

Endogenous Production Networks under Supply Chain Uncertainty*

Alexandr Kopytov
University of Rochester

Bineet Mishra
Cornell University

Kristoffer Nimark
Cornell University

Mathieu Taschereau-Dumouchel
Cornell University

May 22, 2024

Abstract

Supply chain disturbances can lead to substantial increases in production costs. To mitigate these risks, firms may take steps to reduce their reliance on volatile suppliers. We construct a model of endogenous network formation to investigate how these decisions affect the structure of the production network and the level and volatility of macroeconomic aggregates. When uncertainty increases in the model, producers prefer to purchase from more stable suppliers, even though they might sell at higher prices. The resulting reorganization of the network tends to reduce macroeconomic volatility, but at the cost of a decline in aggregate output. The model also predicts that more productive and stable firms have higher Domar weights—a measure of their importance as suppliers—in the equilibrium network. We provide a basic calibration of the model using U.S. data to evaluate the importance of these mechanisms.

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

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

1 Introduction

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

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

To study how supply chain uncertainty affects firms’ sourcing decisions and how these decisions affect the macroeconomy, we construct a model of endogenous network formation that builds on Acemoglu and Azar (2020). In the model, firms produce differentiated goods that can be consumed by a risk-averse representative household or used as intermediate inputs by other producers. Firms can produce their goods using different *production techniques*. A technique is a production function that specifies which intermediate inputs to use and how these inputs are to be combined. Techniques can also differ in terms of productivity. When choosing a technique, a firm can marginally adjust the importance of a supplier or drop that supplier altogether. Consequently, these decisions, when aggregated, lead to changes in the production network along both the intensive and extensive margins.

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

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

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

sentative household's stochastic discount factor and, as such, they inherit the household's attitude toward risk. Consequently, while a firm would generally prefer to purchase from a more productive supplier, it might decide otherwise if this supplier is also riskier. Such a supplier would sell at a lower price on average but it is also more likely to suffer from a large negative productivity shock, in which case the price of its good would rise substantially. Potential customers consider this possibility and balance concerns between average productivity and stability when choosing their production techniques.

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

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

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

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

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

ductivity or a decline in volatility always benefit welfare, their impact on expected GDP can be the opposite of what one would expect. For instance, an *increase* in expected productivity can lead to a *decline* in expected GDP, so that Hulten’s theorem is not a good guide to understanding changes in GDP, even as a first-order approximation. To understand why, imagine a producer with (on average) low but stable productivity. Its high output price makes it unattractive as a supplier. But if its expected productivity increases, its risk-reward profile improves, and other producers might begin to purchase from it. Doing so, they might move away from more productive but riskier producers and, as a result, expected GDP might fall. We show that a similar mechanism also implies that an increase in the volatility of a sector’s productivity can lead to a decline in the variance of aggregate output.

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

The model that we use for our quantitative analysis relies on some simplifying assumptions for tractability reasons. To verify the robustness of our findings, we provide additional empirical evidence that does not rely on the structure of the model. Taking advantage of rich firm-level U.S. data, we find that, as in the model, higher uncertainty leads to a decline in Domar weights, and that network connections involving riskier suppliers are more likely to break down. These results are robust to using different measures of uncertainty and instruments from [Alfaro et al. \(2019\)](#) to tease out exogenous variation in uncertainty.

Our work is related to a large literature that investigates the impact of uncertainty on macroeconomic aggregates ([Bloom, 2009, 2014](#); [Bloom et al., 2018](#)). We propose a novel mechanism through which uncertainty can lower expected GDP. This mechanism operates through a flight to safety process in which firms facing higher uncertainty switch to safer but less productive suppliers, leading to lower but less volatile GDP. In a recent paper, [David et al. \(2022\)](#) argue that uncertainty may lead capital to flow to firms that are less exposed to aggregate risk, rather than to those firms where it would be most productive. In their model, as in ours, uncertainty can lead to lower aggregate output and measured TFP.²

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

²[Fernández-Villaverde et al. \(2011\)](#) investigate the real impact of interest rate volatility for emerging economies. [Jurado et al. \(2015\)](#) provide econometric estimates of time-varying macroeconomic uncertainty. [Baker et al. \(2016\)](#) measure economic policy uncertainty based on newspaper coverage. [Nieuwerburgh and Veldkamp \(2006\)](#) and [Fajgelbaum et al. \(2017\)](#) develop models in which uncertainty can have a long-lasting impact on economic aggregates.

Acemoglu et al. (2012) derive conditions on input-output networks under which idiosyncratic shocks result in aggregate fluctuations, even when the number of producers is large.³ Acemoglu et al. (2017) and Baqaee and Farhi (2019a) describe conditions under which production networks can generate fat-tailed aggregate output distributions. Foerster et al. (2011) and Atalay (2017) study the empirical contributions of sectoral shocks for aggregate fluctuations. Carvalho and Gabaix (2013) argue that the reduction in aggregate volatility during the Great Moderation (and its potential recent undoing) can be explained by changes in the input-output network.⁴

In most of this literature, Hulten’s (1978) theorem applies, so that sales shares are a sufficient statistic to predict the impact of microeconomic shocks on macroeconomic aggregates. In contrast, since firms choose production techniques in the presence of uncertainty, Hulten’s theorem is not a useful guide to how productivity affects expected GDP in our model.⁵ An increase in expected sectoral productivity can even have a negative impact on expected GDP.

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

Several papers in the network literature endow firms with CES production functions, so that

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

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

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

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

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

The next section introduces our model of network formation under uncertainty. In [Section 3](#), we characterize the equilibrium when the network is fixed. We then investigate the firms' technique choice problem in [Section 4](#) and consider the full equilibrium with a flexible network in [Section 5](#). In [Sections 6](#) and [7](#), we describe how the productivity process affects the production network, welfare and GDP. In [Section 8](#), we provide a basic calibration of the model. [Section 9](#) provides additional empirical evidence in support of the mechanisms. The last section concludes. All proofs are in [Appendix A](#) and Supplemental Appendix B in [Kopytov et al. \(2024\)](#).

