

NBER WORKING PAPER SERIES

TAIL RISK IN PRODUCTION NETWORKS

Ian Dew-Becker

Working Paper 30479

<http://www.nber.org/papers/w30479>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

September 2022

Northwestern University and NBER. This paper would not exist without Alireza Tahbaz-Salehi. I appreciate helpful comments from 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, and the Federal Reserve Banks of Chicago and Kansas City. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Ian Dew-Becker. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Tail Risk in Production Networks
Ian Dew-Becker
NBER Working Paper No. 30479
September 2022
JEL No. D24,D57,E13,E32

ABSTRACT

This paper describes the response of the economy to large shocks in a nonlinear production network. While arbitrary combinations of shocks can be studied, it focuses on a sector's tail centrality, which quantifies the effect of a large negative shock to the sector – a measure of the systemic risk of each 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. The paper then uses the results to analyze the determinants of total tail risk in the economy. Increases in interconnectedness in the presence of complementarity can simultaneously reduce the sensitivity of the economy to small shocks while increasing the sensitivity to large shocks. Tail risk is strongest in economies that display conditional granularity, where some sectors become highly influential following negative shocks.

Ian Dew-Becker
Kellogg School of Management
Northwestern University
2001 Sheridan Road
Evanston, IL 60208
and NBER
ian.dewbecker@gmail.com

Tail risk in production networks

Ian Dew-Becker*

August 18, 2022

Abstract

This paper describes the response of the economy to large shocks in a nonlinear production network. While arbitrary combinations of shocks can be studied, it focuses on a sector's *tail centrality*, which quantifies the effect of a large negative shock to the sector – a measure of the systemic risk of each 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. The paper then uses the results to analyze the determinants of total tail risk in the economy. Increases in interconnectedness in the presence of complementarity can simultaneously reduce the sensitivity of the economy to small shocks while increasing the sensitivity to large shocks. Tail risk is strongest in economies that display *conditional granularity*, where some sectors become highly influential following negative shocks.

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, e.g. due to shutdowns at major ports and constraints on trucking, have propagated through the global

*Northwestern University and NBER. This paper would not exist without Alireza Tahbaz-Salehi. I appreciate helpful comments from Nicolas Cruzet, 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, and Caltech.

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 et al. (2017)), 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 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?

The paper’s core contribution is to answer those questions in the context of a general network production model. Its central theoretical tool is a 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

¹Empirically, Barrot and Sauvagnat (2016) and Carvalho et al. (2020) study the effects of large shocks to individual firms due to natural disasters on production. See also related work by Fujiy, Ghose, and Khanna (2021). Liu and Tsyvinski (2021) study the dynamic effects of large shocks in a linear setting.

²For recent models, see Dew-Becker, Tahbaz-Salehi, and Vedolin (2021), Dupraz, Nakamura, and Steinsson (2020), Ilut, Kehrig, and Schneider (2018). Those papers also discuss empirical evidence.

³Jones (2011) and Dew-Becker and Vedolin (2021) 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. That said, in the class of models studied in this paper, the Taylor series has a finite radius of convergence. Outside that radius – i.e. for sufficiently large shocks – it is meaningless as an approximation. See also den Haan and de Wind (2009).

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.

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, e.g. 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 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.

Similarly, the analysis shows that it is the *topology* of the production network, rather than its geometry, 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

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

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 percent 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.

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 exacerbates asymmetry in the distribution of GDP growth: increases in complementarity and in connections running through complementary sectors both create left tail risk.

Second, this section examines the model’s implications for the distribution of GDP under specific distributional assumptions that have appeared in the literature. In a broad range of cases capturing most work, tail risk is determined by the magnitude of a single worst-

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

case scenario. For the case of i.i.d. shocks, that actually becomes particularly simple: what matters is the largest Domar weight (sales share) that any sector can attain for any combination of shocks. That is, whatever sector has the ability to generate the most systemic risk ultimately represents a sufficient statistic for the total tail risk in GDP.

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. As a specific example, in a fully connected and symmetrical network with N sectors, the average transmission of each sector’s shocks to GDP is of order N^{-1} for small shocks, but of order 1 for large shocks. There is no granularity near steady-state, but severe granularity in tail events.

For a more realistic 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,⁶ 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.

Additional related literature

The paper’s framework builds most directly on the literature on production networks, going back to Long and Plosser (1983).⁷ 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.⁸ A contribution

⁶Consumer spending on food services and accommodations fell by 40 percent, or \$403 billion between 2019Q4 and 2020Q2. Spending at movie theaters fell by 99 percent.

⁷That literature is large and work has studied features of networks, e.g. 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).

⁸Liu (2019), Bigio and La’O (2020), and Boehm and Oberfield (2020) study the propagation of distortions in production networks. Costello (2020) and Alfaro, Garcia-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.

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 et al. (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. Section 3 presents the result on approximating output in terms of the exogenous shocks. Sections 4 and 5 analyze the determinants of the tail centrality of individual sectors, while section 6 examines it in the data. Finally, section 7 presents results on the probability of extreme realizations of GDP and section 8 concludes.

2 Structure of the economy

The model is static and frictionless and takes the form of a standard 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, or 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

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

where 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 $1, \dots, N$).⁹

⁹The 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

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 .

$1 - \alpha$ represents labor’s share of income. It is easy to relax the model to allow that to vary across sectors (as it does empirically).

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 I assume a CES specification for simplicity, Appendix D.4 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 Appendix D.4.1 discusses a more general setup from Chodorow-Reich et al. (2022).¹⁰

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

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

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.¹¹

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, e.g. $z_i = \log Z_i$. 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 (Appendix D.5), and allowing for an upward

combined with other inputs with a non-unitary elasticity

¹⁰An example of a model in which the paper’s results do not hold is one where labor cannot be reallocated across sectors and it has an elasticity of substitution with material inputs smaller than 1 (such a model does not have a solution for all levels of productivity).

¹¹One can always add a sector with a non-unitary elasticity of substitution that produces a single final good, with $\beta = 1$ for that sector and equal to zero for all other sectors.

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: $Y_i = C_i + \sum_j X_{j,i}$.*

Since 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

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

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: $p = -(I - \alpha A)^{-1} z$, where A is a matrix collecting the $A_{i,j}$ coefficients, and z is the vector of log productivities.

Equation (3) implies prices do not depend on demand, a “no-substitution” type result (see, e.g., Acemoglu and Azar (2020) and Flynn, Patterson, and Sturm (2022)). Given a solution for p (as a function of z), utility maximization for the consumer yields,

$$gdp = -\beta' p \quad (4)$$

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 nonlinear 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 is a scalar determining magnitude. As examples, $\theta = [\dots, 0, 1, 0, \dots]$ represents a shock to a single sector, while $\theta = [1, 1, \dots] / \sqrt{N}$ represents a common shock to all sectors.

3.1 The large shock limit

Lemma 1. *As $t \rightarrow \infty$, 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. 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 ($\sigma_i > 1$).

Figure 1: Network examples



Notes: The nodes represent sectors and arrows flows of goods. The red nodes and arrows represent a hypothetical tail network following a shock to the darkest red 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 $\sigma < 1$.

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 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. Figure 1 displays four hypothetical networks, with the tail networks for the special case of a shock to a single sector (darkest red) denoted by the red arrows.

An elasticity of substitution less than 1 means that when an input's price rises, its share of expenditures rises, while an elasticity above one means that the share falls. The source of the result in (7) is that in the limit as $t \rightarrow \infty$, expenditure shares are ultimately driven to 0 or 1, depending on the elasticity and the shock, with $\sigma_i = 1$ being the knife-edge case with constant expenditure shares.¹²

There is also a simple recursion for $\mu(\theta)$, which depends on A and σ , but for this paper's

¹²Mathematically, the result comes from the log-sum-exponential that appears in the recursion. Using

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 Appendix D.1), but the remainder of the paper focuses on aggregate output.

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,*

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

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

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

gdp 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, for $t < 0$, formally corresponds to reversing the sign of θ – i.e. 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.


Appendices D.2 and D.3 discuss asymptotic convergence. GDP approaches its asymptote exponentially, so there are no higher-order polynomial terms that appear. When elasticities

$p_i \sim \phi_i t$ asymptotically,

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

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.

Figure 2: Linear, quadratic, and tail approximations



(a) Small-shock approximations

(b) Large-shock approximation

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.

of substitution are closer to 1 or when the units whose shocks are relevant in the tail (in the sense of being the argument of the max/min in (7) for some sector) have smaller production weights, the tail approximation will tend to be less accurate for a given t . 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.

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.

In terms of networks, the result says that what matters in the tail is the topology of the network – the set of edges – rather than the geometry – their weights. In other words, 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 GDP.

Definition. *The left and right tail centralities of unit i are, respectively,*

$$\gamma_i^L \equiv \lim_{\Delta z_i \rightarrow -\infty} \frac{\Delta gdp}{\Delta z_i} \quad (11)$$

$$\gamma_i^R \equiv \lim_{\Delta z_i \rightarrow \infty} \frac{\Delta gdp}{\Delta z_i} \quad (12)$$

where Δ denotes a deviation from steady-state ($z_i = 0 \forall i$)

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

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

$$\gamma_i^L = -\lambda(-e_i) \quad (13)$$

$$\gamma_i^R = \lambda(e_i) \quad (14)$$

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 need not 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 gives greater opportunity to use the output of relatively productive sectors, while complementarity *requires* using all inputs, including the weakest. Since productivity shocks propagate downstream, those are the only elasticities that matter.

Second, 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 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 where as 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 and examples

5.1 Fully complementary production: average closeness

Consider an economy in which $\sigma_i < 1$ for all i (for example, the calibrations of Jones (2011), Baqaee and Farhi (2019), and Rubbo (2020)).¹³ Assume also that $A_{i,i} \in (0, 1)$, which guarantees that every unit uses at least two inputs, one of which is its own output (which is true for 88 percent of sectors according to the BEA). The $A_{i,j}$ are otherwise unconstrained.

Using Theorem 1, it is then immediate that $\gamma_i^R = \beta_i \forall i$. When a unit gets a sufficiently positive shock, it eventually has no downstream impacts, affecting GDP only through its direct effect on consumption.

While positive unit shocks eventually die out, negative unit shocks propagate, since production in all units is complementary. That implies that

$$\gamma_i^L = \frac{1}{1 - \alpha} \sum_{j=1}^n \beta_j \alpha^{d_{\min}(j,i)} \quad (15)$$

where $d_{\min}(j, i)$ is the length of the shortest downstream path from i to j .¹⁴ In the complementary economy, a unit's left tail centrality is measured by its average downstream closeness to final consumption: γ_i^L involves the sum across units of each unit's consumption weight times a term, $\alpha^{d_{\min}(j,i)}$, that decreases in the number of upstream steps from that unit back to i .

Equation (15) answers the question of what types of units have high tail centrality 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 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

¹³See also evidence in Atalay (2017) and Atalay et al. (2018), among others.

¹⁴I.e. 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 .

3. An increase in the share of expenditures on material inputs, α .

On the other side, $\gamma_i^R = \beta_i \forall i$. 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 .

5.1.1 The tail network

Under fully complementary production it is possible to give a fuller description of the tail network that was discussed in section 3.1. At any given productivity, 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 ($z = 0$),

$$D'_{ss} \equiv \beta' (I - \alpha A)^{-1} \tag{16}$$

Proposition 3. *If $\sigma_i \leq 1$ for all i , there is a finite set of $N \times 1$ vectors of asymptotic Domar weights, D_k , such that*

$$\lambda(\theta) = \min_k D'_k \theta \tag{17}$$

It immediately follows that GDP is concave in that $\lambda(\theta) > 0 \iff \lambda(\theta) \leq -\lambda(-\theta)$

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 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.¹⁵ That is, each D'_k is of the form $\beta' (I - \alpha M_k)^{-1}$, where M_k is a matrix representing a particular tail network.

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.¹⁶

Proposition 3 immediately yields an alternative description of tail centrality:

¹⁵The minimization here is reminiscent of the worst-case network analysis in Jiang, Rigobon, and Rigobon (2021).

¹⁶In DeMarzo, Vayanos, and Zweibel’s (2003) 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,

Corollary 2. *In the fully complementary economy,*

$$\gamma_i^L = \max_k D_{k,i} \tag{18}$$

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.

Panels (a)-(c) of Figure 1 plot three hypothetical production networks for $\sigma_i < 1 \forall i$. The arrows represent downstream links. The darkest red nodes are the shocked sectors, and the red arrows represent the tail networks given those shocks. Lighter red nodes are downstream of the shocked sectors.

