\mu_i} = \frac{\partial E[y]}{\partial \mu_i} = \omega_i, \quad \text{and} \quad \frac{dV[y]}{d\Sigma_{ij}} = \frac{\partial V[y]}{\partial \Sigma_{ij}} = \omega_i \omega_j.$$

Proof. When $\Sigma = 0$, the problem of the planner (3) becomes

$$\mathcal{W} = E[y(\alpha^*)] = \max_{\alpha \in \mathcal{A}} E[y(\alpha)] = \max_{\alpha \in \mathcal{A}} \omega(\alpha)^\top (\mu + a(\alpha)).$$

The envelope theorem then implies that $\frac{dE[y]}{d\mu_i} = \omega_i$ which proves the first part of the corollary. The envelope theorem also implies $\frac{dE[y]}{d\Sigma_{ij}} = 0$ which leads to the second part of the corollary when combined with Proposition 5 and Corollary 1. □

A.19 Proofs of Corollaries 7 and 8

These proofs are in Supplemental Appendix D.

B Additional results related to the calibrated economy

In this appendix, we provide additional information about 1) the data used in the calibration of Section 8, 2) our calibration strategy, 3) how well the model fits the data, 4) the quantitative importance of the mechanism, and 5) robustness exercises.

B.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 expenditure accounted for by different sectors. We rely on the harmonized tables constructed by [Vom Lehn and Winberry \(2022\)](#) that provide consistent annual data for $n = 37$ sectors over the period 1948-2020. Table 3 provides the list of the sectors included in this data set.

Table 3: The 37 sectors used in our analysis

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	

Notes: Sectors are classified according to the NAICS-based BEA codes. See [Vom Lehn and Winberry \(2022\)](#) for details of the data construction.

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 \(2022\)](#) 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. We make three departures from [Vom Lehn and Winberry \(2022\)](#) 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.

B.2 Calibration procedure

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 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.

Household preferences

Since the preference parameter β_i corresponds to the household's expenditure share of good i , we pin down its value directly from the data by averaging the consumption share of good i over time. The sectors with the largest consumption shares are "Real estate" (14%), "Retail trade" (12%) and "Health care" (11%). See Appendix J.5 for a version of the calibrated economy with time-varying β 's.

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 specialize the TFP shifter function (2) to

$$\log A_i(\alpha_i) = a_i^\circ - \sum_{j=1}^n \kappa_{ij} (\alpha_{ij} - \alpha_{ij}^\circ)^2 - \kappa_{i0} \left(\sum_{j=1}^n \alpha_{ij} - \sum_{j=1}^n \alpha_{ij}^\circ \right)^2, \quad (69)$$

where the last term can provide a penalty from deviating from an ideal labor share. We denote by κ the matrix with typical element κ_{ij} . This functional form takes as inputs the ideal shares α_{ij}° , the actual shares α_{ij} , 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.³⁷ 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 $2n = 74$ free parameters to estimate.

³⁷We 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.

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,

$$\varepsilon_t = \gamma + \varepsilon_{t-1} + u_t, \quad (42)$$

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

$$\hat{\Theta} = \arg \min_{\Theta} (m(z) - m(\Theta))^{\top} W (m(z) - m(\Theta)),$$

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.³⁸

³⁸To 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 $T - 1$).

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.

B.3 The calibrated economy

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 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.

Evolution of beliefs in the data

Our estimation procedure provides a time-series for μ_t and Σ_t . 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

$$\Delta \bar{\mu}_t = \sum_{j=1}^n \omega_{jt} \Delta \mu_{jt}. \quad (70)$$

We use the Domar weights ω_{jt} in this equation to properly reflect the importance of a sector for GDP, as implied by (13). 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 (14), this can be written as

$$\sigma_{yt} = \sqrt{V[y]} = \sqrt{\omega_t^\top \Sigma_t \omega_t}. \quad (71)$$

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 2007-2009.³⁹

³⁹ σ_{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, $\sigma_{yt} = \sqrt{\omega_t^\top \Sigma_t \omega_t}$. Shaded areas represent NBER recessions.

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 (Corollary 4), 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 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 4).⁴⁰

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 4. 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 sizable portion of the variation in a key moment that characterizes the production network.⁴¹

⁴⁰Since 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 4). 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.

⁴¹One 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.

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.

Table 4: Domar weights, consumption and TFP in the model and in the data

	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)	Corr (ω_{jt}, μ_{jt})	0.08	0.08
(5)	Corr $(\omega_{jt}, \Sigma_{jjt})$	-0.37	-0.31
(6)	Consumption growth volatility	2.65%	2.73%
(7)	TFP growth volatility	1.83%	2.73%

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 volatilities 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 2 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 4, the predictions of the model are borne out in the data. The model is thus able to capture well the impact of beliefs on the structure of

the production network.

Sectoral correlations

The model is also able to replicate features of the correlation between sectoral outputs. We focus on growth rates to accommodate different trends in the data and in the model. For each pair of sectors, we compute the correlation in their output growth in the model and in the data, and plot them in Figure 7. The model reasonably captures cross-sectoral comovements: We find that the correlation between the data- and model-implied values is 0.44. On average, sectoral outputs are positively correlated in the model and in the data, although the model correlation is somewhat weaker on average (see the first column of Table 5).

Figure 7: Cross-sector correlations in the model and in the data



Notes: For each pair of sectors, we compute correlations in the growth rates of sectoral output in the model and in the data. Each dot in the graph shows the value of this correlation in the model (X-axis) and in the data (Y-axis). The solid black line results from the ordinary least square analysis.

Table 5 also reports averages of these correlations during periods of low and high TFP growth and uncertainty growth, as measured by (70) and (71). We see that in the data these correlations are lower in good times, when TFP growth is high and uncertainty growth is low. The model is able to replicate this ranking. Intuitively, in bad times consumption is low and so the household is particularly 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.

B.4 Counterfactual exercises

In this appendix we provide more information about the counterfactual exercises of Section 8.2.

Table 5: Correlations in sectoral sales growth

	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

Notes: For each sector pair (i, j) , we compute correlations in the growth rates of sectoral output in the model and in the data. We then take averages across all sectors. 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.

Long-run moments

Table 6 provides differences in long-run moments between our baseline model and the three alternative economies described in the main text. In the “known ε_t ”, $E[y]$ and \mathcal{W} collapse to realized GDP and $V[y] = 0$. In Table 6, we compute instead these moments *before* ε_t is known but still assuming that the production network is chosen optimally for the realized draw of ε_t

Table 6: Uncertainty, GDP and welfare in the post-war sample

	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, $\sqrt{V[y]}$	+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%

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

Time series under the “known ε_t ” alternative economy

In the “known ε_t ”, 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 8).⁴² 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 6).

The top three panels in the right column of Figure 8 show how the baseline and alternative economies differ in terms of expected log GDP, the standard deviation of log GDP, and (expected)


⁴²Again, here we report the moments *before* ε_t is known but still assuming that the production network is chosen optimally for the realized draw of ε_t .

welfare. Crucially, these measures are evaluated *before* ε is realized.⁴³ 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 6).


Figure 8: The role of uncertainty in the postwar period

Left column: the “as if $\Sigma_t = 0$ ” alternative


Right column: the “known ε_t ” alternative




(a) Difference in expected log GDP [%]




(b) Difference in expected log GDP [%]




(c) Difference in expected st. dev. of log GDP [%]




(d) Difference in expected st. dev. of log GDP [%]




(e) Difference in expected welfare [%]



(f) Difference in expected welfare [%]



(g) Difference in realized log GDP [%]



(h) Difference in realized log GDP [%]

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 $\sqrt{V[y]}$ are evaluated before ε_t is realized.

⁴³Note that *realized* welfare in this economy is simply equal to realized log GDP.

Supplemental appendix (Not for publication)

C Additional derivations and results

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

C.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 10. *The Lagrange multiplier on the budget constraint of the household (4) is*

$$\Lambda = \frac{u'(Y)}{\bar{P}},$$

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

Proof. The household makes decisions after the realization of the state of the world ε . The state-specific 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

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

and the first-order condition with respect to C_i is

$$\beta_i u'(Y) Y = \Lambda P_i C_i. \tag{72}$$

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

$$u'(Y) Y = \Lambda, \tag{73}$$

which, together with (72), implies that

$$P_i C_i = \beta_i. \tag{74}$$

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

$$\begin{aligned} Y &= \prod_{i=1}^n (\beta_i^{-1} C_i)^{\beta_i} = \prod_{i=1}^n \left(\beta_i^{-1} \frac{\beta_i u'(Y) Y}{\Lambda P_i} \right)^{\beta_i} \\ \Lambda &= u'(Y) \prod_{i=1}^n P_i^{-\beta_i} \end{aligned} \quad (75)$$

which, combined with (73), yields

$$Y = \prod_{i=1}^n P_i^{-\beta_i}. \quad (76)$$

This last equation implicitly defines a price index $\bar{P} = \prod_{i=1}^n P_i^{\beta_i}$ such that $\bar{P}Y = 1$. Combining that last equation with (73) yields the result. \square

C.2 Derivation of the unit cost function

The cost minimization problem of the firm is

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

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

$$\begin{aligned} L_i &= \theta \left(1 - \sum_{j=1}^n \alpha_{ij} \right) F(\alpha_i, L_i, X_i), \\ P_j X_{ij} &= \theta \alpha_{ij} F(\alpha_i, L_i, X_i), \end{aligned}$$

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.

C.3 Derivation of the first-order condition of the firm (19)

At an interior solution, the first-order conditions associated with problem (17) are

$$0 = -\frac{da_i(\alpha_i)}{d\alpha_{ij}} + \mathcal{R}_j.$$

We can write these equations in vector form as $f(\alpha_i, \mathcal{R}) = 0$, with element j given by

$$f_j(\alpha_i, \mathcal{R}) = -\frac{da_i(\alpha_i)}{d\alpha_{ij}} + \mathcal{R}_j.$$

A solution α_i to the first-order conditions corresponds to $f(\alpha_i, \mathcal{R}) = 0$. The Jacobian of f with respect to α_i is $-H_i$ where H_i is the Hessian of a_i . The Jacobian of f with respect to the vector \mathcal{R} is the identity matrix. From the implicit function theorem, we can therefore write

$$\frac{\partial \alpha_{ij}}{\partial \mathcal{R}_k} = \left[\frac{\partial \alpha_i}{\partial \mathcal{R}} \right]_{jk} = - \left[\left[\frac{\partial f}{\partial \alpha_i} \right]^{-1} \times \frac{\partial f}{\partial \mathcal{R}} \right]_{jk} = [H_i^{-1}]_{jk}.$$

C.4 Derivation of \bar{a} and $\alpha(\omega)$ under quadratic TFP shifter

In this appendix, we solve the problem (23) at an interior solution when the TFP shifter functions (a_1, \dots, a_n) are of the form (2). From (23), the planner seeks to maximize

$$\sum_i \omega_i \left(\frac{1}{2} \alpha_i^\top H_i \alpha_i - (\alpha_i^\circ)^\top H_i \alpha_i \right),$$

subject to $\omega^\top = \beta^\top \mathcal{L}(\alpha)$ or, since $I - \alpha$ is always invertible for $\alpha \in \mathcal{A}$, $\alpha^\top \omega = \omega - \beta$. We can rewrite this problem as a standard quadratic program

$$\min_{\alpha \in \mathcal{A}} \frac{1}{2} x^\top Q x + c^\top x,$$

subject to $Ex = d$, where

$$\underbrace{Q}_{n^2 \times n^2} = - \begin{bmatrix} \omega_1 H_1 & & 0 \\ & \ddots & \\ 0 & & \omega_n H_n \end{bmatrix},$$

and where the constraint

$$\alpha^\top \omega = \omega - \beta,$$

becomes

$$\underbrace{\begin{bmatrix} I\omega_1, \dots, I\omega_n \end{bmatrix}}_{\substack{E \\ n \times n^2}} \underbrace{\begin{bmatrix} \alpha_{11} \\ \vdots \\ \alpha_{1n} \\ \vdots \\ \alpha_{n1} \\ \vdots \\ \alpha_{nn} \end{bmatrix}}_x = \underbrace{\begin{bmatrix} \omega_1 - \beta_1 \\ \vdots \\ \omega_n - \beta_n \end{bmatrix}}_d,$$

and where $c^\top = \begin{bmatrix} \omega_1 (\alpha_1^\circ)^\top H_1 & \dots & \omega_n (\alpha_n^\circ)^\top H_n \end{bmatrix}$. The solution to this problem is well-known and given by

$$\begin{bmatrix} Q & E^\top \\ E & 0 \end{bmatrix} \begin{bmatrix} x \\ \lambda \end{bmatrix} = \begin{bmatrix} -c \\ d \end{bmatrix}.$$

Proof. Using the block matrix inverse equation we can write

$$\begin{bmatrix} Q & E^\top \\ E & 0 \end{bmatrix}^{-1} = \begin{bmatrix} Q^{-1} - Q^{-1}E^\top (EQ^{-1}E^\top)^{-1}EQ^{-1} & Q^{-1}E^\top (EQ^{-1}E^\top)^{-1} \\ (EQ^{-1}E^\top)^{-1}EQ^{-1} & -(EQ^{-1}E^\top)^{-1} \end{bmatrix},$$

such that the solution to the optimization problem is

$$x = Q^{-1}E^\top (EQ^{-1}E^\top)^{-1}d - \left(Q^{-1} - Q^{-1}E^\top (EQ^{-1}E^\top)^{-1}EQ^{-1} \right) c.$$

Simple matrix algebra implies that

$$EQ^{-1} = - \begin{bmatrix} H_1^{-1} & \dots & H_n^{-1} \end{bmatrix},$$

and

$$EQ^{-1}E^\top = - \sum_i \omega_i H_i^{-1} = -D,$$

where we explicitly define the square matrix D .

It follows that

$$\begin{aligned}
x &= \begin{bmatrix} H_1^{-1} \\ \vdots \\ H_n^{-1} \end{bmatrix} D^{-1} \begin{bmatrix} \omega_1 - \beta_1 \\ \vdots \\ \omega_n - \beta_n \end{bmatrix} + \begin{bmatrix} \omega_1^{-1} H_1^{-1} & & 0 \\ & \ddots & \\ 0 & & \omega_n^{-1} H_n^{-1} \end{bmatrix} \begin{bmatrix} \omega_1 H_1 \alpha_1^\circ \\ \vdots \\ \omega_n H_n \alpha_n^\circ \end{bmatrix} \\
&- \begin{bmatrix} H_1^{-1} \\ \vdots \\ H_n^{-1} \end{bmatrix} D^{-1} \begin{bmatrix} H_1^{-1} & \dots & H_n^{-1} \end{bmatrix} \begin{bmatrix} \omega_1 H_1 \alpha_1^\circ \\ \vdots \\ \omega_n H_n \alpha_n^\circ \end{bmatrix},
\end{aligned}$$

or

$$\alpha_i - \alpha_i^\circ = H_i^{-1} \left(\sum_j \omega_j H_j^{-1} \right)^{-1} \left(\omega - \beta - \sum_j \omega_j \alpha_j^\circ \right).$$

The value function follows immediately from the definition of $\bar{a}(\omega)$ given by (23). \square

C.5 Amplification and dampening

Proposition 8. *Let $\alpha^*(\mu, \Sigma)$ be the equilibrium production network under (μ, Σ) and let $W(\alpha, \mu, \Sigma)$ 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*

$$\underbrace{\mathcal{W}(\mu', \Sigma') - \mathcal{W}(\mu, \Sigma)}_{\text{Change in welfare under a flexible network}} \geq \underbrace{W(\alpha^*(\mu, \Sigma), \mu', \Sigma') - W(\alpha^*(\mu, \Sigma), \mu, \Sigma)}_{\text{Change in welfare under a fixed network}}. \quad (38)$$

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

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

and under the fixed network is

$$W(\alpha^*(\mu, \Sigma), \mu', \Sigma') - W(\alpha^*(\mu, \Sigma), \mu, \Sigma).$$

By Proposition 1, $\alpha^*(\mu', \Sigma')$ maximizes welfare under (μ', Σ') so that

$$W(\alpha^*(\mu', \Sigma'), \mu', \Sigma') \geq W(\alpha^*(\mu, \Sigma), \mu', \Sigma').$$

Combining the two expression gives the result. \square

D Additional proofs

This appendix contains proofs that are not included in the main appendix.