2 A model of endogenous network formation under uncertainty

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

2.1 Firms and production functions

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

As in [Oberfield \(2018\)](#) and [Acemoglu and Azar \(2020\)](#), the representative firm in sector i has access to a set of production techniques \mathcal{A}_i . A technique $\alpha_i \in \mathcal{A}_i$ specifies the set of inputs that are used in production, how these inputs are to be combined, and a productivity shifter $A_i(\alpha_i)$. We model these techniques as Cobb-Douglas technologies that can vary in terms of factor shares and total factor productivity. It is therefore convenient to identify a technique $\alpha_i \in \mathcal{A}_i$ with the intermediate input shares associated with that technique, $\alpha_i = (\alpha_{i1}, \dots, \alpha_{in})$, and to write the

corresponding production function as

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

where L_i is labor and $X_i = (X_{i1}, \dots, X_{in})$ is a vector of intermediate inputs. The term ε_i is the stochastic component of sector i 's total factor productivity. Finally, $\zeta(\alpha_i)$ is a normalization to simplify future expressions.⁷

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

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

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

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

This assumption is both technical and substantial in nature. The strict log-concavity ensures that there exists a unique technique that solves the optimization problem of the firm. It also implies that, for each sector i , there is a unique vector of *ideal* input shares $\alpha_i^\circ \in \mathcal{A}_i$ that maximizes A_i and that represents the most productive way to combine intermediate inputs to produce good i . Without loss of generality, we normalize $A_i(\alpha_i^\circ) = 1$ for all i .⁸

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

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

⁷Namely, $[\zeta(\alpha_i)]^{-1} = \left(1 - \sum_{j=1}^n \alpha_{ij}\right)^{1 - \sum_{j=1}^n \alpha_{ij}} \prod_{j=1}^n \alpha_{ij}^{\alpha_{ij}}$. This normalization is useful to simplify the unit cost expression, given by (8) below. $\zeta(\alpha_i)$ could instead be included in $A_i(\alpha_i)$ without any impact on the model.

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

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

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

In equilibrium, ε will have a direct impact on prices, and its moments (μ, Σ) will affect expectations about the price system. For instance, a sector with a high μ_i will have a low expected unit cost and therefore the price of good i will be low in expectation. Similarly, a high Σ_{ii} implies large productivity shocks and a volatile price of good i . Since production techniques must be chosen before ε is realized, the beliefs (μ, Σ) will affect the sourcing decisions of the firms.

We impose that the representative firm in sector i can only adopt one 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 \mathcal{A}_i . After the realization of the productivity shocks, the firm would only operate the plant that is best suited to the specific ε draw. All the other plants would remain idle. In reality, we think that fixed costs would prevent the firm from setting up all these plants. Information frictions might also impede the reallocation of sectoral demand to the best suited technique. To avoid burdening the exposition of the model, we adopt this restriction in an ad hoc fashion here, but provide a microfoundation for it in Supplemental Appendix C in [Kopytov et al. \(2024\)](#).

2.2 Household preferences

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

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

where $\beta_i > 0$ for all i and $\sum_{i=1}^n \beta_i = 1$. We refer to $Y = \prod_{i=1}^n (\beta_i^{-1} C_i)^{\beta_i}$ as aggregate consumption or, equivalently in this setting, GDP. The utility function $u(\cdot)$ is CRRA with a coefficient of relative risk aversion $\rho \geq 1$.¹⁰ The household makes consumption decisions after uncertainty is resolved

⁹This assumption is common in the literature but implies that an increase in the variance Σ_{ii} of a sector i increases its expected TFP. Through this channel the adverse effect of an increase in uncertainty is mitigated. A common way to correct for this effect is to remove half of the variance of ε from its mean. Supplemental Appendix F in [Kopytov et al. \(2024\)](#) describes why such a correction is problematic in our model and discusses other potential corrections.

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

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

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

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

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

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

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

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

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

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

2.3 Unit cost minimization

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

We begin by solving the second stage problem. Under a given technique α_i , the cost minimization problem of a firm in sector i is

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

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

to maximizing $E[\log Y] - \frac{1}{2}(\rho - 1)V[\log Y]$ such that $\rho \leq 1$ indicates whether the household likes uncertainty in log consumption or not. This is a consequence of the usual increase in the mean of a log-normal variable from an increase in the variance of the underlying normal variable. See Supplemental Appendix F in [Kopytov et al. \(2024\)](#) for a version of the model in which we correct for this term.

¹¹See Supplemental Appendix A in [Kopytov et al. \(2024\)](#) for the derivation of equations (5) and (6).

an important role in our analysis. Since, for a given α_i , the firm operates a constant returns to scale technology, K_i does not depend on the scale of the firm and is only a function of the (relative) prices $P = (P_1, \dots, P_n)$. We show in Supplemental Appendix A in [Kopytov et al. \(2024\)](#) that the production function (1) implies that

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

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

2.4 Technique choice

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

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

where Q_i is the equilibrium demand for good i , and where the profits in different states of the world are weighted by the household's stochastic discount factor Λ . The representative firm takes P , Q_i and Λ as given, and so the only term in (9) over which it has any control is the unit cost $K_i(\alpha_i, P)$. The firm thus selects the technique $\alpha_i \in \mathcal{A}_i$ that minimizes the expected discounted value of the total cost of goods sold $Q_i K_i(\alpha_i, P)$, while taking into consideration that final consumption goods are valued differently across different states of the world, as captured by Λ .¹² Because profits are discounted by Λ , firms effectively inherit the risk attitude of the representative household.

2.5 Equilibrium conditions

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

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

For a given network $\alpha \in \mathcal{A}$, this equation, together with (8), allows us to fully characterize the price system as a function of the random productivity shocks ε .¹³

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

¹³Even without imposing that production techniques are Cobb-Douglas, the system (10) yields a unique price vector P under standard assumptions. But the Cobb-Douglas structure implies that we can write the *distribution* of P in closed form, which allows us to characterize the technique choice problem in a tractable way.

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

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

1. (Optimal technique choice) For each $i \in \{1, \dots, n\}$, the technique choice $\alpha_i^* \in \mathcal{A}_i$ solves (9) given prices P^* , demand Q_i^* and the stochastic discount factor Λ^* given by (5).
2. (Optimal input choice) For each $i \in \{1, \dots, n\}$, factor demands per unit of output L_i^*/Q_i^* and X_i^*/Q_i^* are a solution to (7) given prices P^* and the chosen technique α_i^* .
3. (Consumer maximization) The consumption vector C^* maximizes (3) subject to (4) given prices P^* .
4. (Unit cost pricing) For each $i \in \{1, \dots, n\}$, P_i^* solves (10) where $K_i(\alpha_i^*, P^*)$ is given by (8).
5. (Market clearing) For each $i \in \{1, \dots, n\}$,

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

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

It is straightforward to extend the model along several dimensions without losing tractability. For instance, the model can accommodate disturbances that happen at the link level instead of at the sectoral level. To do so, we can simply think of a link between two producers as a fictitious “transport” sector that is also subject to shocks. It is also straightforward to extend the model to include multiple primary factors or wedges between unit costs and prices. We work out this last extension in Supplemental Appendix I in [Kopytov et al. \(2024\)](#). In Supplemental Appendix J, we also consider additional sources of uncertainty in terms of 1) household preferences, 2) labor supply, and 3) distortions (e.g. due to government policies). We find that these sources of uncertainty either do not matter for the equilibrium network, matter only if they interact with the productivity shocks ε , or have a similar impact to the uncertainty about ε .

On the other hand, certain ingredients are essential to keep the model tractable. Here, the key challenge comes from the fact that technique choices affect equilibrium prices which in turn affect technique choices. The log-linearity implied by the Cobb-Douglas aggregators in (1) and (3) are needed to keep the equilibrium beliefs tractable. While this implies a unit elasticity of substitution

in the production function (1), this elasticity only captures the response of intermediate inputs to realized prices conditional on a chosen production technique. Since firms' expectations affect their technique choice, the model is able to handle richer substitution patterns between expected prices and intermediate input shares, as we explore in more details in Section 6.

