The effects of large shocks in production networks

Ian Dew-Becker*

June 29, 2021

Abstract

This paper studies the response of a production network economy to large shocks. It derives limiting expressions for the effect of arbitrary combinations of shocks on output and prices at both the micro and aggregate levels. Unlike Taylor series, the method has guaranteed accuracy as the size of the shock grows. Combined with a probability distribution over the shocks, the analysis shows how the shock distribution and the structure of the economy combine to determine the magnitude of tail risk and asymmetry in aggregate output. The riskiest economies involve production functions featuring complementarity and have high interconnectedness. Depending on how heavy the tails of the shock distribution are, large deviations in GDP may be guaranteed to come from a single type of shock or they may come from arbitrary combinations.

1 Introduction

Background

If utility is concave, then the events that have the largest individual impact on welfare are large declines in consumption. A large literature studies large movements in GDP, trying to understand their likelihood, their sources, and their effects on welfare and other features of the economy. Acemoglu et al. (2017), for example, show that large movements in GDP are more likely that predicted by the normal distribution, they and Gabaix (2011) suggest that such events could be caused by shocks to influential sectors or firms, and Barro (2006) studies how such movements in GDP and consumption might affect asset prices.

The probability distribution of GDP is a function of the distributions of the fundamental shocks hitting the economy and how the economy's structure maps those shocks into final

*Northwestern University and NBER. This paper would not exist without Alireza Tahbaz-Salehi. I appreciate helpful comments from seminar participants at Northwestern.

output. What makes understanding that mapping particularly difficult is that, except in knife-edge cases, the mapping is nonlinear. The goal of this paper is to understand the determinants of the distribution of GDP in its tail. In particular, how does the structure of the economy determine its sensitivity to different shocks or combinations of shocks, and what types of shocks are most likely to cause crashes?

Contribution

The paper's core contribution is to answer those questions in the context of a network production model. In particular, the paper shows, in closed form, how large microeconomic shocks translate into large movements in GDP in the context of a general network production model. The probability distribution of the shocks and the structure of the economy combine to determine the probability of large deviations in GDP.

As byproducts, the analysis delivers both a novel measure of tail centrality, measuring the ability of a shock to a single economic unit to produce a large (as opposed to small) movement in GDP, and also shows precisely what combinations of micro shocks cause crashes. These results help clarify what factors make a firm or sector systemically risky.

Methods

The paper studies a general network production model of the form studied by Baqaee and Farhi (2019). Economic units – sectors or firms – produce outputs using as inputs both labor and the products of other units. Some of the products are also combined into the final consumption good. The various economic units thus interact through the production network, propagating and potentially amplifying or attenuating shocks.

Importantly, the model allows for arbitrary elasticities of substitution across inputs in the production functions. Since the elasticities are in general different from 1, the model is nonlinear and cannot be solved in closed form.

Theoretical analyses of structural models rely either closed-form solutions available for highly restricted models, or on approximations that are only valid locally. Typically one would analyze a Taylor series, for which accuracy is only guaranteed local to the approximation point. In numerical work sometimes "global" methods are used, but those solutions still only guarantee accuracy in some bounded range. And in both cases, when the exogenous shocks diverge towards $\pm\infty$, the errors in the approximation in general diverge as well.

This paper also studies a series expansion, but one taken at infinity. Formally, for any combination of shocks, the analysis studies a polynomial approximation as the scale of the shocks grow (think of a polar representation for a shock and then sending the radius to ∞). The mathematical ideas are formalized in the paper, but the key point is simply that instead of having guaranteed accuracy local to some specific point, the guarantee holds as the shocks diverge. If one's goal is to understand *large*, rather than small shocks, the methods used here are natural.

To understand the distribution of GDP, the approximation for GDP in terms of the shocks is combined with a description of the tail distribution of the shocks. That allows one to both describe the tails of the distribution of GDP and also answer the question of what types of shocks push the economy into the tail, causing crashes or booms.

Results

The paper's key approximation result – the Taylor series at infinity – is that for any combination of shocks, GDP eventually approaches a (log-)linear asymptote. To see the intuition, think of an economy with just two shocks. Looking along any ray emanating from the origin, log GDP eventually approaches a linear asymptote.