5.1.2 Relationship with other centrality measures

The idea of measuring centrality via average closeness appears elsewhere in the networks literature in the form of harmonic centrality, which is an unweighted average closeness.¹⁷ The concept of the *efficiency* of a network is then measured by the average closeness between pairs of nodes (Marchiori and Latora (2000) and Crucitti et al. (2003)). In the context of complementary production, a network with greater efficiency then also has more tail risk (this is formalized further in section 7).

The difference between average closeness and the usual Bonacich (1987) centrality that appears in a Cobb–Douglas economy 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)).¹⁸

again due to the choice of shortest paths (i.e. the tail network).

The concept of spanning trees (in some cases minimal ones) appears 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)).

¹⁷See Boldi and Vigna (2014) who justify it axiomatically, along with Rochat (2009) and Bloch, Jackson, and Tebaldi (2021)

¹⁸These results also suggest that there might be a relationship with the concept of upstreamness studied in Antras and Chor (2013) and Antras et al. (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

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

$$\text{deg}_i \equiv \sum_{j:i \in S(j)} \beta_j \quad (19)$$

Proposition 4. *Under fully complementary production, left tail centrality satisfies*

$$\frac{1}{1-\alpha} (\beta_i + \alpha \text{deg}_i) \leq \gamma_i^L \leq \frac{1}{1-\alpha} (\beta_i + \alpha \text{deg}_i + \alpha^2 (1 - \text{deg}_i)) \quad (20)$$

Weighted out-degree thus gives upper and lower bounds for tail centrality.¹⁹

5.1.3 Example: fully connected economy

Example 1. *Suppose $\sigma_i < 1$ for all i and every sector uses inputs from itself and every other sector (i.e. $A_{i,j} > 0 \forall i, j$). Then*

$$\phi_i = \theta_i + \frac{\alpha}{1-\alpha} \theta_{\min} \quad (21)$$

$$\lambda(\theta) = \beta' \theta + \frac{\alpha}{1-\alpha} \theta_{\min} \quad (22)$$

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

In the case of a fully connected production network, each sector's ϕ_i 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 the 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 any 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 weight $\beta_i + \alpha / (1 - \alpha)$ when the scale of the shocks, t , is large.

centrality is not just the *composition* of its sales, but also the fraction of final users that it sells to.

¹⁹Out-degree appears frequently in the networks literature, including, recently, Carvalho et al. (2021), Herskovic et al. (2020), Bernard, Moxnes, and Saito (2019), and Mossel, Sly, and Tamuz (2015) among many others.

This again illustrates the idea of conditional granularity. Even if no sector is granular (for large N) 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 Allowing for 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

$$Y_i = \left(X_{i,i'}^{(\sigma_i-1)/\sigma_i} + X_{i,i''}^{(\sigma_i-1)/\sigma_i} \right)^{\sigma_i/(\sigma_i-1)} \quad (23)$$

with $\sigma_i > 1$.²⁰ 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 each individually has a left tail centrality of zero.²¹ 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 the result above. Define $d_{\min}^{-i}(j, i')$ to be the length of the shortest upstream path from j

²⁰Formally, this requires allowing for differential α_i across the production functions in the baseline setup. That is a straightforward extension.

²¹Again, $\sigma_i > 1$ 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 towards steel (regardless of whether total expenditures on iron and steel combined rise or fall).

to i' that does not go through good i . Then

$$\gamma_i^L = \frac{1}{1 - \alpha} \sum_{j=1}^n \beta_j \alpha^{d_{\min}^{-i}(j,i')} \quad (24)$$

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.

6 Tail centrality in the data

This section examines two aspects of tail centralities in the data. First, it gives a simple first-pass estimate of tail centralities and compares them to sales shares in recent data. 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.

6.1 Estimating tail centralities


I study the most recent (2012) sector detail input-output tables reported by the BEA. The tables have 379 private sectors.²² Define an $A_{i,j}$ coefficient to be positive, so that there is an upstream link, if sector i spends at least 0.5 percent of its expenditures on materials for the output of sector j . I set $\alpha = 1/2$ in calculating γ_i^L . The β_i parameters are calculated from the fraction of nominal final expenditure going to each sector.

Figure 3 plots Domar weights (nominal output divided by nominal GDP) against left tail centralities. There is a weak positive correlation of 0.23, 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 centralities close to 0.5. 13 sectors have $\gamma_i^L > 0.8 \max(\gamma_i^L)$, while only two have $D_{ss,i} > 0.8 \max(D_{ss,i})$.

²²For 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.

I exclude customs duties, funds and trusts, real estate sectors, management services, and employment services. Management services are almost entirely offices of holding companies, while employment services represent staffing agencies, so we take both as representing more appropriately relatively generic labor input.

Figure 3: Domar weights and tail centralities



Notes: The x-axis is the Domar weight of each sector. The y-axis is the left tail centrality. The data is the 2012 BEA input-output table. The top four sectors according to both centrality measures are labeled.

One can also see that the top sectors by Domar weight have very different tail centralities – Petroleum Refineries at 0.40, Oil and Gas Extraction at 0.25, and Hospitals at 0.07. Oil and Gas Extraction is lower because it is one more step up the supply chain from refineries. Hospitals are low because they produce essentially only final output.

Table 1 further examines the top sectors sorted by sales and tail centrality. The top sectors for tail centrality are all universal inputs. The first is electricity, which is why it has appeared frequently as an example. The second highest tail centrality is for trucking services – all of final production involves trucking at some phase.

The third-highest tail centrality is for legal services – again, simply because every sector purchases legal services. Does it make sense to claim that a large negative shock to the legal services sector could cause a crash in GDP? There is ample evidence that legal institutions are necessary for the growth of the economy. All aspects of business rely on property rights

Table 1: **Top sectors by left tail centrality and sales share**

Largest by left tail centrality		
<i>Sector</i>	γ_i^L	<i>Sales share</i>
Electric power generation, transmission, and distribution	0.4714	0.0324
Truck transportation	0.4578	0.0211
Legal services	0.4428	0.0224
Advertising, public relations, and related services	0.4404	0.0242
Accounting, tax preparation, bookkeeping, and payroll services	0.4282	0.0112
Services to buildings and dwellings	0.4186	0.0095
Monetary authorities and depository credit intermediation	0.4101	0.0304
Wired telecommunications carriers	0.4049	0.0231
Other nondurable goods merchant wholesalers	0.4040	0.0220
Insurance carriers, except direct life	0.4004	0.0340
Petroleum refineries	0.3997	0.0529
Largest by sales share		
<i>Sector</i>	γ_i^L	<i>Sales share</i>
Hospitals	0.0652	0.0652
Petroleum refineries	0.3997	0.0529
Oil and gas extraction	0.2515	0.0465
Insurance carriers, except direct life	0.4004	0.0340
Electric power generation, transmission, and distribution	0.4714	0.0324
Offices of physicians	0.0316	0.0316
Monetary authorities and depository credit intermediation	0.4101	0.0304
Scientific research and development services	0.0263	0.0269
Other financial investment activities	0.2810	0.0247
Advertising, public relations, and related services	0.4404	0.0242
Wired telecommunications carriers	0.4049	0.0231

Notes: Sales shares and tail centralities calculated from the 2012 BEA input-output tables.

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.

One potential concern with that argument 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 (external) lawyers by firms. That actually illustrates a key advantage of

γ_i^L : measuring it does not require measuring *all* of each sector’s expenditures on each input. All that we need to know is that a sector uses some input. 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 essentially 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.

The bottom section of table 1 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.

6.2 Hospitals and computers

Two prominent sectors that have undergone significant changes in the post-war period are computer equipment and hospitals.²³ The left-hand panel of Figure 4 plots their Domar weights for the period 1963–2020. The Domar weight of hospitals rose by a factor of 5 from 0.02 to 0.10. 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, then, 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 4 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 cutoff is 0.5% of the using sector’s intermediates). Computers, on the other hand, rose from being purchased by 30 to 55 percent of sectors.

²³The computer equipment sector stays consistent in the BEA input-output tables between the 1963–1996 and 1997–2020 versions. For consistency across those two datasets, I use the combined “Hospitals and Nursing and Residential Care” sector.

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 centrality, using the detailed input-output tables as above, the tail centrality 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 centrality of hospitals is always simply equal to its share of final consumption, which is also its Domar weight.

Figure 4: 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 percent 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 $A_{i,j}$).

7 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 first gives a general result and

then studies the behavior of the economy under various specific assumptions that have been used for shocks in the literature.

The key results 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. When $\sigma_i \leq 1 \forall i$, 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.

Additionally, this section generalizes well-known results on extreme realizations of sums of random variables (e.g. Nair, Wierman, and Zwart (2022) and Embrechts, Kluppelberg, and Mikosch (2013), among many others) to a nonlinear setting.

7.1 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 $t > \bar{t}$,

$$\Pr[\theta \in \Theta, t/s(\theta) > x] = m(\Theta) \bar{F}(x) \tag{25}$$

The representation in (25) 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)). Specific parametric examples are studied below. A simple example of a distribution that does not have a representation (25) 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 .²⁴

7.2 General result

Theorem 2. *Given the distribution for z in (25), there exists a function $\varepsilon(x) \geq 0$ with $\lim_{x \rightarrow \infty} \varepsilon(x) = 0$ and an \bar{x} such that for $x > \bar{x}$*

$$\int_{\Theta_-} \bar{F} \left(\frac{x - \mu(\theta) + \varepsilon(x)}{-s(\theta) \lambda(\theta)} \right) dm(\theta) \leq \Pr[gdp < -x] \leq \int_{\Theta_-} \bar{F} \left(\frac{x - \mu(\theta) - \varepsilon(x)}{-s(\theta) \lambda(\theta)} \right) dm(\theta) \quad (26)$$

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

Theorem 2 says that the CDF of log *GDP* is well approximated by

$$\int_{\Theta_-} \bar{F} \left(\frac{x - \mu(\theta)}{-s(\theta) \lambda(\theta)} \right) dm(\theta) \quad (27)$$

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 .

7.2.1 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

²⁴For 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.

(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 – 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 5. *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 $\lambda(\theta) > 0 \iff -\lambda(-\theta) \geq \lambda(\theta)$, $\Pr[gdp < -x] \geq \Pr[gdp > x]$ 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 (26) 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.

7.3 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,

$$\text{var}(\log GDP) \approx D'_{ss} \Sigma D_{ss} \quad (28)$$

Since $D'_{ss} \Sigma D_{ss}$ is continuous in A , any small change in A – i.e. 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 (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, the results above show 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.

7.4 Specific distributions

This section specializes the result in Theorem 2 to specific distributions that appear in the literature. Beyond the specific results, three patterns emerge: (1) for independent shocks

tail risk depends on sector tail centralities (crashes are not caused by simultaneous shocks to multiple sectors); (2) as the tails of the shocks become heavier, more sectors can potentially cause crashes; (3) tail risk arises in the presence of conditional granularity.

7.4.1 Weibull tails

A wide range of distributions, including the normal, gamma, exponential, Gumbel, and Fréchet families all can be said to have Weibull tails, up to asymptotically negligible terms, in the following sense:

Definition. *The shocks have a Weibull-type tail if, for $t > \bar{t}$,*

$$\begin{aligned} \bar{F}(t) &= c \exp(-\eta(t - \bar{t})^\kappa) \\ \text{where } c &= \Pr(t \leq \bar{t}) \end{aligned} \tag{29}$$

for parameters $\kappa > 0$ and $\eta > 0$.

Smaller κ represents heavier tails, with $\kappa = 1$ corresponding to the exponential distribution and $\kappa = 2$ to the normal.²⁵ In addition, all three types of extreme value distributions (Weibull, Gumbel, and Fréchet) have Weibull-type tails. The Weibull family thus covers a broad range of behaviors, including all but the very lightest (e.g. bounded) and very heaviest (Pareto or Cauchy) tails and, as discussed below, has been used prominently in the literature.

Across that entire family, we have a surprisingly simple result. Denote the essential supremum with respect to the measure m over θ of any function $f(\theta)$ by $\|f(\theta)\|_\infty$.²⁶ For example, in the typical case where m has full support, $\|f(\theta)\|_\infty = \max_\theta f(\theta)$ (note that it is *not* the maximum of $|f(\theta)|$). $\|f(\theta)\|_{\infty; \Theta^*}$ denotes the essential supremum on some subset of the sphere Θ^* .

