

Endogenous Production Networks under Supply Chain Uncertainty*

Alexandr Kopytov
University of Hong Kong

Bineet Mishra
Cornell University

Kristoffer Nimark
Cornell University

Mathieu Taschereau-Dumouchel
Cornell University

February 24, 2022

Abstract

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

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

*We thank Ivan Alfaro, Ryan Chahrour, Roger Farmer, Simon Gilchrist, Pablo Kurlat, David Zeke and seminar participants at Cornell University, New York University, the University of Southern California, the University of Pennsylvania, the University of Virginia, Yale University, the Toulouse School of Economics and the University of Hong Kong for useful comments and suggestions. This paper previously circulated under the title “Endogenous Production Networks under Uncertainty”.

1 Introduction

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

Supply chain disruptions are one of the key challenges that managers face in operating their business, and firms devote substantial resources in mitigating these risks. In a survey by [Wagner and Bode \(2008\)](#), business executives in Germany reported that supply chains issues were responsible for significant disruptions to production. Similarly, the [Zurich Insurance Group \(2015\)](#) conducted a global survey of executives in small and medium enterprises and found that, of all the respondent, 39% reported that losing their main supplier would adversely affect their operation, and 14% reported that they would need to significantly downsize their business, require emergency support or shut down. In addition, there is a large literature in operations research that documents the important impact of supply chain risk on firms’ operations (see [Ho et al. \(2015\)](#) for a review).

The COVID-19 pandemic provides a good example of how uncertainty can affect supply relationships. After the onset of the pandemic, many firms realized that their supply chains were exposed to substantially more risk than they thought. In a recent survey of business executives, seventy percent agreed that the pandemic had pushed companies to favor higher supply chain resiliency instead of purchasing from the lowest-cost supplier. Many also reported that they plan to diversify their supply chains across suppliers and geographies to mitigate risk.¹

To investigate whether the concerns expressed by managers translate into actions, we combine data on firm-to-firm input-output relationships in the United States with measures of stock price volatility, which serve as a proxy for uncertainty. 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 1. In column (2), we follow [Alfaro, Bloom, and Lin \(2019\)](#) and address potential endogeneity concerns by instrumenting with industry-level exposure to ten aggregate sources of uncertainty shocks. Finally, in column (3), we use changes in volatility implied by option prices as measure of uncertainty shocks. In all cases, we find a positive and statistically significant relationship between supplier volatility shocks and the end of the supply relationship, 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

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>.

Table 1: 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** (0.010)	0.097*** (0.029)	0.144** (0.064)
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	—	39.0	23.2

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. The 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. 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 levels. F -statistics are Kleibergen-Paap. *, **, *** indicate significance at the 10%, 5%, and 1% levels, respectively.

that a relationship is destroyed, according to the IV estimates.²

Motivated by this evidence, we construct a macroeconomic model of endogenous network formation to investigate how uncertainty affects firms’ sourcing decisions and how, in turn, these decisions affect the macroeconomy. In the model, each firm produces a differentiated good that can be consumed by a representative household or used as an intermediate input by another producer. Firms can produce their good in different ways, which we refer to as production techniques. A technique is a production function that specifies which intermediate inputs to use and how these inputs are to be combined. Techniques can also differ in terms of productivity. When choosing a production technique, a firm can marginally adjust the importance of a supplier or drop that supplier altogether. As a result, these decisions, when aggregated, lead to changes in the production network along both the intensive and extensive margins.

After production techniques have been chosen, firms are subject to random productivity shocks. They can then adjust how much they produce and the quantity of inputs that they use, subject to the constraints imposed by their selected technique. Competitive pressure between producers implies that the productivity shocks, as they affect production costs, are reflected in prices. Since firms are owned by the representative household, they compare profits across states of the world using its stochastic discount factor and inherit the household’s valuation of risk.

Importantly, in our setting firms’ beliefs about the distribution of firm-level productivities can

²The specifications in Table 1 follow [Alfaro et al. \(2019\)](#). See note below the table and Appendix A for more details about the data and this exercise.

influence their choice of production technique and, thus, the structure of the production network. For instance, while a firm would generally prefer to purchase from a more productive firm, it might decide not to do so if this firm is also more risky. A more productive firm would, through competitive pressure, sell at a lower price on average, but if it is also more risky, it is more likely to suffer from a large negative productivity shock, in which case its price would rise substantially. Potential customers take this possibility into account and balance concerns about average productivity and stability when choosing a production technique.

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

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

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

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

decline in the variance of aggregate output.

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

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

To isolate the impact of uncertainty, we also compare our calibrated model to an alternative economy in which firms are unconcerned about risk when making sourcing decisions. While this economy is similar to the calibrated one during normal times, significant discrepancies appear during high-volatility periods, such as the Great Recession. During that episode, we find that firms respond to uncertainty by moving to safer but less productive suppliers. Taken together, these decisions lead to a 2.4% reduction in the volatility of GDP. The added stability comes however at the cost of a 0.25% additional decline in expected GDP. Interestingly, this increase in resilience pays off in our estimation as realized GDP in the economy in which firms disregard risk drops by an additional 2.7% during the recession compared to our baseline model.

Our work is related to a large literature that investigates the impact of uncertainty on macroeconomic aggregates (Bloom, 2009, 2014; Bloom et al., 2018). In this paper, we propose a novel mechanism through which uncertainty can lower expected GDP permanently. That 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. (2021) argue that uncertainty may lead capital to flow to firms that are less exposed to aggregate risk, rather than to those firms where it would be most productive. In their model, as in ours, uncertainty leads to lower aggregate output and measured TFP.³

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

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

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

Our paper is not the first to study the endogenous formation of production networks. Oberfield (2018) considers an economy in which each firm must select one input and studies the emergence of star suppliers. Acemoglu and Azar (2020) build a model of endogenous network formation in which firms have multiple inputs and investigate its implications for growth. These papers focus on the extensive margin of the network, *i.e.* on whether a link between two sectors or firms exists or not. Taschereau-Dumouchel (2020) and Acemoglu and Tahbaz-Salehi (2020) study economies in which the firms’ decisions to operate or not shape the production network. Lim (2018) constructs a model to evaluate the importance of endogenous changes in the network for business cycles fluctuations but does not study how the network adjusts to changes in uncertainty. 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 micro data to study misallocation in the inputs market. In our model, both the intensive and extensive margins are active.⁷ To the best of

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

⁵Other works have looked at the importance of production networks outside of the business cycle literature. Jones (2011) investigate their importance to explain the income difference between countries. Recent work that has studied production networks under distortions include Baqaee (2018), Liu (2019), Baqaee and Farhi (2019b) and Bigio and La’O (2020).

⁶Baqaee and Farhi (2019a) investigate departures from Hulten’s theorem due to higher-order effects of shocks. Barrot and Sauvagnat (2016) and Carvalho et al. (2016) study the propagation of shocks after natural disasters.

⁷Atalay et al. (2011) show that a modified “preferential attachment” model can fit features of the U.S. firm-level production network. Carvalho and Voigtländer (2014) build a rule-based model of network formation to study the

our knowledge, it is also the first model in which uncertainty directly affects the structure of the production network.

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

The remainder of the paper is organized as follows. The next section introduces our model of network formation under uncertainty. In [Section 3](#), we first characterize the equilibrium when the network is kept fixed. We then consider the full equilibrium with a flexible network in [Section 4](#) and prove that an efficient equilibrium always exists. In [Section 5](#), we describe the mechanisms at work in the environment and explain how shocks propagate through the network. In [Section 6](#), we calibrate the model to U.S. data and quantitatively evaluate the importance of uncertainty on the macroeconomy through its impact on the production network. The last section concludes. All proofs are in [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 final consumption. To produce, each firm must choose a production technique, which specifies a set of inputs to use, the factor shares associated with these inputs, and an expected level of productivity. 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 industries, indexed by $i \in \mathcal{N} = \{1, \dots, n\}$, each producing a differentiated good. In each industry, there is a representative firm that behaves competitively so that equilibrium profits are always zero. When this creates no confusion, we use industry i , product i and firm i interchangeably.

Firm 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, the proportions in which these inputs are to be combined,

adoption and diffusion of intermediate inputs.

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 a firm's total factor productivity. Finally, $\zeta(\alpha_i)$ is a normalization to simplify future expressions.⁸

Since a technique α_i corresponds to a vector of factor shares, we define the set of feasible production techniques \mathcal{A}_i for industry i as

$$\mathcal{A}_i = \left\{ \alpha_i \in [0, 1]^n : \sum_{j=1}^n \alpha_{ij} \leq \bar{\alpha}_i \right\}, \quad (2)$$

where the constant $0 < \bar{\alpha}_i < 1$ provides a lower bound on the share of labor in the production of good i .⁹ We denote by $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$ the Cartesian product of the sets $\{\mathcal{A}_1, \dots, \mathcal{A}_n\}$, such that an element $\alpha \in \mathcal{A}$ corresponds to a set of input shares for each firm. As such, it fully characterizes the production network and firms, through their choice of techniques, can influence the structure of this network. Importantly, the set \mathcal{A} allows firms to adjust the importance of a supplier at the margin or to not use a particular input at all by setting the corresponding share to zero. The model is therefore able to capture network adjustments along both the intensive and extensive margins.¹⁰

The choice of technique α_i also influences the total factor productivity of firm i through the term $A_i(\alpha_i)$ in (1). 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 industry i , there is a set of *ideal* input shares α_{ij}^o that maximize A_i and that represent the most efficient way to combine intermediate inputs to produce good i .¹¹ These ideal shares are given by nature and might differ across industries. When deciding on its optimal

⁸The normalization $\zeta(\alpha_i) = \left[\left(1 - \sum_{j=1}^n \alpha_{ij} \right)^{1 - \sum_{j=1}^n \alpha_{ij}} \prod_{j=1}^n \alpha_{ij}^{\alpha_{ij}} \right]^{-1}$ is useful to simplify the expression of the unit cost, given by (9) below. $\zeta(\alpha_i)$ could instead be included in $A_i(\alpha_i)$ without any impact on the model.

⁹We impose $\bar{\alpha}_i < 1$ to make labor essential and rule out pathological cases in which output is infinite.

¹⁰This is in stark contrast with standard network models with CES production. In those models, the share of an input can fluctuate but it can never reach zero. As a result, these models cannot generate the destruction or creation of links observed in the firm-level network data studied in Table 1 and in the sectoral data that we use in Section 6.

¹¹All the proofs go through with a weakly log-concave $A_i(\alpha_i)$ if the covariance Σ of ε is positive definite.

production technique the firm will take these ideal shares into account, but it will also evaluate how high and uncertain each input price is.

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

$$\log A_i(\alpha_i) = a_i^0 - \sum_{j=1}^n \kappa_{ij} (\alpha_{ij} - \alpha_{ij}^{\circ})^2 - \kappa_{i0} \left(\sum_{j=1}^n \alpha_{ij} - \sum_{j=1}^n \alpha_{ij}^{\circ} \right)^2, \quad (3)$$

where $\alpha_i^{\circ} = (\alpha_{i1}^{\circ}, \dots, \alpha_{in}^{\circ})$ represents the ideal TFP-maximizing input shares. The parameter κ_{ij} determines the cost, in terms of productivity, of moving the j^{th} input share α_{ij} away from its ideal share α_{ij}° . The last term captures a productivity penalty of deviating from an ideal labor share.

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

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

Importantly, we impose the restriction that each firm/industry i can only adopt a single production technique α_i . Without this restriction, the firm would set up a continuum of individual plants, each with its own technique to cover the set of available techniques \mathcal{A}_i . After the realization of the productivity shocks ε , the firm would only operate the plant that is best suited to the specific draw ε . All the other plants would remain idle. In reality, we think that fixed costs would prevent firms for setting up all these plants. The restriction that firms can only operate one technique allows us to capture the impact of these costs while keeping the model tractable.

2.2 Household preferences

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

$$u \left(\left(\frac{C_1}{\beta_1} \right)^{\beta_1} \times \dots \times \left(\frac{C_n}{\beta_n} \right)^{\beta_n} \right), \quad (4)$$

where $\beta_i > 0$ for all i and $\sum_{i=1}^n \beta_i = 1$.¹² The utility function u is CRRA with a coefficient of relative risk aversion $\rho \geq 1$.¹³ The household makes consumption decisions after uncertainty is revealed and so in each state of the world it faces the budget constraint

$$\sum_{i=1}^n P_i C_i \leq 1, \quad (5)$$

where P_i is the price of good i and where we use the wage as numeraire so that $W = 1$.¹⁴

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

$$\Lambda = u'(Y) \times 1/\bar{P}, \quad (6)$$

where $Y = \prod_{i=1}^n (\beta_i^{-1} C_i)^{\beta_i}$ is aggregate consumption and $\bar{P} = \prod_{i=1}^n P_i^{\beta_i}$ is the price index.¹⁵ The stochastic discount factor thus captures how much an extra unit of the numeraire contributes to the utility of the household in different states of the world. In our setting, aggregate consumption equals aggregate (real) GDP, so we also refer to Y as GDP in what follows.

From the optimization problem of the household it is straightforward to show (see Appendix C.1 for a derivation) that

$$y = -\beta' p, \quad (7)$$

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

¹²The model can handle $\beta_i = 0$ for some goods at the cost of extra complications in the proofs.

¹³The CRRA assumption is necessary for the stochastic discount factor to be log-normally distributed and, therefore, for the model to remain tractable. The case $0 < \rho < 1$ is straightforward to characterize but is somewhat unnatural since the household then seeks to increase the variance of log consumption. To see this, consider that since $\log Y$, where $Y = \prod_{i=1}^n (\beta_i^{-1} C_i)^{\beta_i}$, is normally distributed, maximizing $E[Y^{1-\rho}]$ amounts to maximizing $E[\log Y] - \frac{1}{2}(\rho - 1) \text{V}[\log Y]$ such that $\rho \leq 1$ indicate 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 $\log Y$ from an increase in the standard deviation of Y .