D.1 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 the moments of log GDP is given by*

$$\frac{\partial \mathbb{E}[y]}{\partial \mu_i} = \omega_i, \quad \text{and} \quad \frac{\partial \mathbb{V}[y]}{\partial \mu_i} = 0.$$

2. *The impact of a change in volatility Σ_{ij} on the moments of log GDP is given by*

$$\frac{\partial \mathbb{E}[y]}{\partial \Sigma_{ij}} = 0, \quad \text{and} \quad \frac{\partial \mathbb{V}[y]}{\partial \Sigma_{ij}} = \omega_i \omega_j.$$

Proof. Equation (14) implies that $\frac{\partial \mathbb{E}[y(\alpha)]}{\partial \mu_i} = \beta^\top \mathcal{L}(\alpha) \mathbf{1}_i$. Since $P^\top C = WL = 1$ by the household's budget constraint, we need to show that $\beta^\top \mathcal{L}(\alpha) \mathbf{1}_i = P_i Q_i$ to complete the proof of the first result. From (74), we know that $P_i C_i = \beta_i$. Using Shepard's Lemma together with the marginal pricing equation (10), we can find the firm's factor demands equations

$$\begin{aligned} P_j X_{ij} &= \alpha_{ij} P_i Q_i, \\ L_i &= \left(1 - \sum_{j=1}^n \alpha_{ij} \right) P_i Q_i. \end{aligned} \tag{77}$$

Using these results, we can write the market clearing condition (11) as

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

Solving the linear system implies

$$\beta^\top \mathcal{L}(\alpha) \mathbf{1}_i = P_i Q_i, \tag{78}$$

which proves the first part of the proposition. For the second part of the result, differentiating (14) with respect to Σ_{ij} and holding Σ symmetric yields

$$\frac{\partial \mathbb{V}[y(\alpha)]}{\partial \Sigma_{ij}} = \frac{1}{2} \beta^\top \mathcal{L}(\alpha) \left[\mathbf{1}_i \mathbf{1}_j^\top + \mathbf{1}_j \mathbf{1}_i^\top \right] [\mathcal{L}(\alpha)]^\top \beta,$$

which is the result. □

D.2 Proof of Corollary 3

Corollary 3. *If $\omega \in \text{int } \mathcal{O}$, then the following holds.*

1. An increase in the expected value μ_i or a decline in the variance Σ_{ii} leads to an increase in ω_j if i and j are global complements, and to a decline in ω_j if i and j are global substitutes.
2. An increase in the covariance Σ_{ij} , $i \neq j$, leads to a decline in ω_k if k is global complement with i and j , and to an increase in ω_k if k is global substitute with i and j .

Proof. Combining (31) and (28), we can write $\frac{d\omega_j}{d\mu_i} = -\mathbf{1}_j^\top \mathcal{H}^{-1} \mathbf{1}_i = -\mathcal{H}_{ji}^{-1} = -\mathcal{H}_{ij}^{-1}$ since \mathcal{H} is symmetric and inversion preserves symmetry. This proves point 1. For point 2, combining (31) and (29) yields $\frac{d\omega_j}{d\Sigma_{ii}} = (\rho - 1) \omega_i \mathbf{1}_j^\top \mathcal{H}^{-1} \mathbf{1}_i = (\rho - 1) \omega_i \mathcal{H}_{ij}^{-1}$, which is the result. For point 3, combining (31) and (29) yields

$$\frac{d\omega_k}{d\Sigma_{ij}} = \frac{1}{2} (\rho - 1) \mathbf{1}_k^\top \mathcal{H}^{-1} (\omega_j \mathbf{1}_i + \omega_i \mathbf{1}_j) = \frac{1}{2} (\rho - 1) (\omega_j \mathcal{H}_{ik}^{-1} + \omega_i \mathcal{H}_{jk}^{-1}),$$

which is the result. \square

D.3 Proof of Lemma 5

Lemma 5. *An increase in the covariance Σ_{ij} induces stronger global substitution between i and j , in the sense that $\frac{\partial \mathcal{H}_{ij}^{-1}}{\partial \Sigma_{ij}} > 0$.*

Proof. Note that

$$\frac{\partial \mathcal{H}_{ij}^{-1}}{\partial \Sigma_{kl}} = \frac{1}{2} (\rho - 1) \mathbf{1}_i^\top \mathcal{H}^{-1} (\mathbf{1}_k \mathbf{1}_l^\top + \mathbf{1}_l \mathbf{1}_k^\top) \mathcal{H}^{-1} \mathbf{1}_j,$$

where, if $k \neq l$, we differentiate with respect to Σ_{kl} and Σ_{lk} to preserve the symmetry of Σ and divide by two. In the special case with $i = k$ and $j = l$,

$$\begin{aligned} \frac{\partial \mathcal{H}_{kl}^{-1}}{\partial \Sigma_{kl}} &= \frac{1}{2} (\rho - 1) \mathbf{1}_k^\top \mathcal{H}^{-1} (\mathbf{1}_k \mathbf{1}_l^\top + \mathbf{1}_l \mathbf{1}_k^\top) \mathcal{H}^{-1} \mathbf{1}_l \\ &= \frac{1}{2} (\rho - 1) \left\{ \mathbf{1}_k^\top \mathcal{H}^{-1} \mathbf{1}_k \mathbf{1}_l^\top \mathcal{H}^{-1} \mathbf{1}_l + \mathbf{1}_k^\top \mathcal{H}^{-1} \mathbf{1}_l \mathbf{1}_k^\top \mathcal{H}^{-1} \mathbf{1}_l \right\} \\ &= \frac{1}{2} (\rho - 1) \left\{ \mathcal{H}_{kk}^{-1} \mathcal{H}_{ll}^{-1} + [\mathcal{H}_{kl}^{-1}]^2 \right\} > 0. \end{aligned}$$

The strict inequality holds because \mathcal{H}^{-1} is a negative definite matrix. \square

D.4 Proof of Lemma 6

Lemma 6. *Suppose that all input shares are (weak) local complements in the production of all goods, that is $[H_i^{-1}]_{kl} \leq 0$ for all i and all $k \neq l$. If $\alpha \in \text{int } \mathcal{A}$, there exists a scalar $\bar{\Sigma} > 0$ such that if $\|\Sigma\| \leq \bar{\Sigma}$, all sectors are global complements, that is $\mathcal{H}_{ij}^{-1} < 0$ for all $i \neq j$.*

Proof. Consider the problem (23). In the internal equilibrium, $\alpha \in \text{int } \mathcal{A}$, the first-order condition with respect to α_{ij} for this problem is

$$\omega_i \frac{\partial a_i}{\partial \alpha_{ij}} - \zeta_j \omega_i = 0 \Leftrightarrow \zeta_j = \frac{\partial a_i}{\partial \alpha_{ij}}, \quad (79)$$

where ζ is the vector of Lagrange multipliers associated with the constraint $\alpha^\top \omega - \omega + \beta = 0$. Applying the envelope theorem to (23), we obtain

$$\nabla \bar{a} = a(\alpha) + (I - \alpha) \zeta.$$

Differentiation of this expression yields

$$\begin{aligned} (\nabla^2 \bar{a})_{ij} &= \frac{d^2 \bar{a}}{d\omega_i d\omega_j} = \sum_{k=1}^n \frac{\partial a_i(\alpha_i)}{\partial \alpha_{ik}} \frac{d\alpha_{ik}}{d\omega_j} + (\mathbf{1}_i - \alpha_i)^\top \frac{d\zeta}{d\omega_j} - \zeta^\top \frac{d\alpha_i}{d\omega_j} \\ &= \sum_{k=1}^n \left(\frac{\partial a_i(\alpha_i)}{\partial \alpha_{ik}} - \zeta_k \right) \frac{d\alpha_{ik}}{d\omega_j} + (\mathbf{1}_i - \alpha_i)^\top \frac{d\zeta}{d\omega_j} \stackrel{(79)}{=} (\mathbf{1}_i - \alpha_i)^\top H_s \frac{d\alpha_s}{d\omega_j}. \end{aligned} \quad (80)$$

Note that from (79), $\frac{d\zeta}{d\omega_j} = H_s \frac{d\alpha_s}{d\omega_j}$ for any sector s . This implies, in particular, that $\frac{d\alpha_k}{d\omega_j} = H_k^{-1} H_s \frac{d\alpha_s}{d\omega_j}$ for any sector pair k, s .

Recall that $\omega^\top = \beta^\top \mathcal{L}(\alpha) \Leftrightarrow \alpha^\top \omega - \omega + \beta = 0$. Differentiating this expression with respect to ω_j , we get

$$\begin{aligned} \sum_{k=1}^n \omega_k \frac{d\alpha_k}{d\omega_j} + \alpha^\top \mathbf{1}_j - \mathbf{1}_j &= 0 \Leftrightarrow \sum_{k=1}^n \omega_k \left(H_k^{-1} H_s \frac{d\alpha_s}{d\omega_j} \right) + \alpha^\top \mathbf{1}_j - \mathbf{1}_j = 0 \Leftrightarrow \\ H_s \frac{d\alpha_s}{d\omega_j} &= \left(\sum_{k=1}^n \omega_k H_k^{-1} \right)^{-1} (I - \alpha^\top) \mathbf{1}_j. \end{aligned} \quad (81)$$

Combining this with (80), we get

$$(\nabla^2 \bar{a})^{-1} = \left(\mathcal{L}^{-1} \left(\sum_{k=1}^n \omega_k H_k^{-1} \right)^{-1} (\mathcal{L}^\top)^{-1} \right)^{-1} = \mathcal{L}^\top \left(\sum_{k=1}^n \omega_k H_k^{-1} \right) \mathcal{L}, \quad (82)$$

where $\mathcal{L}^{-1} = I - \alpha$. Note that all elements of $(\nabla^2 \bar{a})^{-1}$ are negative, $(\nabla^2 \bar{a})^{-1} < 0$, if $[H_i^{-1}]_{kl} \leq 0$ for all i and all $k \neq l$. Indeed, in that case $H_i^{-1} \leq 0$ and H_i^{-1} has a strictly negative diagonal since H_i is negative definite. Furthermore, $\omega_i > 0$ for all i , and all elements of $\mathcal{L} = I + \alpha + \alpha^2 + \dots$ are positive since $\alpha \in \text{int } \mathcal{A}$.

Next, consider $\mathcal{H}^{-1} = (\nabla^2 \bar{a} - (\rho - 1) \Sigma)^{-1}$. For $\Sigma = 0$ we have that $\mathcal{H}^{-1} = (\nabla^2 \bar{a})^{-1} < 0$. Since the function $(\nabla^2 \bar{a} - (\rho - 1) \Sigma)^{-1}$ is continuous in Σ , it follows that there exists a threshold $\bar{\Sigma} > 0$ such that $(\nabla^2 \bar{a} - (\rho - 1) \Sigma)^{-1} < 0$ for $\|\Sigma\| \leq \bar{\Sigma}$, which is the result. \square

D.5 Proof of Lemma 7

Lemma 7. *Suppose that all the TFP shifter functions (a_1, \dots, a_n) take the form 2, with $\alpha_i^\circ = \alpha_j^\circ$ for all i, j , and that H_i^{-1} is of the form (33) for all i . If $\alpha \in \text{int } \mathcal{A}$, there exists a scalar $\bar{\Sigma} > 0$ and a threshold $0 < \bar{s} < 1$ such that if $\|\Sigma\| \leq \bar{\Sigma}$ and $s > \bar{s}$, then all sectors are global substitutes, that is $\mathcal{H}_{ij}^{-1} > 0$ for all $i \neq j$.*

Proof. Imposing $\alpha_i^\circ = l^\circ$ and $H_i = H$ for all i , we can rewrite (24) as

$$\alpha_i(\omega) = l^\circ + \left(\sum_{j=1}^n \omega_j \right)^{-1} \left(\omega - \beta - l^\circ \sum_{j=1}^n \omega_j \right) = \left(\sum_{j=1}^n \omega_j \right)^{-1} (\omega - \beta).$$

Clearly, $\alpha_i(\omega) = l(\omega) > 0 \ \forall i$. Using the Sherman-Morrison formula, we obtain

$$\mathcal{L} = (I - \alpha)^{-1} = \left(I - \mathbf{1} l^\top \right)^{-1} = I + \frac{\mathbf{1} l^\top}{1 - l^\top \mathbf{1}} = I + \frac{1}{1 - \sum_j l_j} \begin{bmatrix} l_1 & l_2 & \dots & l_n \\ l_1 & l_2 & & \vdots \\ \vdots & & \ddots & l_n \\ l_1 & \dots & l_{n-1} & l_n \end{bmatrix}. \quad (83)$$

Plugging H_i^{-1} from (33) into (82), we get

$$(\nabla^2 \bar{a})_{im}^{-1} = \mathbf{1}_i^\top \mathcal{L}^\top \left(\sum_{k=1}^n \omega_k H_k^{-1} \right) \mathcal{L} \mathbf{1}_m = \left(\sum_{k=1}^n \omega_k \right) \sum_{k=1}^n \left(\mathcal{L}_{ki} \left(-\mathcal{L}_{km} + \frac{s}{n-1} \sum_{j \neq k} \mathcal{L}_{jm} \right) \right). \quad (84)$$

Notice that $(\nabla^2 \bar{a})_{im}^{-1}$ is a continuous and strictly increasing function of s (the latter is true because for $\alpha \in \text{int } \mathcal{A}$, all elements of \mathcal{L} are positive). Furthermore, $(\nabla^2 \bar{a})_{im}^{-1} < 0$ if $s = 0$. Using \mathcal{L} given by (83), we get for $i \neq m$

$$\begin{aligned}
(\nabla^2 \bar{a})_{im}^{-1} &= \left(\sum_{k=1}^n \omega_k \right) \left[\mathcal{L}_{mi} \left(-\mathcal{L}_{mm} + \frac{s}{n-1} \sum_{j \neq m} \mathcal{L}_{jm} \right) \right. \\
&\quad \left. + \mathcal{L}_{ii} \left(-\mathcal{L}_{im} + \frac{s}{n-1} \sum_{j \neq i} \mathcal{L}_{jm} \right) + \sum_{k \neq m, i} \left(\mathcal{L}_{ki} \left(-\mathcal{L}_{km} + \frac{s}{n-1} \sum_{j \neq k} \mathcal{L}_{jm} \right) \right) \right] \\
&= \left(\sum_{k=1}^n \omega_k \right) \left[\frac{l_i}{1 - \sum_j l_j} \left(-1 - \frac{(1-s)l_m}{1 - \sum_j l_j} \right) \right. \\
&\quad \left. + \left(1 + \frac{l_i}{1 - \sum_j l_j} \right) \left(\frac{s}{n-1} - \frac{(1-s)l_m}{1 - \sum_j l_j} \right) + \frac{l_i(n-2)}{1 - \sum_j l_j} \left(\frac{s}{n-1} - \frac{(1-s)l_m}{1 - \sum_j l_j} \right) \right] \\
&\xrightarrow{s \rightarrow 1} \left(\sum_{k=1}^n \omega_k \right) \frac{1}{n-1} > 0. \tag{85}
\end{aligned}$$

It follows that we have global (strict) substitution for $\Sigma = 0$ and $s \rightarrow 1$, and global (strict) complementarity for $\Sigma = 0$ and $s = 0$. Furthermore, $(\nabla^2 \bar{a} + J_{\mathcal{E}})^{-1}$ is continuous in s and Σ . Therefore, there exist thresholds $0 < \underline{s} \leq \bar{s} < 1$ and a threshold $\bar{\Sigma} > 0$ such that if $\|\Sigma\| \leq \bar{\Sigma}$ and $s > \bar{s}$ then $\left[(\nabla^2 \bar{a} + J_{\mathcal{E}})^{-1} \right]_{im} > 0$ for $i \neq m$, and if $\|\Sigma\| \leq \bar{\Sigma}$ and $s < \underline{s}$ then $\left[(\nabla^2 \bar{a} + J_{\mathcal{E}})^{-1} \right]_{im} < 0$ for $i \neq m$. \square

D.6 Proof of Proposition 4

Proposition 4. *If $\alpha \in \text{int } \mathcal{A}$, there exists a scalar $\bar{\Sigma} > 0$ such that if $\|\Sigma\| \leq \bar{\Sigma}$ the following holds.*

1. (Complementarity) *Suppose that input shares are local complements in the production of good i , that is $[H_i^{-1}]_{kl} < 0$ for all $k \neq l$. Then a beneficial change to k ($\partial \mathcal{E}_k / \partial \gamma > 0$) increases α_{ij} for all j .*
2. (Substitution) *Suppose that the conditions of Lemma 7 about the TFP shifters (a_1, \dots, a_n) hold. Then there exists a threshold $0 < \bar{s} < 1$ such that if $s > \bar{s}$, a beneficial change to k ($\partial \mathcal{E}_k / \partial \gamma > 0$) decreases α_{ij} for all i and all $j \neq k$, and increases α_{ik} for all i .*

Proof. From (81) we can write

$$\left(\frac{d\alpha_{ij}}{d\omega} \right)^\top = \mathbf{1}_j^\top H_i^{-1} \left(\sum_{l=1}^n \omega_l H_l^{-1} \right)^{-1} (I - \alpha^\top).$$

It follows that

$$\frac{d\alpha_{ij}}{d\gamma} = \left(\frac{d\alpha_{ij}}{d\omega} \right)^\top \underbrace{\frac{d\omega}{d\gamma}}_{\frac{d\alpha_{ij}}{d\omega}^\top} = \mathbf{1}_j^\top H_i^{-1} \left(\sum_{l=1}^n \omega_l H_l^{-1} \right)^{-1} (I - \alpha^\top) \underbrace{\left(-\mathcal{H}^{-1} \frac{\partial \mathcal{E}}{\partial \gamma} \right)}_{\frac{d\omega}{d\gamma}}. \quad (86)$$

If $\Sigma = 0$, then using (32) and (82) we find

$$\frac{d\alpha_{ij}}{d\gamma} = -\mathbf{1}_j^\top H_i^{-1} (I - \alpha)^{-1} \frac{\partial \mathcal{E}}{\partial \gamma}. \quad (87)$$

Recall that $(I - \alpha)^{-1} = I + \alpha + \alpha^2 + \dots > 0$ if $\alpha \in \text{int } \mathcal{A}$. It follows that under local complementarity ($H_i^{-1} < 0$), a beneficial change to \mathcal{E}_k increases α_{ij} for all i, j . Through the same argument as in the proof of Lemma 6, this holds for small $\|\Sigma\|$.