Proposition 6. *If the shocks have Weibull tails,*

$$\lim_{x \rightarrow \infty} \Pr[gdp < -x]^{1/(x^\kappa)} = \exp\left(-\eta \left(\frac{1}{\| -s(\theta) \lambda(\theta) \|_\infty}\right)^\kappa\right) \tag{31}$$

²⁵For the normal distribution, a better approximation for $\bar{F}(t)$ is $c(t - \bar{t})^{-1} \exp(-\eta(t - \bar{t})^2)$. The t^{-1} term is asymptotically dominated by the exponential, and it is straightforward to show that the proposition in this section also holds when any powers of t multiply the exponential in \bar{F} .

²⁶Formally, $\|f(\theta)\|_\infty = \inf\{a \in \mathbb{R} : m(\{\theta : f(\theta) > a\}) = 0\}$.

Furthermore, for any set Θ^* such that $\| -s(\theta) \lambda(\theta) \|_{\infty; \Theta^*} < \| -s(\theta) \lambda(\theta) \|_{\infty}$,

$$\lim_{x \rightarrow \infty} \Pr[\theta \in \Theta^* \mid gdp < -x] = 0 \tag{32}$$

Analogous results hold for $\Pr[gdp > x]$.

Conditional on the distribution of the shocks, the probability that GDP has an extreme decline is determined by a sufficient statistic: the most negative value of $s(\theta) \lambda(\theta)$. That function combines the scale of the shock, $s(\theta)$, and its effects, $\lambda(\theta)$, to determine which θ has the single strongest impact on GDP.²⁷ The exponential form of the distribution for t is what causes only the single most extreme shock to end up mattering, because \bar{F} is effectively infinitely convex as $x \rightarrow \infty$.

The second part of the result says that extreme realizations are driven, in probability, by *only* that single most extreme shock. Any θ with $-s(\theta) \lambda(\theta) < \| -s(\theta) \lambda(\theta) \|_{\infty}$ causes a crash with probability zero asymptotically. That is, when looking across declines in gdp , as those declines become larger, the probability that they are associated with a θ in an arbitrarily small radius around the argmax approaches 1.

The same comparative statics results continue to hold as above, though now really they only matter if they affect $\| -s(\theta) \lambda(\theta) \|_{\infty}$. That is, any change in the model (the elasticities or the inputs used by sectors) that does not affect that supremum does not affect tail risk.

Recall the importance of the results on invariance in Lemma 1 and Theorem 1. Here we have another form of invariance: in the Weibull family, tail risk does not depend on the measure $m(\theta)$ other than through its support, and the combination of shocks that leads to extreme events is invariant to the tail shape parameter κ .

Finally, note that the result in Proposition 6 is a generalization of the well-known fact for sub-exponential distributions (in fact, their definition) that, as $x \rightarrow \infty$, the probability that a sum of i.i.d. subexponentials is greater than x is equal to the probability that their maximum is greater than x . Here instead of a sum we have a nonlinear mixture.

Tail skewness In the Weibull case we can significantly sharpen the conditions for tail asymmetry relative to the general case:

²⁷Under knife-edge conditions, multiple θ 's might achieve that maximum

Corollary 4. *When the shocks have Weibull tails, the following necessary and sufficient conditions for tail asymmetry hold:*

$$\| -s(\theta) \lambda(\theta) \|_\infty > \| s(\theta) \lambda(\theta) \|_\infty \Leftrightarrow \lim_{x \rightarrow \infty} \frac{\Pr [gdp < -x]}{\Pr [gdp > x]} = \infty \quad (33)$$

$$\| -s(\theta) \lambda(\theta) \|_\infty < \| s(\theta) \lambda(\theta) \|_\infty \Leftrightarrow \lim_{x \rightarrow \infty} \frac{\Pr [gdp < -x]}{\Pr [gdp > x]} = 0 \quad (34)$$

When s is symmetrical, a sufficient condition for the first case is $\sigma_i < 1$ for all i , while a sufficient condition for the second case is $\sigma_i > 1$ for all i .

The condition $\| -s(\theta) \lambda(\theta) \|_\infty > \| s(\theta) \lambda(\theta) \|_\infty$ says that the largest negative effects of shocks are larger in magnitude than the largest positive effects. If that is true log GDP is skewed left. When the opposite condition holds, it is skewed right. As in the corollary in section 7.2, a sufficient condition for left tail skewness (in the limiting sense) is that production is complementary, while if production displays substitutability – $\sigma_i > 1$ – then there is right tail skewness. In other words, for a complementary economy with Weibull-type shocks, large booms in GDP are infinitely rare compared to large declines.

In the case of the next two specific examples from the literature, it is possible to further characterize the tail distribution of GDP.

Exponential tailed shocks and the maximum Domar weight Acemoglu, Ozdaglar, and Tahbaz-Salehi (2017) study a model with Cobb–Douglas production and in which shocks are i.i.d. with exponential tails and show that what determines tail risk is the largest Domar weight. This section generalizes that result to a case with complementary production.

Example 3. *Suppose the shocks are i.i.d. with exponential tails so that $s(\theta) = 1/\|x\|_1$ and $m(\theta)$ again has full support. If $\sigma_i \leq 1$ for all i , then*

$$\| -s(\theta) \lambda(\theta) \|_\infty = \max_n \max_j D_{n,j} = \max_j \gamma_j^L \quad (35)$$

$$\Pr [gdp < -x] \rightarrow \exp \left(-\eta \frac{x}{\max_j \gamma_j^L} \right) \quad (36)$$

where $D_{n,j}$ is the j th element of the vector D_n . The shock θ causing the tail event is equal to 1 for the sector with the largest γ_j^L and zero elsewhere.

We thus continue to get the same result as in Acemoglu, Ozdaglar, and Tahbaz-Salehi (2017) that what matters for tail risk with exponential shocks is a maximum Domar weight, but now it is the maximum Domar weight among all possible tail networks.²⁸ 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 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 centrality 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 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 Appendix C.1).

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.

Gaussian shocks A typical starting choice for the distribution of the shocks is normality (e.g. Baqaee and Farhi (2019)).

Example 4. *Suppose $z \sim N(0, \Sigma)$. If $\sigma_i \leq 1$ for all i , then the left tail of GDP is determined*

²⁸In 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.

by

$$\| -s(\theta) \lambda(\theta) \|_\infty = \max_n \sqrt{D'_n \Sigma D_n} \quad (37)$$

$$\Pr [gdp < -x] \rightarrow c \exp \left(-\frac{1}{2} \frac{x^2}{\max_n D'_n \Sigma D_n} \right) \quad (38)$$

as $x \rightarrow \infty$. Denoting the argument for the maximum in (37) by n^* , the shock causing left tail events is $\theta \propto -\Sigma D_{n^*}$.

As above, the tail distribution of GDP remains Gaussian. Compared to the local variance, $D'_{ss} \Sigma D_{ss}$, the left tail probabilities are as though the variance is $\max_n D'_n \Sigma D_n$ – the largest variance of any feasible network.²⁹

Continuing the example of a fully connected network from section 5.1.3:

Example 5. Take the fully connected network with $\sigma_i < 1$ and suppose it is also symmetrical, with $A_{i,j} = \beta_i = 1/N \forall i, j$.³⁰ Then for i.i.d. normal shocks with variance σ^2 ,

1. (Local to steady-state): $D_{ss,i} = N^{-1}/(1 - \alpha)$ and $\text{std}(gdp) \approx N^{-1/2} \sigma / (1 - \alpha)$
2. (Right tail): $\|s(\theta) \lambda(\theta)\|_\infty = N^{-1/2} \sigma / (1 - \alpha)$, which is attained at $\theta_i = N^{-1/2}$ for all i
3. (Left tail): $\| -s(\theta) \lambda(\theta) \|_\infty = \sigma \alpha / (1 - \alpha) + O(N^{-1/2})$, which is attained when $\theta_i \propto N^{-1}$ for all i except a single value, where it is proportional to $N^{-1} + \alpha / (1 - \alpha)$.

The example shows both how diversification works in the model, and also how it differs in the tails. Local to steady-state, i.e. where the model is well described by a first order approximation, output has a standard deviation of order $N^{-1/2}$, and the right tail decays in the same way – i.e. the right-tail is no heavier than would be expected given the variance of small shocks. The shock that causes right tail events is one where all sectors have an equal increase in productivity, which is because of the dependence of output on θ_{\min} . Booms occur when all sectors simultaneously receive positive shocks.

²⁹Note also that in the special case of a linear model ($\sigma_i = 0$ for all i), the tail approximation yields the correct result that the effective variance in both the left and right tails is $D'_{ss} \Sigma D_{ss}$.

As an example, suppose $gdp = \sum_j z_j$ and the z_j are i.i.d. normal. Then the results here say that extreme realizations of GDP are due to $\theta = [1, 1, \dots] N^{-1/2}$. That result in fact holds for Weibull tails more generally as long as $\kappa > 1$. See Nair, Wierman, and Zwart (2020), Proposition 3.1.

³⁰The added perfect symmetry appears in, for example, Jones (2011), and only matters for the steady-state Domar weights.

On the other hand, for output to fall significantly productivity only needs to fall in a single sector. So in this case as N grows there is actually *no* diversification in the tail at all. The left-tail probabilities are as though the volatility of output is simply $\sigma\alpha/(1-\alpha)$. Booms are rare because productivity rising in all sectors simultaneously is relatively unlikely. But productivity falling in one sector is not surprising at all, making crashes much more likely than would be expected either from the local variance or the behavior of the right tail.

The mechanism here is again similar to that of the previous section and Acemoglu, Ozdaglar, and Tahbaz-Salehi (2017), as well as Gabaix (2011) and Acemoglu et al. (2012). Local to the steady-state and in the right tail, the Domar weights are proportional to N^{-1} and each sector is equally important. But in the left tail, the sector with the most negative shock has a weight of $N^{-1} + \alpha/(1-\alpha) = O(1)$. While the model is structurally symmetrical, so that all sectors on average carry the same weight and have the same size, following an extreme shock any given sector can become large and have a major impact on the economy.

7.4.2 Power law tails

In the international trade literature, it is common to assume that productivity has a Pareto distribution.³¹ Because its tail decays polynomially instead of exponentially, the Pareto distribution generates shocks that can be much more extreme than in the Weibull case.

Proposition 7. *Suppose t has a power law tail, $\bar{F}(t) = c(t/\bar{t})^{-\kappa}$, for some c . Then*

1.

$$\lim_{x \rightarrow \infty} \Pr[gdp < -x] / (c\bar{t}^\kappa x^{-\kappa}) = \int_{\Theta_-} (-\lambda(\theta))^\kappa dm(\theta) \quad (39)$$

2.

$$\lim_{x \rightarrow \infty} \Pr[\theta \in \Theta^* \mid gdp < -x] = \frac{\int_{\Theta_-^*} (-\lambda(\theta))^\kappa dm(\theta)}{\int_{\Theta_-} (-\lambda(\theta))^\kappa dm(\theta)} \quad (40)$$

Equation (39) gives two results for the tail of GDP in the case of shocks with power law tails. First, GDP has, in the limit, a power law tail with the same decay rate as the shocks, κ .

Second, the probability of a large deviation in gdp now depends on an average (with respect to the measure m) across all possible shocks of the tail slope, $\lambda(\theta)$. This is because

³¹For discussions, see, for example, Head, Mayer, and Thoenig (2014), Melitz and Redding (2015), Arkolakis et al. (2019), and Allen, Arkolakis, and Takahashi (2020).

even a shock with relatively low impact – small $\lambda(\theta)$ – can still potentially cause a crash simply because the tails are so long. For small κ – a heavier tail – the average gets closer to being equally weighted, while for $\kappa \rightarrow \infty$, the integral depends just on the largest value of $\lambda(\theta)$, consistent with the results for Weibull tails.

The fact that all crashes are, asymptotically, caused by the same shock in the Weibull case is potentially intuitively unappealing. The results here show that that result can be broken, but it requires allowing for the shocks to have extremely heavy tails.

A special case within the general Pareto tail is if the unit shocks are i.i.d.. Then the measure $m(\theta)$ puts mass only on the axes – $\Pr[\theta = e_i] = 1/N$, and $\Pr[\theta \in \Theta^*] = 0$ for any set Θ^* that does not contain one of the unit vectors $e_i = [\dots, 0, 1, 0, \dots]$ (Resnick (2007), section 6.5.1). In that case, the result for GDP specializes to

$$\lim_{x \rightarrow \infty} \Pr[gdp < -x] / (c\bar{c}^\kappa x^{-\kappa}) = N^{-1} \sum_i (-\gamma_i^L)^\kappa \quad (41)$$

Intuitively, when the shocks are i.i.d. with Pareto tails, the probability of two sectors receiving a large negative shock simultaneously is negligible, so tail risk is determined purely by the set of tail centralities. Anything that increases the tail centralities also increases tail risk.