¹⁴There is a different real wage associated with each state of the world (or, equivalently, per realization of ε). However, since P and W are both conditional on the state of the world and only the ratio P/W matters for outcomes, setting $W = 1$ is simply a normalization.

¹⁵See Appendix C.1 for a derivation of Λ .

Equation 7 also shows that it is sufficient to derive the vector of prices to determine GDP.

2.3 Unit cost minimization

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

Under a given production technique α_i , the cost minimization problem of the firm is described by

$$K_i(\alpha_i, P) = \min_{L_i, X_i} \left(L_i + \sum_{j=1}^n P_j X_{ij} \right) \quad (8)$$

subject to $F(\alpha_i, L_i, X_i) \geq 1$,

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

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

$$K_i(\alpha_i, P) = \frac{1}{e^{\varepsilon_i} A_i(\alpha_i)} \prod_{j=1}^n P_j^{\alpha_{ij}}. \quad (9)$$

Equation (9) is the standard unit cost for a Cobb-Douglas production function. It states that the cost of producing one unit of good i is equal to the geometric mean of the individual input prices (weighed by their respective shares) and adjusted for the firm's total factor productivity. The unit cost $K_i(\alpha_i, P)$ therefore rises when inputs become more expensive and declines when the firm becomes more productive.

In equilibrium, competitive pressure from other firms in the same industry will push prices to be equal to unit cost 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 (9), will allow us to fully characterize the price system as a function of the random productivity shocks ε .¹⁶

2.4 Technique choice

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

$$\alpha_i^* \in \arg \max_{\alpha_i \in \mathcal{A}_i} \mathbb{E}[\Lambda Q_i (P_i - K_i(\alpha_i, P))]. \quad (11)$$

Here Q_i is the equilibrium demand for good i and the firm uses the stochastic discount factor Λ of the household to weigh profits in different states of the world. Firms take prices P , demand Q_i and the stochastic discount factor Λ as given and so the only term in (11) over which the firm has any control is the unit cost $K_i(\alpha_i, P)$. The technique choice problem can therefore be written as

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

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

2.5 Equilibrium conditions

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

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

1. (Optimal technique choice) For each $i \in \{1, \dots, n\}$, factor demand L_i^* and X_i^* are a solution to (8), and the technology choice $\alpha_i^* \in \mathcal{A}_i$ solves (11) given prices P^* , demand Q_i^* and the stochastic discount factor Λ^* given by (6).
2. (Consumer maximization) The consumption vector C^* maximizes (4) subject to (5) given prices P^* .
3. (Unit cost pricing) For each $i \in \{1, \dots, n\}$,

$$P_i^* = K_i(\alpha_i^*, P^*), \quad (12)$$

¹⁶Even 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.

where $K_i(\alpha_i^*, P^*)$ is given by (9).

4. (Market clearing) For each $i \in \{1, \dots, n\}$,

$$\begin{aligned} Q_i^* &= C_i^* + \sum_{j=1}^n X_{ji}^*, \\ Q_i^* &= F_i(\alpha_i^*, L_i^*, X_i^*), \\ \sum_{i=1}^n L_i^* &= 1. \end{aligned} \tag{13}$$

Conditions 2, 3 and 4 correspond to the standard competitive equilibrium conditions for an economy with a fixed production network. They imply that firms and the household optimize in a competitive environment and that all markets clear given equilibrium prices. Condition 1 emphasizes that the production techniques, and hence the production network represented by the matrix α^* , are equilibrium objects that depend on the primitives of the economy.

The model studied above is relatively simple, but it is straightforward to extend along several dimensions without losing tractability. For instance, we can generalize the set of techniques \mathcal{A}_i to include lower and upper bounds on specific input shares. These bounds could be used to impose that certain sectors need a given input to produce or, inversely, can never use an input into production. It is also straightforward to extend the model to include multiple types of labor. In this case, we could separate firms between domestic and foreign ones, each using only one type of labor, and use the model to investigate the impact of beliefs and uncertainty on trade networks. Trade costs could then be introduced by imposing that goods that are traded internationally transit through a fictitious “transportation” sector with a productivity less than one.

On the other hand, certain ingredients are essential to keep the model tractable. Here the key challenge comes from a fixed point between the technique choice problem and the beliefs about equilibrium prices. The log-linearity implied by the Cobb-Douglas aggregators in (1) and (4) are needed to keep the equilibrium beliefs tractable. While this requires 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 a firm’s expectations affect its technique choice, the model is able to handle richer substitution patterns between expected prices and intermediate inputs, as we explore in more details in Section 5.

3 Equilibrium prices and GDP in a fixed-network economy

Before analyzing how the equilibrium production network α^* responds to changes in the environment, it is useful to first establish how prices and GDP depend on productivity under a fixed production network. These results will then be useful to characterize how uncertainty affects firms’

choices of production techniques and how these choices, when aggregated, affect the level and volatility of aggregate output.

To this end, we establish a first result that links the vector of firm-level productivities with prices and GDP.

Lemma 1. *For a fixed production network α ,*

$$p(\alpha) = -\mathcal{L}(\alpha)(\varepsilon + a(\alpha)), \quad (14)$$

and

$$y(\alpha) = \beta' \mathcal{L}(\alpha)(\varepsilon + a(\alpha)), \quad (15)$$

where $a(\alpha) = (\log A_i(\alpha_i), \dots, \log A_n(\alpha_n))$ and $\mathcal{L}(\alpha) = (I - \alpha)^{-1}$ is the Leontief inverse.

Proof. All proofs are in Appendix D. □

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

First, consider the impact of the vector of productivities $\varepsilon + a(\alpha)$. Since all elements of β and $\mathcal{L}(\alpha)$ are non-negative, increases in firm-level productivities have a negative impact on prices and a positive impact on GDP.¹⁷ Intuitively, as firms become more productive, their unit costs decline and competition forces them to sell at lower prices. From the perspective of GDP, higher productivity implies that the available labor can be transformed into more consumption goods.

Second, the lemma makes clear that production techniques α matter for prices and GDP through two distinct channels. They have a direct impact on the productivity shifters $a(\alpha)$ because different techniques have different productivities. For instance, if a firm deviates from its ideal input shares, its TFP declines which pushes for higher prices and lower GDP. However, α also affects prices and GDP through its impact on the Leontief inverse. The matrix $\mathcal{L}(\alpha) = (I - \alpha)^{-1} = I + \alpha + \alpha^2 + \dots$ implies that the price of good i depends not only on the productivity of i itself, but also on the productivity of all of its suppliers, and on the productivity of all of *their* suppliers, and so on. These higher-order connections also matter for GDP and the impact of a firm's productivity on aggregate output depends on the firm's importance as a direct and indirect supplier.

To characterize which producers are important suppliers, it is convenient to define the Domar weight of a firm i as $\omega_i(\alpha) = \beta' \mathcal{L}(\alpha) \mathbf{1}_i$, where $\mathbf{1}_i$ is a vector with a 1 as i^{th} element and zeros elsewhere. We can then rewrite (15) as

$$y = \omega(\alpha)'(\varepsilon + a(\alpha)), \quad (16)$$

¹⁷ \mathcal{L} is non-negative since $\mathcal{L}(\alpha) \equiv (I - \alpha)^{-1} = I + \alpha + \alpha^2 + \dots$, and $\alpha \geq 0$. With the assumption that $\sum_{j=1}^n \alpha_{ij} \leq \bar{\alpha}_i$ this ensures that $I - \alpha$ is always strictly diagonally dominant and therefore invertible.

where $\omega(\alpha) = (\omega_1(\alpha), \dots, \omega_n(\alpha))$. As in standard models, ω_i is also equal to the share of firm i 's sales in nominal GDP, so that $\omega_i = \frac{P_i Q_i}{P^i C}$.¹⁸ The Domar weights thus determine the relative importance of sectoral productivity changes in an economy with a fixed production network. As they depend only on β and α , Domar weights are constant in a fixed-network economy but vary when firms are free to adjust their sourcing decisions in response to changes in beliefs. In particular, a change in the production network that would make a given sector a more important supplier would also increase the importance of that sector's productivity for aggregate GDP.

Finally, Lemma 1 also shows that the price vector p and GDP y are linear functions of the productivity vector ε and, as a result inherit the normality of ε . This result is essential for the tractability of the model and allows us to compute the first and second moments of GDP as

$$\mathbb{E}[y(\alpha)] = \omega(\alpha)'(\mu + a(\alpha)), \quad (17)$$

$$\mathbb{V}[y(\alpha)] = \omega(\alpha)' \Sigma \omega(\alpha) \quad (18)$$

It is clear from these equations that the production network α , through its impact on the Domar weights $\omega(\alpha)$, matters for the mean and the variance of GDP. In addition, one important implication of (17) is that the covariance Σ of ε has no impact on expected GDP, except through its influence on the structure of the network. It follows that whenever we discuss the response of expected GDP to a change in uncertainty, the mechanism must operate through the endogenous reorganization of the network.

We conclude this section with a corollary that describes the impact of firm-level shocks on the mean and the variance of GDP. In what follows, we use partial derivatives to emphasize that the network α is kept fixed.

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

1. *The impact of a change in firm-level expected TFP μ_i on expected GDP $\mathbb{E}[y]$ is given by*

$$\frac{\partial \mathbb{E}[y]}{\partial \mu_i} = \omega_i.$$

2. *The impact of a change in firm-level volatility Σ_{ij} on the variance of GDP $\mathbb{V}[y]$ is given by¹⁹*

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

The first part of the lemma demonstrates that for a fixed production network, Hulten's (1978) celebrated theorem also holds in expectational terms. That is, the change in expected GDP follow-

¹⁸See derivation in the proof of Corollary 1.

¹⁹For $i \neq j$, the following derivative simultaneously changes Σ_{ij} and Σ_{ji} to preserve the symmetry of Σ .

ing a change in the expected productivity of an industry i is equal to that industry's sales share ω_i . The second part of the lemma establishes a similar result for changes in volatility.²⁰ In this case, we see that the impact of an increase in the uncertainty of the TFP of a sector on the variance of GDP is equal to the square of that sector's sales share. The corollary also describes how aggregate volatility responds to a change in the correlation between two sectors. In this case, the increase in $V[y]$ is proportional to the product of the two industries' sales shares. Since Domar weights are always positive, an increase in correlation always leads to higher aggregate volatility. Intuitively, positively correlated shocks are unlikely to offset each other, and their expected aggregate impact is therefore larger.

Finally, Corollary 1 emphasizes that for a fixed network, knowing the sales shares of every industry is sufficient to compute the impact of changes to μ and Σ on GDP. In the next section, 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 μ can even lead to a decline in expected GDP.

4 Equilibrium production network

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

4.1 Technique choice

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

Lemma 2. *$\lambda(\alpha^*)$, $k_i(\alpha_i, \alpha^*)$ and $q_i(\alpha^*)$ are normally distributed and the technique choice problem*

²⁰Gabaix (2011) and Acemoglu et al. (2012) emphasize the importance of the Domar weights for the variance of GDP.

of the firm can be written as

$$\alpha_i^* \in \arg \min_{\alpha_i \in \mathcal{A}_i} \mathbb{E} [k_i(\alpha_i, \alpha^*)] + \frac{1}{2} \mathbb{V} [k_i(\alpha_i, \alpha^*)] + \text{Cov} [k_i(\alpha_i, \alpha^*), \lambda(\alpha^*) + q_i(\alpha^*)]. \quad (19)$$

The terms on the right-hand side of (19) capture how beliefs and uncertainty affect the production network. The first term implies that the firm prefers to adopt techniques that provide, in expectation, a lower unit cost of production. Taking the expected value of the log of (9), we can write this term as

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

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

The second term in (19) shows that the firm also prefers production techniques that lower the variance of that unit cost. To understand this, use (9) to write the unit cost variance as

$$\mathbb{V} [k_i(\alpha_i, \alpha^*)] = \sum_{j=1}^n \alpha_{ij}^2 \mathbb{V} [p_j] + \sum_{j \neq k} \alpha_{ij} \alpha_{ik} \text{Cov} [p_j, p_k] + 2 \text{Cov} \left[-\varepsilon_i, \sum_{j=1}^n \alpha_{ij} p_j \right] + \text{const}. \quad (20)$$

The variance of the unit cost can thus be decomposed into three channels, shown on the right-hand side of (20). The first term in (20) implies that the firm prefers inputs that have stable prices. The second term implies that the firm avoids techniques that rely on inputs with correlated prices and instead prefers to *diversify* its set of suppliers and adopt inputs whose variation in prices offset each other. The third term implies that the firm prefers inputs whose prices are positively correlated with its own productivity shocks. When the firm experiences a negative productivity shock, the prices of its inputs are then more likely to also be low, reducing the expected increase in its unit cost.