Suppose now that the TFP shifter functions (a_1, \dots, a_n) take the form 2 with $\alpha_i^\circ = l^\circ$, and $H_i = H$ is given by (33). Then $(I - \alpha)^{-1}$ is given by (83). Plugging those in (87), we obtain

$$\frac{d\alpha_{ij}}{d\gamma} = - \sum_{k \neq j} \left(\frac{s}{n-1} - \frac{(1-s)l_k}{1 - \sum_m l_m} \right) \frac{\partial \mathcal{E}_k}{\partial \gamma} - \left(-1 - \frac{(1-s)l_j}{1 - \sum_m l_m} \right) \frac{\partial \mathcal{E}_j}{\partial \gamma} \xrightarrow{s \rightarrow 1} -\frac{1}{n-1} \sum_{k \neq j} \frac{\partial \mathcal{E}_k}{\partial \gamma} + \frac{\partial \mathcal{E}_j}{\partial \gamma}.$$

The expression above is strictly negative if a beneficial shock hits sector $k \neq j$ ($\partial \mathcal{E}_k / \partial \gamma > 0$, $\partial \mathcal{E}_j / \partial \gamma = 0$), and is strictly positive if a beneficial shock hits sector j ($\partial \mathcal{E}_j / \partial \gamma > 0$, $\partial \mathcal{E}_k / \partial \gamma = 0$). Since the inequalities are strict, the same argument as in the proof of Lemma 7 applies and the results hold for $s > \bar{s} > 0$ and for small $\|\Sigma\|$. \square

D.7 Proof of Corollary 5

Proposition 5. *Let γ denote either the mean μ_i or an element of the covariance matrix Σ_{ij} . The equilibrium response to a change in beliefs γ must satisfy*

$$\underbrace{\frac{d\mathbb{E}[y]}{d\gamma} - \frac{\partial \mathbb{E}[y]}{\partial \gamma}}_{\text{Excess response of } \mathbb{E}[y]} = \frac{1}{2}(\rho - 1) \underbrace{\left(\frac{d\mathbb{V}[y]}{d\gamma} - \frac{\partial \mathbb{V}[y]}{\partial \gamma} \right)}_{\text{Excess response of } \mathbb{V}[y]}. \quad (41)$$

Proof. Suppose that the network is flexible. Then differentiating welfare with respect to γ implies

$$\frac{d\mathcal{W}}{d\gamma} = \frac{d\mathbb{E}[y]}{d\gamma} - \frac{1}{2}(\rho - 1) \frac{d\mathbb{V}[y]}{d\gamma}.$$

Applying the envelope theorem to (21) implies

$$\frac{d\mathcal{W}}{d\gamma} = \frac{\partial \mathbb{E}[y]}{\partial \gamma} - \frac{1}{2}(\rho - 1) \frac{\partial \mathbb{V}[y]}{\partial \gamma},$$

where partial derivatives, as usual, indicate that the network α is held fixed. Combining these two equations yields the result. \square

D.8 Proof of Corollary 7

Corollary 7. *Suppose that $\omega \in \text{int } \mathcal{O}$. There exists a threshold $\bar{\Sigma} < 0$ such that if $\Sigma_{kl} > \bar{\Sigma}$ for all k, l , then the following holds.*

1. *If all sectors are global complements with sector i , that is $\mathcal{H}_{ik}^{-1} < 0$ for $k \neq i$, then*

$$\frac{dE[y]}{d\mu_i} > \omega_i, \quad \text{and} \quad \frac{dV[y]}{d\mu_i} > 0.$$

2. *If all sectors are global complements with sectors i and j , that is $\mathcal{H}_{ik}^{-1} < 0$ and $\mathcal{H}_{jk}^{-1} < 0$ for $k \neq i, j$, then*

$$\frac{dE[y]}{d\Sigma_{ij}} < 0, \quad \text{and} \quad \frac{dV[y]}{d\Sigma_{ij}} < \omega_i \omega_j.$$

Proof. From part 1 of Proposition 7 and using (28),

$$\frac{dE[y]}{d\mu_i} = \omega_i - (\rho - 1) \sum_{jk} \omega_j \Sigma_{jk} \mathcal{H}_{ik}^{-1}, \quad \text{and} \quad \frac{dV[y]}{d\mu_i} = -2 \sum_{jk} \omega_j \Sigma_{jk} \mathcal{H}_{ik}^{-1}.$$

Since \mathcal{H}^{-1} is negative definite, $\mathcal{H}_{ii}^{-1} < 0$. Therefore, if $\mathcal{H}_{ik}^{-1} < 0$ for $k \neq i$, then there exists a threshold $\bar{\Sigma} < 0$ such that if $\Sigma_{jk} > \bar{\Sigma}$ for all j, k , then $\sum_{j,k} \omega_j \Sigma_{jk} \mathcal{H}_{ik}^{-1} < 0$ (recall that $\Sigma_{jj} > 0$ for all j). Therefore, $\frac{dE[y]}{d\mu_i} > \omega_i$ and $\frac{dV[y]}{d\mu_i} > 0$.

Using part 2 of Proposition 7 and (29), we can follow analogous steps to show that $\frac{dE[y]}{d\Sigma_{ii}} < 0$ and $\frac{dV[y]}{d\Sigma_{ii}} < \omega_i^2$ under the same conditions. Finally, from part 2 of Proposition 7 we have

$$\frac{dE[y]}{d\Sigma_{ij}} = \frac{1}{2} (\rho - 1)^2 \sum_{lk} \left(\omega_l \omega_j \Sigma_{lk} \mathcal{H}_{ik}^{-1} + \omega_l \omega_i \Sigma_{lk} \mathcal{H}_{jk}^{-1} \right),$$

and

$$\frac{dV[y]}{d\Sigma_{ij}} = \omega_i \omega_j + (\rho - 1) \sum_{lk} \left(\omega_l \omega_j \Sigma_{lk} \mathcal{H}_{ik}^{-1} + \omega_l \omega_i \Sigma_{lk} \mathcal{H}_{jk}^{-1} \right).$$

Since \mathcal{H}^{-1} is negative definite, $\mathcal{H}_{ii}^{-1}, \mathcal{H}_{jj}^{-1} < 0$. Therefore, if $\mathcal{H}_{ik}^{-1} < 0$ and $\mathcal{H}_{jk}^{-1} < 0$ for $k \neq i$ and $k \neq j$, then there exists a threshold $\bar{\Sigma} < 0$ such that if $\Sigma_{jk} > \bar{\Sigma} \forall j$, then $\sum_{lk} \left(\omega_l \omega_j \Sigma_{lk} \mathcal{H}_{ik}^{-1} + \omega_l \omega_i \Sigma_{lk} \mathcal{H}_{jk}^{-1} \right) < 0$ (recall that $\Sigma_{jj} > 0$ for all j). Therefore, $\frac{dE[y]}{d\Sigma_{ij}} < 0$ and $\frac{dV[y]}{d\Sigma_{ij}} < \omega_i \omega_j$. \square

D.9 Proof of Corollary 8

Corollary 8. *Suppose that $\omega \in \text{int } \mathcal{O}$. Then there exist thresholds $\underline{\Sigma} > 0$ and $\bar{\Sigma} > 0$ such that,*

1. If all sectors are global substitutes with sector i , that is $\mathcal{H}_{ik}^{-1} > 0$ for $k \neq i$, and sector i is not too risky while other sectors are sufficiently risky in the sense that $\Sigma_{ji} < \underline{\Sigma}$ for all j and $\Sigma_{jk} > \bar{\Sigma}$ for all $j, k \neq i$, then

$$\frac{dE[y]}{d\mu_i} < \omega_i, \quad \text{and} \quad \frac{dV[y]}{d\mu_i} < 0.$$

2. If all sectors are global substitutes with sectors i and j , that is $\mathcal{H}_{ik}^{-1} > 0$ and $\mathcal{H}_{jk}^{-1} > 0$ for $k \neq i, j$, and sectors i and j are not too risky while other sectors are sufficiently risky in the sense that $\Sigma_{li} < \underline{\Sigma}$ and $\Sigma_{lj} < \underline{\Sigma}$ for all l , and $\Sigma_{lk} > \bar{\Sigma}$ for all $l, k \neq i$ and $l, k \neq j$, then

$$\frac{dE[y]}{d\Sigma_{ij}} > 0, \quad \text{and} \quad \frac{dV[y]}{d\Sigma_{ij}} > \omega_i \omega_j.$$

Proof. From part 1 of Proposition 7 and using (28),

$$\frac{dE[y]}{d\mu_i} = \omega_i - (\rho - 1) \sum_{jk} \omega_j \Sigma_{jk} \mathcal{H}_{ik}^{-1}, \quad \text{and} \quad \frac{dV[y]}{d\mu_i} = -2 \sum_{jk} \omega_j \Sigma_{jk} \mathcal{H}_{ik}^{-1}.$$

Since \mathcal{H}^{-1} is negative definite, $\mathcal{H}_{ii}^{-1} < 0$. Therefore, if $\mathcal{H}_{ik}^{-1} > 0$ for $k \neq i$, then there exist thresholds $\underline{\Sigma} > 0$ and $\bar{\Sigma} > 0$ such that if $\Sigma_{jk} > \bar{\Sigma}$ for all $j, k \neq i$, and $\Sigma_{ji} < \underline{\Sigma}$ for all j , then $\sum_{j,k} \omega_j \Sigma_{jk} \mathcal{H}_{ik}^{-1} > 0$. Therefore, $\frac{dE[y]}{d\mu_i} < \omega_i$ and $\frac{dV[y]}{d\mu_i} < 0$.

Using part 2 of Proposition 7 and (29), we can follow analogous steps to show that $\frac{dE[y]}{d\Sigma_{ii}} > 0$ and $\frac{dV[y]}{d\Sigma_{ii}} > \omega_i^2$ under the same conditions.

Finally, from part 3 of Proposition 7 we have

$$\frac{dE[y]}{d\Sigma_{ij}} = \frac{1}{2} (\rho - 1)^2 \sum_{lk} \left(\omega_l \omega_j \Sigma_{lk} \mathcal{H}_{ik}^{-1} + \omega_l \omega_i \Sigma_{lk} \mathcal{H}_{jk}^{-1} \right),$$

and

$$\frac{dV[y]}{d\Sigma_{ij}} = \omega_i \omega_j + (\rho - 1) \sum_{lk} \left(\omega_l \omega_j \Sigma_{lk} \mathcal{H}_{ik}^{-1} + \omega_l \omega_i \Sigma_{lk} \mathcal{H}_{jk}^{-1} \right).$$

Since \mathcal{H}^{-1} is negative definite, $\mathcal{H}_{ii}^{-1}, \mathcal{H}_{jj}^{-1} < 0$. Therefore, if $\mathcal{H}_{ik}^{-1} > 0$ and $\mathcal{H}_{jk}^{-1} > 0$ for $k \neq i, j$, then there exist thresholds $\underline{\Sigma} > 0$ and $\bar{\Sigma} > 0$ such that if $\Sigma_{lk} > \bar{\Sigma}$ for all $l, k \neq i$ and $l, k \neq j$, and $\Sigma_{li}, \Sigma_{lj} < \underline{\Sigma}$ for all l , then $\sum_{lk} \left(\omega_l \omega_j \Sigma_{lk} \mathcal{H}_{ik}^{-1} + \omega_l \omega_i \Sigma_{lk} \mathcal{H}_{jk}^{-1} \right) > 0$. Therefore, $\frac{dE[y]}{d\Sigma_{ij}} > 0$ and $\frac{dV[y]}{d\Sigma_{ij}} > \omega_i \omega_j$. \square

E 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 n sectors/goods, but we are now explicit about the firms that operate within a sector. Specifically, in each sector $i \in \{1, \dots, n\}$ 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 from 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.⁴⁴ 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

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

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 $0 \leq \varsigma \leq 1$ of the total surplus. That is to say,

$$\tilde{P}_{il}^{jm} - K_i \left(\alpha_i^l, \left\{ \tilde{P}_k^{il} \right\}_{k \in \{1, \dots, n\}} \right) = \varsigma \left(B_i^{jm} - K_i \left(\alpha_i^l, \left\{ \tilde{P}_k^{il} \right\}_{k \in \{1, \dots, n\}} \right) \right), \quad (88)$$

⁴⁴For notational convenience, let $j = 0$ denote the household and assume that there is a unit mass of “sub-households” indexed by $m \in [0, 1]$.

where $K_i \left(\alpha_i^l, \left\{ \tilde{P}_k^{il} \right\}_{k \in \{1, \dots, n\}} \right)$ 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 (88) 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

$$\mathbb{E} \left[\Lambda \sum_{j=0}^n Q_{ji} dl \int_0^1 \varsigma \left(B_i^{jm} - K_i \left(\alpha_i^l, \left\{ \tilde{P}_k^{il} \right\}_{k \in \{1, \dots, n\}} \right) \right) dm \right],$$

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

$$\mathbb{E} \left[\Lambda Q_i K_i \left(\alpha_i^l, \left\{ \tilde{P}_k^{il} \right\}_{k \in \{1, \dots, n\}} \right) \right], \quad (89)$$

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 $\varsigma \rightarrow 0$, so that the bargaining power of the sellers goes to zero. In that case, (88) 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 (89) 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 $P = (P_1, \dots, P_n)$, and that all prices are equal to unit cost, as in the model in the main text.

F Extension of Proposition 3 to binding constraints

We can straightforwardly extend Proposition 3 to handle the case in which some of the constraints $\omega_i \geq 0$ bind with strictly positive Lagrange multipliers.⁴⁵ Note that these Domar weights ω_i will not respond to a marginal change in beliefs. We still assume, however, that the constraint $1 \geq \omega^\top (\mathbf{1} - \bar{\alpha})$ is slack.

Formally, define a set of indices $\mathcal{I} = \{i : \omega_i \geq \beta_i \text{ binds}\}$. For $j \notin \mathcal{I}$, $\omega_j > \beta_j$ and for $j \in \mathcal{I}$,

⁴⁵Thus, we rule out cases in which $\omega_i = \beta_i$ but the Lagrange multiplier corresponding to the $\omega_i \geq \beta_i$ is also zero. At such points, the derivative of ω_i with respect to a change in beliefs might be not defined.