8 Conclusion

This paper studies large deviations in GDP in the context of a general nonlinear network production model. Its core result is to characterize 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.

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 generally, 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.

References

- Acemoglu, Daron and Pablo D Azar**, “Endogenous production networks,” *Econometrica*, 2020, 88 (1), 33–82.
- , **Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi**, “The network origins of aggregate fluctuations,” *Econometrica*, 2012, 80 (5), 1977–2016.
- Alfaro, Laura, Davin Chor, Pol Antras, and Paola Conconi**, “Internalizing global value chains: A firm-level analysis,” *Journal of Political Economy*, 2019, 127 (2), 508–559.
- , **Manuel García-Santana, and Enrique Moral-Benito**, “On the direct and indirect real effects of credit supply shocks,” *Journal of Financial Economics*, 2021, 139 (3), 895–921.
- Allen, Treb, Costas Arkolakis, and Yuta Takahashi**, “Universal gravity,” *Journal of Political Economy*, 2020, 128 (2), 393–433.
- Antràs, Pol and Davin Chor**, “Organizing the global value chain,” *Econometrica*, 2013, 81 (6), 2127–2204.
- , —, **Thibault Fally, and Russell Hillberry**, “Measuring the upstreamness of production and trade flows,” *American Economic Review*, 2012, 102 (3), 412–16.
- Arkolakis, Costas, Arnaud Costinot, Dave Donaldson, and Andrés Rodríguez-Clare**, “The elusive pro-competitive effects of trade,” *The Review of Economic Studies*,

- 2019, *86* (1), 46–80.
- Atalay, Englin**, “How important are sectoral shocks?,” *American Economic Journal: Macroeconomics*, 2017, *9* (4), 254–80.
- , **Thorsten Drautzburg**, and **Zhenting Wang**, “Accounting for the sources of macroeconomic tail risks,” *Economics Letters*, 2018, *165*, 65–69.
- Baqae, David Rezza and Emmanuel Farhi**, “The macroeconomic impact of microeconomic shocks: beyond Hulten’s Theorem,” *Econometrica*, 2019, *87* (4), 1155–1203.
- Barrot, Jean-Noël and Julien Sauvagnat**, “Input specificity and the propagation of idiosyncratic shocks in production networks,” *The Quarterly Journal of Economics*, 2016, *131* (3), 1543–1592.
- Bernard, Andrew B, Andreas Moxnes, and Yukiko U Saito**, “Production networks, geography, and firm performance,” *Journal of Political Economy*, 2019, *127* (2), 639–688.
- Bigio, Saki and Jennifer La’o**, “Distortions in production networks,” *The Quarterly Journal of Economics*, 2020, *135* (4), 2187–2253.
- Bloch, Francis, Matthew O. Jackson, and Pietro Tebaldi**, “Centrality Measures in Networks,” 2021. Working paper.
- Boehm, Johannes and Ezra Oberfield**, “Misallocation in the Market for Inputs: Enforcement and the Organization of Production,” *The Quarterly Journal of Economics*, 2020, *135* (4), 2007–2058.
- Boldi, Paolo and Sebastiano Vigna**, “Axioms for centrality,” *Internet Mathematics*, 2014, *10* (3-4), 222–262.
- Bonacich, Phillip**, “Power and centrality: A family of measures,” *American journal of sociology*, 1987, *92* (5), 1170–1182.
- Carvalho, Vasco and Alireza Tahbaz-Salehi**, “Production Networks: A Primer,” *Annual Review of Economics*, 2019, *11* (1), 635–663.
- Carvalho, Vasco M, Makoto Nirei, Yukiko U Saito, and Alireza Tahbaz-Salehi**, “Supply chain disruptions: Evidence from the great east japan earthquake,” *The Quarterly Journal of Economics*, 2021, *136* (2), 1255–1321.
- Chodorow-Reich, Gabriel, Xavier Gabaix, and Ralph Koijen**, “The Heterogeneous CES Demand and Production Functions.” Working paper.

- Costello, Anna M**, “Credit market disruptions and liquidity spillover effects in the supply chain,” *Journal of Political Economy*, 2020, *128* (9), 3434–3468.
- Crucitti, Paolo, Vito Latora, Massimo Marchiori, and Andrea Rapisarda**, “Efficiency of scale-free networks: error and attack tolerance,” *Physica A: Statistical Mechanics and its Applications*, 2003, *320*, 622–642.
- DeMarzo, Peter M, Dimitri Vayanos, and Jeffrey Zwiebel**, “Persuasion bias, social influence, and unidimensional opinions,” *The Quarterly journal of economics*, 2003, *118* (3), 909–968.
- den Haan, Wouter J. and Joris de Wind**, “How Well-Behaved are Higher-Order Perturbation Solutions?” Working paper.
- Dew-Becker, Ian, Alireza Tahbaz-Salehi, and Andrea Vedolin**, “Macro skewness and conditional second moments: evidence and theories,” 2021. Working paper.
- Diewert, Erwin**, “Understanding PPPs and PPP-based national accounts: comment,” *American Economic Journal: Macroeconomics*, 2010, *2* (4), 36–45.
- Dupraz, Stephane, Jon Steinsson, and Emi Nakamura**, “A Plucking Model of Business Cycles,” 2020. Working paper.
- Embrechts, Paul, Claudia Klüppelberg, and Thomas Mikosch**, *Modelling extremal events: for insurance and finance*, Vol. 33, Springer Science & Business Media, 2013.
- Flynn, Joel P., Christina Patterson, and John Sturm**, “Fiscal Policy in a Networked Economy,” 2022. Working paper.
- Fujiy, Brian Cevallos, Devaki Ghose, and Gaurav Khanna**, “Production Networks and Firm-level Elasticities of Substitution,” 2021. Working paper.
- Gabaix, Xavier**, “The granular origins of aggregate fluctuations,” *Econometrica*, 2011, *79* (3), 733–772.
- Gofman, Michael, Gill Segal, and Yuchang Wu**, “Production networks and stock returns: The role of vertical creative destruction,” *The Review of Financial Studies*, 2020, *33* (12), 5856–5905.
- Granot, Daniel and Gur Huberman**, “Minimum cost spanning tree games,” *Mathematical programming*, 1981, *21* (1), 1–18.
- Head, Keith, Thierry Mayer, and Mathias Thoenig**, “Welfare and trade without Pareto,” *American Economic Review*, 2014, *104* (5), 310–16.
- Herskovic, Bernard, Bryan Kelly, Hanno Lustig, and Stijn Van Nieuwerburgh**,

- “Firm volatility in granular networks,” *Journal of Political Economy*, 2020, *128* (11), 4097–4162.
- Hill, Robert J**, “Comparing price levels across countries using minimum-spanning trees,” *Review of Economics and Statistics*, 1999, *81* (1), 135–142.
- , “Constructing price indexes across space and time: the case of the European Union,” *American Economic Review*, 2004, *94* (5), 1379–1410.
- Hulten, Charles R**, “Growth accounting with intermediate inputs,” *The Review of Economic Studies*, 1978, *45* (3), 511–518.
- Ilut, Cosmin, Matthias Kehrig, and Martin Schneider**, “Slow to Hire, Quick to Fire: Employment Dynamics with Asymmetric Responses to News,” *Journal of Political Economy*, 2018, *126* (5), 2011–2071.
- Jiang, Bomin, Daniel E. Rigobon, and Roberto Rigobon**, “From Just in Time, to Just in Case, to Just in Worst-Case: Simple models of a Global Supply Chain under Uncertain Aggregate Shocks,” 2021. Working paper.
- Jones, Charles I**, “Intermediate goods and weak links in the theory of economic development,” *American Economic Journal: Macroeconomics*, 2011, *3* (2), 1–28.
- Jr, John B Long and Charles I Plosser**, “Real business cycles,” *Journal of political Economy*, 1983, *91* (1), 39–69.
- La’O, Jennifer and Alireza Tahbaz-Salehi**, “Optimal Monetary Policy in Production Networks,” 2021.
- Liu, Ernest**, “Industrial policies in production networks,” *The Quarterly Journal of Economics*, 2019, *134* (4), 1883–1948.
- and **Aleh Tsyvinski**, “Dynamical Structure and Spectral Properties of Input-Output Networks,” 2021.
- Marchiori, Massimo and Vito Latora**, “Harmony in the small-world,” *Physica A: Statistical Mechanics and its Applications*, 2000, *285* (3-4), 539–546.
- Melitz, Marc J and Stephen J Redding**, “New trade models, new welfare implications,” *American Economic Review*, 2015, *105* (3), 1105–46.
- Mossel, Elchanan, Allan Sly, and Omer Tamuz**, “Strategic learning and the topology of social networks,” *Econometrica*, 2015, *83* (5), 1755–1794.
- Nair, Jayakrishnan, Adam Wierman, and Bert Zwart**, *The Fundamentals of Heavy Tails: Properties, Emergence, and Estimation*, Cambridge University Press, 2022.

- Nehring, Klaus and Clemens Puppe**, “A theory of diversity,” *Econometrica*, 2002, 70 (3), 1155–1198.
- Resnick, Sidney I**, *Heavy-tail phenomena: probabilistic and statistical modeling*, Springer Science & Business Media, 2007.
- Rochat, Yannick**, “Closeness centrality extended to unconnected graphs: The harmonic centrality index,” 2009. Working paper.
- Rubbo, Elisa**, “Networks, phillips curves and monetary policy,” 2020. Working paper.
- Sun, Ning and Zaifu Yang**, “An efficient and incentive compatible dynamic auction for multiple complements,” *Journal of Political Economy*, 2014, 122 (2), 422–466.
- Taschereau-Dumouchel, Mathieu**, “Cascades and Fluctuations in an Economy with an Endogenous Production Network,” 2021. Working paper.
- vom Lehn, Christian and Thomas Winberry**, “The Investment Network, Sectoral Comovement, and the Changing U.S. Business Cycle,” 2021. Working paper.
- Weitzman, Martin L**, “On diversity,” *The Quarterly Journal of Economics*, 1992, 107 (2), 363–405.

A Proofs

A.1 Lemma 1

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

$$GDP = W/P_0 \tag{42}$$

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). I use lower-case letters to denote logs, so $p_0 = \log P_0$, etc. Setting labor to be the numeraire, so that W is normalized to 1, the CES preferences for the consumer along with cost minimization immediately imply

$$p_0 = \sum_{j=1}^N \beta_j p_j \tag{43}$$

$$gdp = -p_0 \tag{44}$$

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

$$p_i = -z_i + \frac{\alpha}{1 - \sigma_i} \log \left(\sum_{j=1}^N A_{ij} \exp((1 - \sigma_i) p_j) \right) \quad (45)$$

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 (43) and (45) gives

$$\lim_{t \rightarrow \infty} t^{-1} gdp = \beta' \phi \quad (46)$$

$$\phi_i = -\theta_i + \alpha_i f_i(\phi) \quad (47)$$

where

$$f_i(\phi) \equiv \begin{cases} \max_{j \in S_i} \phi_j & \text{if } \sigma_i < 1 \\ \sum_j A_{i,j} \phi_j & \text{if } \sigma_i = 1 \\ \min \phi_j & \text{if } \sigma_i > 1 \end{cases} \quad (48)$$

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

$$g_i(\phi) = \theta_i + \alpha f_i(\phi) \quad (49)$$

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 (47) yields

$$\mu_i + \phi_i t = -z_i + \frac{\alpha}{1 - \sigma_i} \log \left(\sum_{j=1}^N A_{ij} \exp((1 - \sigma_i) (\mu_j + \phi_j t)) \right) \quad (50)$$

$$\mu_i = \frac{\alpha}{1 - \sigma_i} \log \left(\sum_{j=1}^N A_{ij} \exp((1 - \sigma_i) (\mu_j + (\phi_j - f_i(\phi)) t)) \right) \quad (51)$$

$$\mu_i = \frac{\alpha}{(1 - \sigma_i)} \log \left(\sum_{j \in j^*(i)} A_{i,j} \exp((1 - \sigma_i) \mu_j) \right) \quad (52)$$

where

$$j^*(i) \equiv \begin{cases} \{j : \phi_j = \max_{k \in S_i} \phi_k\} & \text{if } \sigma_i < 1 \\ \{j : \phi_j = \min_{k \in S_i} \phi_k\} & \text{if } \sigma_i > 1 \end{cases} \quad (53)$$

and S_i is the set of inputs of sector i (i.e. $S_i \equiv \{j : A_{i,j} > 0\}$). (52) follows from the fact that for $\sigma_i > 1$, $\phi_j > f_i(\phi)$ for any $j \notin j^*(i)$ and $j \in S_i$ and $\phi_j < f_i(\phi)$ in the same situation for $\sigma_i < 1$, so that all terms in the summation drop out except $j \in j^*(i)$. The results for the Cobb–Douglas case following using similar analysis. It is again straightforward to confirm that μ is the fixed point of a contraction mapping.