3 Equilibrium prices and GDP in a fixed-network economy

Before analyzing how the equilibrium production network responds to changes in the productivity process, it is useful to first establish how prices and GDP behave under a fixed network. To this end, we first define two objects that will play a central role in our analysis.

The first is the Leontief inverse $\mathcal{L}(\alpha) = (I - \alpha)^{-1}$, which can also be written as the geometric sum $\mathcal{L}(\alpha) = I + \alpha + \alpha^2 + \dots$. An element i, j of $\mathcal{L}(\alpha)$ captures the importance of sector j as an input in the production of good i by taking into account direct and indirect connections between the two sectors in the production network.

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

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

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

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

and log GDP is given by

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

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

Lemma 1 describes how prices and GDP depend on 1) the productivity vector $\varepsilon + a(\alpha)$ and 2) the production network α . Since all the elements of $\omega(\alpha)$ and $\mathcal{L}(\alpha)$ are non-negative, an increase in productivity has a negative impact on log prices and a positive impact on log GDP when the network is fixed. Intuitively, as firms become more productive, their unit costs decline, and competition forces them to sell at lower prices. From the perspective of GDP, higher productivity implies that the available labor can be transformed into more consumption goods.

The lemma makes clear that production techniques α matter for prices and GDP through two distinct channels. They have a direct impact on the productivity shifters $a(\alpha)$ because different techniques have different productivities. In addition, α affects prices and GDP through its impact on the Leontief inverse and the Domar weights. The matrix $\mathcal{L}(\alpha) = I + \alpha + \alpha^2 + \dots$ in (12) implies that the price of good i depends not only on i 's productivity, but also on the productivity of its suppliers, and on the productivity of their suppliers, and so on. These higher-order connections also matter for GDP and thus the impact of sectoral productivity on aggregate output depends on the sector's importance, as captured by its Domar weight.

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

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

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

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

1. *The impact of a change in expected TFP μ_i on the moments of log GDP is given by*

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

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

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

The first part of the corollary demonstrates that for a fixed production network, [Hulten's \(1978\)](#) celebrated theorem also holds in expectational terms. That is, the change in expected log GDP following a change in the expected productivity of a sector i is equal to that sector's sales share ω_i . The second part of the corollary establishes a similar result for a change in Σ . It shows that the impact of an increase in the volatility of a sector's TFP on the variance of log GDP is equal to the square of that sector's sales share. This result also applies to a change in covariance, in which case the increase in $\mathbb{V}[y]$ is equal to the product of the two industries' sales shares. Since Domar weights are always positive, an increase in covariance always leads to higher aggregate volatility. Intuitively, positively correlated shocks are unlikely to offset each other, and their expected aggregate impact is

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

therefore larger. Finally, the Corollary shows that when the network is fixed, the covariance matrix Σ has no impact on $E[y]$. It follows that whenever we discuss the response of expected log GDP to a change in uncertainty, the mechanism must operate through the endogenous reorganization of the network.

Corollary 1 emphasizes that for a fixed network knowing the sales shares of every industry is sufficient to compute the impact of changes in μ and Σ on the moments of log GDP. In Section 7, we show that this is no longer true when firms can adjust their input shares in response to changes in the distribution of sectoral productivity. In fact, when the network is free to adjust, an increase in an element of μ can even lead to a decline in expected log GDP.

4 Firm decisions

In the previous section, we described prices under a given network. Here, we use that information to characterize the problem of the representative firm in sector i that must choose a technique $\alpha_i \in \mathcal{A}_i$. It is convenient to work with the log of the stochastic discount factor $\lambda(\alpha^*) = \log \Lambda(\alpha^*)$ and the log of the unit cost $k_i(\alpha_i, \alpha^*) = \log K_i(\alpha_i, P^*(\alpha^*))$, where α^* denotes the *equilibrium* network. These quantities are normally distributed in equilibrium.

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

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

The objective function in (15) captures how beliefs and uncertainty affect the production network. Its first term implies that the firm prefers to adopt techniques that provide, in expectation, a lower unit cost of production. Taking the expected value of the log of (8), we can write this term as

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

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

The second term in (15) captures the importance of *aggregate risk* for the firm's decision. It implies that the firm prefers to have a low unit cost in states of the world in which the marginal utility of consumption is high. As a result, the coefficient of risk aversion ρ of the household indirectly determines how risk-averse firms are. We can expand this term as $\text{Cov}[\lambda, k_i] = \text{Corr}[\lambda, k_i] \sqrt{V[\lambda]} \sqrt{V[k_i]}$, which implies that the firm tries to minimize the correlation of its unit cost with λ . Furthermore, since prices and GDP tend to move in opposite directions (see Lemma 1), $\text{Corr}[\lambda, k_i]$ is typically positive, and so firms seek to minimize the variance of their

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

unit cost.¹⁶ This has several implications for their choice of suppliers. To see this, we can use (8) to write

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

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

Risk-adjusted prices

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

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

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

where

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

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

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

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

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

¹⁶If i 's productivity shock is strongly negatively correlated with that of the other sectors, it can be that $\text{Corr} [\lambda, k_i] < 0$, in which case i seeks to be *more* volatile to insure the household in states of low consumption.

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

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

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

Whether the increase in \mathcal{R}_k leads to a decline or an increase in the share of other inputs $j \neq k$ depends on whether the shares of j and k are complements or substitutes in the production of good i . If $[H_i^{-1}]_{jk} > 0$ we say that they are *substitutes*, and in that case a higher risk-adjusted price \mathcal{R}_k leads to an increase in α_{ij} . As the firm decreases α_{ik} , the incentives embedded in a_i to increase α_{ij} get stronger, and the firm substitutes α_{ij} for α_{ik} . In contrast, if $[H_i^{-1}]_{jk} < 0$ we say that the shares of j and k are *complements*, and an increase in \mathcal{R}_k leads to a decline in α_{ij} . One sufficient condition for a Hessian matrix H_i to feature complementarities for all sectors is $[H_i]_{jk} \geq 0$ for all $j \neq k$.¹⁷

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

While (19) is only valid at an interior solution of the firm's problem, the forces that it captures are also at work when some of the constraints embedded in $\alpha_i \in \mathcal{A}_i$ bind. But these constraints can also increase the degree of substitution between input shares. Suppose, for instance, that the minimum labor share constraint $\sum_{l=1}^n \alpha_{il} \leq \bar{\alpha}_i$ binds, and that the risk-adjusted price of good j falls. To increase the share of good j , a firm in sector i would have to lower its share of some other input, say k , to avoid violating the constraint. In this case the shares of j and k would behave as substitutes in the production of good i .

