

TAIL RISK IN PRODUCTION NETWORKS

IAN DEW-BECKER

Department of Finance, Northwestern University and NBER

This paper describes the response of the economy to large shocks in a nonlinear production network. A sector's *tail centrality* measures how a large negative shock transmits to GDP, that is, the systemic risk of the sector. Tail centrality is theoretically and empirically very different from local centrality measures such as sales share—in a benchmark case, it is measured as a sector's average downstream closeness to final production. It also measures how large differences in sector productivity can generate cross-country income differences. The paper also uses the results to analyze the determinants of total tail risk in the economy. Increases in interconnectedness can simultaneously reduce the sensitivity of the economy to small shocks while increasing the sensitivity to large shocks. Tail risk is related to *conditional granularity*, where some sectors become highly influential following negative shocks.

KEYWORDS: Tail risk, production networks, large deviations.

1. INTRODUCTION

Background. RECENT EXPERIENCE HAS DEMONSTRATED that dislocations to supply chains can have significant effects on the economy both locally and internationally. Shocks to both the supply of goods, such as semiconductors and natural gas, and also the ability to transport them, for example, due to shutdowns at major ports and constraints on trucking, have propagated through the global supply chain. Over a longer period, research has found that large movements in GDP occur more frequently than predicted by the normal distribution (e.g., Acemoglu, Carvalho, Ozdaglar, and Tahbaz-Salehi (2012)), and a body of work since Gabaix (2011) has developed suggesting how large shocks to influential sectors or firms could cause such events.¹ Additionally, extreme events in the data tend to be negative, so that the distribution of GDP, in both levels and growth rates, is asymmetrical.²

An analysis of large shocks is interesting primarily in a nonlinear setting. In a purely linear model, one immediately knows how the economy responds to large shocks by simply observing its behavior when shocks are small. But when the economy is nonlinear the task of understanding the effects of large shocks becomes much harder—the sectors that are important in normal times need not be the ones that are important in extreme situations. There are some highly specialized cases where nonlinear models can be solved analytically, but in general they are approximated via Taylor series (which need not actually

Ian Dew-Becker: ian@i-db.org

This paper would not exist without Alireza Tahbaz-Salehi. I appreciate helpful comments from three anonymous referees, John Cochrane, Nicolas Crouzet, Joel Flynn, Xavier Gabaix, Stefano Giglio, Francois Gourio, Ernest Liu, Pooya Molavi, Rui Sousa, Fabrice Tourre, Aleh Tsyvinski, and seminar participants at Northwestern, the Triangle Macro-Finance Workshop, the NBER Summer Institute, the Macro Finance Society, Caltech, the Paris School of Economics, and the Federal Reserve Banks of Chicago, Kansas City, and Dallas.

¹Empirically, Barrot and Sauvagnat (2016) and Carvalho, Nirei, Saito, and Tahbaz-Salehi (2021) study the effects of large shocks to individual firms due to natural disasters on production. See also related work by Fujii, Ghose, and Khanna (2021). Liu and Tsyvinski (2021) study the dynamic effects of large shocks in a linear setting. See also Elliott and Golub (2022) for a general review of fragility in economic networks.

²For recent models, see Dew-Becker and Vedolin (2023), Dupraz, Steinsson, and Nakamura (2020), and Ilut, Kehrig, and Schneider (2018). Those papers also discuss empirical evidence.

converge), in which case even allowing for second-order terms can significantly reduce tractability.³

Contribution. This paper asks how the structure of the economy determines the extent to which different sectors create systemic risk. That is, when do large shocks to individual sectors transmit through supply chains to the rest of the economy? And if we know something about that transmission, what does it tell us about the determinants of tail risk in GDP? Relatedly, when can large cross-country differences in sectoral productivity explain differences in income?

The paper's core contribution is to answer those questions in the context of a general network production model. It does so using a theoretical result that gives a closed-form expression for the asymptotic response of GDP to any combination of shocks. That result is first used to understand why large shocks in some sectors propagate and affect the full economy while others may only have local effects. Second, when that result is combined with a probability distribution for the shocks, it is possible to describe the tails of the distribution of GDP. The insights gained from the analysis are significantly different from those from local approximations. The analysis clarifies what factors make a firm or sector systemically risky, and thus also what creates risk for the economy as a whole and which sectors can generate large cross-country income differences.

Methods. In production networks, economic units produce outputs using as inputs both labor and the products of other units. The various units interact, propagating and potentially amplifying or attenuating shocks. Importantly, this paper's model allows for arbitrary elasticities of substitution across inputs in each sector.

Consider a vector of productivity shocks, with a direction and a magnitude. The direction represents a scenario, some mixture of shocks, for example, a positive oil supply shock, or a simultaneous positive oil shock and negative shock to semiconductors. Holding the mixture fixed, the paper asks what happens when the size of the shocks is scaled up. The paper's theoretical tool is a result that shows that for large shocks, GDP and sector prices, and output all converge to linear asymptotes. The analysis can be thought of as giving a first-order *asymptotic*, as opposed to the local, description of the economy.⁴

When combined with an assumption about the distribution of the shocks, the asymptotes also determine the probability of large movements in GDP.

Results. The paper's first application of the limiting approximation is to study what determines whether a large negative shock to a given sector has only local effects or propagates through the economy to GDP. First, consistent with Baqaee and Farhi (2019), it shows that complementarity is key to propagation. A novel finding, though, is that the asymptotic effect does not depend on the precise value of the elasticity of substitution. In the tail, negative shocks propagate through nodes where the elasticity is below 1 and are stopped by nodes where the elasticity is above 1—the distance of the elasticity above or below 1 does not appear. That does not mean the precise elasticity does not actually matter, but rather illustrates that for understanding first-order effects in the tail the sign relative to 1 is all we need to know.

³Jones (2011) and Dew-Becker and Vedolin (2023) study closed-form solutions to nonlinear models (roundabout economies). For a second-order approximation, see, notably, Baqaee and Farhi (2019), for an insightful analysis that simultaneously illustrates the complexity of analyzing a quadratic approximation. See also den Haan and de Wind (2009) for a discussion of the convergence of Taylor series in economic models.

⁴And there are actually no higher-order terms in the Taylor series at infinity.

Similarly, the analysis shows that it is the *topology* of the production network, rather than its geometry—the existence of intersectoral linkages, rather than their intensity—that determines propagation. The importance of a sector depends on how much of GDP is downstream of it. Unlike in a local approximation, the intensity of the use of its output by downstream sectors is (again, to the first order) irrelevant. Another way to put it: the size of a sector in good times does not determine its importance in extreme situations. A sector can be simultaneously small and also systemically important—utilities being the canonical example.

Putting the results on complementarity and downstream propagation together, we can describe how interconnectedness affects tail risk. When a new link is added to the production network whereby a sector has a new input that substitutes for others, that makes the network more robust, while when a new input is added that is a complement, the network becomes more fragile. That fragility can arise even when the new input simultaneously reduces sensitivity to small shocks. That is, the economy can simultaneously become more diversified locally and also face an increased risk of crashes.⁵ As a recent practical example, consider the case of semiconductors. The rise of computer technology has been massively beneficial to the economy, but at the same time it has made essentially every sector sensitive to the supply of semiconductors, making that sector surprisingly influential following a recent negative shock.

Using input–output data for the US, the paper gives a first-pass empirical estimate of tail centrality—the effect on GDP of a large shock to each sector. The basic finding is that tail centrality and sales shares—which measure local centrality—are only about 60% correlated, with numerous sectors with small sales shares having large tail centralities, while many sectors with large sales shares have small tail centralities. The sectors with the highest tail centrality include electricity, trucking, oil, and legal services, with the last being a particularly interesting gut-check, so to speak, to help see the full extent of the model's predictions.

The paper also shows that these tail centralities measure the ability of large productivity differences to explain cross-country income differences. It is precisely the upstream sectors that produce inputs for the entire economy that are most likely to act as bottlenecks.⁶

Finally, but no less importantly, the paper uses the asymptotic expressions for the response of GDP to show how the structure of the economy interacts with the distribution of the shocks to determine the distribution of extreme realizations of GDP.

That analysis first provides comparative statics showing what factors create and exacerbate asymmetry in the distribution of GDP growth: increases in complementarity and in connections running through complementary sectors both create left-tail risk. As an example, for the case of i.i.d. exponential shocks (as in Acemoglu, Ozdaglar, and Tahbaz-Salehi (2017)), tail risk is determined the largest Domar weight (sales share) that any sector can attain for any combination of shocks, rather than just the largest steady-state Domar weight.

The novel idea consistently underlying this paper's results is that what really matters for tail risk is the relative size of the sectors *in extreme scenarios*. Tails are driven not by granularity at steady state, but rather by *conditional* granularity. In the exponential example, what determines tail risk is not whether there is granularity on average, but whether there can ever be granularity—whether a single sector can become pivotal if shocks are large enough.

⁵See Acemoglu and Azar (2020) for related work on changes in interconnectedness in production networks.

⁶For related work on cross-country income differences and weak links in production networks, see Ciccone (2002), Jones (2011), Liu (2019), and Acemoglu and Azar (2020).

For example, take electricity and restaurants. In normal times, those sectors are of similar size, which in a linear approximation would imply that they have similar effects on GDP. But one lesson of Covid was that shutting down restaurants is not catastrophic for GDP,⁷ whereas one might expect that a significant reduction in available electricity would have strongly negative effects—and that those effects would be convex in the size of the decline in available power. Electricity is systemically important not because it is important in good times, but because it *would be* important in bad times. And the paper's analysis shows how to quantify precisely how important.

As noted above, the paper's analysis is based on limits for large shocks. While the results are always useful for understanding the qualitative determinants of tail risk, their quantitative accuracy depends on how large the shocks (or cross-country productivity differences) actually are. Section 6 analyzes how large the shocks need to be for the limits to be quantitatively accurate and compares the magnitude to shocks observed empirically.

Additional Related Literature. The paper's framework builds most directly on the literature on production networks, going back to Long and Plosser (1983).⁸ The closest link is to Baqaee and Farhi (2019), who study higher moments of output in the same nonlinear framework, but studying an explicitly local approximation, which necessarily does not speak specifically to large deviations as it has infinitely large errors in the tails. There are also a number of recent papers on the propagation of shocks and distortions in production networks, both empirical and theoretical.⁹ A contribution of this paper is to potentially give a way for work in those areas to get analytic approximations where they were previously unavailable.

A focus of the analysis is how the network effectively changes as shocks change. Taschereau-Dumouchel (2021) formally studies an endogenous production network and its effects on the distribution of GDP. There is also a related literature in international trade on endogenous value chains (e.g., Alfaro, Chor, Antras, and Conconi (2019)).

The paper's analysis applies to supply shocks to different sectors. There is also work on demand shocks, for which propagation runs upstream through the network, rather than downstream (see the discussion in Carvalho and Tahbaz-Salehi (2019)).

Some of this paper's specific results are related to past work on networks and extreme value theory, and that work is discussed when those results are discussed (e.g., Section 5.1.2).

Outline. The remainder of the paper is organized as follows. Section 2 describes the basic structure of the economy and the main result on approximating output in terms of the exogenous shocks is presented in Section 3. Sections 4 and 5 analyze the drivers of the tail centrality of individual sectors along with examples and extensions, while Section 6 presents results on the determinants of the quantitative accuracy of the approximation.

⁷Consumer spending on food services and accommodations fell by 40%, or \$403 billion between 2019Q4 and 2020Q2. Spending at movie theaters fell by 99%.

⁸That literature is large and work has studied features of networks, for example, what makes a particular sector or firm central and what determines the behavior of GDP. For recent representative work, in addition to other work discussed, see Liu and Tsyvinski (2021), vom Lehn and Winberry (2021), La'O and Tahbaz-Salehi (2021), and Bigio and La'o (2020).

⁹Liu (2019), Bigio and La'o (2020), and Boehm and Oberfield (2020) study the propagation of distortions in production networks. Costello (2020) and Alfaro, García-Santana, and Moral-Benito (2021) study the propagation of credit supply shocks. Gofman, Segal, and Wu (2020) study the propagation of technology shocks and their effects on firm risk.

Section 7 examines tail centrality in the data. Finally, Section 8 presents results on the probability of extreme realizations of GDP and Section 9 concludes.

2. STRUCTURE OF THE ECONOMY

The model is static and frictionless and takes the form of a nested CES production network as studied in Baqaee and Farhi (2019). There are N production units each producing a distinct good. A unit might represent a sector, a firm, or even just part of a sector or firm, though the paper will refer to them as “sectors” as a standard shorthand. Each unit has a CES production function of the form

$$Y_i = \zeta Z_i L_i^{1-\alpha} \left(\sum_j A_{i,j}^{1/\sigma_i} X_{i,j}^{(\sigma_i-1)/\sigma_i} \right)^{\alpha\sigma_i/(\sigma_i-1)}, \tag{1}$$

where ζ is a normalizing constant, Y_i is unit i 's output, Z_i its productivity, L_i its use of labor, and $X_{i,j}$ its use of good j as an input (throughout the paper, summations without ranges are taken over $1, \dots, N$).¹⁰ The parameters $A_{i,j}$, normalized such that $\sum_j A_{i,j} = 1$, determine the relative importance of different inputs. If $A_{i,j} = 0$, unit i does not use good j .

$0 < 1 - \alpha \leq 1$ represents the labor’s share of income. It is easy to relax the model to allow that to vary across sectors (as it does empirically). In that case, α is replaced everywhere with α_i .