A.2 Propositions 1, 2, and 5

The propositions are derived for arbitrary θ , with the tail centralities being special cases. Proposition 5 is for the general case and propositions 1 and 2 are the special cases.

Define $f^0 : \mathbb{R}^N \rightarrow \mathbb{R}^N$ to be the vectorized version of the function in (48). 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 transformation f^1 . First, take the case with σ_i transitioning from below 1 to being equal to 1 or more. Then, necessarily,

$$T^1 \phi \geq T^0 \phi \quad (54)$$

for any ϕ , element-by-element, from which Proposition 1 follows.

The second proposition holds by the same argument. For example, suppose $\sigma_i < 1$ and the set S_i grows. Again, define a T^2 for the model with the larger S_i . We have

$$T^2 \phi \leq T^0 \phi \quad (55)$$

for any ϕ , element-by-element, which establishes Proposition 2.

A.3 Proposition 3

Define a set of $N \times N$ matrices \mathbf{A}_n representing restricted versions of the production network. For each \mathbf{A}_n , each sector is restricted to using just one of its inputs, so that every \mathbf{A}_n 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 n of $\{\mathbf{A}_n\}$ represents every possible restricted network.³² If $\sigma_i = 1$, then sector i always uses the same mix of inputs, and the i th row of \mathbf{A}_n is equal to $A_{i,\cdot}$ for every n .

Now define ϕ^* and n^*

$$n^* = \arg \min_n \beta' (\mathbf{I} - \alpha \mathbf{A}_n)^{-1} \theta \quad (56)$$

$$\phi^* = -(\mathbf{I} - \alpha \mathbf{A}_{n^*})^{-1} \theta \quad (57)$$

where $1_{N \times 1}$ is a vector of 1's. That implies

$$\phi^* = -\theta + \alpha \mathbf{A}_{n^*} \phi^* \quad (58)$$

Suppose \mathbf{A}_{n^*} is not the solution to the recursion from Lemma 1 for ϕ^* . Then, clearly, element-by-element $T\phi^* \geq \phi^*$ (where T is the operator $T\phi \equiv \theta + \alpha f(\phi)$), and 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 Theorem 2

We have

$$gdp(z) = \mu(\theta) + \lambda(\theta)t + \varepsilon(t, \theta) \quad (59)$$

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

Now define

$$\bar{\varepsilon}(x) = \max_{\theta} \max_{t > \frac{x + \mu(\theta)}{-\lambda(\theta)}} |\varepsilon(t, \theta)| \quad (60)$$

³²The index n runs from 1 to the product of the number of inputs used by each each sector.

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

$$\lim_{x \rightarrow \infty} \bar{\varepsilon}(x) = 0 \quad (61)$$

Now note that

$$\Pr [gdp < -x \mid \theta] = \Pr \left[t + \frac{\varepsilon(t, \theta)}{\lambda(\theta)} > \frac{x + \mu(\theta)}{-\lambda(\theta)} \mid \theta \right] \quad (62)$$

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

$$\begin{aligned} \Pr \left[t + \frac{\bar{\varepsilon}(x)}{\lambda(\theta)} > \frac{x + \mu(\theta)}{-\lambda(\theta)} \mid \theta \right] &\leq \Pr \left[t + \frac{\varepsilon(t, \theta)}{\lambda(\theta)} > \frac{x + \mu(\theta)}{-\lambda(\theta)} \mid \theta \right] \leq \Pr \left[t - \frac{\bar{\varepsilon}(x)}{\lambda(\theta)} > \frac{x + \mu(\theta)}{-\lambda(\theta)} \mid \theta \right] \\ \Pr \left[t > \frac{x + \mu(\theta) + \bar{\varepsilon}(x)}{-\lambda(\theta)} \mid \theta \right] &\leq \Pr [gdp < -x \mid \theta] \leq \Pr \left[t > \frac{x + \mu(\theta) - \bar{\varepsilon}(x)}{-\lambda(\theta)} \mid \theta \right] \end{aligned} \quad (63)$$

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

$$\Pr [gdp < -x] = \int_{\Theta} \Pr [gdp < -x \mid \theta] dm(\theta) \quad (64)$$

from which the result follows directly. ■

A.5 Corollary 3

Recall the notation from the proof of Theorem 1 that

$$gdp(\theta t) = \mu(\theta) + \lambda(\theta)t + \varepsilon(\theta, t) \quad (65)$$

and that $|\varepsilon(\theta, t)| < \bar{\varepsilon}(x)$ 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

$$\Pr [gdp < -x] \geq \int_{\theta: \lambda(\theta) < 0} \bar{F} \left(\frac{x - \mu(\theta) + \varepsilon'(x)}{-s(\theta) \lambda(\theta)} \right) dm(\theta) \quad (66)$$

$$\Pr [gdp > x] \leq \int_{\eta: \lambda(\eta) > 0} \bar{F} \left(\frac{x - \mu(\eta) - \varepsilon'(x)}{s(\eta) \lambda(\eta)} \right) dm(\eta) \quad (67)$$

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

$$\frac{x - \mu(-\theta) - \varepsilon'(x)}{s(-\theta)\lambda(-\theta)} - \frac{x - \mu(\theta) + \varepsilon'(x)}{-s(\theta)\lambda(\theta)} \quad (68)$$

$$= \left(\frac{1}{s(-\theta)\lambda(-\theta)} - \frac{1}{-s(\theta)\lambda(\theta)} \right) x + \frac{-\mu(-\theta) - \varepsilon'(x)}{s(-\theta)\lambda(-\theta)} - \frac{-\mu(\theta) + \varepsilon'(x)}{-s(\theta)\lambda(\theta)} \quad (69)$$

So there exists an \bar{x} such that for $x > \bar{x}$, the argument of \bar{F} in the integral for (66) is smaller than that in (67) for any given θ . In addition,

$$m(\{\eta : \lambda(\eta) > 0\}) \leq m(\{\theta : \lambda(\theta) < 0\}) \quad (70)$$

which yields the result.

Online appendix

B Proofs of additional propositions

B.1 Proposition 4

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. ■

B.2 Proposition 6

The statement of Theorem 2 is

$$\int_{\theta:\lambda(\theta)<0} \bar{F}\left(\frac{x - \mu(\theta) + \varepsilon(x)}{-s(\theta)\lambda(\theta)}\right) dm(\theta) \leq \Pr[gdp < -x] \leq \int_{\theta:\lambda(\theta)<0} \bar{F}\left(\frac{x - \mu(\theta) - \varepsilon(x)}{-s(\theta)\lambda(\theta)}\right) dm(\theta) \quad (71)$$

In this case we have

$$\bar{F}(s) = c \exp(-\beta(t - \bar{t})^\kappa) \quad (72)$$

$$\text{where } c = \Pr(t \leq \bar{t}) \quad (73)$$

If the limits of the two integrals in (71) are the same, then that limit is also the limit for $\Pr[gdp < -x]$. This section gives the derivation for the right-hand side limit, with the arguments holding equivalently on the left with the sign of $\varepsilon(x)$ reversed.

We have

$$\left(\int_{\theta:\lambda(\theta)<0} \bar{F}\left(\frac{x - \mu(\theta) - \varepsilon(x)}{-s(\theta)\lambda(\theta)}\right) dm(\theta) \right)^{1/x^\kappa} \quad (74)$$

$$= \left[\int_{\theta \in \Theta} \exp\left(-\left(\frac{1}{-s(\theta)\lambda(\theta)} - \frac{\varepsilon(x) + \mu(\theta)}{x} \frac{1}{s(\theta)\lambda(\theta)} - \frac{\bar{t}}{x}\right)^\kappa\right)^{x^\kappa} dm(\theta) \right]^{1/x^\kappa} \quad (75)$$

Now consider the limit as $x \rightarrow \infty$. I show that the limit of the right-hand side is the essential supremum of $\exp\left(-\left(\frac{1}{-s(\theta)\lambda(\theta)}\right)^\kappa\right)$ with respect to the measure $m(\theta)$ (i.e. the measure of

the set of θ such that $\exp\left(-\left(\frac{1}{s(\theta)\lambda(\theta)}\right)^\kappa\right)$ is above the essential supremum is zero). Denote that by $\left\|\exp\left(-\left(\frac{1}{s(\theta)\lambda(\theta)}\right)^\kappa\right)\right\|_\infty$.

The structure of this proof is from Ash and Doleans-Dade (2000), page 470, with the addition of the convergence of the argument of the integral with respect to x .

Define, for notational convenience,

$$f(\theta) = \exp\left(-\left(\frac{1}{s(\theta)\lambda(\theta)}\right)^\kappa\right) \quad (76)$$

$$f(\theta; x) = \exp\left(-\left(\frac{1}{s(\theta)\lambda(\theta)} - \frac{\varepsilon(x) + \mu(\theta)}{x} \frac{1}{s(\theta)\lambda(\theta)} - \frac{\bar{t}}{x}\right)^\kappa\right) \quad (77)$$

Lemma B1. $\lim_{x \rightarrow \infty} \|f(\theta; x)\|_\infty = \|f(\theta)\|_\infty$.

Proof. $f(\theta; x) \rightarrow f(\theta)$ pointwise trivially. The difference $|f(\theta; x) - f(\theta)|$ is bounded due to the facts that $\varepsilon(x)$ and $\mu(\theta)$ are bounded and that $f(\theta; x)$ is decreasing in $s(\theta)\lambda(\theta)$ (for sufficiently large x), which is bounded from above (and below, by zero). $f(\theta; x)$ then converges uniformly to $f(\theta)$, from which $\|f(\theta; x)\|_\infty \rightarrow \|f(\theta)\|_\infty$ follows, since, using the reverse triangle inequality,

$$\left| \|f(\theta; x)\|_\infty - \|f(\theta)\|_\infty \right| \leq \|f(\theta) - f(\theta; x)\|_\infty \quad (78)$$

■

Lemma B2. $\limsup_{x \rightarrow \infty} \left[\int_{\theta \in \Theta} f(\theta; x)^{x^\kappa} dm(\theta) \right]^{1/x^\kappa} \leq \|f(\theta)\|_\infty$

Proof. We have (except possibly on a set of measure zero)

$$\|f(\theta; x)\|_{x^\kappa} \leq \| \|f(\theta; x)\|_\infty \|_{x^\kappa}$$

Taking limits of both sides

$$\lim_{x \rightarrow \infty} \|f(\theta; x)\|_{x^\kappa} \leq \lim_{x \rightarrow \infty} \| \|f(\theta; x)\|_\infty \|_{x^\kappa} \quad (79)$$

$$= \lim_{x \rightarrow \infty} \|f(\theta; x)\|_\infty \quad (80)$$

$$= \|f(\theta)\|_\infty \quad (81)$$

where the second line follows from the fact that $\|f(\theta; x)\|_\infty$ is constant and the third line uses lemma B1. ■

Lemma B3. $\liminf_{x \rightarrow \infty} \left[\int f(\theta; x)^{x^\kappa} dm(\theta) \right]^{1/x^\kappa} \geq \|f(\theta)\|_\infty$

Proof. Consider some $\eta > 0$, and set $A = \left\{ \theta : \exp\left(-\left(\frac{1}{-s(\theta)\lambda(\theta)}\right)^\kappa\right) \geq \left\| \exp\left(-\left(\frac{1}{-s(\theta)\lambda(\theta)}\right)^\kappa\right) \right\|_\infty - \eta \right\}$. Consider also the set $A' = \left\{ \theta : \exp\left(-\left(\frac{1}{s(\theta)\lambda(\theta)} - \frac{\pm\varepsilon(x) + \mu(\theta)}{x} \frac{1}{\lambda(\theta)} - \frac{\bar{t}}{x}\right)^\kappa\right) \geq \left\| \exp\left(-\left(\frac{1}{s(\theta)\lambda(\theta)}\right)^\kappa\right) \right\|_\infty - \eta \right\}$. For any η such that A has positive measure, there exists an $\bar{x}(\eta)$ sufficiently large that A' has positive measure for all $x > \bar{x}(\eta)$ due to the continuity of $\exp\left(-\left(\frac{1}{s(\theta)\lambda(\theta)} - \frac{\pm\varepsilon(x) + \mu(\theta)}{x} \frac{1}{s(\theta)\lambda(\theta)} - \frac{\bar{t}}{x}\right)^\kappa\right)$ and the fact that $\exp\left(-\left(\frac{1}{s(\theta)\lambda(\theta)} - \frac{\pm\varepsilon(x) + \mu(\theta)}{x} \frac{1}{\lambda(\theta)}\right)^\kappa\right) \rightarrow \exp\left(-\left(\frac{1}{s(\theta)\lambda(\theta)}\right)^\kappa\right)$ as $x \rightarrow \infty$.