$\omega_j = \beta_j$. Define an $\hat{n} \times 1$ vector ω^{nb} that contains elements ω_i such that $i \notin \mathcal{I}$, and an $(n - \hat{n}) \times 1$ vector ω^b that contains elements ω_i such that $i \in \mathcal{I}$, where $0 \leq \hat{n} \leq n$. Then ω^{nb} is implicitly given by the first-order conditions of (26), i.e.

$$\mathcal{E}^{nb} + \frac{d\bar{a}(\omega)}{d\omega^{nb}} = 0,$$

where \mathcal{E}^{nb} is an $\hat{n} \times 1$ vector obtained by deleting elements $k \in \mathcal{I}$ from \mathcal{E} , and $\frac{d\bar{a}(\omega)}{d\omega^{nb}}$ is an $\hat{n} \times 1$ vector obtained by differentiating (23) with respect to ω^{nb} . Applying the implicit function theorem, we can write

$$\frac{d\omega^{nb}}{d\gamma} = -\left(\mathcal{H}^{nb}\right)^{-1} \times \frac{\partial \mathcal{E}^{nb}}{\partial \gamma}. \quad (90)$$

In the expression above,

$$\mathcal{H}^{nb} = \frac{d\bar{a}(\omega)}{d\omega^{nb} d(\omega^{nb})^\top} - (\rho - 1) \Sigma^{nb},$$

where $\frac{d\bar{a}(\omega)}{d\omega^{nb} d(\omega^{nb})^\top}$ is an $\hat{n} \times \hat{n}$ Hessian matrix obtained by differentiating (23) twice with respect to ω^{nb} , and $\frac{\partial \mathcal{E}^{nb}}{\partial \omega^{nb}} = -(\rho - 1) \Sigma^{nb}$, where Σ^{nb} is an $\hat{n} \times \hat{n}$ matrix obtained by deleting columns and rows $k \in \mathcal{I}$ from Σ . Finally, compute $\frac{\partial \mathcal{E}^{nb}}{\partial \gamma}$ for $\gamma = \mu_i$ and Σ_{ij} :

$$\begin{aligned} \frac{\partial \mathcal{E}^{nb}}{\partial \mu_i} &= \mathbf{1}_i^{nb}, \\ \frac{\partial \mathcal{E}^{nb}}{\partial \Sigma_{ij}} &= -\frac{1}{2} (\rho - 1) \left(\omega_j \mathbf{1}_i^{nb} + \omega_i \mathbf{1}_j^{nb} \right), \end{aligned}$$

where $\mathbf{1}_i^{nb}$ is an $\hat{n} \times 1$ vector obtained by deleting elements $k \in \mathcal{I}$ from $\mathbf{1}_i$. In particular, if $i \in \mathcal{I}$ then all elements of $\mathbf{1}_i^{nb}$ are zeros. Consequently, if $i \in \mathcal{I}$ then a change in μ_i or Σ_{ii} has no impact on ω^{nb} .

G More details about the regressions of Section 9

In this section, we provide more details about the regressions presented in Tables 1 and 2. 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 $_j^c = \beta_j^c$ 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:

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

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 1 and 2) to isolate the impact of changes in their second moment alone. Note that we control for year \times customer \times supplier industry (2-digit SIC) fixed effects in Tables 1 and 2. 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.

H 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. Such a change is, however, 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 .

H.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 $y = \omega(\alpha)^\top (\varepsilon + a(\alpha))$, 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

$$\begin{aligned}\mathcal{W} &\equiv \max_{\alpha \in \mathcal{A}} E[y(\alpha)] - \frac{1}{2}(\rho - 1) V[y(\alpha)] \\ &= \max_{\alpha \in \mathcal{A}} \omega(\alpha)^\top (E[\varepsilon] + a(\alpha)) - \frac{1}{2}(\rho - 1) \omega(\alpha)^\top V[\varepsilon] \omega(\alpha).\end{aligned}$$

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 $\varepsilon \sim \mathcal{N}(\mu - \frac{1}{2}B, \Sigma)$ 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

$$\begin{aligned}\mathcal{W} &= \max_{\alpha \in \mathcal{A}} \omega(\alpha)^\top \left(\mu - \frac{1}{2}B + a(\alpha) \right) - \frac{1}{2}(\rho - 1) \omega(\alpha)^\top \Sigma \omega(\alpha) \\ &= \max_{\alpha \in \mathcal{A}} \omega(\alpha)^\top (\mu + a(\alpha)) - \frac{1}{2}\rho \omega(\alpha)^\top \Sigma \omega(\alpha) + \frac{1}{2}\omega(\alpha)^\top (\Sigma \omega(\alpha) - B).\end{aligned}$$

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 $\varepsilon \sim \mathcal{N}(\mu - \frac{1}{2}\Sigma \omega(\alpha), \Sigma)$ 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.

H.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

$$\varepsilon \sim \mathcal{N}(\mu - g(\alpha, \alpha^*, \Sigma), \Sigma), \quad (91)$$

where

$$g(\alpha, \alpha^*, \Sigma) = \frac{1}{2}\Sigma\mathcal{L}(\alpha^*)^\top \beta + \frac{1}{2}(\alpha - \alpha^*)^\top \mathcal{L}(\alpha^*)\Sigma\mathcal{L}(\alpha^*)^\top \beta. \quad (92)$$

The term α^* in this expression is the equilibrium network, so that in equilibrium we have $g(\alpha^*, \alpha^*, \Sigma) = \frac{1}{2}\Sigma\mathcal{L}(\alpha^*)^\top \beta$. When making decisions, the representative firm in sector i chooses $\alpha_i = (\alpha_{i1}, \dots, \alpha_{in})$ 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 (91)–(92) differs from the one discussed above, $\varepsilon \sim \mathcal{N}(\mu - \frac{1}{2}\Sigma\omega(\alpha), \Sigma)$, 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 (91)–(92) 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 (with $E[p(\alpha^*)] = -\mathcal{L}(\alpha^*)(\mu - g(\alpha^*, \alpha^*, \Sigma) + a(\alpha^*))$ in Lemma 2). Furthermore, as derived in Section H.1, the planner's objective is given by

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

Following the same steps as in the main text, we can establish that there exists a unique solution to the planner's problem. One can also establish that there exists a unique equilibrium and it is efficient. The proof is analogous to that of Proposition 1. The key step in that proof is to show that the first-order conditions of the planner's problem and of the firm's problem coincide. This is

indeed the case. For the planner's problem, we have

$$\frac{\partial a_i}{\partial \alpha_i} + \mathcal{L}(\alpha) (\mu + a(\alpha)) - \rho \mathcal{L}(\alpha) \Sigma \mathcal{L}(\alpha)^\top \beta + \chi_i^p - \gamma_i^p = 0. \quad (93)$$

Consider now the firm's problem. Combining (46) with (12), (44) and (45), we find that the problem of the representative firm in sector i can be written as

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

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

$$\begin{aligned} 0 = & \mathbf{1}_j^\top \mathcal{L}(\alpha^*) \Sigma \mathcal{L}(\alpha^*)^\top \beta - \frac{\partial a(\alpha_i)}{\partial \alpha_{ij}} - \mathbf{1}_j^\top \mathcal{L}(\alpha^*) (\mu + a(\alpha^*)) \\ & + \left(\mathbf{1}_j^\top \mathcal{L}(\alpha^*) \right) \Sigma \left((\alpha_i - \mathbf{1}_i - (1 - \rho) \beta)^\top \mathcal{L}(\alpha^*) + \mathbf{1}_i^\top \right)^\top - \chi_{ij}^e + \gamma_i^e, \end{aligned}$$

In equilibrium $\alpha = \alpha^*$ and so the above expression simplifies to

$$\frac{\partial a(\alpha_i^*)}{\partial \alpha_{ij}} + \mathbf{1}_j^\top \mathcal{L}(\alpha^*) (\mu + a(\alpha^*)) - \rho \beta^\top \mathcal{L}(\alpha^*) \Sigma \mathcal{L}(\alpha^*)^\top \mathbf{1}_j + \chi_{ij}^e - \gamma_i^e = 0,$$

which is equivalent to (93). Finally, the results in Sections 6 and 7 remain unchanged, with the only exception that $(\rho - 1)$ should be replaced by ρ .

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

One specification for ε which is used in the literature is $\varepsilon \sim \mathcal{N}(\mu - \frac{1}{2} \text{diag}(\Sigma), \Sigma)$. This adjustment implies that the expected value of firm-level TFP $E[\exp(\varepsilon_i)] = \exp(\mu_i - \frac{1}{2} \Sigma_{ii} + \frac{1}{2} \Sigma_{ii}) = \exp(\mu_i)$ 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. Almost all our analytical results are unaffected by this change in specification. The only differences appear when we take derivatives with respect to Σ . In that case, the results need to be adapted to capture the impact of Σ on the expected value of ε .

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

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


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

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


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 4 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 9. Overall, we find that uncertainty has a larger impact on the economy in this setting than in the baseline model of Section 8. As in Figure 4, 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 9). 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 9: The role of uncertainty when $\varepsilon \sim \mathcal{N}(\mu - \frac{1}{2}\text{diag}(\Sigma), \Sigma)$




(a) Difference in expected log GDP [%]



(b) Difference in expected standard deviation of log GDP [%]



(c) Difference in expected welfare [%]



(d) Difference in realized log GDP [%]

Notes: The differences between the series implied by the baseline model with $\varepsilon \sim \mathcal{N}(\mu - \frac{1}{2}\text{diag}(\Sigma), \Sigma)$ (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.

I Approximated economy

Section I.1 considers the economy in which the cost of deviating from the ideal shares is large. All the proofs are in Section I.2.

I.1 Large costs of deviating from the ideal input shares

In this section, we consider an economy in which the cost of deviating from the ideal input shares α° is large. Let $a_i(\alpha_i) = \bar{\kappa} \times \hat{a}_i(\alpha_i)$, where \hat{a}_i does not depend on $\bar{\kappa}$ and is such that $\hat{a}_i(\alpha_i^\circ) = 0$. Suppose that $\alpha_i^\circ \in \text{int } \mathcal{A}_i$. The parameter $\bar{\kappa} > 0$ captures how costly it is for the firms

to deviate from α° in terms of TFP losses.⁴⁶

Our goal is to characterize the economy when $\bar{\kappa} > 0$ is large. To do so, we use perturbation theory to express the equilibrium production network as a second-order approximation (Judd and Guu, 2001; Schmitt-Grohé and Uribe, 2004). More explicitly, let $\alpha(\bar{\kappa})$ denote the production network under a given cost shifter $\bar{\kappa}$, and consider the expansion

$$\alpha(\bar{\kappa}) = \alpha_{(0)} + \bar{\kappa}^{-1}\alpha_{(1)} + \bar{\kappa}^{-2}\alpha_{(2)} + O(\bar{\kappa}^{-3}), \quad (94)$$

where $\alpha_{(m)}$ denotes the m th order term. Notably, for a sufficiently large $\bar{\kappa}$, $\alpha(\bar{\kappa}) \in \text{int } \mathcal{A}$ because $\alpha^\circ \in \text{int } \mathcal{A}$ by assumption. We will provide expressions for these terms in what follows, but first it is convenient to rewrite some equilibrium quantities using the expansion (94).

Throughout, we will work with variables that are evaluated at the ideal input shares. We use the superscript \circ to denote these quantities. For instance, $\mathcal{L}^\circ = (I - \alpha^\circ)^{-1}$ is the Leontief inverse evaluated at α° , \hat{H}_i° is the Hessian of \hat{a}_i at α_i° , and so on.

The following lemma provides approximate expressions for the Leontief inverse and the TFP shifter.

Lemma 11. *The following holds:*

1. The Leontief inverse \mathcal{L} can be written as $\mathcal{L} = \mathcal{L}^\circ + \bar{\kappa}^{-1}\mathcal{L}_{(1)} + \bar{\kappa}^{-2}\mathcal{L}_{(2)} + O(\bar{\kappa}^{-3})$, where

$$\mathcal{L}_{(1)} = \mathcal{L}^\circ \alpha_{(1)} \mathcal{L}^\circ, \quad (95)$$

$$\mathcal{L}_{(2)} = \mathcal{L}^\circ \alpha_{(2)} \mathcal{L}^\circ + \mathcal{L}_{(1)} \alpha_{(1)} \mathcal{L}^\circ. \quad (96)$$

2. The TFP shifter \hat{a}_i can be written as $\hat{a}_i = \bar{\kappa}^{-2}\hat{a}_{i,(2)} + O(\bar{\kappa}^{-4})$, where

$$\hat{a}_{i,(2)} = \frac{1}{2} \alpha_{i,(1)}^\top \hat{H}_i^\circ \alpha_{i,(1)}, \quad (97)$$

and where $\alpha_{i,(1)}^\top$ is the i th row of $\alpha_{(1)}$.

The first part of the lemma states that to a first order the Leontief inverse can be expressed as a deviation from \mathcal{L}° that is linear in $\alpha_{(1)}$. Naturally, the second-order term is linear in $\alpha_{(2)}$ and quadratic in $\alpha_{(1)}$. Its second part shows that as the production network moves away from α° the TFP loss associated with that move depends on the curvature of the TFP shifter function captured by \hat{H}_i° . The zero-order term in the expression of \hat{a}_i is zero because $\hat{a}_i(\alpha_i^\circ) = 0$, and the first-order term $\nabla \hat{a}_i(\alpha_i^\circ) = 0$ since the ideal shares maximize \hat{a}_i .⁴⁷

⁴⁶Throughout this section we assume that all the third derivatives of \hat{a}_i are zero. Relaxing this assumption is straightforward but requires heavier notation that obscures the exposition of the mechanisms,

⁴⁷The third-order term is zero by assumption that all the third derivatives of \hat{a}_i are zero.

Approximated risk-adjusted prices

With these quantities in hand, we can derive an expression for the risk-adjusted price vector.

Lemma 12. *The risk-adjusted price vector \mathcal{R} can be written as $\mathcal{R} = \mathcal{R}_{(0)} + \bar{\kappa}^{-1}\mathcal{R}_{(1)} + O(\bar{\kappa}^{-2})$, where*

$$\mathcal{R}_{(0)} = \mathcal{R}^\circ = \underbrace{-\mathcal{L}^\circ \mu}_{\mathbb{E}[p^\circ]} + \underbrace{(\rho - 1) \mathcal{L}^\circ \Sigma \omega^\circ}_{\text{Cov}(p^\circ, \lambda^\circ)} \quad (98)$$

is the risk-adjusted price vector (18) evaluated at the ideal input shares α° , and

$$\mathcal{R}_{(1)} = \underbrace{-\mathcal{L}_{(1)}\mu - \mathcal{L}^\circ \hat{a}_{(2)}}_{\text{Change in } \mathbb{E}[p]} + \underbrace{(\rho - 1) \mathcal{L}^\circ \Sigma \mathcal{L}_{(1)}^\top \beta + (\rho - 1) \mathcal{L}_{(1)} \Sigma [\mathcal{L}^\circ]^\top \beta}_{\text{Change in } \text{Cov}(p, \lambda)}, \quad (99)$$

where $\mathcal{L}_{(1)}$ and $\hat{a}_{(2)}$ are given by (95) and (97).

Equation (98) shows that to a first-order approximation \mathcal{R} is simply equal to its value at the ideal input shares α° . The second-order term (99) describes that deviating from α° impacts \mathcal{R} in two ways. First, it changes the importance of different sectors as suppliers, as captured by $\mathcal{L}_{(1)}$, and it modifies the TFP losses associated with the choice of technique, captured by $\hat{a}_{(2)}$. Together, these two terms reflect the impact of α on the expected price vector $\mathbb{E}[p]$. Second, the change in α modifies the covariance between the stochastic discount factor and the price vector. The last two terms in (99) capture that channel.

Approximated production network and Domar weights

The following proposition provides a second-order approximation for the production network α .

Proposition 9. *If $\alpha \in \text{int } \mathcal{A}$, the equilibrium input shares in sector i are approximately given by*

$$\alpha_i = \alpha_i^\circ + \underbrace{\bar{\kappa}^{-1} \left(\hat{H}_i^\circ \right)^{-1} \mathcal{R}_{(0)}}_{\alpha_{i,(1)}} + \underbrace{\bar{\kappa}^{-2} \left(\hat{H}_i^\circ \right)^{-1} \mathcal{R}_{(1)}}_{\alpha_{i,(2)}} + O(\bar{\kappa}^{-3}), \quad (100)$$

where $\mathcal{R}_{(0)}$ and $\mathcal{R}_{(1)}$ are given by (98) and (99).

This expression, together with the two preceding lemmas, provides a closed-form characterization of the equilibrium production network in terms of the primitives of the model as a second-order approximation. To understand its structure, recall from Lemma (2) that the equilibrium can be characterized as the fixed point of the self-map given by the right-hand side of (17). We can iterate on this self-map to find the equilibrium network. This involves starting from an initial risk-adjusted price vector, computing the optimal production network under these prices, computing the new risk-adjusted prices that correspond to that network, and restarting that process again until convergence.