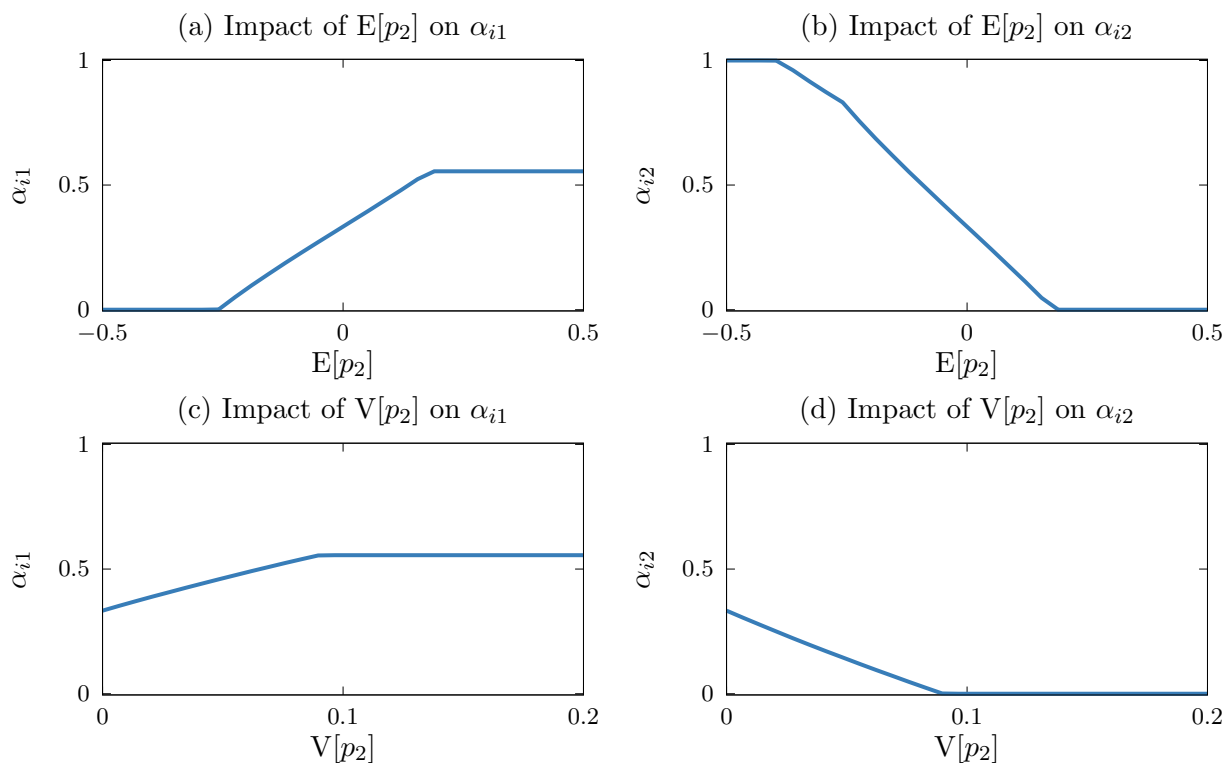
Finally, the third term in (19) captures the importance of *aggregate risk* for the firm's decision. It implies that the firm prefers suppliers whose products are cheap, so that its unit cost is low, in states of the world in which the marginal utility of aggregate consumption is high, or in which demand for the firm's goods is high. As a result, the coefficient of risk aversion ρ of the household indirectly determines how risk averse firms are.

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

Example (Sourcing decisions in partial equilibrium). Consider again the car manufacturer (firm i) that can use steel (good 1) and carbon fiber (good 2) as intermediate inputs in production. The firm must decide on the optimal shares α_{i1} and α_{i2} to pick. Assume that the input prices are

$p = (p_1, p_2) \sim \mathcal{N}(E[p], V[p])$ where the covariance matrix $V[p]$ is diagonal. If we let the penalty function $A_i(\alpha_i)$ take the form (3), it is straightforward to solve the minimization problem (19), and we show in Figure 1 how the solution α_i^* is affected by changes in the mean and the variance of p_2 . We see from panels (a) and (b) that, unsurprisingly, when good 2 is expected to be cheaper, firm i increases α_{i2} and lowers the share of good 1. A similar mechanism is at work when uncertainty about p_2 increases, as seen in Panels (c) and (d). When $V[p_2]$ is large, the firm prefers to use a larger share of the relatively safer good 1. Notice that the share α_{i2} reaches zero when p_2 is expected to be sufficiently large or uncertain. In that case, firm i severs the link with the carbon fiber supplier and an input/output relationship disappears from the production network. In this example, both the intensive and extensive margins of network adjustment are thus active.

Figure 1: Beliefs and input shares



Notes: Parameters: $\rho = 5$, $\kappa_i = (1/3, 1/3, 1/3)$, $a_i^0 = 0$, $\alpha_i^0 = (1/3, 1/3)$, $E[p_1] = 0$ and $V[p_1] = 0$. p_i is uncorrelated with (p_1, p_2) .

4.2 Equilibrium existence and efficiency

The example above demonstrates how an individual firm's technique choice responds to changes in beliefs about input prices. However, prices are equilibrium objects that depend on the production network and, therefore, on the choices made by other firms. Here, we first show that there exists an equilibrium that satisfies the conditions in Definition 1. We then show that there exists a solution

to a social planner’s problem that coincides with a decentralized equilibrium. In this equilibrium, the production network strikes an optimal balance between maximizing the mean level of aggregate output and minimizing its variance.

Existence of an equilibrium

Lemma 2 describes a self-map $\mathcal{K} : \mathcal{A} \rightarrow \mathcal{A}$ that can be used to define an equilibrium network α^* . At a fixed point of this mapping, we have that $\alpha_i^* = \mathcal{K}_i(\alpha^*)$ for all $i \in \mathcal{N}$, where $\mathcal{K}_i(\alpha^*)$ is the right-hand side of (19). Hence, such a fixed point describes an equilibrium network. Lemma 3 establishes that such a fixed point exists.

Lemma 3. *There exists a production network α^* such that $\alpha^* = \mathcal{K}(\alpha^*)$.*

The proof uses that \mathcal{K} is a continuous mapping on the compact set \mathcal{A} . From Brouwer’s fixed point theorem, we then know that there exists at least one element $\alpha^* \in \mathcal{A}$ such that $\alpha^* = \mathcal{K}(\alpha^*)$.²¹

Given an equilibrium network α^* , it is straightforward to compute prices from (15). From there, all other equilibrium quantities can be uniquely determined. The following proposition is then immediate.

Proposition 1. *An equilibrium exists.*

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

Equilibrium efficiency

There is a single household in the economy, and hence finding the set of Pareto efficient allocations amounts to solving the problem of a social planner that maximizes the utility function (4) of the household, subject to the resource constraints (13). The next proposition demonstrates that a solution to the planner’s problem exists, and that that solution corresponds to an equilibrium.

Proposition 2. *There exists an efficient equilibrium.*

From here on, our analysis will focus on the efficient equilibrium.²² This ensures that our results are not consequences of externalities or any other market distortions, and that the forces at work

²¹While the mapping \mathcal{K} is in general not a contraction, iterating on that mapping turns out to be a convenient method for finding a fixed point. When this fails, an equilibrium can be found by solving the planner’s problem, as we explain below. In the appendix, we group the proof of Lemma 3 with that of Proposition 1 below.

²²We have not been able to prove that the equilibrium is unique, although in our numerical computations iterating on the equilibrium mapping from different initial conditions always lead to a unique solution. We discuss conditions under which the solution to the planner’s problem is generically unique in Appendix D.7. In particular, we establish a generic uniqueness result when $A_i(\alpha_i)$ takes the form (3), which we will adopt for our quantitative exercises.

in the decentralized equilibrium are fundamental features of the environment that should not be distorted by policy makers.

Proposition 2 also allows us to investigate the properties of the equilibrium by solving the problem of the social planner directly. This last point implies that we can characterize the equilibrium network as the outcome of a welfare maximization problem, as the following corollary shows.

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

$$\mathcal{W} \equiv \max_{\alpha \in \mathcal{A}} \mathbb{E}[y(\alpha)] - \frac{1}{2}(\rho - 1) \mathbb{V}[y(\alpha)], \quad (21)$$

where \mathcal{W} is the welfare of the representative household and y is GDP as defined in (15).

Corollary 2 follows directly from the fact that, by Proposition 2, the equilibrium network α^* must maximize the expected utility of the representative consumer. It is clear from the objective function (21) that the consumer prefers networks that strike a balance between maximizing expected GDP $\mathbb{E}[y(\alpha)]$ and minimizing aggregate uncertainty $\mathbb{V}[y(\alpha)]$, with the relative risk aversion ρ determining the importance of each term. Another consequence of Corollary 2 is that it casts a complicated network formation problem as a simple optimization problem. We will rely on this result in the next section to characterize how beliefs affect the structure of the production network.

5 Beliefs, the production network and aggregate outcomes

As we have seen, the production techniques chosen by the firms depend on the mean μ and the variance Σ of the productivity vector ε . In this section we show how, through this mechanism, changes in μ and Σ affect the equilibrium structure of the production network, aggregate GDP and welfare.

5.1 Beliefs and Domar weights

In Section 3, we saw that the Domar weights are key objects to understand how changes in μ and Σ affect the expected level and the variance of GDP. In a fixed-network environment, these weights are given and do not respond to changes in beliefs. In contrast, when the network is endogenous the Domar weights are equilibrium objects that also vary with μ and Σ . The next proposition describes the relationship between these quantities.

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

Proposition 3 shows that when the network is endogenous, the Domar weights are increasing in the expected productivity of a firm and decreasing in its variance. This result can be understood both from an individual firm's perspective as well as from the perspective of the social planner.

Individual producers rely more on firms whose prices are low and stable. As a result, these firms become more important suppliers and their Domar weights increase. From the planner's perspective, recall from (16) that the Domar weight of a firm captures the contribution of its productivity to GDP. Since the planner wants to increase and stabilize GDP, it naturally increases the importance of more productive (larger μ_i) or less volatile (smaller Σ_{ii}) sectors in the production network. In Section 5.3 below we show that such an adjustment in the network is welfare-improving. But before doing so, we first discuss how changes in beliefs affect the precise structure of the equilibrium production network α .

5.2 Beliefs and the structure of the production network

Proposition 3 establishes that the Domar weights respond in an intuitive and unambiguous manner to changes in beliefs. The same is not true about the matrix α that describes the complete structure of the production network. In fact, in some cases an increase in the expected productivity of a producer i can even lead some of its customers to lower their usage of input i . In this section, we first describe how α behaves under a weak complementarity property (defined below), in which case the effect of a change in beliefs on the production network can be sharply characterized. We then provide examples of what may occur in a production network when the economy does not satisfy the complementarity property.

Network response when shares are complements

In this section, we consider economies in which the functions (a_1, \dots, a_n) satisfy the following property.

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

Assumption 2 defines a weak complementarity property between the shares that a producer allocates to its suppliers. It states that as a firm increases the share of one input, the marginal benefit of increasing the share of the other inputs weakly increases as well. In the context of the function (3) described in our earlier example, weak complementarity is satisfied if $\kappa_{i0} \leq 0$.

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

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

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

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

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

Parts (ii) and (iii) of Lemma 4 provide similar results for increases in uncertainty and in correlations. As discussed in Section 4.1, firms prefer suppliers with stable and uncorrelated prices. As a result, the additional risk introduced by a higher Σ_{jj} pushes firm j 's direct and indirect customers to reduce their exposure to j . Similarly, an increase in the covariance Σ_{ij} pushes firms to avoid inputs i and j . The complementarity effect is also at work, and so firms overall move toward production techniques that are more labor intensive.²³

Substitution between input shares

Lemma 4 makes sharp predictions about how the structure of the network responds to changes in the productivity processes at the cost of an assumption about the functions (a_1, \dots, a_n) . This restriction imposes some form of complementarity between input shares, but the model can handle much richer substitution patterns. To give an example of what these patterns might involve, we can again go back to our car manufacturer example. Suppose that the price of carbon fiber is expected to decrease (higher μ_{carbon}). The firm might respond by increasing the share of carbon fiber and decrease the share of steel it uses in production. At the same time, it might purchase additional equipment that is needed to handle carbon fiber. The endogenous technique choice thus allows for both substitution between steel and carbon fiber and complementarity between carbon fiber and equipment.

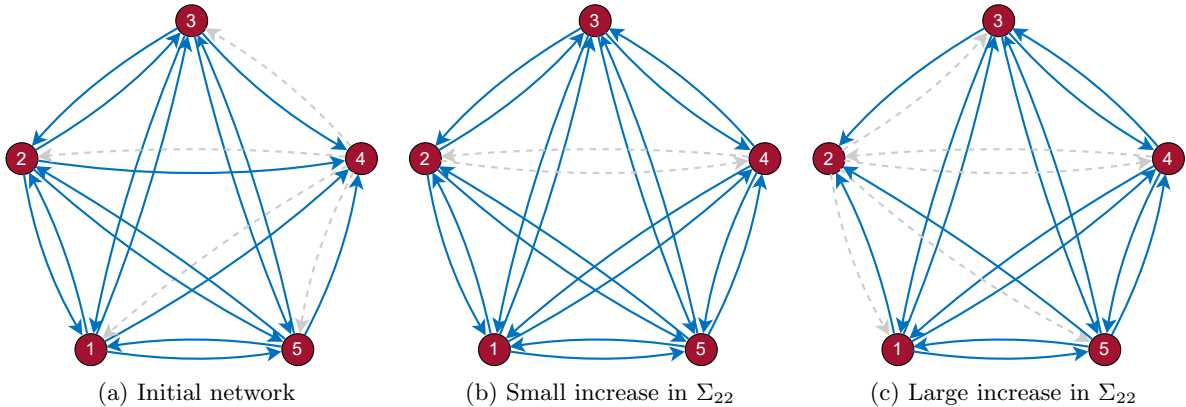
The theory, through the constraints embedded in the set \mathcal{A} and the shape of the functions (a_1, \dots, a_n) , is rich enough to accommodate some inputs that are complements while at the same

²³The assumption $\alpha^* \in \text{int}(\mathcal{A})$ in Lemma 4 is needed to avoid potential substitution patterns between firms. For instance, if $\sum_{k=1}^n \alpha_{ik} = \bar{\alpha}_i$ for a given firm i , an increase in μ_j might lead to a decline in some α_{ik} , $k \neq j$, to accommodate an increase in α_{ij} . The restriction on Σ is needed to prevent a strong uncertainty feedback. For example, if all firms increase their reliance on sector j (e.g., due to an increase in μ_j or a reduction in Σ_{jj}), the economy's exposure to j 's risk may become so large that it will be optimal to reduce α_{kj} for some k . This does not happen when Σ is sufficiently small.

time others are substitutes. For instance, if the constraint $\sum_{k=1}^n \alpha_{ik} \leq \bar{\alpha}_i$ binds, firm i might need to lower the share of another input k to be able to increase α_{ij} after a decline in the expected price of j . In this case the shares of j and k would be substitutes. Similarly, some functional forms for a can generate complementarities between inputs. As an example, consider the function $a(\alpha_i) = -(\alpha_{i1} - \alpha_{i1}^0)^2 - (\alpha_{i1} - \alpha_{i2})^2$. In this case, any increase in α_{i1} will be accompanied by additional incentives to increase α_{i2} . The function a can also be specified to generate substitutabilities through the last term in (3), by setting $\kappa_{i0} > 0$.