It is then the case that for $x > \bar{x}(\eta)$

$$\int \exp\left(-\left(\frac{1}{\lambda(\theta)} - \frac{\pm\varepsilon(x) + \mu(\theta)}{x} \frac{1}{s(\theta)\lambda(\theta)} - \frac{\bar{t}}{x}\right)^\kappa\right) dm(\theta) \quad (82)$$

$$\geq \int_{A'} \exp\left(-\left(\frac{1}{\lambda(\theta)} - \frac{\pm\varepsilon(x) + \mu(\theta)}{x} \frac{1}{s(\theta)\lambda(\theta)} - \frac{\bar{t}}{x}\right)^\kappa\right) dm(\theta) \quad (83)$$

$$\geq \left(\left\| \exp\left(-\left(\frac{1}{\lambda(\theta)}\right)^\kappa\right) \right\|_\infty - \eta \right)^{x^\kappa} \mu(A') \quad (84)$$

Since $\mu(A') > 0$ from the definition of $\left\| \exp\left(-\left(\frac{1}{s(\theta)\lambda(\theta)}\right)^\kappa\right) \right\|_\infty$ (ignoring the trivial case of a constant value for $\exp\left(-\left(\frac{1}{s(\theta)\lambda(\theta)}\right)^\kappa\right)$), and since the above holds for any $\eta > 0$,

$$\liminf_{x \rightarrow \infty} \left[\int \exp\left(-\left(\frac{1}{s(\theta)\lambda(\theta)} - \frac{\pm\varepsilon(x) + \mu(\theta)}{x} \frac{1}{\lambda(\theta)} - \frac{\bar{t}}{x}\right)^\kappa\right)^{x^\kappa} dm(\theta) \right]^{1/x^\kappa} \quad (85)$$

$$\geq \left\| \exp\left(-\left(\frac{1}{s(\theta)\lambda(\theta)}\right)^\kappa\right) \right\|_\infty \quad (86)$$

■

Proof of the proposition: Since both the lim inf and lim sup are equal to $\left\| \exp\left(-\left(\frac{1}{s(\theta)\lambda(\theta)}\right)^\kappa\right) \right\|_\infty$, the limit is also.

For the second part, in the set Θ^* , there exists an η such that $|-s(\theta)\lambda(\theta)| < \|-s(\theta)\lambda(\theta)\|_\infty - \eta$

η . Therefore

$$\frac{\int_{\Theta^*} \exp\left(-\left(\frac{x+\varepsilon(x)-\mu(\theta)}{-s(\theta)\lambda(\theta)} - \bar{t}\right)^\kappa\right) dm(\theta)}{\int \exp\left(-\left(\frac{x-\varepsilon(x)-\mu(\theta)}{-s(\theta)\lambda(\theta)} - \bar{t}\right)^\kappa\right) dm(\theta)} \leq \Pr\left[\begin{array}{c} \theta \in \Theta^* \\ |gdp < -x \end{array}\right] \leq \frac{\int_{\Theta^*} \exp\left(-\left(\frac{x-\varepsilon(x)-\mu(\theta)}{-s(\theta)\lambda(\theta)} - \bar{t}\right)^\kappa\right) dm(\theta)}{\int \exp\left(-\left(\frac{x+\varepsilon(x)-\mu(\theta)}{-s(\theta)\lambda(\theta)} - \bar{t}\right)^\kappa\right) dm(\theta)} \quad (87)$$

Again, we show that both sides of the inequality have the same limit. For a sufficiently large x ,

$$\begin{aligned} \frac{\int_{\Theta^*} \exp\left(-\left(\frac{x\pm\varepsilon(x)-\mu(\theta)}{-s(\theta)\lambda(\theta)} - \bar{t}\right)^\kappa\right) dm(\theta)}{\int \exp\left(-\left(\frac{x\pm\varepsilon(x)-\mu(\theta)}{-s(\theta)\lambda(\theta)} - \bar{t}\right)^\kappa\right) dm(\theta)} &\leq \frac{\int_{\Theta^*} \exp\left(-\left(\frac{x\pm\varepsilon(x)-\mu(\theta)}{(\|s(\theta)\lambda(\theta)\|_\infty - \eta)} - \bar{t}\right)^\kappa\right) dm(\theta)}{\int_{\theta: |\lambda(\theta)| > |\lambda(\theta)| - \eta/2} \exp\left(-\left(\frac{x\pm\varepsilon(x)-\mu(\theta)}{-s(\theta)\lambda(\theta)} - \bar{t}\right)^\kappa\right) dm(\theta)} \\ &\leq \frac{\exp\left(-\left(\frac{x\pm\varepsilon(x)-\mu(\theta)}{-(\|s(\theta)\lambda(\theta)\|_\infty - \eta)} - \bar{t}\right)^\kappa\right)}{\exp\left(-\left(\frac{x\pm\varepsilon(x)-\mu(\theta)}{-(\|s(\theta)\lambda(\theta)\|_\infty - \eta/2)} - \bar{t}\right)^\kappa\right)} \frac{1}{m(\{\theta : |\lambda(\theta)| > \|\lambda(\theta)\|_\infty - \eta/2\})} \quad (88) \\ &\rightarrow 0 \quad (89) \end{aligned}$$

■

B.3 Proof of Proposition 7

We have

$$\bar{F}(s) = c(t/\bar{t})^{-\kappa} \quad (90)$$

$$\text{where } c = \Pr(t \geq \bar{t}) \quad (91)$$

Inserting those into the formula from theorem 1, we again show that the integrals have the same bound. Define $\Theta^- \equiv \{\theta : \lambda(\theta) < 0\}$. The bound is now

$$\int_{\Theta^-} \left(\frac{x + \varepsilon(x) + \mu(\theta)}{-xs(\theta)\lambda(\theta)}\right)^{-\kappa} dm(\theta) \leq \Pr[gdp < -x] / (c\bar{t}^\kappa x^{-\kappa}) \leq \int_{\Theta^-} \left(\frac{x - \varepsilon(x) + \mu(\theta)}{-xs(\theta)\lambda(\theta)}\right)^{-\kappa} dm(\theta) \quad (92)$$

with limit

$$\lim_{x \rightarrow \infty} \int_{\Theta^-} \left(- (s(\theta) \lambda(\theta))^{-1} + x^{-1} \frac{\pm \varepsilon(x) + \mu(\theta)}{-s(\theta) \lambda(\theta)} \right)^{-\kappa} dm(\theta) \quad (93)$$

Again, recall that the $\pm \varepsilon(x)$ term is bounded, as are $\lambda(\theta)$ and $\mu(\theta)$ (since Θ is compact). The argument of the integral therefore converges uniformly, since

$$\begin{aligned} & \left\| \left(- (s(\theta) \lambda(\theta))^{-1} + x^{-1} \frac{\pm \varepsilon(x) + \mu(\theta)}{-s(\theta) \lambda(\theta)} \right)^{-\kappa} - (-s(\theta) \lambda(\theta))^\kappa \right\|_\infty \\ & \leq \left\| \left(-\lambda(\theta)^{-1} + x^{-1} \frac{\pm \varepsilon(x) + \mu(\theta)}{-s(\theta) \lambda(\theta)} \right)^{-1} \right\|_\infty^\kappa + \|-s(\theta) \lambda(\theta)\|_\infty^\kappa \\ & \leq \left\| -s(\theta) \lambda(\theta) \left(1 + \frac{\pm \varepsilon(x) + \mu(\theta)}{x} \right)^{-1} \right\|_\infty^\kappa + \|-s(\theta) \lambda(\theta)\|_\infty^\kappa \quad (94) \end{aligned}$$

$$\leq \|-s(\theta) \lambda(\theta)\|_\infty^\kappa \left\| \frac{x}{x + \inf_{\theta \in \Theta} \{\mu(\theta)\}} \right\|_\infty^\kappa + \|-s(\theta) \lambda(\theta)\|_\infty^\kappa \quad (95)$$

with the last line being bounded. Passing the limit through the integral yields the result from the text. The second claim is again an application of Bayes' rule. \blacksquare

C Derivations for examples

C.1 Example 3

In this case, the probability density in the tail is $\exp(-\|z\|_{1,v}/\eta)$, where

$$\|x\|_{1,v} \equiv \sum_j |x_j|/v_j \quad (96)$$

denotes an l_1 -type norm weighted by a vector v , representing the volatility of each shock. To confirm that $s(\theta) = 1/\|\theta\|_{1,v}$, note that

$$\exp(- (t/s(\theta)) / \eta) = \exp\left(- \left(\|z\| \left\| \frac{z}{\|z\|} \right\|_{1,v}\right) / \eta\right) \quad (97)$$

$$= \exp\left(- \|z\|_{1,v} / \eta\right) \quad (98)$$

as required.

The aim is to find $\max_{\tilde{\theta}: \|\tilde{\theta}\|_2=1} \left\| -s(\tilde{\theta}) \lambda(\tilde{\theta}) \right\|$. Now note that $b\lambda(\tilde{\theta}) = \lambda(b\tilde{\theta})$, and hence $s(\tilde{\theta}) \lambda(\tilde{\theta}) = \lambda(\tilde{\theta}_s(\tilde{\theta}))$. We can then apply a change of variables, with $\theta = \tilde{\theta}_s(\tilde{\theta})$. Note that $\tilde{\theta} = \theta / \|\theta\|$, so we have

$$\max_{\tilde{\theta}: \|\tilde{\theta}\|_2=1} \left\| -s(\tilde{\theta}) \lambda(\tilde{\theta}) \right\| = \max_{\theta: \|\theta/s(\theta/\|\theta\|)\|_2=1} \left\| -\lambda(\theta) \right\| \quad (99)$$

Now in this particular case,

$$\|\theta/s(\theta/\|\theta\|)\| = \left\| \theta \|\theta/\|\theta\|\|_{1,v} \right\| \quad (100)$$

$$= \|\theta\|_{1,v} \quad (101)$$

The objective is then

$$- \max_{\theta} \max_n D'_n \theta = - \max_n \max_{\theta} D'_n \theta \quad (102)$$

subject to the constraint $\|\theta\|_{1,v} = 1$. The inner maximization on the right is a problem with a linear objective and a linear constraint, so it is simply solved at the point that maximizes $D_{n,j} v_j$. We then have

$$- \max_n \max_j D_{n,j} v_j \quad (103)$$

The example in the text is the special case of $v_j = 1 \forall j$.

C.2 Example 4

The aim is to find $\max_{\tilde{\theta}: \|\tilde{\theta}\|_2=1} \left\| -s(\tilde{\theta}) \lambda(\tilde{\theta}) \right\|$. Now note that $b\lambda(\tilde{\theta}) = \lambda(b\tilde{\theta})$, and hence $s(\tilde{\theta}) \lambda(\tilde{\theta}) = \lambda(\tilde{\theta}s(\tilde{\theta}))$. We can then apply a change of variables, with $\theta = \tilde{\theta}s(\tilde{\theta})$. Note that $\tilde{\theta} = \theta / \|\theta\|$, so we have

$$\max_{\tilde{\theta}: \|\tilde{\theta}\|=1} \left\| -s(\tilde{\theta}) \lambda(\tilde{\theta}) \right\| = \max_{\theta: \|\theta/s(\theta)\|=1} \left\| -\lambda(\theta) \right\| \quad (104)$$

Now in this particular case,

$$\|\theta/s(\theta)\| = \left\| \theta (\theta' \Sigma^{-1} \theta)^{1/2} \|\theta\|^{-1} \right\| \quad (105)$$

$$= (\theta' \Sigma^{-1} \theta)^{1/2} \quad (106)$$

The Lagrangian is then

$$\max_{\theta} \max_n (-D'_n \theta) - \frac{\gamma}{2} (\theta' \Sigma^{-1} \theta - 1) \quad (107)$$

where γ is the multiplier on the constraint on θ . Reversing the order of the optimization gives

$$- \min_n \min_{\theta} D'_n \theta + \frac{\gamma}{2} (\theta' \Sigma^{-1} \theta - 1) \quad (108)$$

$$\theta = -\gamma^{-1} \Sigma D_n \quad (109)$$

where the second line uses the first-order condition. Solving for the constraint for a given θ yields

$$\theta = -\frac{1}{\sqrt{D'_n \Sigma D_n}} \Sigma D_n \quad (110)$$

(recall that θ here does not have norm 1 but instead satisfies $\theta' \Sigma^{-1} \theta = 1$).