Proposition 9 mimics this structure. From Equation (19), the inverse Hessian H_i^{-1} captures how risk-adjusted prices affect the input shares chosen by firm i . The term $\alpha_i^\circ + \bar{\kappa}^{-1} \left(\hat{H}_i^\circ \right)^{-1} \mathcal{R}^\circ$ in (100) therefore corresponds to the firms optimal decision under \mathcal{R}° , which acts as the initial vector for the iterations. The following term in (100) captures the next step in the iteration process. The quantity $\mathcal{R}_{(1)}$ corresponds to the response of \mathcal{R} to the first round decision of the firms. Multiplying this quantity by $\left(\hat{H}_i^\circ \right)^{-1}$ then provides the reaction of the firms to this change in \mathcal{R} .

One implication of Proposition 9 is that further iterations of the equilibrium mapping have a decreasing impact on the production network. We see from the second term in (100) that the reaction of the firms to the original risk-adjusted prices is of order $\bar{\kappa}^{-1}$, while the second iteration term is of order $\bar{\kappa}^{-2}$. This suggests that when $\bar{\kappa}$ is large the first few rounds of iteration should provide an accurate picture of the production network.

Combining the expressions for the Leontief inverse (95) and (96) with the expressions for $\alpha_{(1)}$ and $\alpha_{(2)}$ given by (100), it is straightforward to derive a formula for the Domar weights.

Corollary 9. *The equilibrium vector of Domar weights is*

$$\omega = \omega^\circ + \underbrace{\bar{\kappa}^{-1} \mathcal{L}_{(1)}^\top \beta}_{\omega_{(1)}} + \underbrace{\bar{\kappa}^{-2} \mathcal{L}_{(2)}^\top \beta}_{\omega_{(2)}} + O(\bar{\kappa}^{-3}), \quad (101)$$

where $\omega^\circ = (\mathcal{L}^\circ)^\top \beta$, and $\mathcal{L}_{(1)}$ and $\mathcal{L}_{(2)}$ are given by (95) and (96).

Using (100) and (101), one can also explicitly show that the approximate formula (35) shown at the end of Section 6.1 is accurate if $\bar{\kappa}$ is sufficiently large.

Corollary 10. *The approximation of the Domar weight vector (35) is accurate if $\bar{\kappa}$ is large, such that*

$$-[\mathcal{H}^\circ]^{-1} \mathcal{E}^\circ = \bar{\kappa}^{-1} \omega_{(1)} + O(\bar{\kappa}^{-2}). \quad (102)$$

Approximated GDP

Recall that moments of log GDP are given by (14). Using our results above, it is straightforward to derive approximate expressions for $E y$ and $V y$.

Corollary 11. *In equilibrium, the mean and variance of log GDP are*

$$E y = (\omega^\circ)^\top \mu + \bar{\kappa}^{-1} \left(\omega_{(1)}^\top \mu + (\omega^\circ)^\top \hat{a}_{(2)} \right) + \bar{\kappa}^{-2} \left(\omega_{(2)}^\top \mu + \omega_{(1)}^\top \hat{a}_{(2)} \right) + O(\bar{\kappa}^{-3})$$

and

$$V y = (\omega^\circ)^\top \Sigma \omega^\circ + 2\bar{\kappa}^{-1} (\omega^\circ)^\top \Sigma \omega_{(1)} + \bar{\kappa}^{-2} \left(2(\omega^\circ)^\top \Sigma \omega_{(2)} + \omega_{(1)}^\top \Sigma \omega_{(1)} \right) + O(\bar{\kappa}^{-3}).$$

Response of the production network to changes in beliefs

We now provide two results that describe, as closed-form second-order approximation, the response of the production network to changes in beliefs.

The following corollary provides a closed-form expression for $\frac{d\alpha_i}{d\mu_k}$ as a second-order approximation.

Corollary 12. *The impact of an increase in μ_k on the network is given by*

$$\frac{d\alpha_i}{d\mu_k} = \bar{\kappa}^{-1} \left(\hat{H}_i^\circ \right)^{-1} \underbrace{\left(\frac{\partial \mathcal{R}_{(0)}}{\partial \mu_k} + \bar{\kappa}^{-1} \frac{\partial \mathcal{R}_{(1)}}{\partial \mu_k} \right)}_{\text{Response of } \mathcal{R} \text{ with fixed network}} + \bar{\kappa}^{-2} \left(\hat{H}_i^\circ \right)^{-1} \underbrace{\left(\sum_{lm} \frac{d\alpha_{lm,(1)}}{d\mu_k} \times \frac{\partial \mathcal{R}_{(1)}}{\partial \alpha_{lm,(1)}} \right)}_{\text{Impact of the network on } \mathcal{R}} + O(\bar{\kappa}^{-3}), \quad (103)$$

where the impact of μ_k on \mathcal{R} taking the network fixed is given by

$$\frac{\partial \mathcal{R}_{(0)}}{\partial \mu_k} + \bar{\kappa}^{-1} \frac{\partial \mathcal{R}_{(1)}}{\partial \mu_k} = -\mathcal{L}^\circ \mathbf{1}_k - \bar{\kappa}^{-1} \mathcal{L}_{(1)} \mathbf{1}_k, \quad (104)$$

and the change in \mathcal{R} through the response of the network is given by

$$\frac{d\alpha_{lm,(1)}}{d\mu_k} \times \frac{\partial \mathcal{R}_{(1)}}{\partial \alpha_{lm,(1)}} = -\mathbf{1}_m^\top \left(\hat{H}_l^\circ \right)^{-1} \mathcal{L}^\circ \mathbf{1}_k \times (\rho - 1) \mathcal{L}^\circ \Sigma \left(\mathcal{L}^\circ \mathbf{1}_l \mathbf{1}_m^\top \mathcal{L}^\circ \right)^\top \beta. \quad (105)$$

The following corollary provides a similar result for the derivative of the network with respect to Σ_{op} for $o \neq p$. Its results also apply for the case $o = p$ if all the terms $\frac{1}{2} (\mathbf{1}_o \mathbf{1}_p^\top + \mathbf{1}_p \mathbf{1}_o^\top)$ are replaced by $\mathbf{1}_o \mathbf{1}_o^\top$, as in this case there is no need to take two derivatives to preserve the symmetry of Σ .

Corollary 13. *The impact of an increase in Σ_{op} on the network is given by*

$$\frac{d\alpha_i}{d\Sigma_{op}} = \bar{\kappa}^{-1} \left(\hat{H}_i^\circ \right)^{-1} \underbrace{\left(\frac{\partial \mathcal{R}_{(0)}}{\partial \Sigma_{op}} + \bar{\kappa}^{-1} \frac{\partial \mathcal{R}_{(1)}}{\partial \Sigma_{op}} \right)}_{\text{Response of } \mathcal{R} \text{ with fixed network}} + \bar{\kappa}^{-2} \left(\hat{H}_i^\circ \right)^{-1} \underbrace{\left(\sum_{lm} \frac{d\alpha_{lm,(1)}}{d\Sigma_{op}} \times \frac{\partial \mathcal{R}_{(1)}}{\partial \alpha_{op,(1)}} \right)}_{\text{Impact of the network on } \mathcal{R}} + O(\bar{\kappa}^{-3}), \quad (106)$$

where the impact of Σ_{op} on \mathcal{R} taking the network fixed is given by

$$\begin{aligned} \frac{\partial \mathcal{R}_{(0)}}{\partial \Sigma_{op}} + \bar{\kappa}^{-1} \frac{\partial \mathcal{R}_{(1)}}{\partial \Sigma_{op}} &= (\rho - 1) \mathcal{L}^\circ \left[\frac{1}{2} \left(\mathbf{1}_o \mathbf{1}_p^\top + \mathbf{1}_p \mathbf{1}_o^\top \right) \right] \omega^\circ \\ &+ \bar{\kappa}^{-1} \left[(\rho - 1) \mathcal{L}^\circ \left[\frac{1}{2} \left(\mathbf{1}_o \mathbf{1}_p^\top + \mathbf{1}_p \mathbf{1}_o^\top \right) \right] \mathcal{L}_{(1)}^\top \beta + (\rho - 1) \mathcal{L}_{(1)} \left[\frac{1}{2} \left(\mathbf{1}_o \mathbf{1}_p^\top + \mathbf{1}_p \mathbf{1}_o^\top \right) \right] (\mathcal{L}^\circ)^\top \beta \right], \end{aligned}$$

and the change in \mathcal{R} through the response of the network is given by

$$\frac{d\alpha_{lm,(1)}}{d\Sigma_{op}} \times \frac{\partial \mathcal{R}_{(1)}}{\partial \alpha_{lm,(1)}} = (\rho - 1) \left(\hat{H}_i^\circ \right)^{-1} \mathcal{L}^\circ \left[\frac{1}{2} \left(\mathbf{1}_o \mathbf{1}_p^\top + \mathbf{1}_p \mathbf{1}_o^\top \right) \right] \omega^\circ \times (\rho - 1) \mathcal{L}^\circ \Sigma \left(\mathcal{L}^\circ \mathbf{1}_l \mathbf{1}_m^\top \mathcal{L}^\circ \right)^\top \beta. \quad (107)$$

I.2 Proofs related to the approximation

Proof of Lemma 11

Lemma 11. *The following holds:*

1. The Leontief inverse \mathcal{L} can be written as $\mathcal{L} = \mathcal{L}^\circ + \bar{\kappa}^{-1} \mathcal{L}_{(1)} + \bar{\kappa}^{-2} \mathcal{L}_{(2)} + O(\bar{\kappa}^{-3})$, where

$$\mathcal{L}_{(1)} = \mathcal{L}^\circ \alpha_{(1)} \mathcal{L}^\circ, \quad (95)$$

$$\mathcal{L}_{(2)} = \mathcal{L}^\circ \alpha_{(2)} \mathcal{L}^\circ + \mathcal{L}_{(1)} \alpha_{(1)} \mathcal{L}^\circ. \quad (96)$$

2. The TFP shifter \hat{a}_i can be written as $\hat{a}_i = \bar{\kappa}^{-2} \hat{a}_{i,(2)} + O(\bar{\kappa}^{-4})$, where

$$\hat{a}_{i,(2)} = \frac{1}{2} \alpha_{i,(1)}^\top \hat{H}_i^\circ \alpha_{i,(1)}, \quad (97)$$

and where $\alpha_{i,(1)}^\top$ is the i th row of $\alpha_{(1)}$.

Proof. The rules of differentiation for a matrix inverse imply that

$$\frac{d\mathcal{L}(\alpha)}{d\alpha_{kl}} = \frac{d\left((I - \alpha)^{-1}\right)}{d\alpha_{kl}} = (I - \alpha)^{-1} \frac{d\alpha}{d\alpha_{kl}} (I - \alpha)^{-1} = \mathcal{L}(\alpha) \mathbf{1}_k \mathbf{1}_l^\top \mathcal{L}(\alpha),$$

and

$$\left[\frac{d\mathcal{L}(\alpha)}{d\alpha_{kl}} \right]_{ij} = \mathbf{1}_i^\top \mathcal{L}(\alpha) \mathbf{1}_k \mathbf{1}_l^\top \mathcal{L}(\alpha) \mathbf{1}_j = [\mathcal{L}(\alpha)]_{ik} [\mathcal{L}(\alpha)]_{lj}.$$

By Taylor's theorem, and using the notation (94) for the deviation from α° , we find

$$\begin{aligned} \mathcal{L}_{ij} &= \mathcal{L}_{ij}^\circ + \sum_{kl} \mathcal{L}_{ik}^\circ (\bar{\kappa}^{-1} \alpha_{kl,(1)}) \mathcal{L}_{lj}^\circ + \sum_{kl} \mathcal{L}_{ik}^\circ (\bar{\kappa}^{-2} \alpha_{kl,(2)}) \mathcal{L}_{lj}^\circ \\ &\quad + \frac{1}{2} \left\{ \sum_{kl} \left(\sum_{sm} \mathcal{L}_{is}^\circ (\bar{\kappa}^{-1} \alpha_{sm,(1)}) \mathcal{L}_{mk}^\circ \right) (\bar{\kappa}^{-1} \alpha_{kl,(1)}) \mathcal{L}_{lj}^\circ \right. \\ &\quad \left. + \sum_{kl} \mathcal{L}_{ik}^\circ (\bar{\kappa}^{-1} \alpha_{kl,(1)}) \left(\sum_{sm} \mathcal{L}_{ls}^\circ (\bar{\kappa}^{-1} \alpha_{sm,(1)}) \mathcal{L}_{mj}^\circ \right) \right\} + O(\bar{\kappa}^{-3}), \end{aligned}$$

which can be written in the matrix form as

$$\mathcal{L} = \mathcal{L}^\circ + \bar{\kappa}^{-1} \mathcal{L}^\circ \alpha_{(1)} \mathcal{L}^\circ + \bar{\kappa}^{-2} (\mathcal{L}^\circ \alpha_{(2)} \mathcal{L}^\circ + \mathcal{L}^\circ \alpha_{(1)} \mathcal{L}^\circ \alpha_{(1)} \mathcal{L}^\circ) + O(\bar{\kappa}^{-3}).$$

This completes the proof of the first part of the lemma. For the second part, we can write

$$\hat{a}_i(\alpha_i) \approx \hat{a}_i(\alpha_i^\circ) + (\alpha_i - \alpha_i^\circ)^\top \left(\begin{array}{c} \frac{\partial \hat{a}_i}{\partial \alpha_{i1}} \\ \vdots \\ \frac{\partial \hat{a}_i}{\partial \alpha_{in}} \end{array} \right)_{\alpha_i = \alpha_i^\circ} + \frac{1}{2} (\alpha_i - \alpha_i^\circ)^\top \hat{H}_i^\circ (\alpha_i - \alpha_i^\circ) + O((\alpha_i - \alpha_i^\circ)^4).$$

The first term is equal to zero by the definition of \hat{a}_i . The second term is equal to zero since α_i° maximizes \hat{a}_i . The third-order term is zero by assumption that all the third derivatives of \hat{a}_i are zero. Using the notation (94) for the deviation from α° , we can write

$$\hat{a}_i(\alpha_i) \approx \frac{1}{2} \bar{\kappa}^{-2} \alpha_{i,(1)}^\top \hat{H}_i^\circ \alpha_{i,(1)} + O(\bar{\kappa}^{-4}),$$

which is the result. \square

Proof of Lemma 12

Lemma 12. *The risk-adjusted price vector \mathcal{R} can be written as $\mathcal{R} = \mathcal{R}_{(0)} + \bar{\kappa}^{-1} \mathcal{R}_{(1)} + O(\bar{\kappa}^{-2})$, where*

$$\mathcal{R}_{(0)} = \mathcal{R}^\circ = \underbrace{-\mathcal{L}^\circ \mu}_{\mathbb{E}[p^\circ]} + \underbrace{(\rho - 1) \mathcal{L}^\circ \Sigma \omega^\circ}_{\text{Cov}(p^\circ, \lambda^\circ)} \quad (98)$$

is the risk-adjusted price vector (18) evaluated at the ideal input shares α° , and

$$\mathcal{R}_{(1)} = \underbrace{-\mathcal{L}_{(1)} \mu - \mathcal{L}^\circ \hat{a}_{(2)}}_{\text{Change in } \mathbb{E}[p]} + \underbrace{(\rho - 1) \mathcal{L}^\circ \Sigma \mathcal{L}_{(1)}^\top \beta + (\rho - 1) \mathcal{L}_{(1)} \Sigma [\mathcal{L}^\circ]^\top \beta}_{\text{Change in } \text{Cov}(p, \lambda)}, \quad (99)$$

where $\mathcal{L}_{(1)}$ and $\hat{a}_{(2)}$ are given by (95) and (97).

Proof. We can write \mathcal{R} as

$$\mathcal{R} = -\mathcal{L}(\alpha) (\mu + \bar{\kappa} \hat{a}(\alpha)) + (\rho - 1) \mathcal{L}(\alpha) \Sigma [\mathcal{L}(\alpha)]^\top \beta, \quad (108)$$

or

$$\begin{aligned} \mathcal{R} = & -(\mathcal{L}^\circ + \bar{\kappa}^{-1} \mathcal{L}_{(1)} + O(\bar{\kappa}^{-2})) (\mu + \bar{\kappa} \times \bar{\kappa}^{-2} \hat{a}_{(2)} + \bar{\kappa} O(\bar{\kappa}^{-3})) \\ & + (\rho - 1) (\mathcal{L}^\circ + \bar{\kappa}^{-1} \mathcal{L}_{(1)} + O(\bar{\kappa}^{-2})) \Sigma (\mathcal{L}^\circ + \bar{\kappa}^{-1} \mathcal{L}_{(1)} + O(\bar{\kappa}^{-2}))^\top \beta, \end{aligned}$$

where we used the expressions in Lemma 11. Grouping terms yields the result. \square

Proof of Proposition 9

Proposition 9. *If $\alpha \in \text{int } \mathcal{A}$, the equilibrium vector of input shares in sector i is*

$$\alpha_i = \alpha_i^\circ + \underbrace{\bar{\kappa}^{-1} \left(\hat{H}_i^\circ \right)^{-1} \mathcal{R}_{(0)}}_{\alpha_{i,(1)}} + \underbrace{\bar{\kappa}^{-2} \left(\hat{H}_i^\circ \right)^{-1} \mathcal{R}_{(1)}}_{\alpha_{i,(2)}} + O(\bar{\kappa}^{-3}), \quad (100)$$

where $\mathcal{R}_{(0)}$ and $\mathcal{R}_{(1)}$ are given by (98) and (99).