Figure 1 provides a visualization. On the left hand side are standard first- and second-order approximations around shock values of zero. On the right-hand side are the tail approximations, which are the straight lines that log GDP approaches as the shock becomes especially positive or negative. The paper shows that not only is the tail of GDP linear, but true GDP approaches the asymptote exponentially fast. The fact that those two lines are different shows that the approximation is able to capture the nonlinearity and hence asymmetry in GDP.

As a byproduct, the approximation delivers a concept of centrality for each unit that is valid in the tail. In figure 1, note that the effect of a large shock is very different from the local effect of a small shock. When the goal is to understand large movements in GDP, naturally the large shocks are more relevant. The paper gives closed-form expressions for the slopes in the tail. In the case where all inputs are complements, there is a particularly simple result: the left tail centrality of a sector is measured by the average downstream closeness of the sector to every other sector, weighted by their consumption shares. The most influential sectors are those for which much or all of GDP is downstream – i.e. universal inputs.

In addition, the results show that those tail slopes need not have any particular relationship with the local slopes for small shocks. It is easy to construct examples in which sectors or firms that are unimportant for small shocks, think of perhaps electricity or semiconductor production, become dominant when shocks are large. At the same time, sectors that are very large, like restaurants, may have relatively small effects in the tail.

The results described so far describe the mapping from productivity to GDP. To describe the probability distribution of GDP, one must make assumptions about the probability distribution of the shocks. The paper shows what type of assumption one needs about the shocks – really just a description of their tails – and how that then determines the tail distribution of GDP – the likelihood of extreme events, both positive and negative.

Due to the asymptotic linearity of GDP, the tail shape of the distribution of GDP inherits the properties of the tail of the distribution of shocks. It is also possible to determine which units are most prone to cause crashes, which again depends on their relative riskiness and their tail centrality.

An interesting feature of the model's description of the tail risk associated with different firms and sectors is that it maps well into the FSOC's criteria for determining systemically important firms. The FSOC's considerations include size, interconnectedness, lack of substitutes, and various measures of risk (leverage, maturity mismatch, etc.). Those factors turn out to appear in the model endogenously as central determinants of the tail centrality of a sector, and the theory helps clarify their precise roles.

While the majority of the analysis focuses on the behavior of GDP, unit output and relative prices are equally easy to characterize. Those results can then be used to show exactly how extreme shocks propagate through the production network. Intuitively, there is effectively a tail network for every shock, where one input becomes the binding constraint for each unit.

Finally, the last section of the paper measures tail centrality in the data, after adding some additional assumptions. 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.

To summarize, then, the paper's contribution is to develop an approximation for the tail of GDP in terms of micro shocks. That result is used to show how large shocks to units or combinations of units affect aggregate output, thus determining the statistical distribution of GDP and also the likelihood that any particular type of shock creates a crash (or boom).

Related literature

As discussed above, this paper contributes to the literature on large movements in aggregate output, including Rietz (1988), Barro (2006), and Acemoglu et al. (2017). Barrot and Sauvagnat (2016) and Carvalho et al. (2020) study the effects of large shocks to individual firms due to natural disasters on production. Acemoglu et al. (2017) study how heavy tails in GDP can be generated by the production structure of the economy, but in the context of a fully linear (and hence analytically solvable) model. Liu and Tsyvinski (2021) consider the effects of shocks in a dynamic version of a network model.

The paper's framework is builds on the literature on production networks, which in its modern form is traced to the work of Long and Plosser (1983). The closest link is to Baqaee and Farhi (2019), who study higher moments of output in the same nonlinear framework. The key distinction is that that paper analyzes output based on a second-order Taylor series approximation. The analysis is therefore useful for shedding light on how the structure of the economy generates asymmetry in the local response of GDP to small shocks, but it does not speak to large deviations as its approximation 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.[1] A contribution of this paper is to potentially giving a way for work in those areas to get analytic approximations where they were previously unavailable.

Related to this paper and Baqaee and Farhi (2019) also is the work on asymmetry in GDP, including Dew-Becker, Tahbaz-Salehi, and Vedolin (2021), Dupraz, Nakamura, and Steinsson (2020), and others. While those papers primarily focus on asymmetry at the business cycle frequency, and hence for fairly "typical" sorts of shocks, this paper analyzes the determinants of asymmetry in the tails, giving, for example, a potential explanation why there are occasionally extremely large declines in output but never equally large increases.