σ_i is the elasticity of substitution across material inputs for unit i . When $\sigma_i \rightarrow 1$, the production function becomes Cobb–Douglas (with the $A_{i,j}$ becoming the exponents). Though the main analysis assumes a CES specification for simplicity, Supplemental Appendix D.2 (Dew-Becker (2023)) shows that the results also hold under much more general conditions.

As discussed in Baqaee and Farhi (2019), this structure captures arbitrary substitution patterns through nesting of the production functions. For example, if a real-world industry has some inputs that are substitutes and some that are complements, that would be modeled here as two production functions whose outputs are then combined to produce the real-world industry’s output. Section 5.2 gives another example in which substitutability can be modeled as a property of a good instead of a production function, and Supplemental Appendix D.2.1 discusses a more general setup from Chodorow-Reich, Gabaix, and Koijen (2022).¹¹

Last, there is representative consumer whose utility over consumption of the different goods is

$$U(C_1, \dots, C_N) = \prod_i C_i^{\beta_i}, \tag{2}$$

¹⁰ $\zeta = (1 - \alpha)^{-(1-\alpha)} \alpha^{-\alpha}$. The fact that labor in (1) has a unit elasticity of substitution with material inputs is without loss of generality—one can always specify an additional unit that converts labor into labor services, which are then combined with other inputs with a nonunitary elasticity (this requires allowing α to vary across sectors, and the labor services sector has $\alpha_i = 0$).

¹¹An example of a model in which the paper’s results do not hold is one where some input cannot be reallocated across sectors and it has an elasticity of substitution with material inputs smaller than 1 (such a model does not in general have a solution for all levels of productivity).

where $\sum_j \beta_j = 1$ and we define a vector $\beta = [\beta_1, \dots, \beta_N]$. The unitary elasticity of substitution in consumption focuses the analysis on nonlinearity in production, rather than final demand, but it is without loss of generality.¹²

The representative agent purchases C_i units of good i with wages and inelastically supplies a single unit of labor so that $\sum_i L_i = 1$.

Throughout the paper, lower-case letters denote logs, for example, $z_i = \log Z_i$. I also normalize productivity such that $z_i = 0$ represents, informally, the steady-state or average value.

For the main results, I assume labor can be frictionlessly reallocated across sectors. The limits go through identically with fixed labor (Supplemental Appendix C.1), and allowing for an upward sloping aggregate labor supply curve is also straightforward.

Since the economy is frictionless, it can be solved either competitively or from the perspective of a social planner.

DEFINITION: A *competitive equilibrium* is a set of prices, $\{P_i\} \cup W$, and quantities, $\{Y_i\}$, $\{X_{i,j}\}$, $\{C_{i,j}\}$, and $\{L_i\}$, such that each unit i maximizes its profits, $P_i Y_i - W L_i - \sum_j P_j X_{i,j}$, the representative consumer maximizes utility, producers and the consumer take prices as given, and markets clear: $Y_i = C_i + \sum_j X_{j,i}$.

Because there is no government spending or investment, GDP is equal to aggregate consumption. I denote $\log GDP$ by gdp .

The model does not in general have a closed-form solution.

2.1. Cost Minimization

Normalizing the wage to 1, marginal cost pricing along with cost minimization implies that good i 's log price satisfies

$$p_i = -z_i + \frac{\alpha}{1 - \sigma_i} \log \left(\sum_{j=1}^N A_{ij} \exp((1 - \sigma_i) p_j) \right). \tag{3}$$

We have the usual result that shocks propagate downstream: each sector's price depends on its own productivity and the prices of its inputs. In the special case where $\sigma_i = 1$, the recursion is linear and solvable by hand. Stacking the prices and productivities into vectors, $p = -(I - \alpha A)^{-1} z$, where A is a matrix collecting the $A_{i,j}$ coefficients.

Equation (3) implies prices do not depend on demand, a “no-substitution” type result.¹³ With the wage normalized to 1, nominal income and GDP are constant, meaning that real GDP is just the inverse of the price of the consumption good, so that the equilibrium is fully characterized by the solution to (3).

¹²One can always add a sector with a nonunitary elasticity of substitution that produces a single final good, with $\beta = 1$ for that sector and equal to zero for all other sectors. Technically, this violates the restriction of $\alpha_i < 1$, but the results still go through. $\alpha_i < 1$ is a sufficient but not necessary condition—we just need to have that the equilibrium conditions (equation (3) below) are a contraction.

¹³Georgescu-Roegen (1966) and Samuelson (1951). More recently, see Acemoglu and Azar (2020), Flynn, Patterson, and Sturm (2022), and Baqaee and Farhi (2020).

Given a solution for the vector p (as a function of z), utility maximization for the consumer (and the normalization of nominal GDP to 1) implies that real GDP is

$$gdp = -\beta' p \tag{4}$$

showing how the recursion for prices combined with preferences determines gdp .

In the linear case, the analysis is straightforward. For $\sigma_i \neq 1$, the price recursion is non-linear and has no general closed-form solution. If one just wants a quantitative model, it is easy to get a numerical solution even for large N . But for the purposes of characterizing the behavior of the economy theoretically and understanding the forces determining the importance of different sectors and shocks, being able to analyze the model by hand is useful. Even a second-order approximation, though, can become difficult to work with, not only due to the number of terms (quadratic in N), but also due to the fact that the precise values of all the parameters of the model appear.

3. LARGE SHOCK BEHAVIOR

Any vector of log productivities has a polar representation,

$$z = \theta t, \tag{5}$$

where $\theta \in \mathbb{R}^N$, such that $\theta' \theta = 1$ is a unit vector representing a direction in productivity space and $t \geq 0$ is a scalar determining magnitude. As examples, $\theta = [\dots, 0, 1, 0, \dots]$ represents a positive shock to a single sector, while $\theta = [1, 1, \dots] / \sqrt{N}$ represents a common shock to all sectors. Since t is nonnegative, a negative shock to a single sector, rather than being represented by negative t , is represented by $\theta = [\dots, 0, -1, 0, \dots]$. That distinction will matter.

3.1. The Large Shock Limit

LEMMA 1: For each i , there exist unique, continuous, scalar-valued functions $\mu_i(\theta)$ and $\phi_i(\theta)$ such that

$$\lim_{t \rightarrow \infty} p_i(\theta t) - (\mu_i(\theta) + \phi_i(\theta)t) = 0, \tag{6}$$

where

$$\phi_i(\theta) = -\theta_i + \alpha_i \begin{cases} \max_{j \in S_i} \phi_j(\theta) & \text{if } \sigma_i < 1, \\ \sum_j A_{i,j} \phi_j(\theta) & \text{if } \sigma_i = 1, \\ \min_{j \in S_i} \phi_j(\theta) & \text{if } \sigma_i > 1 \end{cases} \tag{7}$$

and $S_i \equiv \{j : A_{i,j} > 0\}$ is the set of inputs used by sector i .

While the recursion for prices (3) is not solvable in closed form, it has a remarkably simple limit as the shocks grow in magnitude. For $\sigma_i < 1$, it involves a maximum upstream, while for $\sigma_i > 1$ a minimum. The result immediately shows how complementarity and substitutability affect shock propagation: negative productivity shocks propagate downstream through parts of the production process that are complementary ($\sigma_i < 1$), while positive productivity shocks propagate through parts that are substitutable ($\sigma_i > 1$).

Since the recursion involves a max/min, it can be interpreted as saying that as $t \rightarrow \infty$, every sector's behavior ends up driven by a single one of its inputs (ignoring the knife-edge




FIGURE 1.—Network examples. *Notes:* The nodes represent sectors and arrows flows of goods. The black and gray nodes and black arrows represent a hypothetical tail network following a shock to the solid black sector (with the shading becoming lighter with distance). All sectors use their own output as an input. For panels (a)–(c), all elasticities are assumed to be less than 1. For panel (d), the two center nodes have elasticities as noted, and the others again have $\sigma < 1$. White nodes are asymptotically unaffected by the shock.

case of $\sigma_i = 1$). In other words, for a given combination of shocks θ , as $t \rightarrow \infty$, there is a *tail network*, which depends on θ , and in which each sector has just a single upstream link.

To illustrate that, Figure 1 displays four hypothetical networks assuming $\sigma_i < 1 \forall i$. Each example illustrates the transmission of a shock to the black node. Arrows represent flows of goods—there is an arrow from i to j if $A_{j,i} > 0$. The solid black arrows show how a large negative shock to the black sector propagates—they represent the tail network. Again, shocks propagate downstream. In the top-left panel, for example, the top node is shocked, and the bottom two are directly affected. The other panels plot richer networks (panels (c) and (d) are discussed further below). The shading of the other nodes shows how strongly they are affected by the shock, with the lightest grays being furthest downstream and, therefore, least affected, with white unaffected.

Since the elasticities are negative in Figure 1, it is negative shocks that propagate. If the black nodes received large *positive* shocks, they would eventually have no marginal impact on the other sectors.

The source of the result in (7) is that in the limit as $t \rightarrow \infty$, each sector’s expenditure shares on inputs that are ultimately driven to 0 or 1, depending on the elasticity and the shock. An elasticity of substitution less than 1 means that when an input’s price rises, its share of expenditures rises (all else equal), an elasticity above 1 means that the share falls, and $\sigma_i = 1$ is the knife-edge case with constant expenditure shares.¹⁴

¹⁴Mathematically, the result comes from the log-sum-exponential that appears in the recursion. Using $p_i \sim \phi_{i,t}$ asymptotically,

$$\phi_i \sim -\theta_i + \frac{\alpha}{1 - \sigma_i} \frac{1}{t} \log \left(\sum_{j=1}^N A_{ij} \exp(\phi_j)^{(1 - \sigma_i)t} \right). \tag{8}$$

There is also a simple recursion for $\mu(\theta)$, which depends on \mathcal{A} and σ , but for this paper’s analysis it will be unimportant (see Appendix A.1). Similarly, sector output follows $y_i \rightarrow \mu_{y,i} - \phi_i t$ for a constant $\mu_{y,i}$ (see Supplemental Appendix C.1), but the remainder of the paper focuses on aggregate output.

Note that the asymptotics here are entirely in terms of the size of the shocks, via the term t . The structure of the economy, including the number of sectors and their relationship, is held fixed. In addition, there is nothing stochastic about the limit—it is describing the economy for given levels of productivity, not saying anything about probability distributions.

3.2. The Behavior of GDP

Using the fact that $gdp = -\beta' p$, we immediately have the paper’s main theoretical tool for calculating the effects of large shocks.

THEOREM 1: *Under the conditions of Lemma 1,*

$$\lim_{t \rightarrow \infty} gdp(\theta t) - (-\beta' \mu(\theta) + \lambda(\theta)t) = 0, \tag{9}$$

$$\text{where } \lambda(\theta) \equiv -\beta' \phi(\theta) \tag{10}$$

and $\mu(\theta)$ and $\phi(\theta)$ are stacked (vector-valued) versions of μ_i and ϕ_i .

Theorem 1 simply says that, in all directions in productivity space, gdp eventually converges to a linear asymptote with slope $\lambda(\theta) \equiv -\beta' \phi(\theta)$.

The panels of Figure 2 plot various approximations for log GDP for some arbitrary value of θ , with t varying along the x-axis. The negative side of the axis formally corre-




FIGURE 2.—Linear, quadratic, and tail approximations. *Notes:* The x-axis is log productivity and the y-axis log aggregate output. The x-axis may represent productivity in a single sector, or it could be the scale of a shock that affects productivity in multiple sectors. The concavity in GDP in this example is consistent with an economy featuring complementarities.

As $t \rightarrow \infty$, the exponent $(1 - \sigma_i)t$ goes to $\pm\infty$, and the log-sum-exp converges to a max or min, except for the case $\sigma_i = 1$, where t drops out.

sponds to reversing the sign of θ , that is, t runs from 0 to ∞ on each side and θ is replaced with $-\theta$ on the left.

When $\sigma_i = 1$ for all i , the model is fully linear with $\lambda(\theta) = \beta'(I - \alpha A)^{-1}\theta$; otherwise it is nonlinear. The nonlinearity can be locally captured by a Taylor series, as is shown in the left-hand panel. The right-hand panel plots the approximation implied by Theorem 1. As t grows both to the left and right, log GDP approaches the two straight lines, with $\lambda(\theta) \neq -\lambda(-\theta)$. That difference is how the tail approximation captures nonlinearity.

Figure 2 intentionally does not include a scale on the x-axis. The general shape of the lines is a robust feature of the model, but the scale at which the nonlinearity appears is parameter dependent. Section 6 examines the determinants of the rate of convergence to the asymptotes in detail. That said, the main reason to use the tail approximation instead of a higher-order Taylor series or numerical solution is its tractability, parameter invariance, and the fact that it is formally describing first-order asymptotic behavior.

At an abstract level, a model is a mapping from shocks to outcomes. Theorem 1 is describing one characteristic of that mapping, which is that it has linear asymptotes in all directions.

3.2.1. Invariance

A significant feature of the results so far is that the asymptotic behavior of the economy is invariant to the specific values of the production parameters. The values of the ϕ_i 's, and hence the limits for prices, do not depend on the exact values of any σ_i or $A_{i,j}$. All that matters is whether the elasticities are above or below 1 and whether the production weights are greater than zero. In the example in Figure 2, changing the exact values of the production parameters (away from $\sigma_i = 1$ or $A_{i,j} = 0$) changes $\mu(\theta)$, and hence the levels of the asymptotes, and it can change the curvature of GDP with respect to productivity, but the slopes of the asymptotes are unaffected.