Proof. Since $\alpha^\circ \in \text{int } \mathcal{A}$, when $\bar{\kappa}$ is large enough the equilibrium network α is also in the interior of \mathcal{A} . From (19), that equilibrium is then fully described by the first-order conditions of the problem (17), such that

$$\frac{\partial \hat{a}_i(\alpha_i)}{\partial \alpha_{ij}} = \bar{\kappa}^{-1} \times \mathcal{R}_j(\alpha), \quad (109)$$

and where $\mathcal{R}_j(\alpha)$ is given by (108). We can write the left-hand side of this equation as

$$\begin{aligned} \frac{\partial \hat{a}_i(\alpha_i)}{\partial \alpha_{ij}} &= \left(\frac{\partial \hat{a}_i}{\partial \alpha_{ij}} \right)_{\alpha_i = \alpha_i^\circ} + (\alpha_i - \alpha_i^\circ)^\top \left(\begin{array}{c} \frac{\partial^2 \hat{a}_i}{\partial \alpha_{i1} \partial \alpha_{ij}} \\ \dots \\ \frac{\partial^2 \hat{a}_i}{\partial \alpha_{in} \partial \alpha_{ij}} \end{array} \right)_{\alpha_i = \alpha_i^\circ} \\ &\quad + \frac{1}{2} (\alpha_i - \alpha_i^\circ)^\top \left(\begin{array}{ccc} \frac{\partial^3 \hat{a}_i}{\partial \alpha_{i1}^2 \partial \alpha_{ij}} & \dots & \frac{\partial^3 \hat{a}_i}{\partial \alpha_{i1} \partial \alpha_{in} \partial \alpha_{ij}} \\ \dots & \dots & \dots \\ \frac{\partial^3 \hat{a}_i}{\partial \alpha_{i1} \partial \alpha_{in} \partial \alpha_{ij}} & \dots & \frac{\partial^3 \hat{a}_i}{\partial \alpha_{in}^2 \partial \alpha_{ij}} \end{array} \right)_{\alpha_i = \alpha_i^\circ} (\alpha_i - \alpha_i^\circ) + O\left((\alpha_i - \alpha_i^\circ)^3\right). \end{aligned}$$

The first term in the right-hand side of this expression is zero since α° maximizes \hat{a}_i . The third term is also zero given our assumption that the third derivatives of \hat{a}_i are zero. Using the notation from (94), we can therefore write

$$\frac{\partial \hat{a}_i(\alpha_i)}{\partial \alpha_{ij}} = (\bar{\kappa}^{-1} \alpha_{i,(1)} + \bar{\kappa}^{-2} \alpha_{i,(2)})^\top \hat{H}_i^\circ \mathbf{1}_j + O(\bar{\kappa}^{-3}).$$

Combing with the right-hand side of (109) and the expressions of Lemma 12, we find

$$(\bar{\kappa}^{-1} \alpha_{i,(1)} + \bar{\kappa}^{-2} \alpha_{i,(2)})^\top \hat{H}_i^\circ \mathbf{1}_j + O(\bar{\kappa}^{-3}) = \bar{\kappa}^{-1} (\mathcal{R}_{j,(0)} + \bar{\kappa}^{-1} \mathcal{R}_{j,(1)} + O(\bar{\kappa}^{-2})).$$

Since this expression must be valid for all $\bar{\kappa}$, we find that

$$\alpha_{i,1} = \left(\hat{H}_i^\circ \right)^{-1} \mathcal{R}^\circ$$

and

$$\alpha_{i,2} = \left(\hat{H}_i^\circ \right)^{-1} \mathcal{R}_{(1)},$$

which is the result. \square

Proof of Corollary 9

Corollary 9. *The equilibrium vector of Domar weights is*

$$\omega = \omega^\circ + \underbrace{\bar{\kappa}^{-1} \mathcal{L}_{(1)}^\top \beta}_{\omega_{(1)}} + \underbrace{\kappa^{-2} \mathcal{L}_{(2)}^\top \beta}_{\omega_{(2)}} + O(\bar{\kappa}^{-3}), \quad (101)$$

where $\omega^\circ = (\mathcal{L}^\circ)^\top \beta$, and $\mathcal{L}_{(1)}$ and $\mathcal{L}_{(2)}$ are given by (95) and (96).

Proof. This corollary immediately follows from the definition of the Domar weights, $\omega = \mathcal{L}(\alpha)^\top \beta$. \square

Proof of Corollary 10

Corollary 10. *The approximation of the Domar weight vector (35) is accurate if $\bar{\kappa}$ is large, such that*

$$-[\mathcal{H}^\circ]^{-1} \mathcal{E}^\circ = \bar{\kappa}^{-1} \omega_{(1)} + O(\bar{\kappa}^{-2}). \quad (102)$$

Proof. Using (18) and (32), we can express the left-hand side of (102) as

$$-[\mathcal{H}^\circ]^{-1} \mathcal{E}^\circ = -[\nabla^2 \bar{a}(\omega^\circ) - (\rho - 1) \Sigma]^{-1} (\mu - (\rho - 1) \Sigma \omega^\circ).$$

Using (82) to express $\nabla^2 \bar{a}(\omega^\circ)$, we get

$$\begin{aligned} -[\mathcal{H}^\circ]^{-1} \mathcal{E}^\circ &= -\left[(\mathcal{L}^\circ)^{-1} \left(\sum_{k=1}^n \omega_k^\circ (H_k^\circ)^{-1} \right)^{-1} \left[(\mathcal{L}^\circ)^\top \right]^{-1} - (\rho - 1) \Sigma \right]^{-1} (\mu - (\rho - 1) \Sigma \omega^\circ) \\ &= -\left[\bar{\kappa} (\mathcal{L}^\circ)^{-1} \left(\sum_{k=1}^n \omega_k^\circ (\hat{H}_k^\circ)^{-1} \right)^{-1} \left[(\mathcal{L}^\circ)^\top \right]^{-1} - (\rho - 1) \Sigma \right]^{-1} (\mu - (\rho - 1) \Sigma \omega^\circ) \\ &= -\bar{\kappa}^{-1} (\mathcal{L}^\circ)^\top \sum_{k=1}^n \left(\omega_k^\circ (\hat{H}_k^\circ)^{-1} \right) \mathcal{L}^\circ (\mu - (\rho - 1) \Sigma \omega^\circ) + O(\bar{\kappa}^{-2}). \end{aligned}$$

From (100), the expression above can be rewritten as

$$-[\mathcal{H}^\circ]^{-1} \mathcal{E}^\circ = \bar{\kappa}^{-1} (\mathcal{L}^\circ)^\top \sum_{k=1}^n \omega_k^\circ \alpha_{k,(1)} + O(\bar{\kappa}^{-2}) = \bar{\kappa}^{-1} \mathcal{L}_{(1)}^\top \beta + O(\bar{\kappa}^{-2}),$$

where $\mathcal{L}_{(1)}$ is given by (95) and $\omega_{(1)} = \mathcal{L}_{(1)}^\top \beta$ by (101). \square

Proof of Corollary 11

Corollary 11. *In equilibrium, the mean and variance of log GDP are*

$$\mathbb{E} y = (\omega^\circ)^\top \mu + \bar{\kappa}^{-1} \left(\omega_{(1)}^\top \mu + (\omega^\circ)^\top \hat{a}_{(2)} \right) + \bar{\kappa}^{-2} \left(\omega_{(2)}^\top \mu + \omega_{(1)}^\top \hat{a}_{(2)} \right) + O(\bar{\kappa}^{-3})$$

and

$$\mathbb{V} y = (\omega^\circ)^\top \Sigma \omega^\circ + 2\bar{\kappa}^{-1} (\omega^\circ)^\top \Sigma \omega_{(1)} + \bar{\kappa}^{-2} \left(2(\omega^\circ)^\top \Sigma \omega_{(2)} + \omega_{(1)}^\top \Sigma \omega_{(1)} \right) + O(\bar{\kappa}^{-3}).$$

Proof. These expressions are straightforward to derive by plugging in (14) expressions (97) and (101). \square

Proof of Corollary 12

Corollary 12. *The impact of an increase in μ_k on the network is given by*

$$\frac{d\alpha_i}{d\mu_k} = \bar{\kappa}^{-1} \left(\hat{H}_i^\circ \right)^{-1} \underbrace{\left(\frac{\partial \mathcal{R}_{(0)}}{\partial \mu_k} + \bar{\kappa}^{-1} \frac{\partial \mathcal{R}_{(1)}}{\partial \mu_k} \right)}_{\text{Response of } \mathcal{R} \text{ with fixed network}} + \bar{\kappa}^{-2} \left(\hat{H}_i^\circ \right)^{-1} \underbrace{\left(\sum_{lm} \frac{\partial \mathcal{R}_{(1)}}{\partial \alpha_{lm,(1)}} \frac{d\alpha_{lm,(1)}}{d\mu_k} \right)}_{\text{Impact of the network on } \mathcal{R}} + O(\bar{\kappa}^{-3}), \quad (103)$$

where the impact of μ_k on \mathcal{R} taking the network fixed is given by

$$\frac{\partial \mathcal{R}_{(0)}}{\partial \mu_k} + \bar{\kappa}^{-1} \frac{\partial \mathcal{R}_{(1)}}{\partial \mu_k} = -\mathcal{L}^\circ \mathbf{1}_k - \bar{\kappa}^{-1} \mathcal{L}_{(1)} \mathbf{1}_k, \quad (104)$$

and the change in \mathcal{R} through the response of the network is given by

$$\frac{d\alpha_{lm,(1)}}{d\mu_k} \times \frac{\partial \mathcal{R}_{(1)}}{\partial \alpha_{lm,(1)}} = -\mathbf{1}_m^\top \left(\hat{H}_l^\circ \right)^{-1} \mathcal{L}^\circ \mathbf{1}_k \times (\rho - 1) \mathcal{L}^\circ \Sigma \left(\mathcal{L}^\circ \mathbf{1}_l \mathbf{1}_m^\top \mathcal{L}^\circ \right)^\top \beta. \quad (105)$$

Proof. Differentiating (94) with respect to μ_k yields

$$\frac{d\alpha_i}{d\mu_k} = \bar{\kappa}^{-1} \frac{d\alpha_{i,(1)}}{d\mu_k} + \bar{\kappa}^{-2} \frac{d\alpha_{i,(2)}}{d\mu_k} + O(\bar{\kappa}^{-3}).$$

Using the expressions for the first and second-order terms given by the Proposition 9, we find

$$\frac{d\alpha_i}{d\mu_k} = \bar{\kappa}^{-1} \left(\hat{H}_i^\circ \right)^{-1} \frac{d\mathcal{R}^\circ}{d\mu_k} + \bar{\kappa}^{-2} \left(\hat{H}_i^\circ \right)^{-1} \left(\sum_{lm} \frac{\partial \mathcal{R}_{(1)}}{\partial \alpha_{lm,(1)}} \frac{d\alpha_{lm,(1)}}{d\mu_k} + \frac{\partial \mathcal{R}_{(1)}}{\partial \mu_k} \right) + O(\bar{\kappa}^{-3}),$$

where, as usual, the partial derivatives imply that the production network is kept constant. From

(98), we have

$$\frac{d\mathcal{R}^\circ}{d\mu_k} = -\mathcal{L}^\circ \mathbf{1}_k,$$

and from (99), we have

$$\frac{\partial \mathcal{R}_{(1)}}{\partial \mu_k} = -\mathcal{L}_{(1)} \mathbf{1}_k.$$

Similarly, we use (99) to compute the partial derivative

$$\begin{aligned} \frac{\partial \mathcal{R}_{(1)}}{\partial \alpha_{ij,(1)}} &= -\mathcal{L}^\circ \mathbf{1}_i \mathbf{1}_j^\top \mathcal{L}^\circ \mu - \mathcal{L}^\circ \mathbf{1}_i \mathbf{1}_j^\top \hat{H}_i^\circ \alpha_{i,(1)} + (\rho - 1) \mathcal{L}^\circ \Sigma \left(\mathcal{L}^\circ \mathbf{1}_i \mathbf{1}_j^\top \mathcal{L}^\circ \right)^\top \beta + (\rho - 1) \left(\mathcal{L}^\circ \mathbf{1}_i \mathbf{1}_j^\top \mathcal{L}^\circ \right) \Sigma [\mathcal{L}^\circ]^\top \beta \\ &= \mathcal{L}^\circ \mathbf{1}_i \mathbf{1}_j^\top \left(-\mathcal{L}^\circ \mu + (\rho - 1) \mathcal{L}^\circ \Sigma [\mathcal{L}^\circ]^\top \beta - \hat{H}_i^\circ \alpha_{i,(1)} \right) + (\rho - 1) \mathcal{L}^\circ \Sigma \left(\mathcal{L}^\circ \mathbf{1}_i \mathbf{1}_j^\top \mathcal{L}^\circ \right)^\top \beta \\ &= \mathcal{L}^\circ \mathbf{1}_i \mathbf{1}_j^\top \left(\mathcal{R}_{(0)} - \hat{H}_i^\circ \alpha_{i,(1)} \right) + (\rho - 1) \mathcal{L}^\circ \Sigma \left(\mathcal{L}^\circ \mathbf{1}_i \mathbf{1}_j^\top \mathcal{L}^\circ \right)^\top \beta. \end{aligned} \tag{110}$$

From Proposition 9, $\alpha_{i,(1)} = \left(\hat{H}_i^\circ \right)^{-1} \mathcal{R}_{(0)}$, hence the first term in the last line of the expression above is zero. Finally, from (100) we find

$$\frac{d\alpha_{lm,(1)}}{d\mu_k} = -\mathbf{1}_m^\top \left(\hat{H}_i^\circ \right)^{-1} \mathcal{L}^\circ \mathbf{1}_k.$$

Grouping terms completes the proof. \square

Proof of Corollary 13

Corollary 13. *The impact of an increase in Σ_{op} on the network is given by*

$$\frac{d\alpha_i}{d\Sigma_{op}} = \bar{\kappa}^{-1} \left(\hat{H}_i^\circ \right)^{-1} \underbrace{\left(\frac{\partial \mathcal{R}_{(0)}}{\partial \Sigma_{op}} + \bar{\kappa}^{-1} \frac{\partial \mathcal{R}_{(1)}}{\partial \Sigma_{op}} \right)}_{\text{Response of } \mathcal{R} \text{ with fixed network}} + \bar{\kappa}^{-2} \left(\hat{H}_i^\circ \right)^{-1} \underbrace{\left(\sum_{lm} \frac{d\alpha_{lm,(1)}}{d\Sigma_{op}} \times \frac{\partial \mathcal{R}_{(1)}}{\partial \alpha_{op,(1)}} \right)}_{\text{Impact of the network on } \mathcal{R}} + O(\bar{\kappa}^{-3}), \tag{106}$$

where the impact of Σ_{op} on \mathcal{R} taking the network fixed is given by

$$\begin{aligned} \frac{\partial \mathcal{R}_{(0)}}{\partial \Sigma_{op}} + \bar{\kappa}^{-1} \frac{\partial \mathcal{R}_{(1)}}{\partial \Sigma_{op}} &= (\rho - 1) \mathcal{L}^\circ \left[\frac{1}{2} \left(\mathbf{1}_o \mathbf{1}_p^\top + \mathbf{1}_p \mathbf{1}_o^\top \right) \right] \omega^\circ \\ &\quad + \bar{\kappa}^{-1} \left[(\rho - 1) \mathcal{L}^\circ \left[\frac{1}{2} \left(\mathbf{1}_o \mathbf{1}_p^\top + \mathbf{1}_p \mathbf{1}_o^\top \right) \right] \mathcal{L}_{(1)}^\top \beta + (\rho - 1) \mathcal{L}_{(1)} \left[\frac{1}{2} \left(\mathbf{1}_o \mathbf{1}_p^\top + \mathbf{1}_p \mathbf{1}_o^\top \right) \right] (\mathcal{L}^\circ)^\top \beta \right], \end{aligned}$$

and the change in \mathcal{R} through the response of the network is given by

$$\frac{d\alpha_{lm,(1)}}{d\Sigma_{op}} \times \frac{\partial \mathcal{R}_{(1)}}{\partial \alpha_{lm,(1)}} = (\rho - 1) \left(\hat{H}_i^\circ \right)^{-1} \mathcal{L}^\circ \left[\frac{1}{2} \left(\mathbf{1}_o \mathbf{1}_p^\top + \mathbf{1}_p \mathbf{1}_o^\top \right) \right] \omega^\circ \times (\rho - 1) \mathcal{L}^\circ \Sigma \left(\mathcal{L}^\circ \mathbf{1}_l \mathbf{1}_m^\top \mathcal{L}^\circ \right)^\top \beta. \tag{107}$$