Figure 2 provides an example of how substitution patterns might arise in equilibrium when we relax Assumption 2. Panel (a) shows the equilibrium network in an economy in which all firms are identical except that firm 4 is less productive and, as a result, does not sell to other firms. Panel (b) shows the same economy except that ε_2 is now more volatile. In response, other producers seek to diversify their set of suppliers and create new supply relationships with firm 4. In panel (c) ε_2 becomes even more volatile. As a result, all producers drop firm 2 as a supplier and reinforce their connection to firm 4. In this example, the substitution comes from the fact that firms do not want to deviate too much from an ideal labor share (last term in (3)).

Figure 2: Uncertainty and the equilibrium network




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 (3) with $a_i^0 = 0$ for all i , $\kappa_{ij} = 1$ for all $i \neq j$, $\kappa_{ii} = \infty$ for all i , $\alpha_{ij}^0 = 1/10$ for all $i \neq j$, and $\alpha_{ii}^0 = 0$ for all i . $\beta_i = 1/n$ for all i . $\mu = 0.1$ except for $\mu_4 = 0.0571$. $\Sigma = 0.3 \times I_{n \times n}$ in Panel (a). Panel (b): same as Panel (a) except $\Sigma_{22} = 0.35$. Panel (c): same as Panel (a) except $\Sigma_{22} = 1$. The risk aversion of the household is $\rho = 5$.

Cascading flight to safety

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

is uncertain (left figure), other producers avoid using it as a supplier. But as Σ_{44} decreases, firm 3, seeking a stable supply of goods, switches to using good 4 as an input. As a result, firm 3's price becomes less volatile which pushes firm 2 to use good 3 in production. The same logic applies to firm 1, which also switches to the less volatile price provided by firm 2. As we can see, a change in the uncertainty of a single firm can lead to a cascading movement to safety. Because uncertainty about each firm's input prices is now lower, aggregate uncertainty also decreases.

Figure 3: Cascading impact of Σ_{44}



Notes: Arrows represent the movement of goods: there is a solid blue arrow from j to i if $\alpha_{ij} > 0$. Dashed gray arrows indicate $\alpha_{ij} = 0$. a is as in (3) with $a_i^0 = 0$ for all i , $\kappa_{ij} = 0$ if there is a potential link between two firms and infinity otherwise. $\alpha_{ij}^0 = 0.5$ if there is a potential link, and 0 otherwise. $\mu = 0$ except for $\mu_4 = 0.1$. In the left figure, Σ is diagonal with $\Sigma_{ii} = 0.1$ for all i except $\Sigma_{44} = 1$. In the right figure $\Sigma_{44} = 0$. The risk aversion of the household is $\rho = 2$. $\beta_i = 1/n$ for all i .

5.3 Implications for GDP and welfare

Above we analyzed how the production network responds to changes in beliefs about the productivity process. What the household ultimately cares about though is the level and variance of consumption, and we now turn to the implications of an endogenous production network for macroeconomic aggregates. We have already established in Corollary 1 how changes to the mean μ and the variance Σ of productivity affect aggregate output when the production network is fixed. Here we generalize these results to our environment with an endogenous network, and further show that some changes to the productivity process can have counterintuitive effects when the network itself responds to changes in the distribution of shocks.

Uncertainty and expected GDP

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

Proposition 4. *Uncertainty lowers the expected value of GDP in equilibrium, such that $E[y]$ is largest when $\Sigma = 0$.*

Proposition 4 follows directly from Corollary 2. When there is no uncertainty ($\Sigma = 0$), the variance $V[y(\alpha)]$ of GDP is zero for all networks $\alpha \in \mathcal{A}$, so that the equilibrium network maximizes only the expected value of GDP. When, instead, the TFP vector is uncertain ($\Sigma \neq 0$), the equilibrium network also seeks to lower $V[y(\alpha)]$, which necessarily leads to a lower expected GDP.

Proposition 4 establishes a novel mechanism through which uncertainty reduces expected GDP. To understand why, consider the technique choice from the firm's perspective. When there is no uncertainty, firms do not worry about risk and move toward cheaper suppliers, which tend to also be the most productive, and toward more productive techniques. As a result, the aggregate economy is particularly productive, and GDP is large. When some suppliers become risky, firms worry that their inputs might become expensive and, to prevent large fluctuations in their own unit cost, start purchasing from more stable but less productive suppliers. As a result, the aggregate economy becomes less productive and expected GDP falls.

The endogenous response of the network is essential for the result of Proposition 4. Indeed, in our model uncertainty affects expected GDP *only* through the endogenous response of the firms' sourcing decisions. As a result, the mechanism through which uncertainty lowers expected GDP is only active when the production network is flexible. If instead the shares α were fixed, uncertainty would have no impact on $E[y]$.

Beliefs and welfare

Proposition 4 establishes that any amount of uncertainty decreases expected GDP when the network can adjust to changes in the productivity process. Here, we investigate how the distribution of shocks affects welfare. As we will see, the endogenous response of the network matters here as well. Throughout this section, we again use partial differentiation to indicate that a derivative is taken keeping the network α fixed.

We begin by establishing that the impact of changes in the productivity process on welfare \mathcal{W} , as defined in (21).

Proposition 5. *When the network α is free to adjust to changes in μ and Σ , the following holds.*

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

$$\frac{d\mathcal{W}}{d\mu_i} = \frac{\partial E[y]}{\partial \mu_i} = \omega_i. \quad (22)$$

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

$$\frac{d\mathcal{W}}{d\Sigma_{ij}} = \begin{cases} -\frac{1}{2}(\rho - 1) \left(\frac{\partial E[y]}{\partial \mu_i} \right)^2 = -\frac{1}{2}(\rho - 1) \omega_i^2 & i = j, \\ -(\rho - 1) \frac{\partial E[y]}{\partial \mu_i} \frac{\partial E[y]}{\partial \mu_j} = -(\rho - 1) \omega_i \omega_j & i \neq j. \end{cases} \quad (23)$$

This proposition follows directly from applying the envelope theorem to (21). Its first part states that the impact of an increase in μ_i on welfare is equal to its marginal impact on expected GDP *taking the network α as fixed*. By Corollary 1, this quantity is also equal to the Domar weight ω_i of firm i . Since Domar weights are positive, it follows that an increase in μ_i always has a positive

impact on welfare. The second part of the proposition provides a similar result for an increase in Σ_{ij} . In this case, the impact of the change is proportional to the product of the Domar weights ω_i and ω_j . Again, (23) implies that an increase in uncertainty must necessarily lower welfare when $\rho > 1$.

Amplification and dampening

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

Proposition 6. *Let $\alpha^*(\mu, \Sigma)$ be the equilibrium production network under (μ, Σ) and let $\mathcal{W}(\alpha, \mu, \Sigma)$ be the welfare of the household under the network α . Then the change in welfare after a change from (μ, Σ) to (μ', Σ') is such that*

$$\underbrace{\mathcal{W}(\alpha^*(\mu', \Sigma'), \mu', \Sigma') - \mathcal{W}(\alpha^*(\mu, \Sigma), \mu, \Sigma)}_{\text{Change in welfare under the flexible network}} \geq \underbrace{\mathcal{W}(\alpha^*(\mu, \Sigma), \mu', \Sigma') - \mathcal{W}(\alpha^*(\mu, \Sigma), \mu, \Sigma)}_{\text{Change in welfare under the fixed network}}. \quad (24)$$

Under a flexible network, the extra margin of adjustment thus allows firms to produce more efficiently and to better mitigate risk, which translates into higher welfare for the household.

Beliefs and GDP

Proposition 5 shows that changes in the mean and the variance of ε have an intuitive impact on welfare when the network is endogenous, qualitatively working in the same direction as in the fixed network economy. The same is not true about the impact of such changes on expected GDP. To see how the endogenous adjustment of the network can lead to counterintuitive implications for output, it is helpful to decompose the impact of a change in beliefs into its direct and indirect impacts. For instance, for a change in μ_i we can write


$$\frac{dE[y]}{d\mu_i} = \underbrace{\frac{\partial E[y]}{\partial \mu_i}}_{\text{direct impact with fixed network}} + \underbrace{\frac{\partial E[y]}{\partial \alpha} \frac{d\alpha}{d\mu_i}}_{\text{network adjustment}}. \quad (25)$$

The first term on the right-hand side of (25) denotes the direct impact, keeping the network α fixed. The second term captures the impact of μ_i on the network α , and the impact of the change in the network α on expected GDP. When the network is fixed, this network adjustment term is zero and the full impact of the change in μ_i is simply equal to its direct impact. This is the situation that we explored in Corollary 1, which states that an increase in μ_i always has a positive impact on $E[y]$. But with an endogenous network the indirect effect can amplify, mitigate or even

overwhelm the direct impact completely, in which case an increase in μ_i can lower expected GDP. When this happens, the Hulten-like result established in Corollary 1 ceases to work, even as a local approximation. A similar mechanism can also flip the impact of changes in uncertainty such that an increase in Σ_{ii} can lower the variance of GDP.

Counterintuitive implications of changes in beliefs We now provide an example to show how the endogenous adjustment of the network can lead to counterintuitive responses to shocks. In the economy depicted in Figure 4, firms 4 and 5 use only labor to produce, while firms 1 to 3 can also use goods 4 and 5 as intermediate inputs. Firm 4 is more productive and volatile than firm 5 ($\mu_4 > \mu_5$ and $\Sigma_{44} > \Sigma_{55}$). Now consider the impact of a positive shock to μ_5 . The solid blue lines in panels (a), (b) and (c) of Figure 5 illustrate the impact of that shock on $E[y]$, $V[y]$, and on welfare. Point O on the graphs represents the economy before the shock. As we can see, the initial increase in μ_5 has a negative impact on expected GDP. To understand why, notice that for a small increase in μ_5 , firm 5 is still less productive (on average) than firm 4, but it now offers a better risk-reward trade-off given its lower variance. As a result, firms 1 to 3 increase their shares of good 5 and reduce their share of good 4. But since $\mu_4 > \mu_5$, this readjustment leads to a fall in expected GDP for a small increase in μ_5 . At the same time, the variance of GDP also declines because firm 5 is less volatile than firm 4. The implied changes in $E[y]$ and $V[y]$ thus have opposite impacts on welfare. By Proposition 5, the overall effect on welfare must be positive though, and this is indeed confirmed in panel (c). Of course, as μ_5 keeps increasing expected GDP eventually starts to increase as well.

Figure 4: The non-monotone impact of shocks



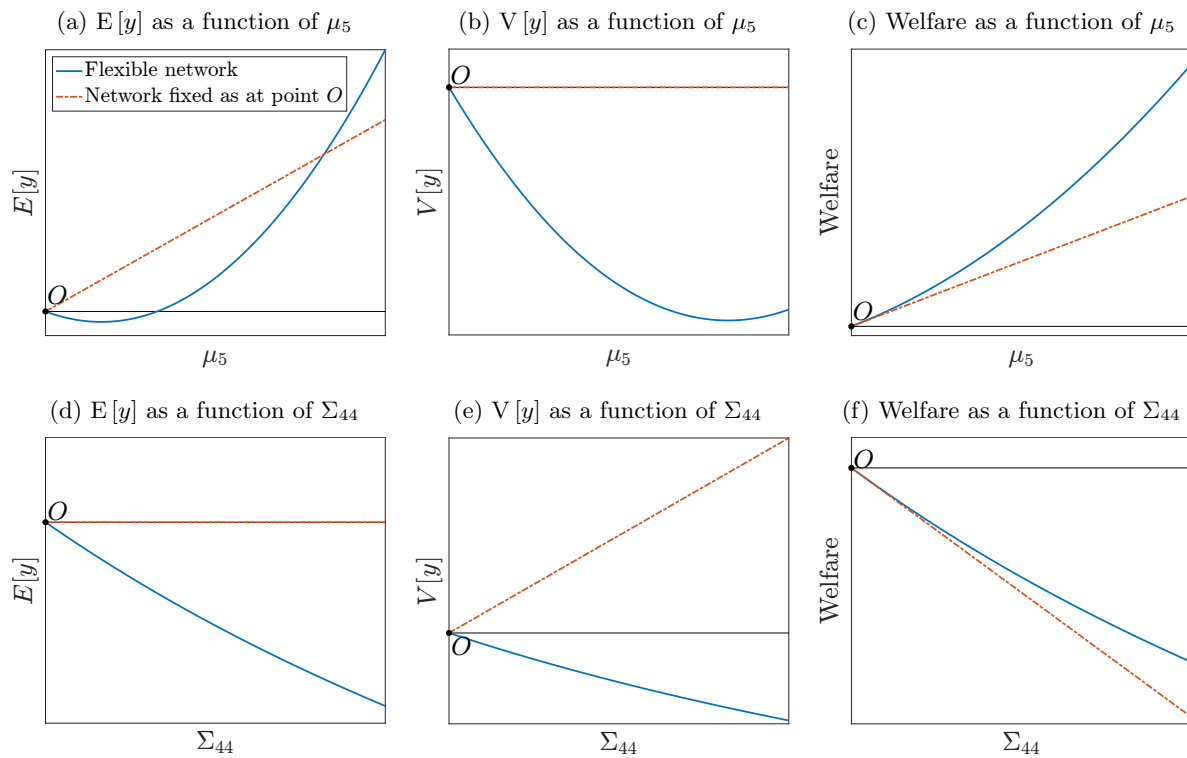
Notes: Arrows represent the movement of goods: there is a solid blue arrow from j to i if $\alpha_{ij} > 0$. a is as in (3) with $a_i^0 = 0$ for all i , $\kappa_{ij} = 0$ if there is a potential link between two firms and very large otherwise. $\alpha_{ij}^0 = 0.5$ if there is a potential link and zero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

To emphasize the importance of the endogenous network for this mechanism, we also show the effect of the same increase in μ_5 when the network is kept fixed (dashed red lines in the same panels). From Corollary 1, the marginal impact of μ_5 on expected GDP is equal to its Domar

weight, and increasing μ_5 has a positive impact on $E[y]$. At the same time, the variance of GDP is unaffected by changes in μ . While an increase in μ_5 is welfare-improving in this case, the effect is less pronounced than in the flexible network economy. Indeed, in the latter case the equilibrium network changes precisely to maximize the beneficial impact of the shock on welfare, as predicted by Proposition 6.