Finally, the value of the objective (which is equal to $\left\| -s(\tilde{\theta}) \lambda(\tilde{\theta}) \right\|$) is

$$\left\| -s(\tilde{\theta}) \lambda(\tilde{\theta}) \right\| = -D'_n \theta \quad (111)$$

$$= \sqrt{D'_n \Sigma D_n} \quad (112)$$

C.3 Example 5

The complete symmetry of the economy, along with the fact that output is homogeneous of degree $1/(1 - \alpha)$ in the vector of productivities immediately implies that $D_{ss,i} = N^{-1}/(1 - \alpha)$ for all i .

It is straightforward to confirm that $\phi_i = \theta_i + \frac{\alpha}{1-\alpha}\theta_{\min}$, where $\theta_{\min} = \min_i \theta_i$. Combining that with the final utility function yields

$$\lambda(\theta) = \sum_i N^{-1}\theta_i + \frac{\alpha}{1-\alpha}\theta_{\min} \quad (113)$$

For the right tail, the Lagrangian is

$$\max_{\theta} \lambda(\theta) - \frac{\gamma}{2} \left(\sum_i \theta_i^2 - 1 \right) = \max_{\theta} \sum_i N^{-1}\theta_i + \frac{\alpha}{1-\alpha}\theta_{\min} - \frac{\gamma}{2} \sum_i \theta_i^2 \quad (114)$$

That problem is nonconvex and is solved at the point $\theta_i = N^{-1/2}$ for all i . That yields

$$\|\sigma\lambda(\theta)\|_{\infty} = \frac{1}{1-\alpha}\sigma N^{-1/2} \quad (115)$$

For the left tail, the optimization is

$$\max_{\theta} - \sum_i N^{-1}\theta_i - \frac{\alpha}{1-\alpha}\theta_{\min} - \frac{\gamma}{2} \left(\sum_i \theta_i^2 - 1 \right) \quad (116)$$

The first-order condition gives

$$\theta_i = \begin{cases} -\frac{N^{-1}}{\sqrt{(N-1)N^{-2} + \left(N^{-1} + \frac{\alpha}{1-\alpha}\right)^2}} & \text{for } i \neq i \text{ min} \\ -\frac{N^{-1} + \frac{\alpha}{1-\alpha}}{\sqrt{(N-1)N^{-2} + \left(N^{-1} + \frac{\alpha}{1-\alpha}\right)^2}} & \text{for } i = i \text{ min} \end{cases} \quad (117)$$

where $i \text{ min} = \arg \min_i \theta_i$. Note that this solution for θ_i is obviously not unique – $i \text{ min}$ can

be equal to any integer between 1 and N . To find $\sigma\lambda(\theta)$, we have

$$\|-\sigma\lambda(\theta)\|_\infty = \sigma \frac{N^{-1} + 2N^{-1} \frac{\alpha}{1-\alpha} + \left(\frac{\alpha}{1-\alpha}\right)^2}{\sqrt{(N-1)N^{-2} + \left(N^{-1} + \frac{\alpha}{1-\alpha}\right)^2}} \quad (118)$$

$$= \sigma \frac{\alpha}{1-\alpha} + O(N^{-1/2}) \quad (119)$$

where $x = O(N^{-1/2}) \Leftrightarrow |x| \leq MN^{-1/2}$ for all x greater than some x_0 and for some constant M .

D Extensions and additional results

D.1 Sector output

To prove the claim from the text, we will show that the following set of limits (along with a third additional result) is consistent with the model's equilibrium conditions.

$$\lim_{t \rightarrow \infty} \frac{-y_j}{t} = \lim_{t \rightarrow \infty} \frac{-c_j}{t} = \lim_{t \rightarrow \infty} \frac{p_j}{t} = \phi_j \quad (120)$$

and the equilibrium conditions are

$$Y_i = \exp(z_i) L_i^{1-\alpha} \left(\sum_j A_{i,j}^{1/\sigma_i} X_{i,j}^{(\sigma_i-1)/\sigma_i} \right)^{\alpha\sigma_i/(\sigma_i-1)} \quad (121)$$

$$Y_j = C_j + \sum_i X_{i,j} \quad (122)$$

$$P_j = P_0 C^{1/\sigma_0} \beta_j^{1/\sigma_0} C_j^{-1/\sigma_0} \quad (123)$$

$$P_j = \alpha P_i \exp(z_i) (Y_i / \exp(z_i))^{\alpha - (\sigma_i-1)/\sigma_i} A_{i,j} X_{i,j}^{-1/\sigma_i} \quad (124)$$

$$1 = (1-\alpha) P_i Y_i / L_i \quad (125)$$

We first prove some small lemmas. Define

$$j^*(i) \equiv \begin{cases} \arg \min_{j \in S(i)} \phi_j & \text{if } \sigma_i > 1 \\ \arg \max_{j \in S(i)} \phi_j & \text{if } \sigma_i < 1 \end{cases} \quad (126)$$

For $\sigma_i = 1$, set $\phi_{j^*(i)} = f_i(\phi)$ (for f_i defined above in (48)).

Lemma D4. $\phi_{j^*(i)} + \sigma_i (\phi_j - \phi_{j^*(i)}) \geq \phi_j$ for all $j \in S(i)$

Proof. This follows from the fact that for all $j \in S(i)$, $\phi_j \leq \phi_{j^*(i)}$ when $\sigma_i > 1$ and $\phi_j \geq \phi_{j^*(i)}$ when $\sigma_i < 1$. It is trivial for $\sigma_i = 1$. ■

Lemma D5. $f_i([\phi_{j^*(i)} + \sigma_i (\phi_j - \phi_{j^*(i)})]) = \phi_{j^*(i)}$

Proof. This follows simply because

$$f_i([\phi_{j^*(i)} + \sigma_i (\phi_j - \phi_{j^*(i)})]) = (1 - \sigma_i) \phi_{j^*(i)} + \sigma_i f_i(\phi_{j^*(i)}) \quad (127)$$

$$= \phi_{j^*(i)} \quad (128)$$

■

To prove the result, we also need the use of inputs. We guess that

$$\lim_{t \rightarrow \infty} \frac{x_{i,j}}{t} = -\phi_{j^*(i)} - \sigma_i [\phi_j - \phi_{j^*(i)}] \quad (129)$$

We need to verify that the above, along with the solution in the proposition, satisfies, in the limit, the equilibrium conditions (121)-(125).

We first take limits of the equilibrium conditions. For any variable g_j , define

$$\phi_{g,j} \equiv \lim_{t \rightarrow \infty} \frac{g_j}{t} \quad (130)$$

Inspection of equation (125) shows that, given the guesses for $\phi_{p,i}$ and $\phi_{y,i}$, we must have $\phi_{l,i} = 0$.

Dividing the equilibrium conditions (equations (121)-(125), respectively) by t and taking limits as $t \rightarrow \infty$ yields

$$\phi_{y,i} = \theta_i + (1 - \alpha) \phi_{l,i} + \alpha f_i([\phi_{x,i,j}]) \quad (131)$$

$$\phi_{y,j} = \max \left\{ \phi_{c,j}, \max_i \phi_{x,i,j} \right\} \quad (132)$$

$$\phi_j = \phi_0 + \sigma_0^{-1} \phi_c - \sigma_0^{-1} (\phi_{j^*(0)} + \sigma_0 [\phi_0 - \phi_{j^*(0)}]) \quad (133)$$

$$\phi_{p,j} = \phi_{p,i} + \theta_i + \frac{\alpha - (\sigma_i - 1)/\sigma_i}{\alpha} (\phi_{y,i} - \theta_i) - \sigma_i^{-1} \phi_{x,i,j} \quad (134)$$

$$0 = \phi_{p,i} + \phi_{y,i} - \phi_{l,i} \quad (135)$$

Equation (131) holds using Lemma D5 and the recursion for ϕ_i . Equation (132) holds trivially using the guesses and Lemma D4. Equations (133)-(135) hold trivially after inserting the various guesses.

D.2 Approximation errors

This section compares Theorem 1 to local approximations.

As a first question, can a Taylor series with sufficiently many terms approximate the response of the economy to large shocks?

Proposition 8. *As $t \rightarrow \infty$, the error from any Taylor series for gdp around any value of z will diverge to $\pm\infty$ for some θ unless $\sigma_i = 1 \forall i$.*

The proposition follows from the fact that GDP has a linear asymptote with a constant that is different from zero. Any finite-order Taylor series necessarily diverges infinitely far from the asymptote unless gdp is actually linear. This is obvious when the order of the approximation is greater than 1, since eventually the higher order terms dominate. A linear approximation also eventually diverges in the nonlinear case since the slope of gdp at $t = 0$ is not the same as at $\pm\infty$.

If the economy has any nonlinearity, the tail approximation is always preferable when the magnitude of shocks is sufficiently large. As $\|z\|$ grows, the error in the tail approximation converges to zero, while it diverges to $\pm\infty$ (in at least some directions) for any Taylor series. That behavior is visualized in Figure 2.³³

While Theorem 1 only guarantees accuracy as $t \rightarrow \infty$, its errors also never diverge:

Corollary 5. *There exists a δ such that, for all θ and t*

$$|gdp(\theta t) - (\mu(\theta) + \lambda(\theta)t)| < \delta \tag{136}$$

In addition, there is a stronger form of the main limit for convergence:

³³And in general a Taylor series for a CES aggregator has a finite radius of convergence, meaning that for t greater than some \bar{t} , as terms are added the Taylor series gets *further* from the truth. Specifically, in logs, consider $\log \sum_j w_j \exp(\gamma \theta_j t)$ for some exponent γ and a unit-norm vector θ . The sum inside the log in general has zeros for complex t , meaning that the function has a pole and hence a finite range of convergence for a given θ .

Remark 1. Equations (6) and (9) in Lemma 1 and Theorem 1, respectively, can be replaced by, for any j

$$\lim_{t \rightarrow \infty} |p_i(\theta t) - (\mu_i(\theta) + \phi_i(\theta) t)| t^j = 0 \quad (137)$$

$$\lim_{t \rightarrow \infty} |gdp(\theta t) - (\mu(\theta) + \lambda(\theta) t)| t^j = 0 \quad (138)$$

In other words, the approximation errors converge to zero exponentially fast as $t \rightarrow \infty$ (i.e. faster than any power of t^{-1}), and further one can show that the rate increases with $|\sigma_i - 1|$. That is, there are no other nonzero terms in the polynomial expansion of prices and GDP as $t \rightarrow \infty$.³⁴

Proof. Consider the next term in a series expansion,³⁵ $p_i = b_i t^{-1} + \mu_i + \phi_i t + o(t^{-1})$, and take limits as $t \rightarrow \infty$,

$$p_i t - (\mu_i + \phi_i t) t = \left(-\theta_i t + \frac{\alpha}{1 - \sigma_i} \log \left(\sum_{j=1}^n A_{i,j} \exp((1 - \sigma_i) p_j) \right) - \mu_i - \phi_i t \right) t \quad (139)$$

$$b_i = \lim_{t \rightarrow \infty} \left(\begin{array}{c} -\theta_i t + \frac{\alpha}{1 - \sigma_i} \log \left(\sum_{j=1}^n A_{i,j} \exp((1 - \sigma_i) (b_j t^{-1} + \mu_j + \phi_j t + o(t^{-1}))) \right) \\ -\mu_i - \phi_i t \end{array} \right) t \quad (140)$$

$$b_i = \lim_{t \rightarrow \infty} \left[\left\{ \begin{array}{c} \frac{\alpha}{1 - \sigma_i} \log \left(\sum_{j=1}^n A_{i,j} \exp((1 - \sigma_i) (b_j t^{-1} + \mu_j + (\phi_j - f_i(\phi)) t + o(t^{-1}))) \right) \\ -\mu_i \end{array} \right\} t \right] \quad (141)$$

The recursion from above for μ_i immediately implies that the limit of the term in braces is zero. Applying L'Hopital's rule then yields the result that the equation is only solved by $b_i = 0 \forall i$. The same analysis goes through for terms of any order. ■ ■

³⁴That is, the term $|gdp(\theta t) - (\mu(\theta) + \lambda(\theta) t)|$ behaves similarly to $\exp(-t)$. If one wanted to add further terms to the approximation, it would be necessary to approximate $\log(gdp(z) - (\mu(\theta) + \lambda(\theta) t))$.