Outline

The remainder of the paper is organized as follows. Section 2 describes the basic structure of the economy. Section 3 presents the main theoretical result on approximating output in terms of the exogenous shocks (i.e. figure 1). That result is then used to develop a novel concept of the importance of economic units – their tail centrality – in section 4, with section 5 presenting examples. Section 6 presents the second main result of the paper, showing how to combine the first result with assumptions about the distribution of micro shocks to obtain the distribution of GDP. Section 7 then presents results on the behavior of unit-level output and prices during extreme events. Finally, section 8 presents simple estimates of tail centrality based on the US input-output matrix. Finally, section 9 concludes.

2 Structure of the economy

This section describes the structure of the economy. The model is static and frictionless and takes the form of a standard 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

[1] Liu (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.

sector or a firm. Each unit has a CES production function of the form

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

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 $1, ..., N$). The fact that labor in (1) has a unit elasticity of substitution with material inputs is without loss of generality – one can always specify an additional unit that converts labor into labor services, which are then combined with other inputs with a non-unitary elasticity (see section 5.2).

σ_i is the elasticity of substitution across material inputs for unit i. In the limit where $\sigma_i \to 1$, the production function becomes Cobb–Douglas (with the $A_{i,j}$ becoming the exponents). As discussed in Baqaee and Farhi (2019), the specification in (1) is general enough to accommodate any nested-CES economy with an arbitrary number of producers, CES nests, elasticities, and intermediate input use by treating each CES aggregator as though it is a distinct producer. Finally, though I assume a CES specification for simplicity, appendix B shows that the results also hold under much more general conditions.

The parameters $A_{i,j}$, normalized such that $\sum_j A_{i,j} = 1$, determine the relative importance of different inputs. If $A_{i,j} = 0$, then unit i does not use good j in production. The set of production weights can be put into a matrix, A, with representative element $A_{i,j}$, which then represents the economy's production network. The network is a weighted directed graph, where there is an edge from unit i to unit j if $A_{i,j} > 0$. An important feature of the graph will be the set of upstream links of each unit, denoted by

$$S_i \equiv \{j : A_{i,j} > 0\} \tag{2}$$

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

$$U(C_1, ..., C_N) = \left(\sum_i A_{0,i}^{1/\sigma_0} C_i^{(\sigma_0 - 1)/\sigma_0} \right)^{\sigma_0 / (\sigma_0 - 1)} \tag{3}$$

where, again, the weights $A_{0,i}$ sum to 1. The subscript 0 (as in σ_0) will be used to denote the consumption aggregator. The representative agent purchases C_i units of good i with wages and inelastically supplies a single unit of labor so that $\sum_i L_i = 1$.

There are two extreme cases for labor that are both straightforward to analyze: a perfectly integrated market in which all units buy labor competitively for some (equilibrium) wage

W, or totally inflexible labor, in which L_i is fixed for all i and thus cannot be reallocated across sectors. The main results are identical in both cases. The paper text focuses on the case with flexible labor because the proofs are simple and transparent, but the main results also hold with fixed labor.

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, $P_iY_i - WL_i - \sum_j P_jX_{i,j}$, the representative consumer maximizes utility, producers and the consumer take prices as given, and markets clear:*

$$Y_i = C_i + \sum_j X_{j,i} \tag{4}$$

Since there is no government spending or investment, GDP is equal to the utility aggregator over consumption, $GDP = U(C_1, ..., C_N)$.

Except for a few knife-edge cases, the model does not have a closed form solution, which has led almost all past work to study first-order approximations local to some point (or, in the case of Baqaee and Farhi (2019), a second-order approximation).

2.1 Hulten's theorem

The analysis in the remainder of the paper focuses primarily on logs of the variables, which are denoted by lowercase letters, e.g. $y_i = \log Y_i$. Without loss of generality, I take the "steady-state" level of productivity to be $z_i = 0 \; \forall \; i$, by which I mean just some average or benchmark level.

Around any point,

$$\mathrm{d}gdp = \sum_i \frac{Y_iP_i}{\sum_j C_jP_j} \mathrm{d}z_i \tag{5}$$

That is, the first-order Taylor approximation for gdp in terms of unit productivity depends only on each unit's sales relative to nominal GDP at that point. Those shares are referred to as sales shares or Domar weights.