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

Figure 5: The non-monotone impact of firm-level shocks on GDP



Notes: The network structure and parameterization are detailed in Figure 4. In panels (a)-(c), μ_5 increases from -0.04 to 0.13 . In panels (d)-(f), Σ_{44} increases from 0.3 to 0.4 .

6 Endogenous network responses in a calibrated model

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

6.1 Data

The Bureau of Economic Analysis (BEA) provides sectoral input-output tables that allow us to compute the intermediate input shares as well as the shares of final consumption accounted for by different sectors. The input-output tables are also needed to compute sectoral TFP. Over the available sample period, the BEA has occasionally changed the number and the definition of the sectors included in the tables. We therefore rely on harmonized tables constructed by [vom Lehn and Winberry \(2021\)](#) that provide consistent annual data for $n = 37$ sectors over the period 1948—2020. Table 2 provides the list of the sectors included in that data set.

Table 2: 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 \(2021\)](#) for details of the data construction.

From these data, we can compute the input shares α_{ijt} of each sector in each year t . We also use the input-output tables to compute total factor productivity for each sector, following the procedure in vom Lehn and Winberry (2021) closely. Sectoral TFP is thus given by 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 (2021) in constructing the TFP series. First, to be consistent with our model, we let the input shares α_{ijt} vary over time. Second, we do not smooth the resulting Solow residuals. Finally, we update the time series to include the years up to 2020.

6.2 Calibration strategy

The three groups of parameters that we need to calibrate are i) the household’s preferences, i.e. the consumption shares β and the risk-aversion ρ , ii) the parameters κ and α° in the TFP shifter function (3), and iii) the processes for the exogenous sectoral productivity shocks, i.e. μ_t and Σ_t . Some of these parameters have counterparts that can be computed directly from the data. The remaining parameters 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

In the household’s utility function the different goods are combined through a Cobb-Douglas aggregator. This implies that an element i of the preference vector β corresponds to the share of good i in total consumption. The final consumption shares from the input-output tables can therefore be used to pin down β directly. The sectors with the largest consumption shares are “Real estate” ($\beta_{\text{real estate}} = 0.141$), “Retail trade” ($\beta_{\text{retail trade}} = 0.124$) and “Health care” ($\beta_{\text{health care}} = 0.108$).

The CRRA 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

The procedure to compute sectoral Solow residuals described by vom Lehn and Winberry (2021) gives us a measure of sectoral total factor productivity which, in our model, corresponds to the product $e^{\varepsilon_{it}} A_i(\alpha_{it}) \zeta(\alpha_{it})$. To extract the stochastic component ε_{it} we must first remove the endogenous productivity shifter $A_i(\alpha_{it})$ and the normalizing term $\zeta(\alpha_{it})$ from measured TFP. To do so, we need an explicit function for $A_i(\alpha_{it})$, and we adopt the form (3) introduced earlier.²⁴

²⁴We set the constant a_i^0 in (3) to the average TFP of sector i .

The functional form (3) takes as inputs the ideal shares α_{ij}° , the actual shares α_{ijt} and the coefficients κ_{ij} . The ideal shares α_{ij}° are set to the time average of the input shares observed in the data. The coefficients κ_{ij} , which determine how costly it is in terms of productivity to deviate from the ideal shares, 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 in κ^i then scales the cost for producer k of changing the share of any of its input, 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 $2n$ free parameters to estimate.

Exogenous productivity process

The source of uncertainty in the model is the vector of exogenous 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 extract the endogenous productivity shifter $A_i(\alpha_{it})$ and the normalization $\zeta(\alpha_{it})$ from measured sectoral TFP. We then estimate the evolution of μ_t and Σ_t from the remaining component.

The vector of exogenous sectoral productivity shocks is assumed to follow a random walk with drift,

$$\varepsilon_t = \gamma + \varepsilon_{t-1} + u_t, \quad (26)$$

where γ is an $n \times 1$ vector of deterministic sectoral 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 innovations $\varepsilon_t - \varepsilon_{t-1}$ over time.²⁵

When making decisions in period t , firms know the past realizations of ε_t so that the conditional mean μ_t of ε_t is given by $\gamma + \varepsilon_{t-1}$. 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 $\hat{\Sigma}_{ijt} = \sum_{s=1}^{t-1} \lambda^{t-s-1} u_{is} u_{js}$, where $0 < \lambda < 1$ is a parameter that determines the weight of more recent observations relative to those further back into the past. 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 $\lambda = 0.47$. Note that the time series for ε_t depends on the parameters in the function (3). The estimation of the stochastic process for sectoral productivity thus has to be done jointly with the estimation of κ .

²⁵We adopt this random-walk specification since our sector-level productivity data shows clear signs of non-stationarity that cannot be accounted for by a simple time-trend.

Matching model and data moments

The parameters that we estimate via indirect inference are collected in the vector $\Theta \equiv \{\rho, \kappa\}$. We then choose Θ to minimize

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

where $m(y)$ is a vector of moments computed from the data, and $m(\Theta)$ is the vector of the corresponding model-implied moments conditional on the parameters Θ . The moments that we use are the $n^2 \times T$ 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 for the trade-off that firms face when choosing production techniques. To 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 share moments and a weight of $(T - 1)^{-1}$ to the consumption growth moment.

The number of moments that we match is $n^2 \times T + T - 1$, while the number of free parameters that we estimate is only $2n + 1$. The model is thus strongly over-identified. We use particle swarm optimization (Kennedy and Eberhart, 1995) to find the global minimizer $\hat{\Theta}$. The estimated value of the risk-aversion parameter is $\hat{\rho} = 4.27$ which is similar to commonly found values in the macroeconomics literature.

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. Appendix B.1 provides the values of $\hat{\kappa}^i$ and $\hat{\kappa}^j$ for all the sectors. 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, the average sectoral TFP loss is 0.06%.

6.3 Model implied moments and the data

The mechanisms of the model make the Domar weights dependent on the conditional mean μ_t and the covariance Σ_t of the sectoral TFP process. 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 . Doing so, we verify that some of the key mechanisms of the model are at work in the data.

Sectoral total factor productivity

The stochastic process for exogenous sectoral TFP is parameterized by the vector of sectoral drifts γ and the covariance matrix Σ_t . The estimated drift vector $\hat{\gamma}$ features a lot of 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 on average 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 a lot of 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%.

Aggregate variation in beliefs

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

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

We use the Domar weights ω_{jt} in this equation to properly reflect the importance of a sector for GDP, as in (16). The solid blue line in Figure 6 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 perceived standard deviation of current GDP. From (18), this can be computed as


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

The red dashed line in Figure 6 represents the evolution of σ_{yt} between 1948 and 2020. 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.

Unconditional Domar weights

We want our model to fit key features of the data that relate to the structure of the production network, how that network changes in response to shocks, and how these changes affect macroe-

Figure 6: Domar-weighted uncertainty and TFP 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, $\sigma_{yt} = \sqrt{\omega_t' \Sigma_t \omega_t}$. Shaded areas represent NBER recessions.

conomic aggregates. As we have seen earlier, the Domar weights play a central role for these mechanisms and we now describe how the model fits these weights and their relationship with the beliefs (μ_t, Σ_t) .


Figure 7 shows the average Domar weight of each sector in the data (blue bars) and in the model (black line). The cross-sector 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 due to the fact that 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. In the calibrated model, the standard deviation of consumption growth is 2.73%, which is very close to its data counterpart (2.65%). It is reassuring to see that the model is able to match that moment well given its importance for the impact of uncertainty on the economy.

Figure 7 also shows that the sectors with the highest Domar weights in the data are “Real estate”, “Food and beverage”, “Retail trade”, “Finance and insurance” and “Health care”. According to our theory (Proposition 5), changes in the expected level and variance of productivity in those sectors will have the largest effects on welfare.

The mechanism can also account for about 40% of the observed average standard deviation of the Domar weights over time, as shown in row 2 of Table 3. 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 62% of the observed variation in Domar weights. This suggests that the model is able to generate a realistic amount of variation in a key moment that characterizes the production network.²⁶

²⁶One reason why the Domar weights are less volatile in the model is that we assume that the $A_i(\alpha_i)$ function is time invariant. In reality, technological changes might affect the shape of that function which would translate into additional variation in Domar weights.

Figure 7: 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 3: Domar weights 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 $\sigma(\omega_j) / \bar{\omega}_j$	0.11	0.07
(4) $\text{Corr}(\omega_{jt}, \mu_{jt})$	0.08	0.08
(5) $\text{Corr}(\omega_{jt}, \Sigma_{jjt})$	-0.37	-0.31

Notes: For each sector, we compute the time series of its Domar weights ω_{jt} , as well as their standard deviation σ_{ω_j} and their mean $\bar{\omega}_j$. Rows (1), and (2) report cross-sectional averages of these statistics. Row (3) is the ratio of rows (1) and (2). 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.

Domar weights and beliefs

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, implying a decline in its Domar weight. Proposition 3 makes this point formally for a single shock to μ_i or Σ_{ii} . Of course, in the data multiple changes in μ_t and Σ_t occur at the same time, and so it is not possible to isolate the impact of a single shock 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 the last two rows of Table 3, the model predictions are borne out in the data. The model can not only explain these correlations qualitatively but is also doing a good job quantitatively. The model is thus able

to capture well the impact of beliefs on the structure of the production network.²⁷

6.4 The production network, welfare and output

To evaluate the importance of an endogenous production network for welfare and GDP, we now compare the calibrated model to two alternative economies. In the first alternative, we keep the input shares in the production network completely fixed. This exercise therefore informs us about the overall impact of changes in the structure of the production network. In the second alternative, we reduce the risk aversion of the household so that it becomes indifferent to changes in the variance of sectoral productivities Σ_t . This exercise therefore informs us about the impact that uncertainty alone has on the production network and, through that channel, on macroeconomic aggregates.

Fixed vs endogenous production networks

We start by comparing the baseline model to the alternative economy in which the production network cannot respond to changes in μ_t and Σ_t . Specifically, we fix the input share matrix α to its time average in the data and recompute all other equilibrium quantities. We find that the economy with a fixed network is on average 2.12% less productive than the economy with a flexible network (see Table 4). The intuition for this difference is straightforward. As some sectors of the economy become more productive, firms would like to take advantage of their cheaper inputs by relying more on them in production. With a flexible network this is possible, and the aggregate economy becomes more productive as a result. In contrast, when the network is fixed, firms are stuck with less productive suppliers and the economy is overall less efficient. Perhaps surprisingly, the fixed network economy is also slightly more stable than its flexible counterpart with a standard deviation of GDP that is 0.13% smaller. To understand why, recall that the planner balances increasing the expected value of GDP against lowering its variance. In the calibrated model, the planner is willing to suffer a slight increase in variance, compared to the fixed network alternative, for large gains in expected value. Overall, this comparison with a fixed network suggests that policy interventions that impede or slow down the reorganization of supply chains might have a sizable impact on welfare.

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

²⁷While the data Domar weights are directly observable and independent of the model, the data moments in rows 4 and 5 of Table 3 also depend on the processes for μ_t and Σ_t which are extracted using the baseline values for κ .

The role of uncertainty

These differences between the flexible and fixed network economies come from variations in both μ_t and Σ_t . To isolate the role of uncertainty in shaping the production network and GDP, we consider an economy in which the household has a relative risk aversion of $\rho = 1$. In this case, firms' sourcing decisions are independent of Σ_t , and since these decisions are the only ones taken under uncertainty, the only *direct* impact of this change in risk aversion is on the shape of the network.²⁸ We then compare the allocation in this economy with the baseline calibration. Table 4 reports long-run moments related to GDP and welfare, which is evaluated from the perspective of the more risk-averse representative household of our baseline model. We see that, in line with the theory, the baseline economy is slightly less productive and slightly less volatile than the alternative. When $\rho > 1$, it is worthwhile from the firms' perspective to use suppliers that are less productive but safer, which translates into the observed differences in $E[y]$ and $V[y]$. However, the differences are fairly small. The reason for this is that for most of the sample period, uncertainty is relatively low so that firms simply buy their inputs from the most productive suppliers without much regard for any risk involved.²⁹ We explore the sensitivity of our results to different values of ρ in Appendix B.3.

Table 4: Uncertainty, GDP and welfare over the long run

	Baseline model compared to...	
	Fixed network	$\rho = 1$
Expected GDP	+2.122%	-0.008%
Std. dev. of GDP	+0.131%	-0.105%
Welfare \mathcal{W}	+2.109%	+0.010%

The Great Recession

While the average effect of uncertainty is quite small, it can be substantial when uncertainty increases sharply. During the Great Recession, we did see such a spike in uncertainty and firms then move toward safer but less productive suppliers to avoid potentially disastrous increases in costs. Such an adjustment can have large effects on output and welfare.

Figure 8 shows how the baseline economy compares to the $\rho = 1$ alternative (denoted with tildes in the figure) over our whole sample period. As we can see, the biggest differences occur during periods of high uncertainty such as the oil shock of the 70's, the dot-com bubble, and the Great Recession. As we can see from Figure 6, uncertainty was particularly large during the Great

²⁸By combining (17) and (21), it is clear that the equilibrium production network is unaffected by Σ when $\rho = 1$.


²⁹As in Lucas (1987), the utility cost of business cycles fluctuations is, on average, small in our model and the planner does not want to sacrifice much in terms of the level of GDP for a reduction in volatility.