³⁵Formally this is a Laurent series since it has both positive and negative powers of t . The terms in the expansion for t^j with $j > 1$ must be zero because otherwise the first limit in this section would not converge.

One way to think about this is that it is an expansion in $\tau = t^{-1}$ around $\tau = 0$. There is a pole at $\tau = 0$ since log GDP and all log prices go to $\pm\infty$. The τ^{-1} terms remove the pole, at which point we just have a standard Taylor series in τ . The remark says that all terms in the Taylor series of order higher than 0 (i.e. everything but the constant) are equal to zero.

D.3 Which is the right approximation to use?

The usual Taylor approximation is around $z = 0$, while this paper focuses on $z \rightarrow \infty$. As z grows, the tail approximation is eventually superior, so for any statements about limiting probabilities as $gdp \rightarrow \pm\infty$, it is the correct representation. But at what point does that transition happen? To shed light on that question, first note that $gdp(0) = 0$. So to know the size of the error from using the tail approximation when $z = 0$, we need to know the constants $\mu(\theta)$.

Recall that the constant in the tail approximation is $-\beta'\mu$ where the vector μ solves the recursion

$$\mu_i = \frac{\alpha}{(1 - \sigma_i)} \log \left(\sum_{j \in j^*(i)} A_{i,j} \exp((1 - \sigma_i) \mu_j) \right)$$

and

$$j^*(i) \equiv \begin{cases} \{j : \phi_j = \max_{k \in S_i} \phi_k\} & \text{if } \sigma_i < 1 \\ \{j : \phi_j = \min_{k \in S_i} \phi_k\} & \text{if } \sigma_i > 1 \end{cases} \quad (142)$$

The constant, $\mu(\theta)$, thus increases when the elasticity of substitution is closer to 1 and when the upstream source of shocks is units that are relatively small (have small $A_{i,j}$). Those factors cause the tail approximation to have a relatively larger error as $t \rightarrow 0$.

The concave case

In the case where gdp is globally concave in the shocks $-\sigma_i \leq 1 \forall i$ – a stronger result is available. The error for the tail approximation then is smaller than for the first-order Taylor series when

$$t > \frac{\mu(\theta)}{D'_{ss}\theta - \lambda(\theta)} \quad (143)$$

The tail approximation is superior if t is sufficiently large – larger when the constant $\mu(\theta)$ is larger or the gap between the local and tail approximations, $D'_{ss}\theta - \lambda(\theta)$, is smaller. That immediately implies that when any elasticity gets closer to 1, the cutoff point gets larger, since σ_i has no impact on λ and D_{ss} . The closer are the various elasticities to 1, the larger the shocks have to be in order for the tail approximation to be superior to a local approximation.

It is less clear what the effects of the $A_{i,j}$ parameters on the cutoff is because they affect both μ and D_{ss} . Note, though, that (in the concave case), when $\lambda(\theta) < 0$ – i.e. when thinking about shocks that reduce GDP – the tail approximation cannot possibly be the

better of the two until $\mu(\theta) + \lambda(\theta)t < 0$, and the point where that happens necessarily increases as the A parameters for the minimizing units (i.e. the units $j \in j^*(i)$ for some i) decline.

D.4 Relaxing the CES assumption

This section extends the baseline result to a broader class of production functions and shows that theorem 1 holds with no modification.

Consider the same competitive economy as in the main analysis, with the only difference that each sector's production need not be CES. Rather, just assume that in each sector has constant returns to scale. Again, without loss of generality, assume that labor and materials are combined with a unit elasticity of substitution. Those assumptions imply that, in competitive equilibrium, the price of good i is given by

$$P_i = \frac{1}{Z_i} W^{1-\alpha} (C_i(P_1, \dots, P_n))^\alpha \quad (144)$$

where Z_i is the productivity shock to industry i , C_i is a homogenous function of degree one, and $\alpha < 1$. In addition to the intermediate-input-producing industries, there is also an industry with cost function C_0 that produces a final good, which is then sold to the representative consumer. Therefore, the final good price, P_0 , also satisfies equation (144), with the convention that $\alpha_0 = 1$ and $Z_0 = 1$.

The only additional assumption imposed on C_i is that

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log C_i(\exp(\phi_l t), \exp(\phi_1 t), \dots, \exp(\phi_n t)) = \tilde{f}_i(\phi_l, \phi_1, \dots, \phi_n) \quad (145)$$

for some function \tilde{f}_i . A sufficient condition for that limit to exist is that

$$\lim_{t \rightarrow \infty} \frac{d}{dt} C_i(\exp(\phi_l t), \exp(\phi_1 t), \dots, \exp(\phi_n t)) \quad (146)$$

exists. That is, it is sufficient that the gradients of the cost functions have limits, but even that is not strictly necessary. The restriction of C_i to the CES family leads to the set of functions f_i that appear in theorem 1.

Theorem 3. *Under the assumptions of this section, and with $z = \theta t$,*

$$\lim_{t \rightarrow \infty} (gdp(z) - \lambda(\theta) t) t^{-1} = 0 \quad (147)$$

where $\lambda(\theta) = \phi_0$ and $\phi \in \mathbb{R}^{N+1}$ is a function of θ that is implicitly defined by the system of equations

$$\phi_i = \theta_i + \alpha \tilde{f}_i(\phi) \text{ for } i \in \{0, 1, \dots, N\} \quad (148)$$

This result shows that what ultimately determines the behavior of GDP for extreme shocks is the limiting slope of the sector-level cost functions.

Proof. The price of good i is

$$p_i = -\log z_i + \alpha_i \log C_i(\exp(p)) \quad (149)$$

Let

$$\phi_i = -\lim_{t \rightarrow \infty} t^{-1} p_i \quad (150)$$

we maintain for the moment that this limit exists and is finite and verify that later. Then

$$t^{-1} p_i = -\theta + \alpha_i t^{-1} \log C_i(\exp(p)) \quad (151)$$

$$\phi_i = -\theta + \alpha_i \lim_{t \rightarrow \infty} t^{-1} \log C_i(\exp(p)) \quad (152)$$

$$= -\theta + \alpha_i f_i(\phi) \quad (153)$$

where the second line takes the limit as $t \rightarrow \infty$ and the third line uses the definition of f_i along with the continuity of C_i and the price function.

Note also that the price of the final good is

$$\log GDP = -\log P = f_0(\phi) t + o(t) \quad (154)$$

Finally, to show that a solution to the system exists, define

$$\hat{g}_i(\phi) = \theta_i + \alpha_i f_i(\phi) \quad (155)$$

This has a unique solution if \hat{g} is a contraction. To see why that is true, we just check

Blackwell's sufficient conditions of monotonicity and discounting. Monotonicity holds simply because the cost function itself is assumed to be monotone. Constant returns in the function C_i also imply that $f(\phi + a) = f(\phi) + a$. Since $\alpha_i < 1$, \hat{g}_i has the discounting property, making it a contraction, so we can then apply the Banach fixed point theorem. ■

D.4.1 The heterogeneous CES setup of Chodorow-Reich, Gabaix, and Koijen (2022)

Chodorow-Reich et al. (2022) study an aggregator of the form

$$\sum_i \phi_i \frac{(X_i/Y)^{(\sigma_i-1)/\sigma_i} - 1}{(\sigma_i - 1)/\sigma_i} + \phi_0 = 0 \quad (156)$$

where the X_i are uses of inputs, The ϕ_i are parameters, and Y is output, which is an implicit function of the inputs. They show that the unit cost function for this case is solved by

$$\sum_i \frac{\sigma_i}{\sigma_i - 1} (P_i/\mu)^{1-\sigma_i} + \phi_0 = 0 \quad (157)$$

$$C = \mu \sum_i (P_i/\mu)^{1-\sigma_i} \quad (158)$$

Now suppose the prices all have limits $\log P_i \rightarrow g_i t$ as $t \rightarrow \infty$. It is then the case that if all $\sigma_i < 1$, $C \rightarrow (\max_i g_i) t$, while if $\sigma_i > 1$, $C \rightarrow (\min_i g_i) t$. That is, in this more general case, the precise value of the elasticity of substitution for each good continues to play no role, as long as all of the elasticities (within a given sector) are above or below 1. In the case where elasticities are mixed within a sector in this model, the analysis, for general g_i , becomes much more difficult and does not yield a simple solution.

D.5 Fixed labor

Assume labor is normalized to 1 and the elasticity of substitution at the household level is 1. Then the production function, resource constraint, and FOCs are

$$Y_i = \exp(z_i) \left(\sum_j a_{i,j} x_{i,j}^{\gamma_i} \right)^{\alpha/\gamma_i} \quad (159)$$

$$Y_j = c_j + \sum_i x_{i,j} \quad (160)$$

$$p_j = b_j c_j^{-1} \quad (161)$$

$$p_j = \alpha p_i \exp(z_i) (y_i / \exp(z_i))^{(\alpha - \gamma_i) / \alpha} a_{i,j} x_{i,j}^{\gamma_i - 1} \quad (162)$$

We assume productivity in each sector is

$$\log z_i = \theta_i t \quad (163)$$

with $t \rightarrow \infty$. Define ϕ to be the solution to the recursion

$$\phi_i = \theta_i + \alpha f_i(\phi) \quad (164)$$

Proposition 9. *We have the following limits,*

$$\lim_{t \rightarrow \infty} \frac{\log y_j}{t} = \lim_{t \rightarrow \infty} \frac{\log c_j}{t} = \lim_{t \rightarrow \infty} \frac{-\log p_j}{t} = \phi_j \quad (165)$$

Proof. The result can be proven by simply verifying that it satisfies the equilibrium conditions. An additional result we need is the use of inputs, $x_{i,j}$. The limit is

$$\lim_{t \rightarrow \infty} \frac{\log x_{i,j}}{t} = \phi_{j^*(i)} + \frac{1}{(1 - \gamma_i)} [\phi_j - \phi_{j^*(i)}] \quad (166)$$

where

$$j^*(i) = \begin{cases} \arg \min_{j \in S(i)} \phi_j & \text{if } \gamma_i < 0 \\ \arg \max_{j \in S(i)} \phi_j & \text{if } \gamma_i > 0 \end{cases} \quad (167)$$

We need to verify that the above, along with the solution in the proposition, satisfies the limits of the two FOCs, the resource constraint, and the production function.

We first take limits of the equilibrium conditions. For any variable g_j , define

$$\phi_{g,j} \equiv \lim_{t \rightarrow \infty} \frac{\log g_j}{t} \quad (168)$$

First, a small lemma:

$$\phi_{x,i,j} \leq \phi_j \quad (169)$$

To see why, first suppose $\gamma_i < 0$. Then $\phi_j - \phi_{j^*(i)} \geq 0$ and $\frac{1}{(1-\gamma_i)} < 1$, from which the result immediately follows. To see the result for $\gamma_i > 1$, note that

$$\phi_{x,i,j} = \phi_{j^*(i)} + \frac{1}{(1-\gamma_i)} (\phi_j - \phi_{j^*(i)}) \quad (170)$$

$$= \phi_j + \frac{\gamma_i}{(1-\gamma_i)} (\phi_j - \phi_{j^*(i)}) \quad (171)$$

Since $\frac{\gamma_i}{1-\gamma_i} > 0$ and $\phi_j - \phi_{j^*(i)} \leq 0$ in this case, the result again follows. It holds trivially for $\gamma_i = 0$. Furthermore, note that

$$f_i(\phi_{x,i,j}) = \phi_{j^*(i)} \quad (172)$$

Then the limits of the three equilibrium conditions and the production function are (equations (159) to (162), respectively)

$$\phi_{y,i} = \zeta_i + \alpha_i f_i(\phi_{x,i,j}) \quad (173)$$

$$\phi_{y,j} = \max \left\{ \phi_{c,j}, \max_i \phi_{x,i,j} \right\} \quad (174)$$

$$\phi_{p,j} = -\phi_{c,j} \quad (175)$$

$$\phi_{p,j} = \phi_{p,i} + \zeta_i + \frac{\alpha_i - \gamma_i}{\alpha_i} (\phi_{y,i} - \zeta_i) + (\gamma_i - 1) \phi_{x,i,j} \quad (176)$$

where the first equation uses equation (169). The first and second equations hold because of (172). The third is trivial. The fourth holds by substituting in the various ϕ terms and using the recursion defining ϕ . ■