Example: Substitutability and complementarity in partial equilibrium

To show how the substitution patterns embedded in a_i affect technique choices, we can revisit the car manufacturer example from the introduction. Suppose that this manufacturer primarily uses steel (input 1) to produce cars, and that it relies on equipment (input 2) such as milling machines and lathes to transform raw steel into usable components. As before, the manufacturer also has the option to purchase carbon fiber (input 3) to replace steel components if needed. It

¹⁷In this case $-H_i$ is an M-matrix and therefore inverse-positive. Intuitively, $[H_i]_{jk} \geq 0$ implies that a higher α_{ij} increases the TFP benefit of raising α_{ik} .

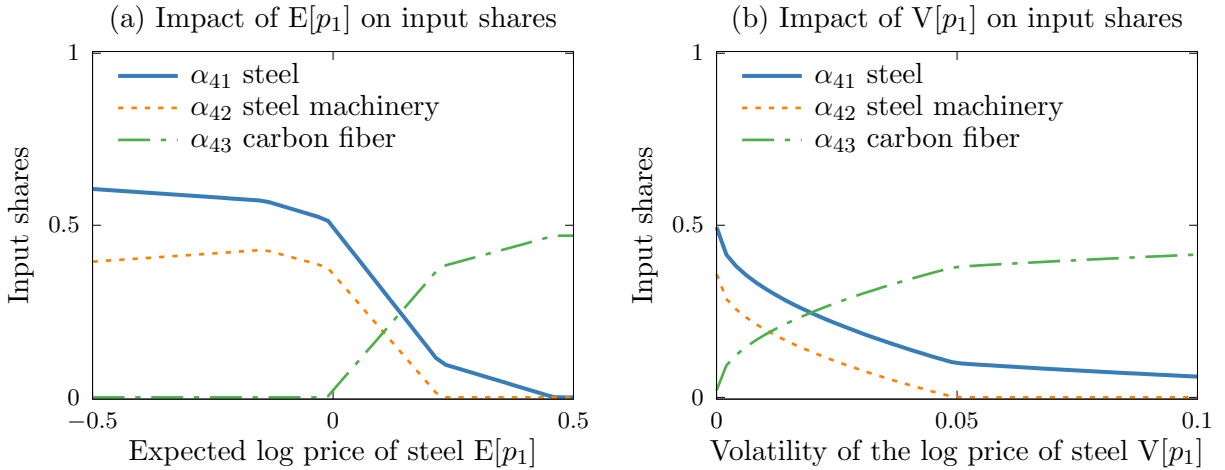
would be natural to endow this manufacturer (sector $i = 4$) with a TFP shifter function of the form

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

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

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

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



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

5 Equilibrium existence, uniqueness and efficiency

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

5.1 The efficient allocation

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

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

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

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

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

is welfare under a given network α .¹⁸

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

Recasting household welfare in terms of Domar weights

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

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

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

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

Formally, consider the optimization problem

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

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

Example. We can solve explicitly for $\bar{a}(\omega)$ and $\alpha(\omega)$ under the quadratic TFP shifter function specified in (2). At an interior solution $\alpha \in \text{int } \mathcal{A}$, the optimal production network $\alpha(\omega)$ that solves (23) for a given vector of Domar weights ω is¹⁹

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

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

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

From (24), it is straightforward to show that the gradients ∇a_i of the TFP shifter functions are all equal to each other such that $\nabla a_i = \nabla a_j$ for all i, j .²⁰ It follows that at an interior solution, input shares must be such that the marginal TFP benefit $[\nabla a_i]_k$ of increasing α_{ik} is equal across all sectors i .

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

¹⁹See Supplemental Appendix A in Kopytov et al. (2024) for the full derivation.

²⁰It is clear from (24) that $H_i(\alpha_i - \alpha_i^\circ) = H_j(\alpha_j - \alpha_j^\circ)$ for all i, j . Furthermore, since $a_i(\alpha)$ is a quadratic function given by (2), we have that $\nabla a_i = H_i(\alpha_i - \alpha_i^\circ)$, where $\nabla a_i = \frac{da_i}{d\alpha_i}$ denotes the gradient vector of a_i .

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

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

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

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

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

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

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

5.2 Fundamental properties of the equilibrium

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

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

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

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

6 Beliefs and the production network

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

6.1 Domar weights

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

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

This proposition can be understood from both the perspective of an individual producer and from the perspective of the social planner. Individual producers rely more on sectors whose prices are low and stable. As a result, these sectors are more important suppliers and their Domar weights are larger. From the planner's perspective, recall from (13) that the Domar weight of a sector captures its contribution to log GDP. Since the planner wants to increase and stabilize GDP, it naturally increases the importance of more productive (larger μ_i) and less volatile (smaller Σ_{ii}) sectors in the production network.

Risk-adjusted productivity shocks

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

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

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

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

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

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

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

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

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

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

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

Response of the Domar weight vector to changes in beliefs

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

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

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

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

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

and where the matrix $\nabla^2 \bar{a}$ is the Hessian of the aggregate TFP shifter function \bar{a} , and $\frac{d\mathcal{E}}{d\omega} = -\frac{d\text{Cov}[\varepsilon, \lambda]}{d\omega} = -(\rho - 1)\Sigma$ is the Jacobian matrix of the risk-adjusted TFP vector \mathcal{E} .²¹

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

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

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

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

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

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

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

Σ and global substitution patterns

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

²¹This proposition focuses on an interior equilibrium, such that $\omega \in \text{int } \mathcal{O}$, but this restriction can be satisfied even if some shares α_{ij} are equal to zero. Indeed, for $\omega_j \geq \beta_j$ to bind, it must be that $\alpha_{ij} = 0$ for all i . Furthermore, Proposition 3 can be extended to include some binding constraints. When $\omega_i \geq \beta_i$ binds, ω_i is not affected by a marginal change in beliefs. We can therefore exclude these constrained Domar weights from the application of the implicit function theorem. It follows that a version of (31) holds for unconstrained Domar weights, as we show in Supplemental Appendix D in Kopytov et al. (2024).

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

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

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

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

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

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

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

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

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

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

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

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

An approximate equation for the equilibrium Domar weights

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

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

This approximation is accurate if, for instance, the cost of deviating from the ideal shares embedded in the local TFP shifters is large. We work out that case formally in Supplemental Appendix G in Kopytov et al. (2024).

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

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

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

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

This proposition provides an approximate expression for the equilibrium Domar weights in terms of the global substitution patterns embedded in $[\mathcal{H}^\circ]^{-1}$ and the expected attractiveness of all

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

6.2 The production network

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

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

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

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

An approximate equation for the equilibrium production network

As for the Domar weights, one must in general use numerical methods to find the equilibrium network α . We can, however, derive an approximation for the equilibrium production network when the cost of deviating from the ideal shares α° is large. Specifically, let $a_i(\alpha_i) = \bar{\kappa} \times \hat{a}_i(\alpha_i)$, where \hat{a} does not depend on κ , and suppose that $\alpha_i^\circ \in \text{int } \mathcal{A}_i$. The parameter $\bar{\kappa} > 0$ captures how costly it is for the firms to deviate from α° in terms of TFP loss. When $\bar{\kappa}$ is large, we can use perturbation theory to derive an approximate equation for α (Judd and Guu, 2001; Schmitt-Grohé and Uribe, 2004).