A production system can be thought of as a weighted directed network, where the edges represent use of a good by a sector, and their weights correspond to the importance of the good in production, measured by the $A_{i,j}$. But here the exact values of the $A_{i,j}$ play no role. In that sense, the result says that what matters in the tail is the topology of the network—the set of edges or links between sectors—rather than the geometry—their weights or usage intensity.

So, when thinking about the supply-chain risks associated with large shocks, what is important is not how large a given supplier is on average, but rather how many sectors it supplies (the link to out-degree is formalized below). Unlike the usual analysis for small shocks or a Cobb–Douglas economy, this result implies that for large shocks, the economy is analyzed as an *unweighted* network. The second-order Taylor series in Figure 2, on the other hand, depends on the precise value of every parameter of the model.

4. SECTOR TAIL CENTRALITY

This section studies how large shocks to individual sectors affect the GDP.

DEFINITION: The *left-tail centrality* and **tail elasticity** of unit i are, respectively,

$$\gamma_i^L \equiv \lim_{\Delta z_i \rightarrow -\infty} \frac{\Delta \text{gdp}}{\Delta z_i}, \tag{11}$$

$$\delta_i^L \equiv \lim_{\Delta z_i \rightarrow -\infty} \frac{\Delta \text{gdp}}{\Delta y_i}, \tag{12}$$

where Δ denotes a deviation from steady state ($z_i = 0 \forall i$). Right centralities and elasticities are the same but for $\Delta z_i \rightarrow +\infty$.

The usual local approximation takes $\Delta z_i \rightarrow 0$; here, we study $\Delta z_i \rightarrow \pm\infty$.

Which is preferred between the tail centrality and elasticity and centrality depends on context. γ_i^L is more fundamental theoretically—mapping between exogenous shocks and gdp —and it is more closely related to the Domar weights studied in past work, so the theoretical results focus on it somewhat more. $\Delta gdp/\Delta y_i$, on the other hand, involves objects more easily observable in the data and will turn out to be somewhat better behaved empirically. Both will be discussed in what follows.

COROLLARY 1: *Let e_i denote a vector equal to 1 in element i and zero otherwise. Then in the notation of Theorem 1,*

$$\gamma_i^L = -\lambda(-e_i), \tag{13}$$

$$\gamma_i^R = \lambda(e_i). \tag{14}$$

4.1. Comparative Statics

Because of the simplicity of Theorem 1, it is straightforward to characterize how the parameters of the model affect tail centralities.

Sector i has a direct downstream link to sector j if $A_{j,i} > 0$, and sector j is *downstream* of sector i if there is a path via direct downstream links from i to j . Note that it is possible for i and j to both be downstream of each other—the economy does not need to have a strict hierarchy.

Complementarity magnifies negative shocks and attenuates positive shocks.

PROPOSITION 1: *γ_i^L weakly increases and γ_i^R weakly decreases when σ_j transitions from above to below 1 for any j downstream of i .*

Intuitively, substitutability means it is possible to primarily use the outputs of relatively productive sectors in production, while complementarity *requires* using all inputs, including the weakest. Since productivity shocks propagate downstream, it is downstream elasticities that matter.

Second, for similar reasons, interconnectedness in the network increases tail risk under complementarity and reduces it under substitutability.

PROPOSITION 2: *When the set of inputs used by sector i , S_i , grows, in the sense that $S_i \rightarrow S_i \cup j$ for some $j \notin S_i$, γ_k^L weakly increases and γ_k^R weakly decreases for all k if $\sigma_i > 1$ and decreases if $\sigma_i < 1$.*

One way to state that result makes it seem obvious: if the number of inputs needed to produce output grows, then the supply chain is more fragile. On the other hand, if there are more options for production, it becomes less fragile. Just like in the previous result, $\sigma_i < 1$ is a situation whereas the sector effectively needs all of its inputs, while $\sigma_i > 1$ is a situation where it can use just a single input.

There is a less obvious implication of this result, though: if a sector discovers an input that strongly increases the marginal product of all of its other inputs, then production is more delicate, with all left-tail centralities (weakly) rising. Obviously, such a discovery will increase output, but it also will make output in the future sensitive to more shocks, since

now shocks to the new input will matter, where they did not previously. Take electricity, for example, obviously we are better off for having it, but at the same time the economy is now sensitive to the risk of electricity being cut off.

Panels (b) and (c) of Figure 1 give an example of the effect of adding a link to the network. When the top-left sector is shocked, adding a single link (the thick arrow) causes the shock to now propagate to the entire network.

5. SPECIAL CASES, EXAMPLES, AND EXTENSIONS

5.1. Fully Complementary Production: Average Closeness

A number of papers, including Jones (2011), Baqaee and Farhi (2019), and Rubbo (2020), study economies characterized by complementarity, with $\sigma_i < 1$ for all i .¹⁵

DEFINITION: A *complementary economy* is one in which $\sigma_i < 1$ for all i .

PROPOSITION 3: In the complementary economy,

$$\delta_i^L = \sum_{j=1}^n \beta_j \alpha^{d_{\min}(j,i)}, \tag{15}$$

where $d_{\min}(j, i)$ is the length of the shortest downstream path from i to j .¹⁶ If, additionally, $A_{i,i} \in (0, 1)$, then

$$\gamma_i^L = \frac{1}{1 - \alpha} \sum_{j=1}^n \beta_j \alpha^{d_{\min}(j,i)}. \tag{16}$$

In the complementary economy, the asymptotic effect of a shock to sector i on *gdp* is proportional to the average downstream closeness, measured by $\alpha^{d_{\min}(j,i)}$, of i to *gdp*. Since $d_{\min}(j, i)$ is the shortest path from i to j and $\alpha < 1$, $\alpha^{d_{\min}(j,i)}$ measures how *close* i and j are.¹⁷ If more of *gdp* is downstream of, and close to, i , then its shocks are asymptotically more influential. The $1/(1 - \alpha)$ term in γ_i^L is the asymptotic effect of z_i on y_i .

Proposition 3 answers the question of what types of units create systemic risk under complementarity: those that are direct suppliers to producers of a large fraction of GDP (and that do not have substitutes). That also implies that tail centralities increase when the economy is more connected.

More generally, all of the following will increase δ_i^L and γ_i^L :

1. An increase in the number of units downstream of i or an increase in their share of GDP.
2. A decrease in the number of steps between unit i and the units downstream of it.
3. An increase in the share of expenditures on material inputs, α .

¹⁵See also evidence in Atalay (2017) and Atalay, Drautzburg, and Wang (2018), among others.

¹⁶That is, if $i \neq j$ and $A_{i,j} > 0$, $d_{\min}(j, i) = 1$. If $A_{i,j} = 0$, but there exists a k such that $A_{i,k} > 0$ and $A_{k,j} > 0$, then $d_{\min}(i, j) = 2$, etc.

The assumption that $A_{i,i} \in (0, 1)$ ensures that the shocked sector is directly downstream of itself, which determines its ϕ_i .

¹⁷It is somewhat intriguing to note that the matrix formed by $\alpha^{d_{\min}(j,i)}$ can be obtained as a power series, like the Leontief inverse one gets in the Cobb–Douglas case, but under an alternative algebra. See Butkovic (2010) and Joswig and Schröter (2021).

On the other side, $\delta_i^R = \gamma_i^R = \beta_i \forall i$. That is, positive shocks do not propagate, so their only asymptotic effect is from their direct impact. When $\sigma_i \geq 1 \forall i$, the results for γ_i^L and γ_i^R are switched—right-tail centrality is equal to average downstream closeness to GDP and left-tail centrality is simply β_i (and the same holds for the tail elasticities).

5.1.1. *The Tail Network and Conditional Granularity*

In a complementary economy, it is possible to give a fuller description of the tail network that was discussed in Section 3.1. For any given vector of productivities z , there is a vector of Domar weights, D , with $dgdp/dz = D$ (which, by Hulten’s (1978) theorem, are nominal sales shares). D measures the importance of each sector in a given state. In steady-state ($z = 0$),

$$D'_{ss} \equiv \beta'(I - \alpha A)^{-1}. \tag{17}$$

PROPOSITION 4: *Conditional on the parameters of the model, as $t \rightarrow \infty$, the Domar weights converge to a finite number of possible limits (across the values of θ), denoted by the set $\{D_k\}$. In a complementary economy,*

$$\lambda(\theta) = \min_k D'_k \theta. \tag{18}$$

It immediately follows that GDP is concave in that $\lambda(\theta) \leq -\lambda(-\theta)$.

In a linear model, where the production network is fixed, the Domar weights are constant so that there is a single slope determining the response to any θ , $\lambda(\theta) = D'_{ss} \theta$. In a nonlinear model, the Domar weights vary depending on productivity, but the proposition says that in the limit they only take on a finite set of values. That is, there are sets, say Θ_k , such that for all $\theta \in \Theta_k$, the Domar weights always converge to the same D_k as $t \rightarrow \infty$. That fact follows from the recursion defining ϕ —for $\sigma_i \neq 1$, every sector’s price just depends on that of a single upstream input in the tail, and there are only a finite set of possible upstream sectors.¹⁸ Each vector D'_k is of the form $\beta'(I - \alpha M_k)^{-1}$, where M_k is a matrix representing a particular tail network.¹⁹

In the language of graph theory, the tail network is a minimal spanning tree over the sectors downstream of i , rooted at i , where a spanning tree connects all downstream nodes back to i and it is minimal in that it uses the fewest possible links.²⁰ Analysis of persuasion, the influence of node i depends on the number of spanning trees rooted at i . Here, on the other hand, all that matters is the *minimal* tree, again due to the choice of shortest paths (i.e., the tail network).

Spanning trees appear elsewhere in economics including in the analysis of diversity (Weitzman (1992) and Nehring and Puppe (2002)), price indexes (Hill (1999), Hill (2004), and Diewert (2010)), game theory (Granot and Huberman (1981), and auctions (Sun and Yang (2014))).

Finally, Proposition 4 immediately yields an alternative description of tail centrality.

¹⁸The minimization here is reminiscent of the worst-case network analysis in Jiang, Rigobon, and Rigobon (2021).

¹⁹Proposition 4 also gives a way to visualize the tail approximation more richly in the complementary case: it is the minimum of a set of hyperplanes (which is a convex polytope), with each hyperplane representing one particular tail network, defined by D_k .

²⁰In DeMarzo, Vayanos, and Zwiebel’s (2003)

COROLLARY 2: *In a complementary economy,*

$$\gamma_i^L = \max_k D_{k,i}. \tag{19}$$

That is, a sector’s left-tail centrality is measured by the *largest value* that its Domar weight can take on for *any* feasible tail network structure. This is the paper’s first view of the importance of **conditional granularity**. A sector need not be granular in steady state to be able to significantly damage the economy. What matters is whether it can *ever* be granular.

5.1.2. *Relationship With Other Centrality Measures*

The idea of measuring centrality via average closeness, as in Proposition 3, appears elsewhere in the networks literature in the form of harmonic centrality, which is an unweighted average closeness.²¹ The concept of the *efficiency* of a network is then measured by the average closeness between all pairs of nodes (Marchiori and Latora (2000) and Crucitti, Latora, Marchiori, and Rapisarda (2003)). In the context of complementary production, a network with greater efficiency then also has more tail risk (this is formalized further in Section 8).

The difference between average closeness and the usual Bonacich (1987) centrality that appears in a Cobb–Douglas economy (with $\sigma_i = 1 \forall i$) is that the latter measures centrality by looking across every possible path through the network, while average closeness is measured based only on shortest paths (see Carvalho and Tahbaz-Salehi (2019)).²²

Intuitively, the result on closeness suggests that out-degree of a unit—the number of units directly downstream of it—would be closely linked to tail centrality. Define weighted out-degree to be

$$\text{deg}_i \equiv \sum_{j:i \in S(j)} \beta_j \tag{20}$$

deg_i measures what fraction of final consumption demand is accounted for by the sectors directly downstream of i .

PROPOSITION 5: *Under fully complementary production, left-tail centrality satisfies*

$$\frac{1}{1 - \alpha} (\beta_i + \alpha \text{deg}_i) \leq \gamma_i^L \leq \frac{1}{1 - \alpha} (\beta_i + \alpha \text{deg}_i + \alpha^2 (1 - \text{deg}_i)). \tag{21}$$

Weighted out-degree thus gives upper and lower bounds for tail centrality.²³

²¹See Boldi and Vigna (2014) who justify it axiomatically, along with Rochat (2009) and Bloch, Jackson, and Tebaldi (2021).

²²These results also suggest that there might be a relationship with the concept of upstreamness studied in Antràs and Chor (2013) and Antràs, Chor, Fally, and Hillberry (2012). However, the normalization here is different. For those papers, a sector is fully downstream if it sells only to final users. Here, though, what determines a sector’s centrality is not just the *composition* of its sales, but also the fraction of final users that it sells to.

²³Out-degree appears frequently in the networks literature, including recently Carvalho et al. (2021), Her-skovic, Kelly, Lustig, and Van Nieuwerburgh (2020), Bernard, Moxnes, and Saito (2019), and Mossel, Sly, and Tamuz (2015) among many others.