Proof. Differentiating (94) with respect to Σ_{op} yields

$$\begin{aligned}\frac{d\alpha_i}{d\Sigma_{op}} &= \bar{\kappa}^{-1} \frac{d\alpha_{i,(1)}}{d\Sigma_{op}} + \bar{\kappa}^{-2} \frac{d\alpha_{i,(2)}}{d\Sigma_{op}} + O(\bar{\kappa}^{-3}) \\ &= \bar{\kappa}^{-1} \left(\hat{H}_i^\circ \right)^{-1} \frac{d\mathcal{R}^\circ}{d\Sigma_{op}} + \bar{\kappa}^{-2} \left(\hat{H}_i^\circ \right)^{-1} \left(\sum_{lm} \frac{\partial \mathcal{R}_{(1)}}{\partial \alpha_{lm,(1)}} \frac{d\alpha_{lm,(1)}}{d\Sigma_{op}} + \frac{\partial \mathcal{R}_{(1)}}{\partial \Sigma_{op}} \right) + O(\bar{\kappa}^{-3}),\end{aligned}$$

where we have used the expressions for the first and second-order terms given by the Proposition 9. From (98), we have

$$\frac{d\mathcal{R}^\circ}{d\Sigma_{op}} = (\rho - 1) \mathcal{L}^\circ \left[\frac{1}{2} \left(\mathbf{1}_o \mathbf{1}_p^\top + \mathbf{1}_p \mathbf{1}_o^\top \right) \right] \omega^\circ,$$

and from (99), we have

$$\frac{\partial \mathcal{R}_{(1)}}{\partial \Sigma_{op}} = (\rho - 1) \mathcal{L}^\circ \left[\frac{1}{2} \left(\mathbf{1}_o \mathbf{1}_p^\top + \mathbf{1}_p \mathbf{1}_o^\top \right) \right] \mathcal{L}_{(1)}^\top \beta + (\rho - 1) \mathcal{L}_{(1)} \left[\frac{1}{2} \left(\mathbf{1}_o \mathbf{1}_p^\top + \mathbf{1}_p \mathbf{1}_o^\top \right) \right] [\mathcal{L}^\circ]^\top \beta.$$

From (110), we know that

$$\frac{\partial \mathcal{R}_{(1)}}{\partial \alpha_{ij,(1)}} = (\rho - 1) \mathcal{L}^\circ \Sigma \left(\mathcal{L}^\circ \mathbf{1}_i \mathbf{1}_j^\top \mathcal{L}^\circ \right)^\top \beta.$$

Finally, from (100), we can compute

$$\frac{d\alpha_{lm,(1)}}{d\Sigma_{op}} = \left(\hat{H}_i^\circ \right)^{-1} \frac{d\mathcal{R}^\circ}{d\Sigma_{op}} = (\rho - 1) \left(\hat{H}_i^\circ \right)^{-1} \mathcal{L}^\circ \left[\frac{1}{2} \left(\mathbf{1}_o \mathbf{1}_p^\top + \mathbf{1}_p \mathbf{1}_o^\top \right) \right] \omega^\circ.$$

□

J Additional information about the calibrated economy

We provide here additional information about the calibrated economy.

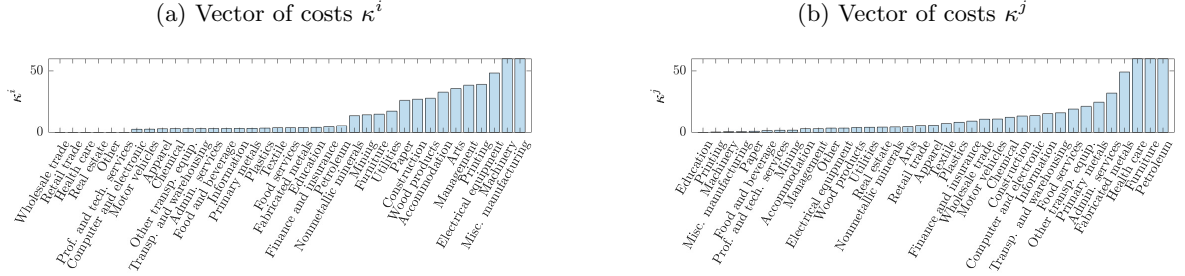
J.1 Cost of deviating from the ideal input shares

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 10 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 10: The calibrated costs of deviating from the ideal input shares



J.2 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.

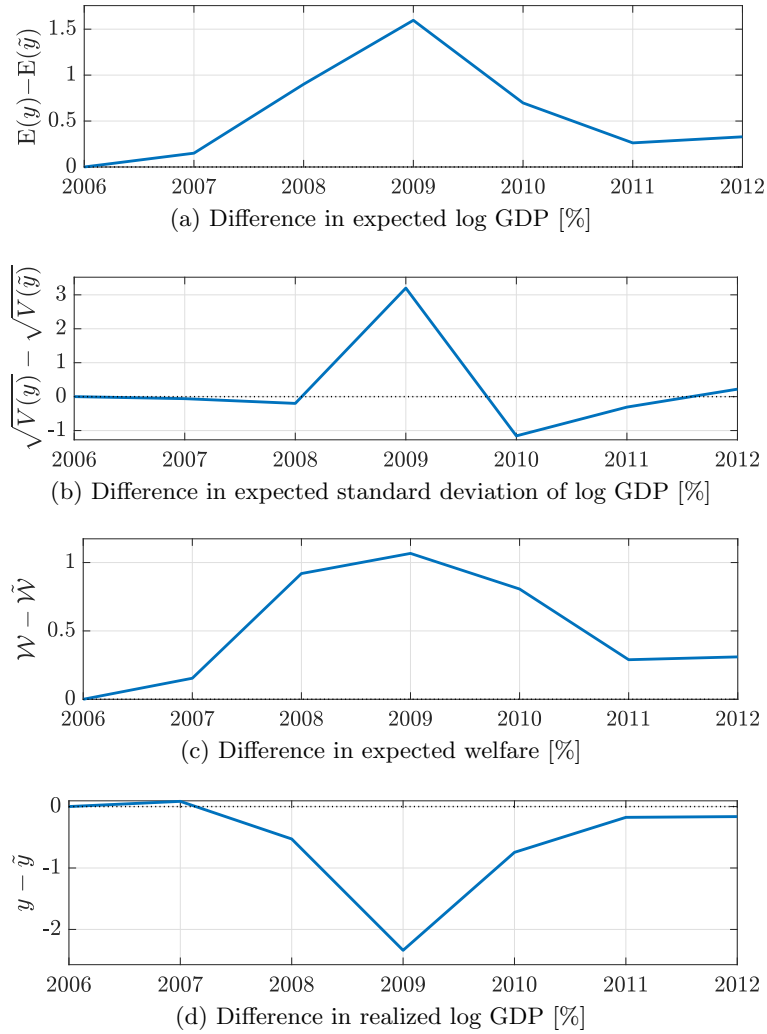
Similarly, the estimated covariance matrix $\hat{\Sigma}_t$ suggests that there is also substantial dispersion in uncertainty across sectors. The most volatile productivity is found in “Electrical equipment” with an average $\sqrt{\hat{\Sigma}_{iit}}$ of 38.0%, and the least volatile sector is “Real 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%.

J.3 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 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.

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.

J.4 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 8: 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 7 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.

Table 7: Domar weights and input shares in the model and in the data


Statistic	Version of the model			
	Data	Baseline	$\Sigma_t = 0$	Known ε_t
(1) Average Domar weight $\bar{\omega}_j$	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

Notes: For each sector, we compute the time series of its Domar weight ω_{jt} , as well as 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 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 the planner selects the network as if $\Sigma = 0$. The “Known ε_t ” 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 12 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.

Figure 12: Average of absolute differences in Domar weight growths in the postwar period



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.

J.5 Robustness exercises

In this appendix, we provide two robustness exercises. First, we investigate what happens with the coefficient of risk aversion varies. Second, we explore how the calibrated economy behaves when β can change every period.

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).

The economy also responds to the spike in uncertainty during the Great Recession much more for $\rho = 10$ (Figure 13). 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.

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

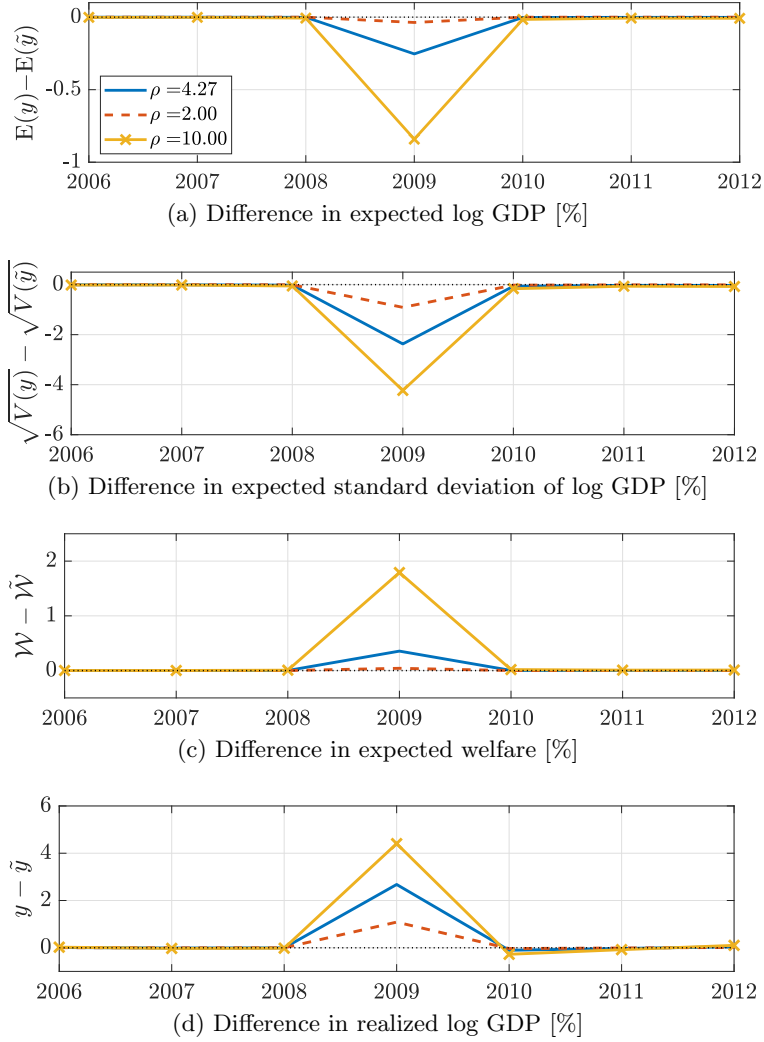
	Comparison with $\Sigma = 0$ model		
	$\rho = 2$	$\rho = 4.27$ (baseline)	$\rho = 10$
Expected log GDP $E[y(\alpha)]$	+0.001%	+0.008%	+0.033%
Std. dev. of log GDP $\sqrt{V[y(\alpha)]}$	+0.038%	+0.105%	+0.208%
Welfare \mathcal{W}	-0.001%	-0.010%	-0.057%

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

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 14 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 4 (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 13: 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.

Figure 14: 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.

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 along the lines of [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

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

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 it follows that

$$\lambda = (\rho - 1) \beta^\top p - \rho \log(1 + T), \quad (111)$$

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 (10) becomes

$$P_i = (1 + \tau_i) K_i(\alpha, P), \quad (112)$$

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 ,

$$p = -\mathcal{L}(\alpha) (\varepsilon + a(\alpha) - \log(1 + \tau)), \quad (113)$$

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 (11) for good i must also be adjusted for the potential loss in resources. It becomes

$$Q_i \left(1 - (1 - \zeta_i) \frac{\tau_i}{1 + \tau_i} \right) = C_i + \sum_j X_{ji}. \quad (114)$$

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 (112) and (114), we get

$$T(\alpha) = \begin{pmatrix} \zeta_1 \tau_1 \\ \zeta_2 \tau_2 \\ \vdots \\ \zeta_n \tau_n \end{pmatrix}^\top \left[\begin{pmatrix} 1 + \zeta_1 \tau_1 & 0 & \dots & 0 \\ 0 & 1 + \zeta_2 \tau_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 + \zeta_n \tau_n \end{pmatrix} - \beta \begin{pmatrix} \zeta_1 \tau_1 \\ \zeta_2 \tau_2 \\ \vdots \\ \zeta_n \tau_n \end{pmatrix}^\top - \alpha^\top \right]^{-1} \beta, \quad (115)$$

and so we can fully characterize the stochastic discount factor (111) 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

$$y = \beta^\top \mathcal{L}(\alpha) (\varepsilon + a(\alpha) - \log(1 + \tau)) + \log(1 + T(\alpha, \zeta, \tau)). \quad (116)$$

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

Proof. Denote $\chi_i = \zeta_i \tau_i$. Rewrite (115) as

$$T(\alpha, \chi) = \chi^\top \left[I + \text{diag}(\chi) - \beta \chi^\top - \alpha^\top \right]^{-1} \beta,$$

where $\text{diag}(\chi)$ is a diagonal matrix with the vector χ on its main diagonal. Note that $I + \text{diag}(\chi) - \beta \chi^\top - \alpha^\top$ is a diagonally dominant Z-matrix with a positive main diagonal. Therefore, it is an M-matrix whose inverse has all nonnegative inputs. Therefore, $\left[I + \text{diag}(\chi) - \beta \chi^\top - \alpha^\top \right]^{-1}$ is a nonnegative matrix as well. Differentiating $T(\alpha, \chi)$ with respect to α_{ij} yields

$$\frac{\partial T}{\partial \alpha_{ij}} = \chi^\top \left[I + \text{diag}(\chi) - \beta \chi^\top - \alpha^\top \right]^{-1} \left(\mathbf{1}_j \mathbf{1}_i^\top \right) \left[I + \text{diag}(\chi) - \beta \chi^\top - \alpha^\top \right]^{-1} \beta.$$

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

Firm decision

In the presence of wedges, the technique choice by firms is described in the following lemma. It is a direct analogue of Lemma 2 from the main text. The only difference is that expected equilibrium log prices $p^*(\alpha)$ in 118 needs to be adjusted by $\mathcal{L}(\alpha^*) \log(1 + \tau)$, as follows from Equation (113).

Lemma 14. *In the presence of wedges, the equilibrium technique choice problem of the representative firm in sector i is*

$$\alpha_i^* \in \arg \max_{\alpha_i \in \mathcal{A}_i} a_i(\alpha_i) - \sum_{j=1}^n \alpha_{ij} \mathcal{R}_j(\alpha^*), \quad (117)$$

where

$$\mathcal{R}(\alpha^*) = \mathbb{E}[p(\alpha^*)] + \text{Cov}[p(\alpha^*), \lambda(\alpha^*)], \quad (118)$$

is the equilibrium risk-adjusted price of good j , and where

$$\mathbb{E}[p(\alpha^*)] = -\mathcal{L}(\alpha^*) (\mu - \log(1 + \tau) + a(\alpha^*)) \text{ and } \text{Cov}[p(\alpha^*), \lambda(\alpha^*)] = (\rho - 1) \mathcal{L}(\alpha^*) \Sigma [\mathcal{L}(\alpha^*)]^\top \beta.$$

Proof. The proof is analogous to that of Lemma 2, with log prices given by (113) instead of (12). \square

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. 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

$$\mathcal{W}^d \equiv \max_{\alpha \in \mathcal{A}} \mathbb{E} \left[y^d(\alpha) \right] - \frac{1}{2} (\rho - 1) \mathbb{V} \left[y^d(\alpha) \right], \quad (119)$$

where

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

is what log GDP *would be* if nothing was rebated to the household.