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

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

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


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

Recall from (19) that H_i^{-1} describes how a marginal change in \mathcal{R} affects α_i in the problem of firm i . The approximation (36) captures the same forces. It shows that the deviation of α_i from α_i° depends, approximately, on the vector of risk-adjusted prices \mathcal{R} evaluated at the ideal shares α° . Intuitively, when it is costly for firms to deviate from α° , we can evaluate the equilibrium prices *as if* firms chose α° and use these prices to compute the sourcing decisions of the firm. By Lemma 9, these decisions provide a first-order approximation of the true equilibrium network.²²

Example: cascading link destruction

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

Figure 2. Cascading impact of a change in Σ_{44}



Notes: Arrows represent the movement of goods: there is a solid blue arrow from j to i if $\alpha_{ij} > 0$. Dashed gray arrows indicate $\alpha_{ij} = 0$. a is as in (69) in the appendix with $\kappa_{ij} = 0$ if there is a potential link between two firms and infinity otherwise. $\alpha_{ij}^\circ = 0.5$ if there is a potential link, and 0 otherwise. $\mu = 0$ except for $\mu_4 = 0.1$. In the left figure, $\Sigma = 0$. In the right figure $\Sigma = 0$ except $\Sigma_{44} = 1$. The risk aversion of the household is $\rho = 2$. $\beta_i = 1/n$ for all i .

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

²²See Supplemental Appendix G in Kopytov et al. (2024) for more details.

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

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

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

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

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

Instead, if we consider the response of α_{15} , the direct effect is absent because sector 5 does not rely on sector 4 ($\mathcal{L}_{54}^\circ = 0$). Since sector 2 is 1's only other possible connection, only the indirect effect through that sector remains. The relevant term here is $[(H_1^\circ)^{-1}]_{52} \mathcal{L}_{24}^\circ$, which is positive because $\mathcal{L}_{24}^\circ > 0$, and the shares of goods 5 and 2 are substitutes in the production of good 1, $[(H_1^\circ)^{-1}]_{52} > 0$. Therefore, an increase in Σ_{44} leads to a larger α_{15} . The same logic applies to the responses of firms 2 and 3, thus explaining the cascading effect illustrated in Figure 2.²³

²³Lemma 9 assumes that $\alpha \in \text{int } \mathcal{A}$, which is not the case in Figure 2, but it still captures the main forces that push the shares in response to changes in (μ, Σ) and is therefore informative about the response of the network.

7 Implications for GDP and welfare

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

7.1 Beliefs and welfare

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

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

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

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

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

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

This result follows directly from the fact that a flexible network provides an extra margin of adjustment to the planner and thus cannot leave the household worse off than under a fixed network.²⁴

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

²⁴We provide a proof of this result in Supplemental Appendix A in Kopytov et al. (2024).

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

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

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

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

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

7.2 Beliefs and GDP

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

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

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

Proposition 6 establishes a novel mechanism through which uncertainty reduces expected log GDP. To understand that mechanism, consider the technique choice problem from the firm's perspective. When there is no uncertainty, firms do not worry about risk and move toward cheaper suppliers, which tend to be the most productive ones, and toward techniques with higher TFP. As a result, the aggregate economy is maximally productive, and $E[y]$ is large. When some suppliers become risky, customers worry about a possible increase in input costs and start purchasing from more stable but less productive suppliers. As a result, the aggregate economy becomes less productive on average and expected log GDP falls.

The endogenous response of the network is essential for the result of Proposition 6. Indeed, in our model uncertainty affects expected log GDP *only* through the endogenous response of the network. If the shares α were fixed, uncertainty would have no impact on $E[y]$.

Response of GDP to a marginal change in beliefs

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

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

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

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

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

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

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

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

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

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

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

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

$$\frac{dE[y]}{d\mu_i} - \frac{\partial E[y]}{\partial \mu_i} \propto \frac{dV[y]}{d\mu_i} - \frac{\partial V[y]}{\partial \mu_i} = \underbrace{2\omega^\top \Sigma}_{\frac{dV[y]}{d\omega^\top}} \times \underbrace{\left(-\mathcal{H}^{-1} \frac{\partial \mathcal{E}}{\partial \mu_i}\right)}_{\frac{d\omega}{d\mu_i}} - 0,$$

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

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

The role of risk and of the global substitution patterns

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

with i , the increase in ω_i leads to a decline in all the other Domar weights. If the variances of those sectors are large relative to Σ_{ii} , this decline in Domar weights leads to a substantial decrease in $V[y]$. By the logic of Proposition 7, this implies that $E[y]$ must increase by less than its fixed-network term ω_i . Through a similar mechanism, an increase in Σ_{ii} leads to an increase in $V[y]$ that is larger than under a fixed network. In this case, $E[y]$ *increases* in response to the higher Σ_{ii} , such that uncertainty can be beneficial to expected log GDP at the margin.²⁵

Counterintuitive implications of changes in beliefs

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

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

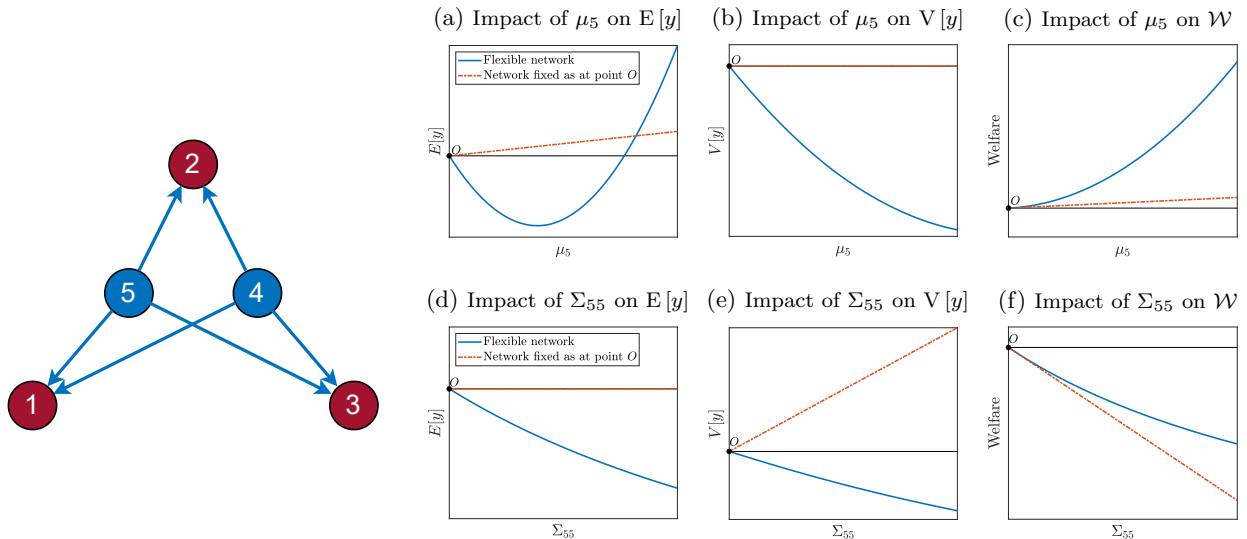
Consider the impact of a positive shock to μ_5 when inputs 4 and 5 are substitutes. The solid blue lines in panels (a) to (c) illustrate the impact of this change, and point O represents the economy before the change. As we can see, the initial *increase* in μ_5 has a *negative* impact on expected log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

²⁵This does not contradict Proposition 6 as Corollary 8 only applies at the margin when $\Sigma_{jk} > \bar{\Sigma} > 0$ for all $j, k \neq i$. Eliminating uncertainty altogether would still lead to an increase in $E[y]$.