5.1.3. *Example: Fully Connected Economy*

EXAMPLE 1: Consider a complementary economy in which every sector uses inputs from itself and every other sector (i.e., $A_{i,j} > 0 \forall i, j$). Then

$$\phi_i(\theta) = -\theta_i - \frac{\alpha}{1-\alpha} \theta_{\min}, \tag{22}$$

$$\lambda(\theta) = \beta' \theta + \frac{\alpha}{1-\alpha} \theta_{\min}, \tag{23}$$

where $\theta_{\min} \equiv \min_i \theta_i$. The tail centrality of any sector i is $\gamma_i^L = \beta_i + \alpha/(1 - \alpha)$ and the tail elasticity is $\delta_i^L = (1 - \alpha)\gamma_i^L$.

In the case of a fully connected production network, each sector’s sensitivity, $\phi_i(\theta)$, is a linear combination of its own productivity and that of the weakest sector, and GDP then depends on both a linear combination of the θ ’s and also their minimum. So even if, for example, the economy is fully symmetric, with each good used in equal amounts so that all sectors have identical Domar weights in steady state, the effect of a shock on GDP in the tail depends additionally on the productivity of the weakest sector

Note again the invariance: the results in this example do not depend on the exact value of most of the production parameters. A sector can be large or small on average, but if, given θ , it has the minimal value of θ_i , it will have Domar weight $\beta_i + \alpha/(1 - \alpha)$ when the scale of the shocks, t , is large.

This again illustrates the idea of conditional granularity. Even if no sector is granular when shocks are small, as the shocks become large, the sector with the most negative shock becomes granular in the sense that it becomes a uniquely important determinant of GDP. It is thus possible for the economy to diversify, with the vector β having smaller average values, while tail risk stays large, simply because in this economy a large negative shock to any single sector has the power to significantly impact GDP. Tail centrality is thus independent of diversification, the number of units, and steady-state Domar weights.

Panel (a) of Figure 1 represents the tail network for a version of this economy with three sectors.

5.2. *Extension: Allowing for Goods to Be Substitutes*

In the description of the economy in equation (1), substitutability is a characteristic of a sector. But it is also possible to treat substitutability as a characteristic of a good. For example, for some goods i' and i'' to be substitutes, they can be combined into are combined into a bundle i via the function

$$Y_i = \left(X_{i,i'}^{(\sigma_i-1)/\sigma_i} + X_{i,i''}^{(\sigma_i-1)/\sigma_i} \right)^{\sigma_i/(\sigma_i-1)} \tag{24}$$

with $\sigma_i > 1$.²⁴ If goods i' and i'' are used only in production of good i , that is, i' and i'' are substitutes for each other and they never appear individually—then $\gamma_{i'}^L = \gamma_{i''}^L = 0$, regardless of any other elasticities or production weights. For example, it might be that iron and steel are substitutes for each other in all uses (if imperfect ones), in which case

²⁴Formally, this requires allowing for differential α_i across the production functions in the baseline setup. That is, a straightforward extension and again is allowable as long as the equilibrium conditions in (3) remain a contraction.

each individually has a left-tail centrality of zero.²⁵ This is the formalization of the idea described in the Introduction that what determines tail centrality is having a large fraction of GDP downstream and having no close substitutes.

To generalize further, one could imagine a situation where good i' is used both in a bundle with i'' and also separately on its own. Then, if $\sigma_j < 1 \forall j \neq i$, we have a modified version of Proposition 3. Define $d_{\min}^{-i}(j, i')$ to be the length of the shortest upstream path from j to i' that does not go through good i . Then

$$\gamma_i^L = \frac{1}{1 - \alpha} \sum_{j=1}^n \beta_j \alpha^{d_{\min}^{-i}(j, i')} \tag{25}$$

That is, if a good has substitutes for some uses but not others, then its tail centrality is calculated based on its closeness to final production only via paths where it cannot be substituted. Panel (d) of Figure 1 gives an example of this situation.

5.3. Extension: The Neoclassical Growth Model

Jones (2011) embeds a simple production network in a neoclassical growth model and shows how it can help explain income differences across countries. The analysis here can easily be extended to fit into that same framework.

Suppose each sector’s productivity is fixed on all dates τ at some $Z_{i,\tau} = Z_i$. To incorporate the neoclassical growth mechanism, we add capital so that the production function on date τ is

$$Y_{i,\tau} = \zeta Z_i (K_{i,\tau}^\eta L_{i,\tau}^{1-\eta})^{1-\alpha} \left(\sum_j A_{i,j}^{1/\sigma_i} X_{i,j,\tau}^{(\sigma_i-1)/\sigma_i} \right)^{\alpha\sigma_i/(\sigma_i-1)} \tag{26}$$

Supplemental Appendix D.1 then shows that after completing the model with standard assumptions, we have the following.

PROPOSITION 6: *Steady-state GDP per capita in the neoclassical growth model with sector production functions (26) is*

$$\left[(\beta^{-1} - 1 + \delta)^{-1} \eta \right]^{\eta/(1-\eta)} \exp\left(\frac{-1}{1-\eta} \beta' p \right), \tag{27}$$

$$\text{where } p \text{ solves (3) given } Z_i. \tag{28}$$

That is, we have the same formula for GDP as in the baseline case except that $-\beta' p$ is now multiplied by the usual factor $1/(1 - \eta)$ for long-run responses in a neoclassical growth model. That is, simply the textbook result that a decline in productivity feeds into a reduction in investment, thus shrinking the capital stock and further reducing output.²⁶

Theorem 1 and the examples and extensions discussed so far thus also have implications for income differences across countries. The sectors with the largest tail centralities have

²⁵Again, $\sigma_i > 1$ implies that if the price of good i' rises, then expenditures on it fall relative to those on i'' —if iron gets more expensive, then expenditures shift relatively toward steel (regardless of whether total expenditures on iron and steel combined rise or fall).

²⁶Jones (2011) emphasizes an additional multiplier on productivity coming through the use of intermediates. That appears here as the term $1/(1 - \alpha)$ appearing, for example, in equation (16).

the greatest potential to cause large cross-country income differences. While the analysis in Jones (2011) is symmetrical across goods, the discussion and intuition there often focuses on universal inputs, like electricity, and the analysis here shows that intuition is correct: productivity in universal inputs with no substitutes has the strongest effects on steady-state income.

6. HOW FAST IS THE CONVERGENCE TO THE LIMIT?

This section presents results describing the determinants of how large shocks need to be in order for the model to be “close” in some sense to the limit. The two primary determinants of the convergence rate are the difference between the elasticities of substitution and 1, and how flexibly inputs are reallocated across sectors.

6.1. *The Elasticity of Substitution*

Dividing the recursion for prices, (3), by t yields

$$p_i/t = -\theta_i + \frac{\alpha}{(1 - \sigma_i)t} \log \left(\sum_{j=1}^N A_{ij} \exp((1 - \sigma_i)t(p_j/t)) \right). \quad (29)$$

A restatement of Lemma 1 is that as $t \rightarrow \infty$, p_i/t converges to a constant. Note, though, that in the recursion for p_i/t , t is always multiplied by $1 - \sigma_i$. If $1 - \sigma_i$ is divided by some factor for all i , then to be equally close to the limit, t has to be multiplied by exactly the same factor. For example, if σ_i rises from 0.5 to 0.75 for all sectors, then the convergence to the limit happens exactly half as fast. And that is not an asymptotic statement—it is true along the entire path, simply because what appears in (29) is the product $(1 - \sigma_i)t$. The model is fully linear when $\sigma_i = 1 \forall i$, so this result also shows that when σ_i approaches 1, the convergence to the max/min limit of Lemma 1 becomes arbitrarily slow.

6.2. *Flexibility of Inputs*

The results so far assume that inputs, both labor and materials, are perfectly flexible across sectors, meaning that when a sector receives a negative shock resources can be reallocated to it, dampening its decline in output. In reality, though, inputs cannot be instantaneously reallocated—physical capital is subject to time to build, worker flows face numerous frictions, and firms often produce outputs from inventories of inputs that are purchased ahead of time. For example, inventories of materials and supplies in manufacturing industries represent about a month of production, according to the Census M3 survey.²⁷

The results in Theorem 1 in fact continue to hold in cases where inputs are not perfectly flexible. Supplemental Appendix C.1 describes one version of the result, where labor cannot be adjusted. To see how inflexibility in material inputs affects convergence to the limit, Supplemental Appendix C.2 develops a simple dynamic version of a special case of the model. This section describes the key results.

²⁷Liu and Tsyvinski (2021) study adjustment of inputs in a production network in detail and Jones (2011) also discusses the effects of optimal versus suboptimal allocation of inputs.

Assume that on each date τ there is a single final good, Y_τ , produced as an aggregate over sector outputs,

$$Y_\tau = \left(\sum_i a_i^{1/\sigma} Y_{i,\tau}^{(\sigma-1)/\sigma} \right)^{\sigma/(\sigma-1)}, \tag{30}$$

$$\text{where } Y_{i,\tau} = Z_{i,\tau} X_{i,\tau-1}. \tag{31}$$

$X_{i,\tau-1}$ represents inputs in sector i that must be purchased one period ahead of time. The resource constraint is $C_\tau + \sum_i X_{i,\tau} = Y_\tau$.

Now suppose we normalize $Z_{i,\tau} = 1$ for all $\tau < 0$, and then there is an unexpected shock on date 0 where each sector receives a new $Z_{i,0}$, which is then fixed forever going forward. Supplemental Appendix C.2 then shows that effective productivity, measured as output per unit of inputs, is

$$Y_\tau / \sum_i X_{i,\tau-1} = \begin{cases} 1, & \text{for } \tau < 0, \\ \left(\sum_i a_i Z_{i,0}^{(\sigma-1)/\sigma} \right)^{\sigma/(\sigma-1)}, & \text{for } \tau = 0, \\ \left(\sum_i a_i Z_{i,0}^{\sigma-1} \right)^{1/(\sigma-1)}, & \text{for } \tau > 0. \end{cases} \tag{32}$$

On the impact of the shock at $\tau = 0$, because there is no reallocation (since inputs available on date 0 depend on decisions from date -1), final output depends on a CES aggregate over the sector productivities with exactly the same elasticity as in production itself. As an example, if $\sigma = 0$, then aggregate productivity is just the minimum of the sector productivities, exactly as in the limit—the limit holds locally. For $\tau > 0$, though, inputs can be reallocated, and productivity becomes a less concave (or more convex, for $\sigma > 1$) function of the sector productivities. For $\sigma = 0$, aggregate productivity with full reallocation is a harmonic mean over sector productivities and, therefore, much less sensitive to the most extreme draw than the $\tau = 0$ case without reallocation, meaning that larger shocks will be required to approach the limiting behavior.

6.3. Numerical Examples

The left-hand panels of Figure 3 plot the level of GDP varying productivity in a single sector in the dynamic model above. They are for an economy with 30 sectors, each with the same initial share of GDP, $a_i = 1/30$. The lines plot log GDP for different elasticities of substitution, σ , ranging between 0.01 and 0.5. The top-left panel reports the short-run response ($\tau = 0$) to a productivity shock, and the bottom-left panel the long-run response ($\tau \geq 1$), which also corresponds (up to a scaling factor) to the perfectly flexible case studied in the earlier sections.

To see the limit derived in the theoretical analysis, just look at the lines for the smallest elasticities in the top-left panel. The theory formalizes the observation that GDP approaches a straight line (on both sides) and shows how to calculate the slope. One can see that the limits for the slopes are all the same.

There are two additional important features of the plots that help understand comparative statics. First, there is quantitatively far more nonlinearity for small shocks in the initial period ($\tau = 0$ in the top panel) with no adjustment than with full reallocation of inputs ($\tau > 0$ bottom panel). Note that in the bottom panel the scale runs to ± 5 , compared to ± 1 in the top panel. With no reallocation, the rate of convergence to the limit depends on

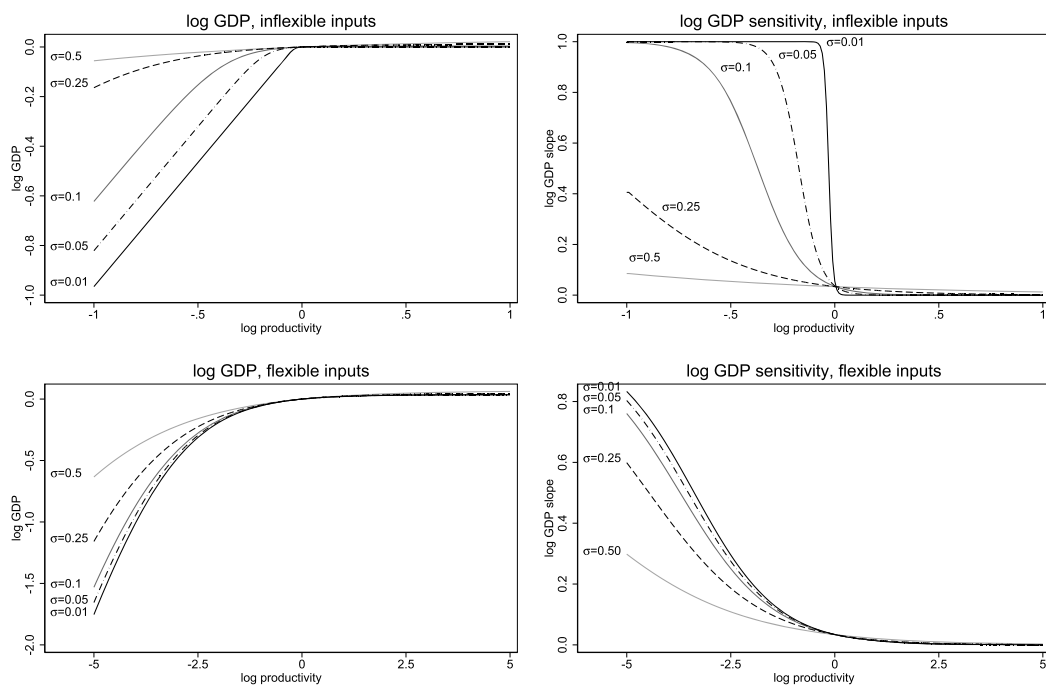


FIGURE 3.—Numerical examples. *Notes:* The panels report the effect of shocking log productivity in a single sector, holding all others fixed. The right-hand panels report the local response of log GDP to the sector’s productivity (its Domar weight).