Recession and, accordingly, this is also when the differences between the two models are the largest. The top panel of Figure 8 shows that expected GDP in the baseline economy is about 0.25% lower in 2009 than in the $\rho = 1$ economy. Because of the large increase in uncertainty, firms are worried about the possibility of important increases in costs and move toward safer but more expensive suppliers. The result, in terms of aggregate volatility, is visible in the second panel where we see that GDP is expected to be about 2.37% less volatile in 2009 in the baseline economy. While the implied reduction in expected GDP in the baseline economy is quite small, the reduction in variance is important enough to lead to an increase in welfare of about 0.36%, when compared to the $\rho = 1$ economy (third panel). Interestingly, realized GDP, shown in the last panel, is substantially higher in the baseline economy. Essentially, firms were worried about bad draws from the TFP processes and opted for safer suppliers, and then their fears were realized. The year 2009 saw bad TFP draws (as evident from Figure 6), and so the baseline economy fared about 2.67% better in terms of realized GDP.

Overall, our findings in this section are that allowing the production network to reorganize itself in response to changes in the productivity process, can lead to large welfare gains, but that changes in uncertainty are, most of the time, not a major factor in shaping the production network. The situation is quite different however when uncertainty spikes. In this case, its impact on the structure of the network and, through that channel, on GDP and welfare can be sizable. Our results therefore highlight the importance for firms of reorganizing supply lines during turbulent periods. It also suggests that policies, such as trade barriers, that would slow down this reorganization might have significant side effects.

While the calibrated model suggests that the effect of uncertainty on welfare and output may occasionally be substantial, there are also reasons to believe that the model estimates are conservative. In reality, many firms are owned by individuals whose consumption is likely to be more volatile than that of the representative household. It is also likely that many entrepreneurs have a less diversified portfolio than the representative household and earn most of their income from a single firm. More uncertain income that covaries more strongly with firm profits would make entrepreneurs more sensitive to risk, hence making them more likely to take action to mitigate uncertainty about input prices. In addition, the model probably underestimates the risk associated with input prices. While we only have 37 sectors in our analysis, in reality each sector is made up of heterogeneous firms that, in addition to sector-specific shocks, are also subject to firm-specific disturbances. To the extent that individual firms buy intermediate inputs from other individual firms, this will add to the input price uncertainty that they are facing.

Figure 8: The role of uncertainty in the postwar period



Notes: The differences between the series implied by the baseline model and the $\rho = 1$ alternative. 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. Welfare is always computed according to the preferences of the household in the baseline economy.

7 Conclusion

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

The proposed model is tractable and can serve as a framework to study other related questions. For instance, while we study a closed economy, uncertainty related to international supply chains have recently been at the forefront of the economic news. It would be straightforward to extend our model to analyze the impact of uncertainty on international trade networks. It would also be natural to extend the analysis in this paper to a model calibrated to firm-level data and to allow firms to enter and exit. However, such an extension would be more involved, as it might necessitate moving away from the perfect competition framework proposed here.

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*.
- ALFARO, I., N. BLOOM, AND X. LIN (2019): "The Finance Uncertainty Multiplier," .
- 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," *The 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,” *The 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,” *The Quarterly Journal of Economics*, 131, 1543–1592.
- BIGIO, S. AND J. LA’O (2020): “Distortions in Production Networks,” *The 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, J. AND E. OBERFIELD (2020): “Misallocation in the Market for Inputs: Enforcement and the Organization of Production,” *The Quarterly Journal of Economics*, 135, 2007–2058.
- 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,” Tech. rep.
- CARVALHO, V. M., M. NIREI, Y. SAITO, AND A. TAHBAZ-SALEHI (2016): “Supply chain disruptions: Evidence from the great east japan earthquake,” *Columbia Business School Research Paper*.
- COX, G. (2020): “Almost sure uniqueness of a global minimum without convexity,” *The Annals of Statistics*, 48, 584–606.
- DAVID, J. M., L. SCHMID, AND D. ZEKE (2021): “Risk-adjusted capital allocation and misallocation,” .
- DHYNE, E., A. K. KIKKAWA, M. MOGSTAD, AND F. TINTELNOT (2021): “Trade and domestic production networks,” *The Review of Economic Studies*, 88, 643–668.
- DUPOR, B. (1999): “Aggregation and irrelevance in multi-sector models,” *Journal of Monetary Economics*, 43, 391–409.
- FAJGELBAUM, P. D., E. SCHAAL, AND M. TASCHEREAU-DUMOUCHEL (2017): “Uncertainty Traps,” *The Quarterly Journal of Economics*, 132, 1641–1692.
- 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.
- 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,” *The Review of Economic Studies*, 45, 511–518.
- JONES, C. I. (2011): “Intermediate Goods and Weak Links in the Theory of Economic Development,” *American Economic Journal: Macroeconomics*, 3, 1–28.
- JURADO, K., S. C. LUDVIGSON, AND S. NG (2015): “Measuring Uncertainty,” *American Economic Review*, 105, 1177–1216.
- KENNEDY, J. AND R. EBERHART (1995): “Particle swarm optimization,” in *Proceedings of ICNN’95-international conference on neural networks*, IEEE, vol. 4, 1942–1948.
- LIM, K. (2018): “Endogenous Production Networks and the Business Cycle,” working paper.
- LIU, E. (2019): “Industrial Policies in Production Networks,” *The Quarterly Journal of Economics*, 134, 1883–1948.
- LONG, J. B. AND C. I. PLOSSER (1983): “Real Business Cycles,” *Journal of Political Economy*, 91, 39–69.
- (1987): *Models of business cycles*, vol. 26, Basil Blackwell Oxford.
- NIEUWERBURGH, S. V. AND L. VELDKAMP (2006): “Learning asymmetries in real business cycles,” *Journal of Monetary Economics*, 53, 753–772.
- OBERFIELD, E. (2018): “A Theory of Input-Output Architecture,” *Econometrica*, 86, 559–589.
- TASCHEREAU-DUMOUCHEL, M. (2020): “Cascades and Fluctuations in an Economy with an Endogenous Production Network,” *Working Paper*.
- VOM LEHN, C. AND T. WINBERRY (2021): “The Investment Network, Sectoral Comovement, and the Changing U.S. Business Cycle,” *The Quarterly Journal of Economics*.
- WAGNER, S. M. AND C. BODE (2008): “An Empirical Examination of Supply Chain Performance Along Several Dimensions of Risk,” *Journal of Business Logistics*, 29, 307–325.
- ZURICH INSURANCE GROUP (2015): “Potential business impact of loss of the main supplier for small and medium enterprises (SMEs) in 2015,” .

Appendices

A Motivational evidence from the introduction

In this section, we provide more details about the regressions presented in Table 1. 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 details, 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 US 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:

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

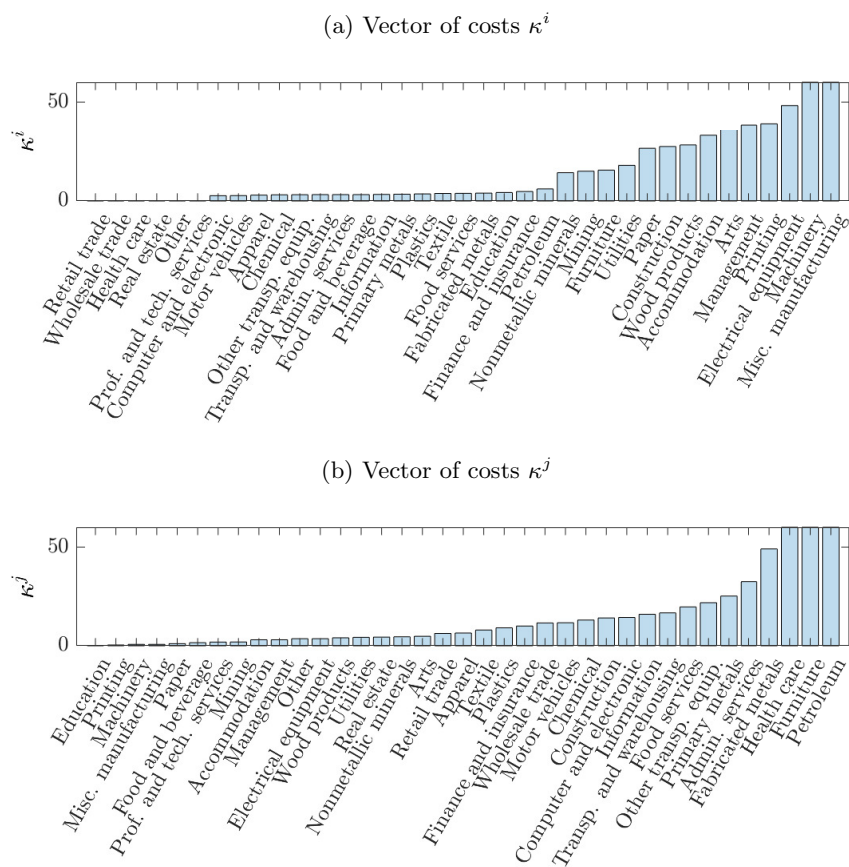
where $\Delta\sigma_{t-1}^c$ denotes the volatility of the aggregate variable c . As a result, instruments vary on the 3-digit SIC industry-by-year level. As in Alfaro et al. (2019), we also include in the IV regressions the first moments associated with each aggregate series c (“1st moment $10IV_{i,t-1}$ ” in Table 1) 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 Table 1. 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.

B Additional results related to the calibrated economy

B.1 Matrix κ

The overall mean of the elements of the calibrated cost matrix $\hat{\kappa}$ is 194 with a standard deviation of 447. To better understand the structure of the $\hat{\kappa}$ matrix, Figure 9 shows for each sector the elements of the vectors $\hat{\kappa}^i$ and $\hat{\kappa}^j$ (recall that $\hat{\kappa} = \hat{\kappa}^i \hat{\kappa}^j$). 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 expensive 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 shares of these sectors.

Figure 9: The calibrated costs of deviating from the ideal input shares



B.2 Great Recession: Flexible vs fixed network

In this section, we explore the role of network flexibility during the Great Recession—the period in which the economy was hit by large adverse shocks (see Figure 6). 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 10 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 TFPs 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.

B.3 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 . Not surprisingly, we find that ignoring uncertainty becomes costlier for higher values of ρ (Table 5).

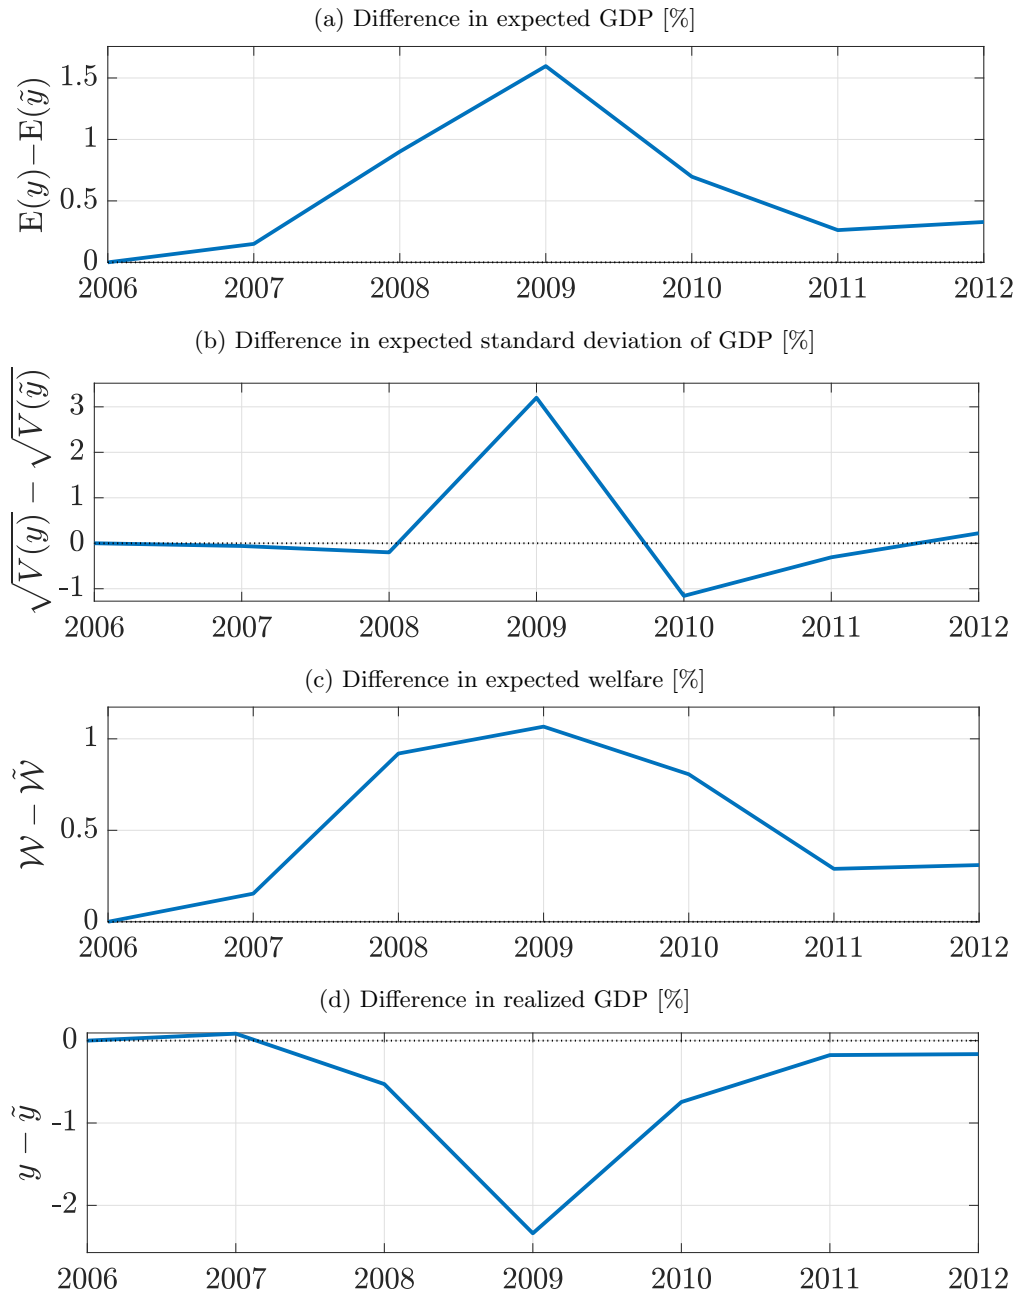
Table 5: Uncertainty and GDP over the long run

	Model with $\rho = 1$ compared to...		
	$\rho = 2$	$\rho = 4.27$ (baseline)	$\rho = 10$
Expected GDP $E[y(\alpha)]$	+0.001%	+0.008%	+0.033%
Std. dev. of GDP $\sqrt{V[y(\alpha)]}$	+0.038%	+0.105%	+0.208%
Welfare \mathcal{W}	-0.001%	-0.010%	-0.057%

Notes: The differences between the series implied by the models featuring various degrees of risk aversion and the $\rho = 1$ alternative. Welfare is always computed according to the preferences of the household in the baseline economy.