To emphasize the role of the endogenous network for this mechanism, Figure 3 also shows the effect of the same increase in μ_5 when the network is kept fixed (dashed red lines). From Corollary 1, the marginal impact of μ_5 on expected log GDP is equal to its Domar weight, and increasing μ_5 has a positive impact on $E[y]$. At the same time, $V[y]$ is unaffected by changes in μ . While an increase in μ_5 is welfare-improving in this case, the effect is less pronounced than in the flexible network economy. Indeed, in the latter case the equilibrium network adjusts precisely to maximize the beneficial impact of the change in beliefs on welfare, as implied by (38).

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

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



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

8 A basic calibration of the model

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

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

8.1 Parametrization

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

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

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

where γ is an $n \times 1$ vector of deterministic drifts and $u_t \sim \text{iid } \mathcal{N}(0, \Sigma_t)$ is a vector of shocks. We further assume that firms know γ and ε_{t-1} at time t , so that the conditional mean and the covariance of beliefs are given by $\mu_t = \gamma + \varepsilon_{t-1}$ and Σ_t . Importantly, we allow uncertainty Σ_t to vary over time and estimate it from TFP data using a rolling window that puts more weight on more recent observations.

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

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

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

We next assess how well the calibrated model fits key moments in the data. As we have seen above, the Domar weights, and how they react to changes in μ_t and Σ_t , are central for the mechanisms of the model. The model is able to roughly replicate features of the empirical Domar weights, with a cross-sectional correlation between the time-averaged Domar weights in the model and in the data of 0.96. However, the average Domar weight in the model (0.03) is lower than its data counterpart (0.05).²⁶ Overall, the model can account for about 40% of the over-time standard deviation of Domar weights, which indicates that other mechanisms, such as technological progress that might expand the set of available techniques, might be at work in reality.

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

8.2 The production network, welfare and output

To evaluate the quantitative potential of an endogenous production network for welfare and GDP, we compare the calibrated model to two sets of alternative economies. First, we compare our baseline model to an economy in which the network is kept completely fixed at its sample average. This exercise therefore informs us about the overall impact of changes in the structure of the production network. We then investigate the role of uncertainty alone in shaping the production network. We do so by considering 1) an economy in which production techniques are chosen *as if* $\Sigma_t = 0$,²⁷ and 2) a perfect-foresight economy in which firms observe the realization of ε_t before making technique choices (the “known ε_t ” economy).²⁸ In both cases, uncertainty is irrelevant for decisions, and so these exercises allow us to isolate the impact of uncertainty on the production network and, through that channel, on macroeconomic aggregates.

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

²⁶We explain in Appendix B that this discrepancy can be explained by our choice to target consumption growth instead of GDP growth in the estimation.

²⁷Specifically, we set $\Sigma = 0$ when solving the problem of the social planner (21) for the equilibrium network α^* . We then reintroduce uncertainty when computing the moments of GDP and welfare.

²⁸One interpretation is that adopting a new technique is immediate, so that firms can wait to pick the best technique for a particular ε_t draw. Techniques and intermediate input choices are thus made simultaneously and conditional on observed prices.

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

When we isolate the role of uncertainty, however, these numbers become smaller. In line with the theory, the baseline economy is on average less productive and less volatile than under the “as if $\Sigma_t = 0$ ” alternative but the numbers are small, on the order of 0.01% for $E[y]$ and 0.10% for $V[y]$. This suggests that, for most of the sample period, uncertainty is sufficiently low that firms simply buy their inputs from the most productive suppliers without much concern for any risk involved.²⁹

The differences between our calibrated economy and the “no uncertainty” alternatives are however larger during high-uncertainty episodes like the Great Recession.³⁰ The top row of Figure 4 shows that expected log GDP in the baseline economy is about 0.25% lower in 2009 than in the alternative “as if $\Sigma_t = 0$ ” economy. Because of the large increase in uncertainty, firms adjust their production techniques toward safer but less productive suppliers to avoid potentially large increases in costs. The result in terms of aggregate volatility is visible in the top-right panel, where we see that log GDP is about 2.4% less volatile in 2009 in the baseline economy. Interestingly, realized log GDP, shown in the left-bottom panel, is substantially higher in the baseline economy than in the “as if $\Sigma_t = 0$ ” alternative. Essentially, firms took out an insurance against particularly bad TFP draws and opted for safer suppliers. When these fears were realized, this insurance policy paid off so that the baseline economy fared about 2.7% better in terms of realized log GDP compared to the alternative.

The right-bottom panel provides the same information for the “known ε_t ” alternative. In this case, beliefs (μ_t, Σ_t) , and in particular uncertainty, play no role in shaping the network and, from the planner’s problem, the optimal network is simply the one that maximizes (realized) consumption. It follows that realized consumption (or GDP) is always larger than in the baseline model. Unsurprisingly, the difference is particularly pronounced during episodes of high uncertainty, when knowing ε_t provides a larger advantage, and reaches a high of 3% during the Great Recession.³¹

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

²⁹As in Lucas (1987), the utility cost of business cycles is on average small in our model and the planner does not want to sacrifice much in terms of the level of GDP for a reduction in its volatility. We provide the same moments for the “known ε_t ” economy in Appendix B.4.


³⁰The differences between our calibrated and fixed-network economies are also particularly large during volatile periods, when adjusting the network is most beneficial. In Supplemental Appendix H in Kopytov et al. (2024), we show that allowing the network to adjust leads to large gains in expected GDP during the Great Recession.

³¹Since ε_t is known in this exercise, $E[y] = \mathcal{W} = y$ and $V[y] = 0$ and, so we do not report these moments in Figure 4. Alternatively, one can compute $E[y]$, $V[y]$ and \mathcal{W} before ε_t is known but still assuming that the production network is chosen optimally for the given realized draw of ε_t . We report these moments in Appendix B.4.


progress that would expand the set of production techniques. Similarly, the production function might not be Cobb-Douglas in reality, in which case changes in prices would affect Domar weights. We also made the implicit assumption that it takes one year (the frequency of our data) for firms to change production techniques. While this assumption might be reasonable for some sectors, it is likely that the time it takes to retool a factory varies significantly by industry, or even depending on what the new and the old techniques are.³² While we believe that the mechanisms that we explore in this paper would still be present in a richer model, more work would be needed to fully assess their importance.