If we specialize the above to take derivatives around the steady-state, it can be compactly written in terms of the production network. Define the matrix A to have representative element $A_{i,j}$. All indexing or vectors and matrices throughout the paper use zero-based indexing, with the first element indexed by 0 (see also Baqaee and Farhi (2019)). Then

$$\mathrm{d}gdp = e_0 (I - A)^{-1} \mathrm{d}z \quad (6)$$

$$= D_{ss,i} dz \quad (7)$$

where e_0 is a vector equal to 1 in element 0 and 0 elsewhere (i.e. $e_0 = [1, 0, ...]'$), z is the vector of log productivities, and D_{ss} is the vector of steady-state Domar weights, with $D_{ss,i} = \frac{P_i Y_i}{\sum_j C_j P_j}$ with prices and quantities measured at steady-state. In a local sense, then, there are two ways to think about centrality. The first is just that big units, with large Domar weights, $D_{ss,i}$, are central. To the first order, if you know each unit's size, in terms of its sales, you do not need to know anything else. Sales are a function of the production network, though, so the second way is in terms of $(I - A)^{-1}$ and the consumption weights. The sales shares are sufficient statistics for the first-order approximation, and the sales shares themselves depend on all of the parameters of the model.

But the economy here is nonlinear, so the first-order approximation does not fully capture its response to shocks, especially large ones, and the steady-state sales shares are not sufficient statistics. The second-, third- or in fact any order derivatives are not sufficient statistics, either. Finding a way to calculate the effects of large shocks is the focus of this paper, and the next section shows how to do it.

3 Large shock limit

This section presents the paper's key result characterizing output as a function of productivity. It is a limiting result for large shocks, taking the vector of log productivities to infinity in some direction.

An important first observation is that if the economy exhibits any nonlinearity, the impact of various shocks cannot be studied independently: nonlinearities in production and/or consumption result in non-trivial interactions between shocks' impacts. Those interactions can be particularly strong when shocks are large.

For a given vector of productivities, $z = [z_1, ..., z_n]'$, we can always write z in a polar form,

$$z = \theta t \quad (8)$$

$\theta \in \mathbb{R}^N$ is a vector representing a direction in productivity space, while t is a scalar controlling magnitude. While it would be typical in a polar representation to normalize θ to have unit length, it will prove useful later on to avoid that restriction. The representation $z = \theta t$ is thus not unique (and it need not be) – θ can be scaled up or down by any factor. For the purpose of developing intuition, though, nothing will be lost for now in just thinking about the case where $\|\theta\| = 1$.

As two simple examples, studying the effects of a large shock to a single unit would correspond to setting θ equal to zero except for a single entry, while studying the effect of simultaneous equal shocks to all units would correspond to setting θ to have the same value for all entries (e.g. $1/\sqrt{N}$). θ determines the mixture of shocks to different units embodied in z.

3.1 Main result

Theorem 1. *Write* $z = \theta t$.

Part 1: *There exist unique scalar-valued functions* $\lambda(\theta)$ *and* $\mu(\theta)$ *independent of* t *such that*

$$\lim_{t\to\infty} |gdp(z) - (\mu(\theta) + \lambda(\theta)t)| = 0 \tag{9}$$

Part 1 of theorem 1 says that log GDP has a formal linear approximation as $t \to \infty$. For any vector of shocks, as the shocks grow in scale, holding the direction fixed, *gdp* eventually converges to a linear asymptote. That formalizes the graphical intuition in figure 1.[2]

Taylor series are accurate local to some point (holding the order fixed), i.e. guaranteeing convergence as $t \to 0$. The approximation here, on the other hand, guarantees convergence as $t \to \infty$. That is exactly the form of convergence that is relevant when one is concerned with accuracy for large rather than small shocks.

Since the parameter t is non-negative, the approximation is piecewise linear. That is, $\lambda(\theta) \neq \lambda(-\theta)$ and $\mu(\theta) \neq \mu(-\theta)$. That idea is illustrated in figure 1, where the slope to the left is different from the slope to the right.