The economy also responds to the spike in uncertainty during the Great Recession much more for $\rho = 10$ (Figure 11). 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.

Figure 10: The role of network flexibility during the Great Recession



Notes: The differences between the series implied by the full model and the model in which the network is fixed at its 2006 level. All differences are expressed in percentage terms.

Figure 11: 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 and the $\rho = 1$ alternative. All differences are expressed in percentage terms. Welfare is always computed according to the preferences of the household in the baseline economy.

C Additional derivations

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 5. *The Lagrange multiplier on the budget constraint of the household (5) 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

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

and the first-order condition with respect to C_i is

$$\beta_i u'(Y) Y = \Lambda P_i C_i. \tag{27}$$

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

$$u'(Y) Y = \Lambda, \tag{28}$$

which, together with (27), implies that

$$P_i C_i = \beta_i. \tag{29}$$

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

$$\begin{aligned} Y &= \prod_{i=1}^n (\beta_i^{-1} C_i)^{\beta_i} = \prod_{i=1}^n \left(\beta_i^{-1} \frac{\beta_i u'(Y) Y}{\Lambda P_i} \right)^{\beta_i} \\ \Lambda &= u'(Y) \prod_{i=1}^n P_i^{-\beta_i} \end{aligned} \tag{30}$$

which, combined with (28), yields :

$$Y = \prod_{i=1}^n P_i^{-\beta_i}. \quad (31)$$

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 (28) yields the result. \square

C.2 Derivation of the unit cost function

The cost minimization problem of the firm is

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

subject to $F(\alpha_i, L_i, X_i) \geq 1$,

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

$$L_i = \theta \left(1 - \sum_{j=1}^n \alpha_{ij} \right) F(\alpha_i, L_i, X_i),$$

$$P_j X_{ij} = \theta \alpha_{ij} F(\alpha_i, L_i, X_i),$$

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

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

which is the result.

D Proofs

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

D.1 Proof of Lemma 1

Lemma 1. *For a fixed production network α ,*

$$p(\alpha) = -\mathcal{L}(\alpha) (\varepsilon + a(\alpha)), \quad (14)$$

and

$$y(\alpha) = \beta' \mathcal{L}(\alpha) (\varepsilon + a(\alpha)), \quad (15)$$

where $a(\alpha) = (\log A_1(\alpha_1), \dots, \log A_n(\alpha_n))$ and $\mathcal{L}(\alpha) = (I - \alpha)^{-1}$ is the Leontief inverse.

Proof. Combining the unit cost equation (9) with the equilibrium condition (12) and taking the log we find that, for all i ,

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

where $a_i(\alpha_i) = \log(A_i(\alpha_i))$. This is a system of linear equations whose solution is (14). The log price vector is also normally distributed since it is a linear transformation of normal random variable. Combining with (7) yields (15). \square

D.2 Proof of Corollary 1

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

1. *The impact of a change in firm-level expected TFP μ_i on expected GDP $\mathbf{E}[y]$ is given by*

$$\frac{\partial \mathbf{E}[y]}{\partial \mu_i} = \omega_i.$$

2. *The impact of a change in firm-level volatility Σ_{ij} on the variance of GDP $\mathbf{V}[y]$ is given by³⁰*

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

Proof. (17) implies that $\frac{\partial \mathbf{E}[y(\alpha)]}{\partial \mu_i} = \beta' \mathcal{L}(\alpha) 1_i$. Since $P'C = WL = 1$ by the household's budget constraint, we need to show that $\beta' \mathcal{L}(\alpha) 1_i = P_i Q_i$ to complete the proof of the first result. From (29), we know that $P_i C_i = \beta_i$. Using Shepard's Lemma together with the marginal pricing equation (12), we can find the firm's factor demands equations

$$\begin{aligned} P_j X_{ij} &= \alpha_{ij} P_i Q_i \\ L_i &= \left(1 - \sum_{j=1}^n \alpha_{ij} \right) P_i Q_i. \end{aligned} \quad (32)$$

³⁰For $i \neq j$, the derivative simultaneously changes Σ_{ij} and Σ_{ji} to preserve the symmetry of Σ .

Using these results, we can write the market clearing condition (13) as

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

Solving the linear system implies

$$\beta' \mathcal{L}(\alpha) \mathbf{1}_i = P_i Q_i, \quad (33)$$

which proves the first part of the proposition. For the second part of the result, differentiating (18) with respect to Σ_{ij} and holding Σ symmetric yields

$$\frac{\partial V[y(\alpha)]}{\partial \Sigma_{ij}} = \begin{cases} \beta' \mathcal{L}(\alpha) \mathbf{1}_i \mathbf{1}_i' \mathcal{L}(\alpha)' \beta & i = j, \\ \beta' \mathcal{L}(\alpha) [\mathbf{1}_i \mathbf{1}_j' + \mathbf{1}_j \mathbf{1}_i'] \mathcal{L}(\alpha)' \beta & i \neq j \end{cases} = \begin{cases} \omega_i^2 & i = j, \\ 2\omega_i \omega_j & i \neq j, \end{cases}$$

which is the result. \square

D.3 Proof of Lemma 2

Lemma 2. $\lambda(\alpha^*)$, $k_i(\alpha_i, \alpha^*)$ and $q_i(\alpha^*)$ are normally distributed and the technique choice problem of the firm can be written as

$$\alpha_i^* \in \arg \min_{\alpha_i \in \mathcal{A}_i} \mathbb{E}[k_i(\alpha_i, \alpha^*)] + \frac{1}{2} \text{V}[k_i(\alpha_i, \alpha^*)] + \text{Cov}[k_i(\alpha_i, \alpha^*), \lambda(\alpha^*) + q_i(\alpha^*)]. \quad (19)$$

Proof. We first consider the stochastic discount factor. (31) shows that aggregate consumption can be written as a function of prices. Combining that equation with (6) we can write $\lambda = \log(\Lambda)$ as

$$\lambda(\alpha^*) = -(1 - \rho) \sum_{i=1}^n \beta_i p_i(\alpha^*) \quad (34)$$

Taking the log of (9) yields

$$k_i(\alpha_i, \alpha^*) = -(\varepsilon_i + a(\alpha_i)) + \sum_{j=1}^n \alpha_{ij} p_j(\alpha^*). \quad (35)$$

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 problem of the firm, we can write (11) as

$$\alpha_i^* \in \arg \min_{\alpha_i \in \mathcal{A}_i} \mathbb{E}[\Lambda Q_i K_i(\alpha_i, P)], \quad (36)$$

or, taking the logs

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

where $q_i (\alpha^*) = \log Q_i (\alpha^*)$ and where we emphasize that λ and q_i depend only on the equilibrium technique choice α^* . From (33), q_i is normally distributed and so are all the terms in the exponential. We can therefore use the expression for the expected value of a lognormal distribution and write

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

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

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

We can expand that expression as

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

The term $\mathbb{V} [\lambda (\alpha^*) + q_i (\alpha^*)]$ can be dropped as it does not affect the optimization and we find (19).

For later derivations, it is also convenient to write (36) in terms of P_i as

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

where we have used (33). We can drop $\beta' \mathcal{L} (\alpha^*) 1_i \geq 0$ since it is deterministic and does not depend on α_i . Going through the same steps as above, the firm's problem becomes

$$\alpha_i^* \in \arg \min_{\alpha_i \in \mathcal{A}} \mathbb{E} [k_i (\alpha_i, \alpha^*)] + \frac{1}{2} \mathbb{V} [\lambda (\alpha^*) - p_i (\alpha^*) + k_i (\alpha_i, \alpha^*)]. \quad (37)$$

□

D.4 Proof of Proposition 1

Proposition 1. *An equilibrium exists.*

Proof. We group here the proofs of Lemma 3 and Proposition 1. We proceed in three steps. First we show that there is a unique technique α_i that solves the problem of the firm, i.e. \mathcal{K}_i is a function. Second, we show that that function is continuous. Finally, we use a fixed-point theorem to show the existence of an equilibrium.

Step 1. We show that the right-hand side of (37) is a strictly concave function. First, note that from (35) we can write

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

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

Similarly, combining (34) and (35) we can write

$$\frac{1}{2} \mathbb{V}[\lambda(\alpha^*) - p_i(\alpha^*) + k_i(\alpha_i, \alpha^*)] = \frac{1}{2} \mathbb{V}\left[-(\varepsilon_i + a(\alpha_i)) - p_i(\alpha^*) + \sum_{j=1}^n (\alpha_{ij} - (1 - \rho)\beta_j) p_j(\alpha^*)\right].$$

We can remove the term $a(\alpha_i)$ from the variance as it is not stochastic. Combining with the equilibrium price equation (14), we get

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

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

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

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

$$\mathbb{V}[\lambda(\alpha^*) - p_i(\alpha^*) + k_i(\alpha_i, \alpha^*)] = B(\alpha_i, \alpha^*) \Sigma B(\alpha_i, \alpha^*)',$$

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

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

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

$$\mathbb{E}[k_i(\alpha_i, \alpha^*)] + \frac{1}{2} \mathbb{V}[\lambda(\alpha^*) - p_i(\alpha^*) + k_i(\alpha_i, \alpha^*)]$$

is strictly convex in the vector α_i .

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

Step 2. We now show that κ_i is continuous. To simplify the notation, define

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

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

We will first show that $g_i(\alpha, \alpha^*)$ is continuous. From (34) and (35), λ and k are continuous functions of α and linear functions of $p(\alpha^*)$. It therefore suffices to show that $p(\alpha^*)$ is continuous. From (14), we see that $p(\alpha^*)$ is continuous since $\mathcal{L}(\alpha^*)$, as a matrix inverse, is continuous and $a(\alpha^*)$ is continuous by Assumption 1. So $g_i(\alpha, \alpha^*)$ is continuous.

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

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

D.5 Proof of Proposition 2

Proposition 2. *There exists an efficient equilibrium.*

Proof. Since we only have one agent in the economy, any Pareto efficient allocation must maximize the utility of the representative household. Under a given network and a given productivity shock ε the first welfare theorem applies and the equilibrium is efficient. The consumption chosen by the planner is therefore given by (15). Taking a step back, the efficient production network must

therefore solve

$$\begin{aligned}
\max_{\alpha \in \mathcal{A}} \mathbb{E}[u(Y)] &= \max_{\alpha \in \mathcal{A}} \frac{1}{1-\rho} \mathbb{E}[\exp((1-\rho) \log Y)] \\
&= \max_{\alpha \in \mathcal{A}} \frac{1}{1-\rho} \exp\left((1-\rho) \mathbb{E}[\log Y] + \frac{1}{2} (1-\rho)^2 \text{V}[\log Y]\right) \\
&= \max_{\alpha \in \mathcal{A}} \mathbb{E}[\log Y] - \frac{1}{2} (\rho-1) \text{V}[\log Y]
\end{aligned} \tag{38}$$

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

First-order conditions of the planner. Using (17) and (18), we can write (38) as

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

The first-order conditions are

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

where $\underline{\mu}_{ij}$ are the Lagrange multipliers on the constraints on $\alpha_{ij} \geq 0$ and γ_i is the Lagrange multiplier on the constraint $\sum_j \alpha_{ij} \leq \bar{\alpha}_i$.³¹ Now,

$$\frac{\partial}{\partial \alpha_{ij}} \mathcal{L}(\alpha) = \frac{\partial}{\partial \alpha_{ij}} (I - \alpha)^{-1} = -(I - \alpha)^{-1} \left[\frac{\partial}{\partial \alpha_{ij}} (I - \alpha) \right] (I - \alpha)^{-1} \tag{39}$$

$$= (I - \alpha)^{-1} [O_{ij}] (I - \alpha)^{-1} = \mathcal{L}(\alpha) O_{ij} \mathcal{L}(\alpha) \tag{40}$$

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

$$\begin{aligned}
0 &= \beta' \mathcal{L}(\alpha) \mathbf{1}_i \mathbf{1}'_j \mathcal{L}(\alpha) [\mu + a(\alpha)] + \beta' \mathcal{L}(\alpha) \mathbf{1}_i \frac{\partial}{\partial \alpha_{ij}} a_i(\alpha) \\
&\quad + (1-\rho) \beta' \mathcal{L}(\alpha) \mathbf{1}_i \mathbf{1}'_j \mathcal{L}(\alpha) \Sigma \mathcal{L}(\alpha)' \beta + \underline{\mu}_{ij} - \gamma_i
\end{aligned}$$

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

³¹Note that $\sum_j \alpha_{ij} \leq \bar{\alpha}_i < 1$ implies that $\alpha_{ij} < 1$, so we do not need to explicitly consider a constraint $\alpha_{ij} \leq 1$.