Figure 4. The role of uncertainty in the postwar period

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




(a) Difference in expected log GDP [%]




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

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



(c) Difference in realized log GDP [%]

Right column: “known ε_t ” as alternative



(d) Difference in realized log GDP [%]

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

9 Model-free evidence for the mechanisms

The model proposed in this paper relies on simplifying assumptions for tractability. In this section, we present additional evidence in support of the main mechanisms of the model that does not rely on this structure. Through firm-level regressions that closely follow [Alfaro, Bloom, and Lin \(2019\)](#) we document that 1) higher uncertainty about a firm leads to a decline in its Domar weight, and 2) network connections involving riskier suppliers are more likely to break down. We test these predictions at the firm level to take advantage of the abundance of data and of instrumental variables that are available at this level of aggregation. Supplemental Appendix E in [Kopytov et al.](#)

³²In the car industry, General Motors took about one year to retool a factory for electric vehicle production ([Lutz, 2021](#)), but it took Ford eight weeks to switch from using steel to aluminum for the body of the F-150 ([Dean, 2015](#)).

(2024) describes the data and the instruments in detail.

9.1 Uncertainty and Domar weights

We first test the model’s prediction that Domar weights decrease with uncertainty. We use annual U.S. data from 1963 to 2016 provided by Compustat. Our main variables of interest are a firm’s Domar weight, constructed by dividing its sales by nominal GDP, and a measure of its stock price volatility, which we use as a proxy for uncertainty.³³ We then regress the change in Domar weight on the change in stock price volatility. The results are presented in the first column of Table I. In column (2), we follow Alfaro et al. (2019) and address potential endogeneity concerns by instrumenting stock price volatility with industry-level exposure to ten aggregate sources of uncertainty shocks. In column (3), we use option prices to back out an implied measure of future volatility. In all cases, we find a negative and significant relationship between uncertainty and Domar weights. The effect is also economically large with a decline in Domar weight of about 18% following a doubling in firm-level volatility (roughly a 3.3 standard deviation volatility shock), according to the IV estimates. Overall, these results provide evidence that higher uncertainty leads to lower Domar weights, in line with the predictions of our theoretical model.

TABLE I. Domar weights and uncertainty

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

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

³³Ersahin et al. (2022) use textual analysis of earning conference calls to measure firm-level supply chain risk, and find that it is positively correlated with stock price volatility. They also find that firms respond to higher supply chain risks by switching to a wider range of less risky suppliers.

9.2 Uncertainty and link destruction

We conduct a similar exercise, this time at the firm-to-firm relationship level, to investigate whether higher supplier uncertainty is associated with a higher likelihood of link destruction. We proceed by combining the uncertainty data described above with data from 2003 to 2016 about firm-level supply relationships provided by Factset. We then regress a dummy variable that equals one in the last year of a relationship on the change in the supplier’s stock price volatility. The results are presented in column (1) of Table II. As in the last exercise, column (2) uses industry-level sensitivity to aggregate shocks as instruments, and column (3) uses implied volatility from option prices as a measure of uncertainty. In all cases, we find a positive and statistically significant relationship between supplier volatility and the end of supply relationships, which is consistent with buyers moving away from riskier suppliers. The effect is also economically large with a doubling in volatility associated with a 12 percentage point increase in the likelihood that a relationship is destroyed, according to the IV estimates.

TABLE II. Link destruction and supplier volatility

	Dummy for last year of supply relationship		
	(1): OLS	(2): IV	(3): IV
$\Delta\text{Volatility}_{t-1}$ of supplier	0.026** (0.012)	0.097*** (0.035)	0.144** (0.063)
1st moment 10IV $_{t-1}$ of supplier	No	Yes	Yes
Type of volatility	Realized	Realized	Implied
Fixed effects	Yes	Yes	Yes
Observations	35,629	35,620	26,195
F -statistic	—	22.9	10.4

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

10 Conclusion

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

uncertainty, through its effect on the network, unambiguously lowers expected log GDP. When calibrated to the United States economy, the model predicts that the impact of uncertainty on the network can potentially have a sizable effect on GDP and welfare during periods of high uncertainty such as the Great Recession.

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

References

- ACEMOGLU, D. AND P. D. AZAR (2020): “Endogenous Production Networks,” *Econometrica*, 88, 33–82.
- ACEMOGLU, D., V. M. CARVALHO, A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2012): “The Network Origins of Aggregate Fluctuations,” *Econometrica*, 80, 1977–2016.
- ACEMOGLU, D., A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2017): “Microeconomic Origins of Macroeconomic Tail Risks,” *American Economic Review*, 107, 54–108.
- ACEMOGLU, D. AND A. TAHBAZ-SALEHI (2020): “Firms, Failures, and Fluctuations: The Macroeconomics of Supply Chain Disruptions,” Working paper.
- ALESSANDRIA, G., S. YAR KHAN, A. KHEDERKARIAN, C. MIX, AND K. J. RUHL (2022): “The Aggregate Effects of Global and Local Supply Chain Disruptions: 2020-2022,” Working paper.
- ALFARO, I., N. BLOOM, AND X. LIN (2019): “The Finance Uncertainty Multiplier,” Working paper.
- ATALAY, E. (2017): “How Important Are Sectoral Shocks?” *American Economic Journal: Macroeconomics*, 9, 254–280.
- ATALAY, E., A. HORTACSU, J. ROBERTS, AND C. SYVERSON (2011): “Network Structure of Production,” *Proceedings of the National Academy of Sciences*, 108, 5199–5202.
- BAKER, S. R., N. BLOOM, AND S. J. DAVIS (2016): “Measuring Economic Policy Uncertainty,” *Quarterly Journal of Economics*, 131, 1593–1636.
- BAQAEE, D. R. (2018): “Cascading Failures in Production Networks,” *Econometrica*, 86, 1819–1838.