We can rewrite the distorted planner's problem in terms of Domar weights.

$$\mathcal{W}^d \equiv \max_{\omega \in \mathcal{O}} \omega^\top (\mu - \log(1 + \tau)) + \bar{a}(\omega) - \frac{1}{2} (\rho - 1) \omega^\top \Sigma \omega, \quad (120)$$

where

$$\begin{aligned} \bar{a}(\omega) &= \max_{\alpha \in \mathcal{A}} \omega^\top a(\alpha) \\ \text{s.t. } \omega^\top &= \beta^\top \mathcal{L}(\alpha). \end{aligned} \quad (121)$$

Similar to the main model, the distorted planner's problem has a unique solution because the distorted planner's objective function is concave in ω , and there exists unique $\alpha = \alpha(\omega)$ solving (121). Furthermore, following the same steps as in the proof of Proposition 1, we can find that any solution to the firm's problem is also a solution to the distorted planner's problem.

Proposition 10. *There exists a unique equilibrium.*

Proof. The proof is analogous to that of Proposition 1. □

The unique equilibrium is no longer efficient because the distorted planner does not maximize the true welfare of the representative household. Specifically, the distorted planner does not take into account the fact that part of the distortion income is rebated to the household.

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 11. *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 2. Recall that

$$\mathcal{W}^d = \max_{\omega \in \mathcal{O}} \omega^\top (\mu - \log(1 + \tau)) + \bar{a}(\omega) - \frac{1}{2} (\rho - 1) \omega^\top \Sigma \omega.$$

By the envelope theorem,

$$\frac{d\mathcal{W}^d}{d\mu_i} = -\frac{d\mathcal{W}^d}{d\log(1+\tau_i)} = \omega_i \text{ and } \frac{d\mathcal{W}^d}{d\Sigma_{ij}} = -\frac{1}{2}(\rho-1)\omega_i\omega_j.$$

With these derivatives in hand, we can follow analogous steps to those in the proof of Proposition 2 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. \square

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

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

where $\alpha^* = \alpha^*(\omega^*)$ solves (121) and ω^* solves (120).

Proposition 12. *In the distorted equilibrium, the following holds.*

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

$$\frac{d\mathcal{W}}{d\mu_i} = \omega_i + \frac{1}{1+T(\alpha^*)} \sum_{k,l} \frac{\partial T}{\partial \alpha_{kl}} \frac{d\alpha_{kl}^*}{d\mu_i}.$$

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

$$\frac{d\mathcal{W}}{d\Sigma_{ij}} = -\frac{1}{2}(\rho-1)\omega_i\omega_j + \frac{1}{1+T(\alpha^*)} \sum_{k,l} \frac{\partial T}{\partial \alpha_{kl}} \frac{d\alpha_{kl}^*}{d\Sigma_{ij}}.$$

3. *If $\zeta = 0$, then the results of Proposition 4 hold in the distorted equilibrium.*

Proof. Taking derivative of \mathcal{W} with respect to μ_i yields

$$\frac{d\mathcal{W}}{d\mu_i} = \frac{d\mathcal{W}^d(\alpha^*)}{d\mu_i} + \frac{1}{1+T(\alpha^*)} \sum_{k,l} \frac{\partial T}{\partial \alpha_{kl}} \frac{d\alpha_{kl}^*}{d\mu_i} = \omega_i + \frac{1}{1+T(\alpha^*)} \sum_{k,l} \frac{\partial T}{\partial \alpha_{kl}} \frac{d\alpha_{kl}^*}{d\mu_i}.$$

Similar steps yield the expression for the derivative with respect to Σ_{ij} . Note that $T = 0$ when $\zeta = 0$ and so $\frac{\partial T}{\partial \alpha_{kl}} = 0$ for all k, l in this case. We then get the expressions of Proposition 4. \square

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 4. 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, if conditions of part 1 of Proposition 4 are satisfied, an increase in μ_i leads to an increase in all shares, that is, $\frac{d\alpha_{kl}^*}{d\mu_i} > 0$ for all k, l . This in turn increases the rebated amount T by Lemma 13 and, as a result, $\frac{d\mathcal{W}}{d\mu_i} > \omega_i$.

L Other forms of uncertainty

In this appendix, we discuss how our model can be extended to handle other types of uncertainty. We assume that when picking techniques, firms not only face uncertainty about the productivity shocks ε but are also uncertain about 1) household's preferences; 2) labor supply; 3) distortions. As in the baseline model, all uncertainties are realized before firms pick quantities and the household picks its consumption basket. We start by describing how these additional types of uncertainty can be introduced in our framework.

Household's preferences The household's demand for different good types is determined by the vector of preference parameters β . In this appendix, we assume that β is unknown when firms pick production techniques. We maintain the restrictions that $\beta_i > 0$ for all i and $\sum_{i=1}^n \beta_i = 1$. One natural distribution that satisfies these restrictions is the Dirichlet distribution.

Labor supply In the baseline model, the labor supply is fixed at one. In this appendix, we assume instead that the total labor supply is given by $\bar{L} > 0$ and that \bar{L} is unknown when firms pick production techniques. Uncertainty in \bar{L} might stem from, for example, changes in immigration, retirement, or health (e.g., stay-at-home orders or mandated shutdowns) regulations that affect the size of the labor force.

Distortions We introduce distortions in the same way as in Appendix K. Specifically, the price of good i exceeds the unit cost for firms, namely, $P_i = (1 + \tau_i) K_i(\alpha_i, P)$, where $K_i(\alpha, P)$ is given by (8). We assume that a fraction $1 - \zeta_i \in [0, 1]$ of the distortion revenue is pure waste. Note that these distortions can be driven by various sources, including government interventions and markups.⁴⁸ To study the role of distortion uncertainty, we assume that when firms pick their production techniques, they are uncertain about τ .

To focus on how different types of uncertainty affect equilibrium network, we are going to make the following assumption.

Assumption 2. *The shocks to productivity (ε), preferences (β), distortions (τ) and labor supply (\bar{L}) are independent.*

Analysis

After all uncertainties are realized, the representative household chooses its consumption basket and firms pick quantities of production factors given their production techniques α . Several parts

⁴⁸Our modeling of distortions follows Acemoglu and Azar (2020). Liu (2019) considers a richer structure of government interventions. Importantly, these papers do not study the role of distortion uncertainty.

of the model remain the same as in the main text. In particular, the objective function of the household remains unchanged. Its budget constraint must however be adjusted to

$$\sum_{i=1}^n P_i C_i \leq \bar{L} + T,$$

where the first term on the right-hand side is the household's labor income (recall that we can normalize wage to one without loss of generality), and the second term on the right-hand side $T = \sum_{i=1}^n \zeta_i \frac{\tau_i P_i Q_i}{1+\tau_i}$ is the amount of distortion revenue rebated to the household. A different form of the right-hand side of the budget constraint implies a different stochastic discount factor Λ . From the first-order conditions of the household we can write the log of the stochastic discount factor as

$$\lambda = (\rho - 1) \beta^\top p - \rho \log(\bar{L} + T). \quad (122)$$

On the firms' side, the cost minimization problem (7) is unaffected, 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 now turn to the market clearing conditions and the pricing equations. Those are not the same as in our main model. In particular, the pricing equation (10) becomes

$$P_i = (1 + \tau_i) K_i(\alpha, P). \quad (123)$$

We can combine this equation with (8) to write an expression for log prices as a function of the network α ,

$$p = -\mathcal{L}(\alpha) (\varepsilon + a(\alpha) - \log(1 + \tau)), \quad (124)$$

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

The market clearing condition (11) for good i must also be adjusted for the potential loss in resources. It becomes

$$Q_i \left(1 - (1 - \zeta_i) \frac{\tau_i}{1 + \tau_i} \right) = C_i + \sum_j X_{ji}. \quad (125)$$

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 (123) and (125), we get

$$T = \bar{L} \sum_{i=1}^n \zeta_i \tau_i \hat{\omega}_i, \quad (126)$$

where

$$\hat{\omega} = \left[\begin{pmatrix} 1 + \zeta_1 \tau_1 & 0 & \dots & 0 \\ 0 & 1 + \zeta_2 \tau_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 + \zeta_n \tau_n \end{pmatrix} - \beta \begin{pmatrix} \zeta_1 \tau_1 \\ \zeta_2 \tau_2 \\ \vdots \\ \zeta_n \tau_n \end{pmatrix}^\top - \alpha^\top \right]^{-1} \beta \quad (127)$$

is a vector of sectoral weights in the total factor income, $\hat{\omega}_i = \frac{P_i Q_i}{WL} = \frac{P_i Q_i}{L}$. In the absence of taxes ($\tau = 0$) or if taxes are purely wasteful ($\zeta = 0$), the rebate amount is zero, $T = 0$, and $\hat{\omega} = \omega = \mathcal{L}^\top \beta$ is the Domar weight vector.

In the main model with only ε uncertainty, firms in sector i prefer techniques that are more productive, in the sense of increasing $a_i(\alpha_i)$, and have low risk-adjusted prices \mathcal{R}_j (Lemma 2). The following lemma shows how this result can be extended in the model with multiple sources of uncertainty.

Lemma 15. *The technique choice of the representative firm in sector i solves*

$$\alpha_i^* \in \arg \min_{\alpha_i \in \mathcal{A}_i} -a_i(\alpha_i) + \sum_{j=1}^n \alpha_{ij} \hat{\mathcal{R}}_{ij}(\alpha^*), \quad (128)$$

where

$$\hat{\mathcal{R}}_{ij} = \mathbb{E}[p_j] + \frac{\text{Cov}[p_j, \Lambda \hat{\omega}_i]}{\mathbb{E}[\Lambda \hat{\omega}_i]}. \quad (129)$$

Here Λ is the marginal utility of the representative household with respect to changes in income, and $\hat{\omega}_i$ is the weight of sector i in the total factor income.

Proof. The objective function of firm in sector i is given by (9). The log stochastic discount factor is (122), where T is given by (126); demand for sector i 's good is $Q_i = \frac{\hat{\omega}_i \bar{L}}{P_i}$, where $\hat{\omega}$ is given by (127) and $\log P$ is given by (124); and the unit cost function is (8). Combining them together, the objective function (9) becomes

$$\alpha_i^* \in \arg \min_{\alpha_i \in \mathcal{A}_i} \mathbb{E} \exp \left(\underbrace{(\rho - 1) \sum_i \beta_i p_i - \rho \log \bar{L} - \rho \log \left(1 + \sum_{i=1}^n \zeta_i \tau_i \hat{\omega}_i \right)}_{=\log \Lambda} \right. \\ \left. + \underbrace{\log \bar{L} + \log \hat{\omega}_i + \log(1 + \tau_i) - p_i - \varepsilon_i - a_i(\alpha_i)}_{=\log Q_i} + \underbrace{\sum_j \alpha_{ij} p_j}_{=\log K_i} \right). \quad (130)$$

Taking the first-order condition with respect to α_{ij} and imposing that in equilibrium $\alpha = \alpha^*$, we

find

$$\mathbb{E} \left[\Lambda \bar{L} \hat{\omega}_i \left(-\frac{\partial a(\alpha_i^*)}{\partial \alpha_{ij}} + p_j \right) \right] + \chi_{ij}^e - \gamma_i^e = 0. \quad (131)$$

Notice that we can write $\Lambda(\bar{L}) = \bar{L}^{-\rho} \Lambda(\bar{L} = 1)$. Therefore, under Assumption 2, the equilibrium network is unaffected by \bar{L} . Then (131) can be rewritten as

$$\mathbb{E} [\Lambda \hat{\omega}_i] \left(-\frac{\partial a(\alpha_i^*)}{\partial \alpha_{ij}} + \mathbb{E} [p_j] \right) + \text{Cov} [\Lambda \hat{\omega}_i, p_j] + \tilde{\chi}_{ij}^e - \tilde{\gamma}_i^e = 0, \quad (132)$$

where we redefine Lagrange multipliers by dividing them by $\mathbb{E} \bar{L}$. This is the first-order condition of (128) with respect to α_{ij} . \square

From (8), a marginal change in α_{ij} changes the log unit cost of firms in sector i by $-\frac{\partial a_i}{\partial \alpha_{ij}} + p_j$. Firms care not only about this effect but also how it comoves with the household's marginal utility with respect to changes in income Λ , adjusted by the relative importance of sector i , i.e. $\hat{\omega}_i$. Under normality of the productivity shocks ε , (129) can be rewritten as⁴⁹

$$\hat{\mathcal{R}}_{ij} = \mathbb{E}_{\beta, \tau} [\mathcal{R}_j] + \frac{\text{Cov}_{\beta, \tau} [\mathcal{R}_j, (\mathbb{E}_\varepsilon \Lambda) \hat{\omega}_i]}{\mathbb{E}_{\beta, \tau} [(\mathbb{E}_\varepsilon \Lambda) \hat{\omega}_i]}, \quad (133)$$

where $\mathbb{E}_x(\cdot)$ means that we take expectations with respect to the random variable x , and where \mathcal{R} is given by

$$\mathcal{R} = \underbrace{-\mathcal{L}(\alpha^*) (\mu + a(\alpha^*) - \log(1 + \tau))}_{\mathbb{E}_\varepsilon [p]} + \underbrace{(\rho - 1) \mathcal{L}(\alpha^*) \Sigma [\mathcal{L}(\alpha^*)]^\top \beta}_{\text{Cov}_\varepsilon [p, \lambda]},$$

which is a direct analogue of (18). When picking their production techniques, firms adjust their beliefs about suppliers' prices not only for risk in ε but also for other risk types, as captured by the second term on the right-hand side of (133). Notably, the labor supply uncertainty and, in fact, the level of \bar{L} do not affect firms' production technique choices. As shown in the proof of Lemma 15, \bar{L} enters the firm's objective (8) as a scaling factor of sectoral output Q_i and stochastic discount factor Λ . Therefore, under Assumption 2, the labor supply \bar{L} does not affect firms' technique choices, and thus can be normalized to 1 without loss of generality.

Proposition 13. *The following statements about the impacts of different types of uncertainty on equilibrium network hold.*

1. *Labor supply uncertainty does not matter for the production network;*
2. *If distortions are purely wasteful, $\zeta = 0$, then uncertainty about $\log(1 + \tau)$ has an equivalent impact on the production network as uncertainty about ε ;*

⁴⁹In particular, we use the law of total covariance and the fact that if x_1 and x_2 are normal variables, $\text{Cov}(x_1, X_2) = \text{Cov}(x_1, x_2) \mathbb{E}(X_2)$, where $X_2 = \exp(x_2)$.

3. *Uncertainty about the household's preferences do not matter for the production network if ε and τ are nonrandom.*

Proof. Part 1 has been already established in the proof of Lemma 15. To prove part 2, impose $\zeta = 0$ and denote $\hat{\varepsilon} = \varepsilon - \log(1 + \tau)$. Then $\hat{\omega}$ does not explicitly depend on τ and ε , and the price function (124) and the firm's objective function (130) depend only on $\hat{\varepsilon}$, not on ε and τ separately. Finally, part 3 follows from the fact that the price (124) does not depend on β . Therefore, if ε and τ are nonrandom, (132) can be written as

$$-\frac{\partial a(\alpha_i^*)}{\partial \alpha_{ij}} + p_j + \tilde{\chi}_{ij}^e - \tilde{\gamma}_i^e = 0,$$

where we redefine Lagrange multipliers by dividing them by $E[\Lambda \hat{\omega}_i]$. It is then clear that α^* does not depend on β . \square

Proposition 13 justifies our focus on uncertainty in productivity ε in the main text. First, as discussed above, labor supply uncertainty does not matter for the production network. Second, shocks to productivity and to taxes have analogous effects on firms' decisions. In particular, a high distortion τ_i makes sector i an expensive supplier, which for other firms is equivalent to firms in industry i being unproductive. Notably, shocks to ε and to $\log(1 + \tau)$ are no longer equivalent if distortions are not purely wasteful. In that case, some of the distortion revenues are rebated to the household, affecting the stochastic discount factor. Finally, preference uncertainty matters *only if* there is uncertainty about ε or τ . To understand this result, notice that the price vector (124) does not depend on the household's preferences because it is determined by firms' quantity choices after uncertainty is realized. In the absence of uncertainty in ε and τ , the price vector p is then a constant. When choosing their production techniques, firms do not need to consider how it covaries with the household's stochastic discount factor, that is, the second term on the right-hand side of (129) is zero.