$\frac{\sigma-1}{\sigma}$, while with full reallocation it depends on $\sigma - 1$. For small σ , those give very different rates—with no reallocation, convergence to the limit is infinitely fast as $\sigma \rightarrow 0$.

Second, the curvature is also clearly sensitive to the elasticity of substitution, falling as the elasticity rises, especially with no reallocation. In the Cobb–Douglas case ($\sigma = 1$), the model is linear, and GDP would be a straight (and nearly flat) line.

The right-hand panels plot the derivative of log GDP with respect to the single sector’s shock for different values of the shock. That derivative is, by Hulten’s theorem, also equal to the sector’s share of nominal GDP. So, the right-hand panels measure conditional granularity—they show whether and to what extent a sector may become granular following a shock.

The top-right panel shows that the jump from zero to full influence can happen in a small range of productivity, depending on σ . With $\sigma = 0.1$, far larger than the calibration of Baqaee and Farhi (2019), a shock of only about -30% is enough for the slope to get quantitatively close to the limit, and with $\sigma = 0.01$ the shift between the two limits happens about ten times faster. For larger elasticities, the convergence is clearly much slower. In terms of granularity, the sectors have steady-state shares of $1/30$. But for a shock of -30% , a sector can conceivably (in this highly stylized example) account for anything up to even 100% of GDP. And, again, this is *conditional*—the sector that is large is the one that was shocked, so it changes depending on what shock hits the economy. In other words, when the elasticity of substitution is low, *every* sector has the potential to be granular in this example.

The bottom-right panel, on the other hand, shows that with full reallocation much larger shocks are required in order to get quantitatively close to the limits, even with

low elasticities of substitution. So, conditional granularity is much less of an issue until the shocks become very large.

But what is a plausible upper end for the magnitude of shocks? Outside of food, the largest year-over-year changes in producer prices in the US are in energy prices—shifts on the order of a factor of 2–3 (0.7 to 1.1 in logs) in crude oil prices have occurred a number of times in the US, and in Europe the price of natural gas futures rose by a factor of 50 (3.3 in logs) between 2020 and 2022.

A second example is the shifts in expenditures during 2020 due to Covid. While these, like price changes, are also not productivity shocks, and thus map at best imperfectly into the model, the declines include 96% (−3.24 in logs) for air transportation, 43% (−0.56 in logs) for vehicles purchases, and 33% for health care (−0.4 in logs).²⁸

Finally, in cross-country comparisons, Duarte and Restuccia (2020) report 90/10 cross-country sectoral productivity ratios as large as 81 (4.2 in logs) at a high level of aggregation (see also Herrendorf and Valentinyi (2012)), and at the plant level there is evidence of far larger divergences (see Hsieh and Klenow (2009) and also the calibration of Jones (2011)). Overall, log shocks or productivity differences on the order of 3–5 or more appear to be a reasonable description of the most extreme scenarios in the data.

The theoretical results above always have qualitative value in helping to understand what makes a sector important in the tail and what sectors will become more important as their productivity falls. This section shows that the *quantitative* accuracy of the formulas depends on the elasticities of substitution, the extent to which low productivity in a sector can be alleviated by allocating more inputs to it, and the size of the shocks. When the elasticities are further from 1, and in the short-run or under other circumstances where reallocation of inputs is difficult, smaller shocks are needed for the model to approach the limits. For $\sigma \approx 0$ and no reallocation, the limits are hit essentially immediately. A canonical example is electricity—it has no close substitutes, and when there is a shock that restricts its supply, like a blackout, there is no reallocation of inputs that immediately solves the problem.

6.4. Additional Factors

In addition to the elasticity of substitution and flexibility of inputs, Supplemental Appendix C.3 discusses some additional factors determining the accuracy of the tail approximation. It shows that in general when the shock is to a sector that starts out relatively small, a larger shock is required to reach the asymptote (since the economy effectively starts out relatively far from it). It also studies how large the shocks have to be for the tail approximation to be more accurate than a Taylor series, and shows that that depends on the same factors that determine convergence to the asymptote.

More generally, we cannot say that there is some specific region in productivity space where either the Taylor series is “valid” or “invalid.” As $t \rightarrow 0$, the errors in a Taylor series approach zero, while as $t \rightarrow \infty$, they grow (for an n th-order series, they grow with t^n), but the error bounds from Taylor’s theorem always hold.²⁹ The tail approximation has the opposite features: its errors are largest for $t = 0$ (though they remain bounded), and shrink to zero as $t \rightarrow \infty$. The exact point where one becomes better than the other is parameter dependent.

²⁸There were also declines of 84% in hotels, 55% in gasoline, 50% in clothing, and 48% in food services. These sectors and those in the text combined account for 30% of personal consumption expenditures.

²⁹Note that Taylor series errors grow *faster* as the order, n , grows. Relatedly, the Taylor series for this model has a finite domain of convergence—for sufficiently large t , increasing the order increases rather than decreases the approximation errors.

7. EMPIRICAL ILLUSTRATION

This section examines two aspects of tail responses in the data. First, it compares local and tail responses of output to sector shocks. Second, it studies two sectors that have had significant changes in sales shares over time and examines how those changes relate to their out-degrees, and hence tail centralities.

7.1. *Local and Tail Responses*

I study the most recent (2012) sector detail input–output tables reported by the BEA. The tables have 379 private sectors.³⁰ The $A_{i,j}$ coefficients are set to be positive, so that there is an upstream link, if sector i spends at least 0.5% of its expenditures on materials for the output of sector j . That choice is made for two reasons. On a practical level, it stops the input–output matrix from being too dense and implying that many sectors are equally highly influential. Economically, it can be thought of as assuming, perhaps unrealistically in some cases, that inputs with very small expenditure shares are not strictly necessary for production. That said, the results are not terribly sensitive to varying the cut-off.

The β_i parameters are calculated from the fraction of nominal final expenditure going to each sector. I calculate α_i for each sector based on expenditures on materials relative to value added. I assume $\sigma_i < 1$, which allows us to use the results in Section 5.1.1. The tail elasticities, δ_i^L , can then be calculated from Proposition 3. I focus on the tail elasticity instead of tail centrality because it does not require the assumption that $A_{i,i} > 0$ (which is not true of all sectors) and it does not involve the $1/(1 - \alpha_i)$ term, which in the data occasionally becomes very large. Unlike in the main analysis, the α_i here are allowed to vary by sector. The local elasticities are extremely close empirically to Domar weights (sales shares). Appendix B describes the data and the details of the calculations of the various objects.

Figure 4 plots local against left-tail elasticities. There is a weak positive correlation of 0.20, but the figure makes apparent that the distributions are very different. There are a few sectors, such as petroleum refineries, that have sales shares noticeably higher than most others. But there are numerous sectors with tail elasticities close to 0.5. Twenty-one sectors have $\delta_i^L > 0.8 \max(\delta_i^L)$, while only two have local elasticities that are at least 80% of the maximum.

One can also see that the top sectors by local influence have very different tail influence—petroleum refineries at 0.40, oil and gas extraction at 0.34, and hospitals at 0.06. Oil and gas extraction is lower than petroleum refineries because it is one more step up the supply chain from refineries. Hospitals are low because they produce essentially only final output.

Table I further examines the top sectors sorted by local and tail elasticities. The top sectors for tail elasticity are all universal inputs. The first is electricity, which is why it has appeared frequently as an example. The second highest tail elasticity is for trucking services—all of final production involves trucking at some phase.³¹

The third-highest tail elasticity is for legal services—again, simply because every sector purchases legal services. Does it make sense that a large negative shock to the legal

³⁰For this paper's purposes, it is important to use a detailed version of the input–output tables because at higher levels of aggregation, the sectors become very strongly connected. The disaggregated table has much more sparse links. See Appendix B for a complete description of the data.

³¹The top-tail elasticities are clustered near 0.5 because an upper bound for δ_i^L is $\beta_i + \alpha(1 - \beta_i)$, and β_i is small while α is near 0.5.




FIGURE 4.—Tail and local output elasticities. *Notes:* The x-axis is the local elasticity of *gdp* with respect to y_i for each sector. The y-axis is the left-tail elasticity (see equation (15)). The data is the 2012 BEA input–output table. See Appendix B for details.

TABLE I
TOP SECTORS BY LEFT-TAIL ELASTICITY AND SALES SHARE.

Sector	δ_i^L	γ_i^L	Sales Share
<i>Largest by left-tail elasticity</i>			
Electric power generation, transmission, and distr.	0.4459	0.7194	0.0309
Truck transportation	0.4404	1.0287	0.0203
Legal services	0.4138	0.6999	0.0210
Advertising, public relations, and related services	0.4078	0.6757	0.0091
Accounting, tax prep., bookkeeping, and payroll serv.	0.4061	0.5702	0.0214
Services to buildings and dwellings	0.4045	0.5701	0.0290
Monetary authorities and depository credit intermed.	0.4033	0.6080	0.0231
Wired telecommunications carriers	0.3989	2.2404	0.0505
Other nondurable goods merchant wholesalers	0.3977	1.1949	0.0070
Insurance carriers, except direct life	0.3973	0.6295	0.0113
Petroleum refineries	0.3948	0.5114	0.0107
<i>Largest by sales share</i>			
Hospitals	0.0622	0.0622	0.0622
Petroleum refineries	0.3989	2.2404	0.0505
Oil and gas extraction	0.3389	0.5812	0.0444
Insurance carriers, except direct life	0.0358	0.0551	0.0363
Electric power generation, transmission, and distribution	0.3750	0.7277	0.0325
Offices of physicians	0.4459	0.7194	0.0309
Monetary authorities and depository credit intermediation	0.0302	0.0302	0.0302
Scientific research and development services	0.4045	0.5701	0.0290
Other financial investment activities	0.2537	0.5226	0.0236
Advertising, public relations, and related services	0.4033	0.6080	0.0231
Wired telecommunications carriers	0.3507	0.6306	0.0220

Note: Sales shares and tail elasticities calculated from the 2012 BEA input–output tables. See Appendix B for details.

services sector could cause a crash in GDP, or that large differences in “productivity” in legal services could explain cross-country income differences? There is ample evidence that legal institutions are necessary for the growth of the economy. All aspects of business rely on property rights and contract enforcement. If, for some reason, the legal system literally shut down and legal services were actually no longer available to firms, it is entirely plausible that there would be significant declines in output. In addition, even if one does not believe that there are exactly TFP shocks to sectors like legal services, the results here are still useful for formalizing how differences in the availability of legal services, broadly speaking, can explain income differences across countries, as in Section 5.3.

One potential concern with that argument, though, is that the input–output tables do not actually measure things like enforcement of property rights or the use of courts; they just measure expenditures on lawyers by firms. That actually illustrates a key advantage of δ_i^L and γ_i^L : measuring them does not require measuring *all* of each sector’s use of each input. All that we need to know is that a sector uses some input at all—again, we need to know links between sectors, not their intensity. And the input–output tables are certainly correct that all sectors directly use legal services.

In addition to utilities and professional services like lawyers and accountants, the last major category of sectors that appears repeatedly among the top sources of tail risk is financial institutions. Just as with legal services, all firms use financial services in one way or another (as do all households). The analysis here thus helps explain why the financial sector would be a relevant source of crashes throughout history—when financial services are disrupted, every firm in the economy faces more difficulty in production.

There is past work examining, both in models and in the data, the effects of shocks to the energy sector, financial services, and legal and accounting institutions. The analysis here shows how those shocks are linked: they all represent shocks to universal inputs, where tail centralities are far larger than steady-state sales shares. As in the previous section, the tail approximation here will be most accurate when the shocks are large or when reallocation of inputs is inhibited. So, the identification of important sectors here will be most useful either for very large high-frequency shocks or in cross-country comparisons.

Table I also reports the tail centrality, γ_i^L . When $A_{i,i} > 0$, $\gamma_i^L = \delta_i^L / (1 - \alpha_i)$. The tail centrality is always larger, and since α_i is sometimes near 1, some values of γ_i^L are very high, showing less clustering than δ_i^L .

The bottom section of Table I reports the top sectors sorted by sales share. Again, not all have particularly high-tail centralities. In many cases, they only produce final goods, like hospitals.

Last, note that in a complementary economy, responses to large shocks, γ_i^L and δ_i^L , are not terribly interesting, both being equal to β_i .

7.2. Hospitals and Computers

Two prominent sectors that have undergone significant changes in the post-war period are computer equipment and hospitals.³² The left-hand panel of Figure 5 plots their Domar weights (sales shares, which are nearly identical to local output elasticities) for the period 1963–2020. The Domar weight of hospitals rose by a factor of 5 from 0.02 to 0.10.

³²The computer equipment sector stays consistent in the BEA input–output tables between the 1963–1996 and 1997–2020 versions. For consistency across those two data sets, I use the combined “hospitals and nursing and residential care” sector.