- BAQAEE, D. R. AND E. FARHI (2019a): “The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten’s Theorem,” *Econometrica*, 87, 1155–1203.
- (2019b): “Productivity and Misallocation in General Equilibrium,” *Quarterly Journal of Economics*, 135, 105–163.
- BARROT, J.-N. AND J. SAUVAGNAT (2016): “Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks,” *Quarterly Journal of Economics*, 131, 1543–1592.
- BERNARD, A. B., E. DHYNE, G. MAGERMAN, K. MANOVA, AND A. MOXNES (2022): “The Origins of Firm Heterogeneity: A Production Network Approach,” *Journal of Political Economy*, 130, 000–000.
- BIGIO, S. AND J. LA’O (2020): “Distortions in Production Networks,” *Quarterly Journal of Economics*, 135, 2187–2253.
- BLOOM, N. (2009): “The Impact of Uncertainty Shocks,” *Econometrica*, 77, 623–685.
- (2014): “Fluctuations in Uncertainty,” *Journal of Economic Perspectives*, 28, 153–76.
- BLOOM, N., M. FLOETOTTO, N. JAIMOVICH, I. SAPORTA-EKSTEN, AND S. J. TERRY (2018): “Really Uncertain Business Cycles,” *Econometrica*, 86, 1031–1065.
- BOEHM, C. E., A. FLAAEN, AND N. PANDALAI-NAYAR (2019): “Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tōhoku Earthquake,” *Review of Economics and Statistics*, 101, 60–75.
- BOEHM, J. AND E. OBERFIELD (2020): “Misallocation in the Market for Inputs: Enforcement and the Organization of Production,” *Quarterly Journal of Economics*, 135, 2007–2058.
- CALIENDO, L., F. PARRO, AND A. TSYVINSKI (2022): “Distortions and the Structure of the World Economy,” *American Economic Journal: Macroeconomics*, 14, 274–308.
- CARVALHO, V. AND X. GABAIX (2013): “The Great Diversification and its Undoing,” *American Economic Review*, 103, 1697–1727.
- CARVALHO, V. AND N. VOIGTLÄNDER (2014): “Input Diffusion and the Evolution of Production Networks,” Working paper.
- CARVALHO, V. M., M. NIREI, Y. U. SAITO, AND A. TAHBAZ-SALEHI (2021): “Supply Chain Disruptions: Evidence from the Great East Japan Earthquake,” *Quarterly Journal of Economics*, 136, 1255–1321.
- DAVID, J. M., L. SCHMID, AND D. ZEKE (2022): “Risk-Adjusted Capital Allocation and Misallocation,” *Journal of Financial Economics*, 145, 684–705.
- DEAN, J. (2015): “How Ford’s largest truck factory was completely overhauled in 8 weeks,” *Popular Mechanics*.
- DUPOR, B. (1999): “Aggregation and Irrelevance in Multi-Sector Models,” *Journal of Monetary Economics*, 43, 391–409.
- ELLIOTT, M., B. GOLUB, AND M. V. LEDUC (2022): “Supply Network Formation and Fragility,” *American Economic Review*, 112, 2701–47.

- ERSAHIN, N., M. GIANNETTI, AND R. HUANG (2022): “Supply Chain Risk: Changes in Supplier Composition and Vertical Integration,” Working paper.
- FAJGELBAUM, P. D., E. SCHAAL, AND M. TASCHEREAU-DUMOUCHEL (2017): “Uncertainty Traps,” *Quarterly Journal of Economics*, 132, 1641–1692.
- FERNALD, J. (2014): “A Quarterly, Utilization-Adjusted Series on Total Factor Productivity,” Working paper.
- FERNÁNDEZ-VILLAYERDE, J., P. GUERRÓN-QUINTANA, J. F. RUBIO-RAMÍREZ, AND M. URIBE (2011): “Risk Matters: The Real Effects of Volatility Shocks,” *American Economic Review*, 101, 2530–2561.
- FOERSTER, A. T., P.-D. G. SARTE, AND M. W. WATSON (2011): “Sectoral versus Aggregate Shocks: A Structural Factor Analysis of Industrial Production,” *Journal of Political Economy*, 119, 1–38.
- GABAIX, X. (2011): “The Granular Origins of Aggregate Fluctuations,” *Econometrica*, 79, 733–772.
- GROSSMAN, G. M., E. HELPMAN, AND A. SABAL (2023): “Resilience in Vertical Supply Chains,” Tech. rep., National Bureau of Economic Research.
- HO, W., T. ZHENG, H. YILDIZ, AND S. TALLURI (2015): “Supply Chain Risk Management: A Literature Review,” *International Journal of Production Research*, 53, 5031–5069.
- HORVATH, M. (2000): “Sectoral Shocks and Aggregate Fluctuations,” *Journal of Monetary Economics*, 45, 69–106.
- HULTEN, C. R. (1978): “Growth Accounting with Intermediate Inputs,” *Review of Economic Studies*, 45, 511–518.
- HUNEUS, F. (2018): “Production Network Dynamics and the Propagation of Shocks,” Working paper.
- JONES, C. I. (2011): “Intermediate Goods and Weak Links in the Theory of Economic Development,” *American Economic Journal: Macroeconomics*, 3, 1–28.
- JUDD, K. L. AND S.-M. GUU (2001): “Asymptotic methods for asset market equilibrium analysis,” *Economic theory*, 18, 127–157.
- JURADO, K., S. C. LUDVIGSON, AND S. NG (2015): “Measuring Uncertainty,” *American Economic Review*, 105, 1177–1216.
- KENNEDY, J. AND R. EBERHART (1995): “Particle Swarm Optimization,” in *Proceedings of ICNN’95-international conference on neural networks*, IEEE, vol. 4, 1942–1948.
- KOPYTOV, A. (2023): “Booms, Busts, and Common Risk Exposures,” *Journal of Finance*, 78, 3299–3341.
- KOPYTOV, A., K. NIMARK, B. MISHRA, AND M. TASCHEREAU-DUMOUCHEL (2024): “Supplement to “Endogenous Production Networks under Supply Chain Uncertainty.”,” Working paper.
- LIM, K. (2018): “Endogenous Production Networks and the Business Cycle,” Working paper.

- LIU, E. (2019): “Industrial Policies in Production Networks,” *Quarterly Journal of Economics*, 134, 1883–1948.
- LONG, J. B. AND C. I. PLOSSER (1983): “Real Business Cycles,” *Journal of Political Economy*, 91, 39–69.
- LUCAS, R. E. (1987): *Models of Business Cycles*, vol. 26, Basil Blackwell Oxford.
- LUTZ, H. (2021): “Retooling plants for evs will save GM Billions,” *Automotive News*.
- NIEUWERBURGH, S. V. AND L. VELDKAMP (2006): “Learning Asymmetries in Real Business Cycles,” *Journal of Monetary Economics*, 53, 753–772.
- OBERFIELD, E. (2018): “A Theory of Input-Output Architecture,” *Econometrica*, 86, 559–589.
- SCHMITT-GROHÉ, S. AND M. URIBE (2004): “Solving dynamic general equilibrium models using a second-order approximation to the policy function,” *Journal of economic dynamics and control*, 28, 755–775.
- TASCHEREAU-DUMOUCHEL, M. (2020): “Cascades and Fluctuations in an Economy with an Endogenous Production Network,” Working paper.
- VOM LEHN, C. AND T. WINBERRY (2022): “The investment network, sectoral comovement, and the changing US business cycle,” *The Quarterly Journal of Economics*, 137, 387–433.
- WAGNER, S. M. AND C. BODE (2008): “An Empirical Examination of Supply Chain Performance Along Several Dimensions of Risk,” *Journal of Business Logistics*, 29, 307–325.
- ZURICH INSURANCE GROUP (2015): “Potential Business Impact of ‘Loss of the Main Supplier’ for Small and Medium Enterprises (SMEs) in 2015,” Global survey report.