$$0 = 1'_j \mathcal{L}(\alpha) [\mu + a(\alpha)] + \frac{\partial}{\partial \alpha_{ij}} a_i(\alpha) + (1 - \rho) 1'_j \mathcal{L}(\alpha) \Sigma \mathcal{L}(\alpha)' \beta \quad (41)$$

$$+ (\beta' \mathcal{L}(\alpha) 1_i)^{-1} (\underline{\mu}_{ij} - \gamma_i) \quad (42)$$

Firms first-order conditions. We can repeat similar steps for the equilibrium. Combining (37) with (14), (34) and (35), we find that firm i 's problem can be written as

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

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

$$\begin{aligned} 0 &= -\frac{\partial a(\alpha_i)}{\partial \alpha_{ij}} - 1'_j \mathcal{L}(\alpha^*) (\mu + a(\alpha^*)) + \frac{1}{2} (1'_j \mathcal{L}(\alpha^*)) \Sigma ((\alpha_i - 1_i - (1 - \rho) \beta)' \mathcal{L}(\alpha^*) + 1'_i)' \\ &+ \frac{1}{2} ((\alpha_i - 1_i - (1 - \rho) \beta)' \mathcal{L}(\alpha^*) + 1'_i) \Sigma \mathcal{L}(\alpha^*)' 1_j + \underline{\mu}_{ij}^e - \gamma_i^e \end{aligned}$$

or

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

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

$$-\frac{\partial a(\alpha_i^*)}{\partial \alpha_{ij}} - 1'_j \mathcal{L}(\alpha^*) (\mu + a(\alpha^*)) + ((\alpha_i^* - 1_i - (1 - \rho) \beta)' \mathcal{L}(\alpha^*) + 1'_i) \Sigma \mathcal{L}(\alpha^*)' 1_j + \underline{\mu}_{ij}^e - \gamma_i^e = 0.$$

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

$$-\frac{\partial a(\alpha_i^*)}{\partial \alpha_{ij}} - 1'_j \mathcal{L}(\alpha^*) (\mu + a(\alpha^*)) - (1 - \rho) \beta' \mathcal{L}(\alpha^*) \Sigma \mathcal{L}(\alpha^*)' 1_j + \underline{\mu}_{ij}^e - \gamma_i^e = 0.$$

Notice that these are the same first-order conditions (up to a normalization of the Lagrange multipliers) as the planner's (equation 41). The complementary slackness conditions are also the same in both problems. As a result, any equilibrium allocation also satisfied the planner's first-order conditions and vice versa. \square

D.6 Proof of Corollary 2

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

$$\max_{\alpha \in \mathcal{A}} E[y(\alpha)] - \frac{1}{2}(\rho - 1)V[y(\alpha)], \quad (21)$$

where \mathcal{W} is the welfare of the representative household and y is GDP as defined in (15).

Proof. This is an intermediate result that was proven at (38) in the proof of Proposition 2. \square

D.7 Generic uniqueness of the efficient equilibrium

Consider the planner's objective function from (21): $\mathcal{W}(\alpha; z) = \beta' \mathcal{L}(\alpha) (\mu + a(\alpha, z)) + \frac{1}{2}(1 - \rho) \beta' \mathcal{L}(\alpha) \Sigma \mathcal{L}(\alpha)' \beta$, where z is a vector of parameters, which includes μ, Σ, β, ρ and any additional parameters of the $a(\alpha, z)$ function. Define a space \mathcal{Z} on the set of parameters z . We endow this space with an absolutely continuous probability measure \mathbb{P} . We will call the solution to that problem generically unique if the set \mathcal{Z}^* for which \mathcal{W} has multiple maximizers is almost surely empty, i.e. $\mathbb{P}(z \in \mathcal{Z}^*) = 0$.

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

Proposition A1. *Suppose that $A_i(\alpha_i)$ takes the form (3) and all elements of the κ matrix are positive.³² Then the Pareto efficient equilibrium is generically unique.*

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

1. The set $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$, where \mathcal{A}_i is the set of feasible production technique for firm i given by (2), must be a disjoint union of finitely or countably many second-countable Hausdorff manifolds, possibly with boundary or corner. This assumption is satisfied in our case since \mathcal{A} is a manifold in \mathbb{R}^{n^2} of dimension $n^2 - n$.
2. We need $\mathcal{W}(\alpha, z)$ to be differentiable with respect to z and the derivative to be continuous with respect to α and z . This is satisfied in our case given the form (3).
3. It must be that for all $\alpha_1, \alpha_2 \in \mathcal{A}$ such that $\alpha_1 \neq \alpha_2$ we have $\frac{d\mathcal{W}(\alpha_1, z)}{dz} \neq \frac{d\mathcal{W}(\alpha_2, z)}{dz}$, where the derivative here indicates the gradient. We prove this by contraposition. For that purpose, take $\alpha_1, \alpha_2 \in \mathcal{A}$ such that $\frac{d\mathcal{W}(\alpha_1, z)}{dz} = \frac{d\mathcal{W}(\alpha_2, z)}{dz}$. We are going to show that it implies that $\alpha_1 = \alpha_2$. From Proposition 5, it must be that $\frac{d\mathcal{W}(\alpha_1, z)}{d\mu_i} = \omega_i(\alpha_1, z) = \omega_i(\alpha_2, z) = \frac{d\mathcal{W}(\alpha_2, z)}{d\mu_i}$. Since this is true for all i , it follows that the vector of Domar weights must be the same, that

³²This is the functional form we use for the quantitative analyses. In our calibrated model, all elements of the κ matrix are positive.

it $\omega(\alpha_1, z) = \omega(\alpha_2, z) > 0$. Next, differentiate $\mathcal{W}(\alpha, z)$ with respect to α_{il}^0 to write

$$\frac{d\mathcal{W}(\alpha, z)}{d\alpha_{il}^0} = 2\omega_i \left[\kappa_{il}(\alpha_{il} - \alpha_{il}^0) + \kappa_{i0} \left(\sum_{j=1}^n \alpha_{ij} - \sum_{j=1}^n \alpha_{ij}^0 \right) \right].$$

Suppose by contradiction that $\alpha_1 \neq \alpha_2$. Then there exists a pair i, l such that $(\alpha_{il})_1 \neq (\alpha_{il})_2$. Without loss of generality, suppose that $(\alpha_{il})_1 > (\alpha_{il})_2$. Then it must be that $\sum_{j=1}^n (\alpha_{ij})_1 < \sum_{j=1}^n (\alpha_{ij})_2$ for $\frac{d\mathcal{W}(\alpha_1, z)}{d\alpha_{il}^0} = \frac{d\mathcal{W}(\alpha_2, z)}{d\alpha_{il}^0}$ to hold. Therefore, there exists l' such that $(\alpha_{il'})_1 < (\alpha_{il'})_2$. But then it must be $\frac{d\mathcal{W}(\alpha_1, z)}{d\alpha_{il'}^0} < \frac{d\mathcal{W}(\alpha_2, z)}{d\alpha_{il'}^0}$. Therefore, we have a contradiction and $\alpha_1 = \alpha_2$.

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

D.8 Proof of Proposition 3

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

Proof. Fix the initial mean and variance-covariance matrix at μ^0 and Σ^0 , and denote the optimal network by $\alpha^*(\mu^0, \Sigma^0)$. Now, consider an increase in μ_i from μ_i^0 to μ_i^1 (holding other elements of μ^0 and Σ^0 fixed). The welfare changes from $\mathcal{W}(\alpha^*(\mu_i^0, \mu_{-i}^0, \Sigma^0); \mu_i^0, \mu_{-i}^0, \Sigma^0)$ to $\mathcal{W}(\alpha^*(\mu_i^1, \mu_{-i}^0, \Sigma^0); \mu_i^1, \mu_{-i}^0, \Sigma^0)$, which, by Proposition 5, can be written as

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

Now suppose instead that the network is fixed at its original value $\alpha^*(\mu_i^0, \mu_{-i}^0, \Sigma^0)$. From Equations (17) and (18), a change in μ_i affects welfare only through its impact on expected GDP. By Lemma 1, the change in welfare can be written as

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

Since the initial network $\alpha^*(\mu_i^0, \mu_{-i}^0, \Sigma^0)$ is attainable at $(\mu_i^1, \mu_{-i}^0, \Sigma^0)$, it must be that $\mathcal{W}(\alpha^*(\mu_i^1, \mu_{-i}^0, \Sigma^0); \mu_i^1, \mu_{-i}^0, \Sigma^0) \geq \mathcal{W}(\alpha^*(\mu_i^0, \mu_{-i}^0, \Sigma^0); \mu_i^1, \mu_{-i}^0, \Sigma^0)$. Because μ_i^1 can be picked arbitrary close to μ_i^0 , it must therefore be that $\omega_i(\mu_i^1, \mu_{-i}^0, \Sigma^0) \geq \omega_i(\mu_i^0, \mu_{-i}^0, \Sigma^0)$, or $\frac{d\omega_i}{d\mu_i} \geq 0$.

For the second part of the proposition, recall that $\frac{d\mathcal{W}}{d\Sigma_{ii}} = (1 - \rho)\omega_i^2$ by Proposition 5. Using analogous steps, we then can establish the second part of this proposition. \square

D.9 Proof of Lemma 4

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

(i) *an increase in μ_j leads to an increase in the shares α_{kl}^* for all k, l ;*

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

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

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

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

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

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

Then

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

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

Next, differentiate F_{jk} with respect to α_{lm} to get

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

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

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

$$\begin{pmatrix} \left(\frac{\partial \alpha_1}{\partial \mu_i} \right)' \\ \left(\frac{\partial \alpha_2}{\partial \mu_i} \right)' \\ \dots \\ \left(\frac{\partial \alpha_n}{\partial \mu_i} \right)' \end{pmatrix} = -(A + B)^{-1} \frac{\partial F}{\partial \mu_i}. \quad (43)$$

We will now show that when $\Sigma = 0$ (and so $B = 0$), all the elements on the left-hand side of (43) are positive. Since the right-hand side of (43) is continuous in the elements of Σ , the left-hand side will remain positive for small Σ .

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

$$\lambda_{\min}(A_i^{-1}) \leq \frac{x' A_i^{-1} x}{x' x} \leq \lambda_{\max}(A_i^{-1}),$$

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

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

Now, in the case of no uncertainty, $\Sigma = 0$, $B = 0$ and the right-hand side of 43 must be strictly positive and so is the vector of $\frac{\partial \alpha_{kl}}{\partial \mu_i}$. In this case, both parts of the Lemma hold. If there is uncertainty ($\Sigma \neq 0$), the result still holds if all the elements of Σ are sufficiently close to zero. Indeed, $-(A + B)^{-1}$ is continuous in Σ and, thus there exists $\bar{\Sigma} > 0$ such that if $|\Sigma_{ij}| < \bar{\Sigma}$ for all $i, j \in \mathcal{N}^2$ then elements of $-(A + B)^{-1} \frac{\partial F}{\partial \mu_i}$ have the same signs as the corresponding elements of

$-A^{-1} \frac{\partial F}{\partial \mu_i}$.³³

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

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

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

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

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

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

D.10 Proof of Proposition 4

Proposition 4. *Uncertainty lowers the expected value of GDP, such that $E[y]$ is largest when $\Sigma = 0$.*

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

D.11 Proof of Proposition 5

Proposition 5. *When the network α is free to adjust to changes in μ and Σ , the following holds.*

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

$$\frac{dW}{d\mu_i} = \frac{\partial E[y]}{\partial \mu_i} = \omega_i. \quad (22)$$

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

$$\frac{dW}{d\Sigma_{ij}} = \begin{cases} -\frac{1}{2} (\rho - 1) \left(\frac{\partial E[y]}{\partial \mu_i} \right)^2 = -\frac{1}{2} (\rho - 1) \omega_i^2 & i = j, \\ -(\rho - 1) \frac{\partial E[y]}{\partial \mu_i} \frac{\partial E[y]}{\partial \mu_j} = -(\rho - 1) \omega_i \omega_j & i \neq j. \end{cases} \quad (23)$$

³³Note that $(A + B)^{-1}$ exists for small Σ because A the eigenvalues of A are strictly negative (and so $\det(A) \neq 0$ and A is invertible) and that the determinant of $A + B$ is a continuous function of Σ . Note also that as we move Σ away from 0 the optimal matrix α changes and so do A and \mathcal{L} . But these changes are continuous so the strict inequality $-(A + B)^{-1} \frac{\partial F}{\partial \mu_i} > 0$ is preserved for small enough Σ .

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

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

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

$$\frac{d\mathcal{W}}{d\mu_i} = \frac{\partial \mathbb{E}[y]}{\partial \mu_i} = \beta' \mathcal{L}(\alpha) \mathbf{1}_i = \omega_i,$$

and

$$\frac{d\mathcal{W}}{d\Sigma_{ij}} = -\frac{1}{2}(\rho - 1) \frac{\partial \text{V}[y(\alpha)]}{\partial \Sigma_{ij}} = (1 - \rho) \beta' \mathcal{L}(\alpha) (\mathbf{1}_i \mathbf{1}_j') \mathcal{L}(\alpha)' \beta = (1 - \rho) \omega_i \omega_j,$$

where we used the expressions for the expectation and the variance of output given by (17) and (18). \square

D.12 Proof of Proposition 6

Proposition 6. *Let $\alpha^*(\mu, \Sigma)$ be the equilibrium production network under (μ, Σ) and let $\mathcal{W}(\alpha, \mu, \Sigma)$ be the welfare of the household under the network α . Then the change in welfare after a change from (μ, Σ) to (μ', Σ') is such that*

$$\underbrace{\mathcal{W}(\alpha^*(\mu', \Sigma'), \mu', \Sigma') - \mathcal{W}(\alpha^*(\mu, \Sigma), \mu, \Sigma)}_{\text{Change in welfare under the flexible network}} \geq \underbrace{\mathcal{W}(\alpha^*(\mu, \Sigma), \mu', \Sigma') - \mathcal{W}(\alpha^*(\mu, \Sigma), \mu, \Sigma)}_{\text{Change in welfare under the fixed network}}. \quad (24)$$

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

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

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

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

Combining the two expression gives the result. \square