FIGURE 5.—Time series of Domar weights and out-degree. *Notes:* The left-hand panel plots Domar weights for the two sectors calculated from BEA annual input–output tables. The right-hand panel plots, for each year, the fraction of sectors that spent at least 0.5% of expenditures on material inputs on the industry’s output (note this is measuring computers as a material input; investment expenditures are not counted in measuring the production network parameters $A_{i,j}$).

Computer equipment rose from about 0.03 to a peak of 0.07 and then fell back to nearly where it started. According to the standard local analysis, hospitals have become progressively more important, while the importance of computers to the economy peaked around 2000 and has subsequently fallen by half.

The right-hand panel of Figure 5 plots their out-degrees, measured here as the fraction of sectors that purchase output from those same two sectors. Hospitals never sell output to more than one other sector (where again the cut-off is 0.5% of the using sector’s intermediates). Computers, on the other hand, rose from being purchased by 30 to 55% of sectors. The rise in the Domar weight of the computer-producing sector can thus be said to be driven by the extensive margin; its Domar weight increases by the same factor as the number of sectors using its output—whereas the rise in the Domar weight of hospitals is driven by the intensive margin—the share of final expenditures going to them has risen.

In terms of tail elasticities, using the detailed input–output tables as above, the tail elasticity of the semiconductor-producing sector (the figure uses “computer equipment”—a higher level of aggregation—because it is available at the annual frequency) rose from 0.18 to 0.31 between 1963 and 2012, while the tail elasticity of hospitals is always simply equal to its share of final consumption, which is also its Domar weight. In addition, recall that Proposition 5 shows that there is a close link between the out-degree plotted in Figure 5 and left-tail centrality.

8. THE RISK OF LARGE DEVIATIONS IN GDP AND THEIR SOURCE

The results so far describe how the economy responds to a given shock to productivity. This section combines Theorem 1 with assumptions about the probability distribution for shocks to describe the probability distribution of GDP. It gives a general result for the determinants of the tails of GDP and discusses the implications and then examines one particular example from the literature.

The key results from this section are as follows:

1. The tail approximation from Theorem 1 is sufficient for characterizing the tail of GDP (meaning that the invariance that holds for Theorem 1 also holds for the determinants of GDP tail risk).
2. In a complementary economy, increases in interconnectedness increase tail risk.
3. Whereas past work has studied the riskiness of the steady-state production network, tail risk is in general driven by the riskiest of the *tail* networks, as in Section 5.1.1.

8.1. Shock Distributions

I assume that there is a positive function $s(\theta)$ that determines the scale of the shocks in direction θ . Specifically, for t greater than some \bar{t} , $t/s(\theta)$ has a cumulative distribution function F , with complementary CDF $\bar{F} \equiv 1 - F$ (note \bar{F} is positive and decreasing). So, for example, if $s(\theta) = ks(\theta')$, then the n th percentile of z in direction θ is k times that in direction θ' . For the purposes of this paper, consistent with the analysis so far, it is only necessary to choose the distribution of z for large t (i.e., when $\|z\|$ is large), with its behavior for $t \leq \bar{t}$ left unrestricted.

I assume θ has a probability measure m . Since $z = \theta t$ is a unique decomposition, we can write its probability distribution equivalently over z or θ and t (with $t = \|z\|$ and $\theta = z/\|z\|$). To formalize the above assumptions, we set, for $t > \bar{t}$,

$$\Pr[\theta \in \Theta, t/s(\theta) > x] = m(\Theta)\bar{F}(x). \tag{33}$$

The representation in (33) accommodates standard distributions studied in the literature such as the multivariate normal, elliptical distributions more generally, transformations of Laplace distributed vectors, and Pareto-tailed distributions (Resnick (2007)). A simple example of a distribution that does not have a representation (33) is the case with $N = 1$ so that z is a scalar and z is distributed normally conditional on being positive but exponentially conditional on being negative. Intuitively, the restriction, which can easily be relaxed, is that the tail shape (as distinct from the scale) is the same for all θ .³³

8.2. General Result

THEOREM 2: *Given the distribution for z in (33), there exists a function $\varepsilon(x) \geq 0$ with $\lim_{x \rightarrow \infty} \varepsilon(x) = 0$ and an \bar{x} such that for $x > \bar{x}$,*

$$\begin{aligned} \int_{\Theta_-} \bar{F}\left(\frac{x - \mu(\theta) + \varepsilon(x)}{-s(\theta)\lambda(\theta)}\right) dm(\theta) &\leq \Pr[\text{gdp} < -x] \\ &\leq \int_{\Theta_-} \bar{F}\left(\frac{x - \mu(\theta) - \varepsilon(x)}{-s(\theta)\lambda(\theta)}\right) dm(\theta), \end{aligned} \tag{34}$$

where $\Theta_- = \{\theta : s(\theta)\lambda(\theta) < 0\}$.

Theorem 2 says that the CDF of $\log \text{GDP}$ is well approximated by

$$\int_{\Theta_-} \bar{F}\left(\frac{x - \mu(\theta)}{-s(\theta)\lambda(\theta)}\right) dm(\theta) \tag{35}$$

³³For practical purposes, if the tail decays significantly faster in some direction ($z > 0$ in this example), then that can be analyzed by just setting the measure m to zero in that direction.

and in fact the $\mu(\theta)$ term is also irrelevant since x eventually dominates. Intuitively, this says that the CDF of GDP, in the tail, depends on the average across all shocks ($\int dm(\theta)$), of the probability that each shock (θ) creates a large decline in GDP, where $(x - \mu(\theta))/\lambda(\theta)$ is the size of a shock needed in direction θ to generate a decline of size x .

8.2.1. *General Properties of the Tail of GDP*

Even without further specialization, there are general results that follow from Theorem 2.

Determinants of Tail Risk. First, the probability of large deviations in GDP depends on the probability of large deviations in productivity, scaled by the limiting slope, $\lambda(\theta)$, showing that the tail approximation is the correct way to analyze tail risk in this setting. Other aspects of the economy—such as the steady-state Domar weights, the precise values of the elasticities of substitution, or terms in a Taylor expansion—are irrelevant. The invariance results for the function λ thus also hold for tail risk—it is unaffected by the exact values of the production parameters and only depends on the topology of the production network (which $A_{i,j} > 0$) along with whether σ_i is above or below 1.

A second observation is that the volatility of the shocks in different directions, captured by $s(\theta)$, interacts with $\lambda(\theta)$ to determine tail risk. When the shocks are more volatile, s is larger and tail risk is greater.

Comparative Statics. Generalized versions of the comparative statics in Section 4.1 are useful here for showing what makes the economy riskier.

PROPOSITION 7: *For sufficiently large x , any factor that weakly increases $\lambda(\theta)$ for all θ weakly reduces tail risk in the limiting sense of Theorem 1. In particular,*

1. *when any σ_i transitions from below to above 1;*
2. *when the set of inputs used by any sector i grows if $\sigma_i > 1$ or shrinks if $\sigma_i < 1$.*

The second part of the proposition shows how changes in interconnectedness affect tail risk—interconnectedness reduces tail risk when it increases the number of substitutes and increases tail risk when it increases the number of complements.

Skewness. We also obtain a general result on skewness in the tail. It is an asymptotic form of skewness, as opposed to the scaled third moment.

COROLLARY 3: *If the distribution of z is symmetrical ($s(\theta) = s(-\theta)$ and $m(\theta) = m(-\theta)$), then when GDP is concave in the sense that $\lambda(\theta) \leq -\lambda(-\theta)$, $\Pr[\text{gdp} < -x] \geq \Pr[\text{gdp} > x]$ for sufficiently large x . In particular, that holds when $\sigma_i < 1$ for all i .*

So, under very general (but still only sufficient) conditions, as long as the elasticities are all below 1, the left tail of GDP is heavier than the right. Concavity in production thus robustly generates left skewness in GDP, in the limiting sense of the corollary. This is a formal tail version of results that are intuitively described and studied in a local approximation by Baqaee and Farhi (2019).

Finally, Theorem 2 shows how nonlinearity in the economy generates increases in tail risk. If the economy were linear, the argument of \bar{F} in (34) would be $\frac{x}{-s(\theta)D_{ss}\theta}$. When $\lambda(\theta)$ is larger in magnitude than $D'_{ss}\theta$, there is a larger chance of a large movement in GDP.

8.3. *Interconnectedness and Risk in the Economy*

As discussed above and in Section 4.1, when a sector sells to a new downstream sector, left-tail risk weakly increases if the new downstream sector has an elasticity of substitution less than 1. In other words, complementarity and interconnectedness combine to increase left-tail risk (and at the same time reduce the probability of large booms in GDP).

But obviously the tail probabilities in Theorem 2 are not the only way to evaluate the risk of the economy. Another interesting question is how the economy responds to small shocks, or, equivalently, what the variance of $\log GDP$ is in a first-order Taylor approximation.

If Σ is the covariance matrix of z , we have from a first-order approximation that

$$\text{var}(\log GDP) \approx D'_{ss} \Sigma D_{ss}. \tag{36}$$

Since $D'_{ss} \Sigma D_{ss}$ is continuous in A , any small change in A , that is, a change in some $A_{i,j}$ from zero to a small positive number will cause only a small change in $D'_{ss} \Sigma D_{ss}$, even though it can cause a discrete shift in the values of the function λ , and hence in tail risk. In other words, local risk is always affected smoothly by A , but tail risk is affected discretely by it.

In addition, an increase in interconnectedness, even though it cannot reduce tail risk when $\sigma_i < 1 \forall i$, can certainly reduce the sensitivity of GDP to small shocks. Since the sum of the Domar weights, $D_{ss,i}$, is always equal to $(1 - \alpha)^{-1}$, we have the following simple example.

EXAMPLE 2: Suppose the shocks are uncorrelated so that Σ is diagonal. A marginal increase in the sales share of any sector starting from zero, if it (weakly) reduces the sales shares of all other sectors, will reduce $D'_{ss} \Sigma D_{ss}$.

The example gives simple sufficient—and far from necessary—conditions for when adding a new sector diversifies the economy. At the same time, though, Proposition 7 shows that adding a new sector will weakly increase tail risk (weakly reduce $\lambda(\theta)$ for all θ) when the elasticity of substitution in production is less than 1. This section thus shows that in the model increases in interconnectedness—measured here by the number of links in the production network ((i, j) pairs such that $A_{i,j} > 0$)—can diversify the economy, making it less sensitive to small shocks, while at the same time increasing the probability of an extreme negative realization of GDP.

8.4. *Example: Exponential Tailed Shocks*

EXAMPLE 3: Suppose the sector productivity shocks are i.i.d. with exponential tails, implying that $s(\theta) = 1/\|\theta\|_1$ and $m(\theta)$ has full support. Then as $s \rightarrow \infty$,

$$\Pr[gdp < -x] \rightarrow \exp\left(-\eta \frac{x}{\max_n \max_j D_{n,j}}\right). \tag{37}$$

If, in addition, $\sigma_i \leq 1 \forall i$, then

$$\Pr[gdp < -x] \rightarrow \exp\left(-\eta \frac{x}{\max_j \gamma_j^L}\right). \tag{38}$$

The shock θ causing the tail event is equal to 1 for the sector with the largest γ_j^L and zero elsewhere.

Acemoglu, Ozdaglar, and Tahbaz-Salehi (2017) study a model with Cobb–Douglas production so that the Domar weights are constant (the model is log-linear) and in which shocks are i.i.d. with exponential tails, as in this example. They show in their model that what determines tail risk is the largest Domar weight. The first part of example 3 says that we get a very similar result, but here it is the maximum Domar weight among all possible tail networks that determines tail risk.³⁴ So it need not be the case that $\max_j D_{ss,j}$ is large for there to be significant tail risk. Rather, under complementarity there just needs to be *some* Domar weight that can be large in some situation. The second part of the result follows from Corollary 2—each sector’s tail centrality is in fact equal to its maximum Domar weight across all networks when production is complementary.

The fact that extreme events are caused by a shock to a single sector—the one with the highest left-tail centrality—is again due to the importance of conditional granularity in the model. Crashes appear not necessarily because of granularity local to steady state, but because there *can be* granularity in an extreme event. If the model is such that granularity cannot occur—the maximum tail centrality (which is the maximum possible Domar weight among all tail networks) is small—then tail risk will also be small.

As an example, the steady-state Domar weight of electricity is not particularly large empirically—it is certainly not the largest sector in the economy—but its tail elasticity is highest. One can imagine a scenario in which electricity or some other energy sector receives a large negative shock, becomes a limiting input in production, and then becomes much more expensive. That is the type of scenario that these limits show is important for driving the largest declines in GDP in this model, and it is a very different scenario from the model of Acemoglu, Ozdaglar, and Tahbaz-Salehi (2017), in which tail risk arises only when there is a big sector at the steady state.

When this example is generalized so that the shocks are exponential but with different scales, then the sector that causes crashes is the one with the highest product of its tail centrality with its volatility (see Supplemental Appendix E).

Note also that in Acemoglu, Ozdaglar, and Tahbaz-Salehi (2017), when the shocks are distributed symmetrically, tail risk is also symmetrical. Here, on the other hand, tail risk is in general asymmetrical even for symmetrical shocks.

9. CONCLUSION

This paper studies large deviations in GDP in the context of a general nonlinear network production model. Its core result characterizes the asymptotic response of GDP to arbitrary combinations of shocks. That result yields a description of the determinants of tail risk and a measure of the risk associated with large shocks to individual sectors. In addition, when combined with a probability distribution for shocks, it yields a description of the tail of the probability distribution of GDP.