It is straightforward to characterize the slope function, $\lambda(\theta)$:

Theorem 1. ***Part 2:*** *The slope in equation (9) is* $\lambda(\theta) = \phi_0$*, where* $\phi \in \mathbb{R}^{N+1}$ *is a function of* θ *that is implicitly defined by the system*

$$\phi_i = \theta_i + \alpha_i f_i(\phi) \;\; \text{for } i \in \{0, 1, ..., N\} \tag{10}$$

[2] That result also immediately implies that μ is homogenous of degree zero and λ homogenous of degree one ($\mu(k\theta) = \mu(\theta)$ and $\lambda(k\theta) = k\lambda(\theta)$ for a positive scalar k).

with $\theta_0 \equiv 0$ *and* $f_i : \mathbb{R}^{N+1} \to \mathbb{R}$ *is defined as*

$$f_i(\phi) = \begin{cases} \max_{j \in S_i} \phi_j \text{ if } \sigma_i > 1 \\ \sum_j A_{i,j} \phi_j \text{ if } \sigma_i = 1 \\ \min \phi_j \text{ if } \sigma_i < 1 \end{cases} \tag{11}$$

for $i \geq 0$*.* ϕ *is unique and continuous in* θ*, as is* $\lambda(\theta)$*.*

The response of *gdp* to a productivity shock in the direction θ depends on the solution to a system of equations. Naturally, the slope depends on the direction. That is, shocks to some units affect GDP more than others. $\lambda(\theta)$ will be larger than if θ represents a shock to an important unit than to an unimportant unit. The concept of importance will be further expanded below. Note also that λ is continuous in θ, so that two shocks θ and θ' that are similar in the sense that θ is close to θ' will also have similar impacts on the economy, with $\lambda(\theta)$ close to $\lambda(\theta')$.

What is most interesting about the system is what does *not* affect it: the exact values of the production weights and the elasticities of substitution. The production weights only matter to the extent that they are zero versus positive, determining what inter-industry links exist. Similarly, the elasticities only matter for being greater than, less than, or equal to 1. Another way to say it is that all that matters in the production network is the set of edges – not their weight – and how they are combined – linearly (in logs), through a concave aggregator, or through a convex aggregator.

Intuitively, the form of f is due to the fact that the CES aggregator behaves like a maximum or minimum as the scale of the inputs diverges (it is referred to in the neural networks literature as the softmax/softmin function for precisely this reason).

The recursion thus gives a formal and simple way to see how complementarities amplify and propagate negative shocks and attenuate and isolate positive shocks (with substitutability doing the opposite). A negative shock in a unit will propagate downstream to any unit with $\sigma_j \leq 1$, and when the inequality is strict it will eventually dominate any positive shocks from other units.

3.2 Approximation errors

Proposition 1. *As* $t \to \infty$*, the error from a finite order Taylor series for gdp around any value of* z *will diverge to* $\pm\infty$ *unless* $\sigma_i = 0 \ \forall \ i \geq 0$*.*

That fact holds due to the fact that *gdp* has a linear asymptote – no Taylor series can match that behavior unless the true function is linear. For the tail approximation, on the other hand, we have the following,

Corollary 1.1. *The errors in the tail approximation for gdp are globally bounded in the sense that there exists a δ such that*

$$|gdp(\theta t) - (\mu(\theta) + \lambda(\theta) t)| < \delta \tag{12}$$

for all θ and $t \geq 0$.

The tail approximation is not just accurate in the tails: in fact its errors are globally bounded. The next section discusses the bound on the errors further, but for now note that this is a property that does not hold for any Taylor series around steady-state in this model – quite the opposite: any Taylor series will have errors that eventually diverge – and in fact does not in general hold in approximations typically used in economics.

Beyond the bound, though, it is valuable to get a more formal quantitative statement about how fast the approximation errors converge to zero as t grows. We have the following stronger form of theorem 1 part 1:

Theorem 1. ***Part 3:*** *Equation (9) in part 1 can be replaced by*

$$\lim_{t\to\infty} |gdp(z) - (\mu(\theta) + \lambda(\theta) t)| \, t^j = 0 \tag{13}$$

for any j.

In other words, it is not only the case that a linear approximation for *gdp* exists as $t \to \infty$, but in fact no other powers of t appear in the series expansion. That is, the errors in the approximation converge to zero faster than any polynomial grows as $t \to \infty$.[3]

Remark. *The above can be pushed further to show that the convergence is exponential at a rate that increases with $|\sigma_i - 1|$.*

The next section gives another way to see how $|\sigma_i - 1|$ determines the accuracy of the approximation.

3.3 When is the tail approximation the right approximation?

The usual approximation is around $z = 0$, while this paper focuses on $z \to \infty$. As z grows, the approximation from theorem 1 is eventually superior. But at what point does that transition happen?

[3] Formally, there is a singularity at $t = \infty$, so that the series approximation at ∞ is not convergent. It is similar to looking at a Taylor series for $\lambda x + \exp\left(-1/\left(x^2\right)\right)$ at $x = 0$. The derivatives at all orders above 1 are equal to zero.