The simple statement of the core idea is that what determines tail risk is the structure of the economy in the tail. For example, while granularity near steady state affects the dynamics of the economy near steady state, what determines behavior in the tail is whether the economy displays granularity in the tail. The paper shows how that can easily happen even in a perfectly symmetrical economy where all sectors are of equal size at steady state.

³⁴In the case where $\sigma_i = 1$ for all $i \geq 0$, $\{D_n\}$ is just the singleton D_{ss} and we recover their original result.

A closely related point is that to understand the systemic risk of a sector—whether a large shock to it will spill over into the rest of the economy—one needs to understand the importance of the sector not on average but rather conditional on the occurrence of a large shock. The analysis shows that it is upstream sectors that produce inputs for a large fraction of GDP that are most systemically risky, while sectors that exclusively produce final outputs do not produce systemic risk.

More broadly, the paper provides a general theoretical foundation for analyzing tail risk. It shows how to construct an approximation for the dynamics of the economy that, rather than being valid only for small shocks, is valid explicitly for large shocks. That approximation can then be combined with assumptions about the shape of the tail of the shock distribution to yield a description of the tail behavior of the full economy.

APPENDIX A: PROOFS

A.1. Lemma 1

The assumption that aggregate labor supply is inelastic and normalized to one implies that real GDP is

$$GDP = W/P_0, \tag{39}$$

where W is the wage and P_0 is the price of the consumption bundle. The index 0 indicates consumption (P_0 might be called a pseudo price, since it is the cost of the consumption bundle, but not of an actual individual good). The CES preferences for the consumer along with cost minimization and the normalization $W = 1$ immediately imply

$$p_0 = \sum_{i=1}^N \beta_i p_i, \tag{40}$$

$$gdp = -p_0. \tag{41}$$

Similarly, marginal cost pricing by the producers implies that the log price of good i is

$$p_i = -z_i + \frac{\alpha}{1 - \sigma_i} \log \left(\sum_{j=1}^N A_{ij} \exp((1 - \sigma_i)p_j) \right). \tag{42}$$

Now define $\phi_i = -\lim_{t \rightarrow \infty} p_i/t$ and set the vector $\phi \equiv [\phi_1, \dots, \phi_N]$. If that limit exists and is finite (a claim established below), then dividing by t and taking limits of both sides of equations (40) and (42) give

$$\lim_{t \rightarrow \infty} t^{-1} gdp = \beta' \phi, \tag{43}$$

$$\phi_i = -\theta_i + \alpha_i f_i(\phi), \tag{44}$$

where

$$f_i(\phi) \equiv \begin{cases} \max_{j \in \mathcal{S}_i} \phi_j & \text{if } \sigma_i < 1, \\ \sum_j A_{i,j} \phi_j & \text{if } \sigma_i = 1, \\ \min \phi_j & \text{if } \sigma_i > 1. \end{cases} \tag{45}$$

To show that the system has a unique solution (guaranteeing that ϕ is also finite), define a mapping $g : \mathbb{R}^N \rightarrow \mathbb{R}^N$ such that the i th element of the vector $g(\phi)$ is

$$g_i(\phi) = \theta_i + \alpha f_i(\phi). \tag{46}$$

The set of solutions for ϕ is the set of fixed points for g , so we must just show that g has a unique fixed point. That follows from the Banach fixed-point theorem if g_i is a contraction. It is straightforward to confirm the Blackwell’s sufficient conditions hold here, giving the result. The continuity of the solution follows from the continuity of g in θ .

To get the constant $\mu(\theta)$, consider a series expansion, $p_i = \mu_i + \phi_i t + o(1)$ (as $t \rightarrow \infty$). Inserting that into (3) taking limits, and using (44) yields a recursion for μ .

A.2. Propositions 1, 2, and 7

Define $f^0 : \mathbb{R}^N \rightarrow \mathbb{R}^N$ to be the vectorized version of the function in (45). Define a transformation $T^0 \phi = -\theta + \alpha f^0(\phi)$, with $\phi^0 = T^0 \phi^0$ the fixed point of that transformation.

After changing some σ_i , we have a new f^1 (analogous to f^0) and associated T^1 . First, take the case with σ_i transitioning from above 1 to being equal to 1 or below. Then, necessarily,

$$T^1 \phi \geq T^0 \phi \tag{47}$$

for any ϕ , element-by-element. That means that the fixed-point $\phi^1 \geq \phi^0$ elementwise, from which Proposition 1 follows.

Proposition 2 by the same argument. For example, suppose $\sigma_i < 1$ and the set S_i grows. Again, define an f^2 and T^2 for the model with the larger S_i . We have

$$T^2 \phi \geq T^0 \phi \tag{48}$$

for any ϕ , elementwise, so that $\phi^2 \geq \phi$ elementwise, establishing Proposition 2.

Since both of those statements hold for arbitrary θ , they also establish Proposition 7

A.3. Proposition 4

Define a set of $N \times N$ matrices \mathbf{A}_k representing restricted versions of the production network. For each \mathbf{A}_k , each sector is restricted to using just one of its inputs, so that every \mathbf{A}_k has a single value of 1 in each row and is otherwise equal to zero, with links (1’s) only appearing where $A_{i,j} > 0$. The set over all k of $\{\mathbf{A}_k\}$ represents every possible restricted network.³⁵ If $\sigma_i = 1$, then sector i always uses the same mix of inputs, and the i th row of \mathbf{A}_k is equal to $A_{i,\cdot}$ for every k .

For each \mathbf{A}_k , there is an associated vector of Domar weights,

$$D'_k = \beta'(\mathbf{I} - \alpha \mathbf{A}_k)^{-1}. \tag{49}$$

Now define ϕ^* and k^* :

$$k^* \equiv \arg \min_k \beta'(\mathbf{I} - \alpha \mathbf{A}_k)^{-1} \theta = \arg \min_k D'_k \theta, \tag{50}$$

$$\phi^* \equiv -(\mathbf{I} - \alpha \mathbf{A}_{k^*})^{-1} \theta. \tag{51}$$

³⁵The index k runs from 1 to the product of the number of inputs used by each each sector.

That implies

$$\phi^* = -\theta + \alpha \mathbf{A}_{k^*} \phi^*. \tag{52}$$

As above, define $f : \mathbb{R}^N \rightarrow \mathbb{R}^N$ to be the vectorized version of the function in (45). Now suppose \mathbf{A}_{k^*} is not the true tail network, in that,

$$\mathbf{A}_{k^*} \phi^* \neq f(\phi^*). \tag{53}$$

Then, clearly, element-by-element $T\phi^* \geq \phi^*$, where T is again the operator $T\phi \equiv -\theta + \alpha f(\phi)$. Then whatever the solution is for ϕ in Lemma 1, it will be, element-by-element, weakly greater than ϕ^* . But that solution is always of the form $-(\mathbf{I} - \alpha \mathbf{A}_n)^{-1} \theta$, leading to a contradiction with the original construction of ϕ^* . So, ϕ^* must be the solution to the recursion with $T\phi^* = \phi^*$. The result for GDP then follows immediately.

A.4. Proposition 5

The left-hand inequality follows from assuming that the sectors immediately downstream of i have no other downstream users (except final output). The right-hand inequality follows from assuming that the remainder of GDP that is not immediately downstream of sector i 's users is a single step further downstream.

A.5. Theorem 2

We have

$$gdp(z) = \mu(\theta) + \lambda(\theta)t + \varepsilon(t, \theta), \tag{54}$$

where $\varepsilon(t, \theta)$ is an error that converges to 0 as $t \rightarrow \infty$ (from Theorem 1).

Now define

$$\bar{\varepsilon}(x) = \max_{\theta} \max_{t > \frac{x + \mu(\theta)}{-\lambda(\theta)}} |\varepsilon(t, \theta)|. \tag{55}$$

Consider its limit as $x \rightarrow \infty$. Since the right-hand side is bounded and continuous in θ , the limit can be passed through the maximum and we have

$$\lim_{x \rightarrow \infty} \bar{\varepsilon}(x) = 0. \tag{56}$$

Now note that

$$\Pr[gdp < -x \mid \theta] = \Pr\left[t + \frac{\varepsilon(t, \theta)}{\lambda(\theta)} > \frac{x + \mu(\theta)}{-\lambda(\theta)} \mid \theta\right], \tag{57}$$

where $\lambda(\theta) < 0$. In addition,

$$\begin{aligned} \Pr\left[t + \frac{\bar{\varepsilon}(x)}{\lambda(\theta)} > \frac{x + \mu(\theta)}{-\lambda(\theta)} \mid \theta\right] &\leq \Pr\left[t + \frac{\varepsilon(t, \theta)}{\lambda(\theta)} > \frac{x + \mu(\theta)}{-\lambda(\theta)} \mid \theta\right] \\ &\leq \Pr\left[t - \frac{\bar{\varepsilon}(x)}{\lambda(\theta)} > \frac{x + \mu(\theta)}{-\lambda(\theta)} \mid \theta\right], \\ \Pr\left[t > \frac{x + \mu(\theta) + \bar{\varepsilon}(x)}{-\lambda(\theta)} \mid \theta\right] &\leq \Pr[gdp < -x \mid \theta] \leq \Pr\left[t > \frac{x + \mu(\theta) - \bar{\varepsilon}(x)}{-\lambda(\theta)} \mid \theta\right]. \end{aligned} \tag{58}$$

By integrating over the measure for θ (i.e., applying Fubini’s theorem),

$$\Pr[gdp < -x] = \int_{\Theta} \Pr[gdp < -x \mid \theta] dm(\theta) \tag{59}$$

from which the result follows directly.

A.6. Corollary 3

Recall the notation from the proof of Theorem 2 that

$$gdp(\theta t) = \mu(\theta) + \lambda(\theta)t + \varepsilon(\theta, t) \tag{60}$$

and that $|\varepsilon(\theta, t)| \leq \bar{\varepsilon}(x)$ for $t > \frac{x + \mu(\theta)}{-\lambda(\theta)}$. We want to compare $\Pr[gdp < -x]$ with $\Pr[gdp > x]$. Define $\varepsilon'(x) = \max(\bar{\varepsilon}(x), \bar{\varepsilon}(-x))$. We have the bounds

$$\Pr[gdp < -x] \geq \int_{\theta: \lambda(\theta) < 0} \bar{F}\left(\frac{x - \mu(\theta) + \varepsilon'(x)}{-s(\theta)\lambda(\theta)}\right) dm(\theta), \tag{61}$$

$$\Pr[gdp > x] \leq \int_{\eta: \lambda(\eta) > 0} \bar{F}\left(\frac{x - \mu(\eta) - \varepsilon'(x)}{s(\eta)\lambda(\eta)}\right) dm(\theta). \tag{62}$$

Now first note that, for θ such that $\lambda(\theta) < 0$,

$$\frac{x - \mu(-\theta) - \varepsilon'(x)}{s(-\theta)\lambda(-\theta)} - \frac{x - \mu(\theta) + \varepsilon'(x)}{-s(\theta)\lambda(\theta)} \tag{63}$$

$$= \left(\frac{1}{s(-\theta)\lambda(-\theta)} - \frac{1}{-s(\theta)\lambda(\theta)}\right)x + \frac{-\mu(-\theta) - \varepsilon'(x)}{s(-\theta)\lambda(-\theta)} - \frac{-\mu(\theta) + \varepsilon'(x)}{-s(\theta)\lambda(\theta)}. \tag{64}$$

So, there exists an \bar{x} such that for $x > \bar{x}$, the argument of \bar{F} in the integral for (61) is smaller than that in (62) for any given θ . In addition,

$$m(\{\eta : \lambda(\eta) > 0\}) \leq m(\{\theta : \lambda(\theta) < 0\}) \tag{65}$$

which yields the result.

APPENDIX B: ESTIMATES OF LOCAL AND TAIL ELASTICITIES

B.1. Data

The estimates are based on the 2012 input–output tables from the BEA. The specific table is the 405-industry table after redefinitions.³⁶ Each sector’s α_i (discussed further below) is constructed as total intermediate expenditures divided by total intermediate expenditures plus value added. The β_i ’s are constructed as shares of final use, which includes consumption, private and public investment (excluding inventories) and exports. Imports and inventories are excluded because they do not represent final uses of domestically produced commodities.

³⁶As of the writing of the paper, the tables were located at https://apps.bea.gov/industry/xls/io-annual/IOUse_After_Redefinitions_PRO_DET.xlsx.

I keep all commodities except for 4200ID (customs duties), 525000 (funds and trusts), 531HSO, 531HST, 531ORE (real estate), 550000 (management of companies; mostly offices of holding companies), 561300 (employment services; e.g., temp agencies), 811400, 812100, 812200, 812300, and 812900 (miscellaneous personal services).

For Figure 3, the input-out matrix is *before* instead of after redefinitions because the BEA does not produce an after-redefinitions file for the period 1963–1996.

B.2. Constructing Local Elasticities

To get the local elasticity, $dgdp/dy_i$, I first get $dgdp/dz_i$ based on Hulten’s theorem, which says that it is equal to the sector’s nominal sales divided by nominal GDP (where GDP here is calculated based on the sum of final uses, as in the construction of the β_i ’s). I then divide $dgdp/dz_i$ by dy_i/dz_i . To get the latter, I use the result from Carvalho and Tahbaz-Salehi (2019) that in a Cobb–Douglas economy—which is first-order equivalent to the general CES economy— dy_i/dz_i is equal to the i, i element of the Leontief inverse matrix. That matrix is constructed from the input–output table and the α_i ’s described above.

B.3. Constructing Tail Elasticities

As discussed in Section 2, the paper’s analysis goes through identically if α is sector specific. To construct δ_i^L and γ_i^L , I calculate $\phi(-e_i)$, where e_i is equal to 1 in element i and 0 elsewhere, by iterating on the recursion for ϕ , (7). We then have that $\gamma_i^L = -\beta' \phi(-e_i)$ and $\delta_i^L = \beta' \phi(-e_i)/\phi_i(-e_i)$.

REFERENCES

- ACEMOGLU, DARON, AND PABLO D. AZAR (2020): “Endogenous Production Networks,” *Econometrica*, 88 (1), 33–82. [2091,2094]
- ACEMOGLU, DARON, VASCO M. CARVALHO, ASUMAN OZDAGLAR, AND ALIREZA TAHBAZ-SALEHI (2012): “The Network Origins of Aggregate Fluctuations,” *Econometrica*, 80 (5), 1977–2016. [2089]
- ACEMOGLU, DARON, ASUMAN OZDAGLAR, AND ALIREZA TAHBAZ-SALEHI (2017): “Microeconomic Origins of Macroeconomic Tail Risks,” *American Economic Review*, 107 (1), 54–108. [2091,2116]
- ALFARO, LAURA, DAVIN CHOR, POL ANTRAS, AND PAOLA CONCONI (2019): “Internalizing Global Value Chains: A Firm-Level Analysis,” *Journal of Political Economy*, 127 (2), 508–559. [2092]
- ALFARO, LAURA, MANUEL GARCÍA-SANTANA, AND ENRIQUE MORAL-BENITO (2021): “On the Direct and Indirect Real Effects of Credit Supply Shocks,” *Journal of Financial Economics*, 139 (3), 895–921. [2092]
- ANTRÁS, POL, AND DAVIN CHOR (2013): “Organizing the Global Value Chain,” *Econometrica*, 81 (6), 2127–2204. [2102]
- ANTRÁS, POL, DAVIN CHOR, THIBAUT FALLY, AND RUSSELL HILLBERRY (2012): “Measuring the Upstreamness of Production and Trade Flows,” *American Economic Review*, 102 (3), 412–416. [2102]
- ATALAY, ENGHIN (2017): “How Important Are Sectoral Shocks?” *American Economic Journal: Macroeconomics*, 9 (4), 254–280. [2100]
- ATALAY, ENGHIN, THORSTEN DRAUTZBURG, AND ZHENTING WANG (2018): “Accounting for the Sources of Macroeconomic Tail Risks,” *Economics Letters*, 165, 65–69. [2100]
- BAQAEI, DAVID, AND EMMANUEL FARHI (2020): “Entry vs. Rents,” Working paper. [2094]
- BAQAEI, DAVID REZZA, AND EMMANUEL FARHI (2019): “The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten’s Theorem,” *Econometrica*, 87 (4), 1155–1203. [2090,2092,2093,2100,2107,2114]
- BARROT, JEAN-NOËL, AND JULIEN SAUVAGNAT (2016): “Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks,” *The Quarterly Journal of Economics*, 131 (3), 1543–1592. [2089]
- BERNARD, ANDREW B., ANDREAS MOXNES, AND YUKIKO U. SAITO (2019): “Production Networks, Geography, and Firm Performance,” *Journal of Political Economy*, 127 (2), 639–688. [2102]
- BIGIO, SAKI, AND JENNIFER LA’O (2020): “Distortions in Production Networks,” *The Quarterly Journal of Economics*, 135 (4), 2187–2253. [2092]

- BLOCH, FRANCIS, MATTHEW O. JACKSON, AND PIETRO TEBALDI (2021): "Centrality Measures in Networks," Working paper. [2102]
- BOEHM, JOHANNES, AND EZRA OBERFIELD (2020): "Misallocation in the Market for Inputs: Enforcement and the Organization of Production," *The Quarterly Journal of Economics*, 135 (4), 2007–2058. [2092]
- BOLDI, PAOLO, AND SEBASTIANO VIGNA (2014): "Axioms for Centrality," *Internet Mathematics*, 10 (3–4), 222–262. [2102]
- BONACICH, PHILLIP (1987): "Power and Centrality: A Family of Measures," *American journal of sociology*, 92 (5), 1170–1182. [2102]
- BUTKOVIC, PETER (2010): *Max-Linear Systems: Theory and Algorithms*. Springer. [2100]
- CARVALHO, VASCO, AND ALIREZA TAHBAZ-SALEHI (2019): "Production Networks: A Primer," *Annual Review of Economics*, 11 (1), 635–663. [2092,2102,2121]
- CARVALHO, VASCO M., MAKOTO NIREI, YUKIKO U. SAITO, AND ALIREZA TAHBAZ-SALEHI (2021): "Supply Chain Disruptions: Evidence From the Great East Japan Earthquake," *The Quarterly Journal of Economics*, 136 (2), 1255–1321. [2089,2102]
- CHODOROW-REICH, GABRIEL, XAVIER GABAIX, AND RALPH KOIJEN (2022): "The Heterogeneous CES Demand and Production Functions," Working paper. [2093]
- CICCONE, ANTONIO (2002): "Input Chains and Industrialization," *The Review of Economic Studies*, 69 (3), 565–587. [2091]
- COSTELLO, ANNA M. (2020): "Credit Market Disruptions and Liquidity Spillover Effects in the Supply Chain," *Journal of Political Economy*, 128 (9), 3434–3468. [2092]
- CRUCITTI, PAOLO, VITO LATORA, MASSIMO MARCHIORI, AND ANDREA RAPISARDA (2003): "Efficiency of Scale-Free Networks: Error and Attack Tolerance," *Physica A: Statistical Mechanics and its Applications*, 320, 622–642. [2102]
- DEMARZO, PETER M., DIMITRI VAYANOS, AND JEFFREY ZWIEBEL (2003): "Persuasion Bias, Social Influence, and Unidimensional Opinions," *The Quarterly journal of economics*, 118 (3), 909–968. [2101]
- DEN HAAN, WOUTER J., AND JORIS DE WIND (2009): "How Well-Behaved Are Higher-Order Perturbation Solutions?" Working paper. [2090]
- DEW-BECKER, IAN (2023): "Supplement to 'Tail Risk in Production Networks'," *Econometrica Supplemental Material*, 91, <https://doi.org/10.3982/ECTA21064>. [2093]
- DEW-BECKER, IAN, AND ANDREA VEDOLIN (2023): "Macro Skewness and Conditional Second Moments: Evidence and Theories," Working paper. [2089,2090]
- DIEWERT, ERWIN (2010): "Understanding PPPs and PPP-Based National Accounts: Comment," *American Economic Journal: Macroeconomics*, 2 (4), 36–45. [2101]
- DUARTE, MARGARIDA, AND DIEGO RESTUCCIA (2020): "Relative Prices and Sectoral Productivity," *Journal of the European Economic Association*, 18 (3), 1400–1443. [2108]
- DUPRAZ, STEPHANE, JON STEINSSON, AND EMI NAKAMURA (2020): "A Plucking Model of Business Cycles," Working paper. [2089]
- ELLIOTT, MATTHEW, AND BENJAMIN GOLUB (2022): "Networks and Economic Fragility," *Annual Review of Economics*, 14, 665–696. [2089]
- FLYNN, JOEL P., CHRISTINA PATTERSON, AND JOHN STURM (2022): "Fiscal Policy in a Networked Economy," Working paper. [2094]
- FUJII, BRIAN CEVALLOS, DEVAKI GHOSE, AND GAURAV KHANNA (2021): "Production Networks and Firm-Level Elasticities of Substitution," Working paper. [2089]
- GABAIX, XAVIER (2011): "The Granular Origins of Aggregate Fluctuations," *Econometrica*, 79 (3), 733–772. [2089]
- GEORGESCU-ROEGEN, NICHOLAS (1966): "Some Properties of a Generalized Leontief Model," in *Analytical Economics*. Harvard University Press. [2094]
- GOFMAN, MICHAEL, GILL SEGAL, AND YOUCHANG WU (2020): "Production Networks and Stock Returns: The Role of Vertical Creative Destruction," *The Review of Financial Studies*, 33 (12), 5856–5905. [2092]
- GRANOT, DANIEL, AND GUR HUBERMAN (1981): "Minimum Cost Spanning Tree Games," *Mathematical programming*, 21 (1), 1–18. [2101]
- HERRENDORF, BERTHOLD, AND AKOS VALENTINYI (2012): "Which Sectors Make Poor Countries so Unproductive?" *Journal of the European Economic Association*, 10 (2), 323–341. [2108]
- HERSKOVIC, BERNARD, BRYAN KELLY, HANNO LUSTIG, AND STIJN VAN NIEUWERBURGH (2020): "Firm Volatility in Granular Networks," *Journal of Political Economy*, 128 (11), 4097–4162. [2102]
- HILL, ROBERT J. (1999): "Comparing Price Levels Across Countries Using Minimum-Spanning Trees," *Review of Economics and Statistics*, 81 (1), 135–142. [2101]
- (2004): "Constructing Price Indexes Across Space and Time: The Case of the European Union," *American Economic Review*, 94 (5), 1379–1410. [2101]

- HSIEH, CHANG-TAI, AND PETER J. KLENOW (2009): "Misallocation and Manufacturing TFP in China and India," *The Quarterly journal of economics*, 124 (4), 1403–1448. [2108]
- HULTEN, CHARLES R. (1978): "Growth Accounting With Intermediate Inputs," *The Review of Economic Studies*, 45 (3), 511–518. [2101]
- ILUT, COSMIN, MATTHIAS KEHRIG, AND MARTIN SCHNEIDER (2018): "Slow to Hire, Quick to Fire: Employment Dynamics With Asymmetric Responses to News," *Journal of Political Economy*, 126 (5), 2011–2071. [2089]
- JIANG, BOMIN, DANIEL E. RIGOBON, AND ROBERTO RIGOBON (2021): "From Just in Time, to Just in Case, to Just in Worst-Case: Simple Models of a Global Supply Chain Under Uncertain Aggregate Shocks," Working paper. [2101]
- JONES, CHARLES I. (2011): "Intermediate Goods and Weak Links in the Theory of Economic Development," *American Economic Journal: Macroeconomics*, 3 (2), 1–28. [2090,2091,2100,2104,2105,2108]
- JOSWIG, MICHAEL, AND BENJAMIN SCHRÖTER (2021): "Parametric Shortest-Path Algorithms via Tropical Geometry," Working paper. [2100]
- LA'O, JENNIFER, AND ALIREZA TAHBAZ-SALEHI (2021): "Optimal Monetary Policy in Production Networks." [2092]
- LIU, ERNEST (2019): "Industrial Policies in Production Networks," *The Quarterly Journal of Economics*, 134 (4), 1883–1948. [2091,2092]
- LIU, ERNEST, AND ALEH TSYVINSKI (2021): *Dynamical Structure and Spectral Properties of Input-Output Networks*. [2089,2092,2105]
- LONG JR, JOHN B., CHARLES I. PLOSSER (1983): "Real Business Cycles," *Journal of political Economy*, 91 (1), 39–69. [2092]
- MARCHIORI, MASSIMO, AND VITO LATORA (2000): "Harmony in the Small-World," *Physica A: Statistical Mechanics and its Applications*, 285 (3–4), 539–546. [2102]
- MOSSEL, ELCHANAN, ALLAN SLY, AND OMER TAMUZ (2015): "Strategic Learning and the Topology of Social Networks," *Econometrica*, 83 (5), 1755–1794. [2102]
- NEHRING, KLAUS, AND CLEMENS PUPPE (2002): "A Theory of Diversity," *Econometrica*, 70 (3), 1155–1198. [2101]
- RESNICK, SIDNEY I. (2007): *Heavy-Tail Phenomena: Probabilistic and Statistical Modeling*. Springer Science & Business Media. [2113]
- ROCHAT, YANNICK (2009): "Closeness Centrality Extended to Unconnected Graphs: The Harmonic Centrality Index," Working paper. [2102]
- RUBBO, ELISA (2020): "Networks, Phillips Curves and Monetary Policy," Working paper. [2100]
- SAMUELSON, PAUL A. (1951): "Abstract of a Theorem Concerning Substitutability in Open Leontief Models," in *Activity Analysis of Production and Allocation*, ed. by Tjalling C. Koopmans. [2094]
- SUN, NING, AND ZAIFU YANG (2014): "An Efficient and Incentive Compatible Dynamic Auction for Multiple Complements," *Journal of Political Economy*, 122 (2), 422–466. [2101]
- TASCHEREAU-DUMOUCHEL, MATHIEU (2021): "Cascades and Fluctuations in an Economy With an Endogenous Production Network," Working paper. [2092]
- VOM LEHN, CHRISTIAN, AND THOMAS WINBERRY (2021): "The Investment Network, Sectoral Comovement, and the Changing U.S. Business Cycle," Working paper. [2092]
- WEITZMAN, MARTIN L. (1992): "On Diversity," *The Quarterly Journal of Economics*, 107 (2), 363–405. [2101]

Co-editor Charles I. Jones handled this manuscript.

Manuscript received 15 August, 2022; final version accepted 18 August, 2023; available online 1 September, 2023.

The replication package for this paper is available at <https://doi.org/10.5281/zenodo.8247084>. The Journal checked the data and codes included in the package for their ability to reproduce the results in the paper and approved online appendices.