To shed light on that question, first note that $gdp\left(0\right) = 0$. So to know the size of the error from using the tail approximation when $z = 0$, we need to know the constants, $\mu\left(\theta\right)$. We have

Theorem 1. ***Part 4:*** *The constant in equation (9) is $\mu\left(\theta\right) = \mu_0$, where the vector μ solves the recursion*

$$\mu_i = \frac{\alpha_i}{1-\sigma_i} \log\left(\sum_{j\in j^*(i)} A_{i,j} \exp\left(\left(1-\sigma_i\right)\mu_j\right)\right) \tag{14}$$

with

$$j^*\left(i\right) \equiv \begin{cases} \{j : \phi_j = \min_{k\in S_i} \phi_k\} & \text{if } \sigma_i < 1 \\ \{j : \phi_j = \max \phi_k\} & \text{if } \sigma_i > 1 \\ \{j\} & \text{if } \sigma_i = 1 \end{cases} \tag{15}$$

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

In the case where *gdp* is globally concave in the shocks, we can do even better. That happens when $\sigma_i \leq 1 \; \forall \; i$. In that case, the error for the tail approximation is smaller when

$$t > \frac{\mu\left(\theta\right)}{D'_{ss}\theta - \lambda\left(\theta\right)} \tag{16}$$

The tail approximation is superior if t is sufficiently large – larger when the constant $\mu\left(\theta\right)$ is larger or the gap between the local and tail approximations, $D'_{ss}\theta - \lambda\left(\theta\right)$, 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.

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 $\lambda\left(\theta\right) < 0$ – i.e. when thinking about shocks that reduce GDP – the tail approximation cannot possibly be the better of the two until $\mu\left(\theta\right) + \lambda\left(\theta\right) t < 0$, and the point where that happens necessarily increases as the A parameters for the minimizing units (i.e. the units $j \in j^*\left(i\right)$ for some i) decline.

Overall, then, when elasticities of substitution are closer to 1, or when the units that are relevant in the tail (in the sense of being in the set $j^*\left(i\right)$ for some i) have smaller production weights, the tail approximation will tend to be less accurate for small t, and when the elasticities are further from 1, the tail approximation will outperform the linear approximation for smaller values of t.

4 Tail centrality

A key question in the networks literature, both in economics and more broadly, is what nodes of the network – which might represent sectors or firms, in this case – are most influential. That issue touches on one of the motivating questions of this paper, which is what makes a firm or sector systemically risky, in the sense that a shock to that unit of the production network may be prone to propagate to the rest of the economy.

As discussed above, Hulten's theorem says that local centrality – the impact of a small shock to a unit on GDP – is exactly the sales share of the unit. This section measures the importance of units based on how *large* shocks to their productivity affect GDP, which is typically closer to the spirit of the question being asked when trying to evaluate systemic risk.

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

$$\gamma_i^L \equiv \lim_{\Delta z_i \to -\infty} \frac{\Delta gdp}{\Delta z_i} \tag{17}$$

$$\gamma_i^R \equiv \lim_{\Delta z_i \to \infty} \frac{\Delta gdp}{\Delta z_i} \tag{18}$$

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

The tail centralities measure how systemic a large shock to a given unit is in the sense of passing through to GDP. They are defined without any reference to theorem 1. A unit with a large left tail centrality is one where a large negative shock has a large effect on GDP. Right tail centrality, instead of measuring systemic risk, measures the extent to which a large boom in a particular unit – say due to a new invention that massively increases productivity – will benefit the rest of the economy.

As discussed above, Hulten's theorem says that the Domar weights – nominal sales shares – determine the effect of a small shock to a given unit on GDP. They are in fact the same limit as in definition 1, but when $\Delta z \to 0$. The tail centralities, on the other hand, measure the effect of large shocks. One way to interpret them is that they are the limiting values of the Domar weights as $z_i \to \pm\infty$. The gap between the Domar weights and tail centralities can be interpreted as a measure of the extent to which nonlinearities in the economy affect the response to shocks to a particular unit.

From theorem 1, it is trivial to obtain the tail centralities:

Corollary 1.2. *Let e_i denote a vector equal to 1 in element i and zero otherwise. Then in* *the notation of theorem 1,*

$$\gamma_i^L = \lambda(-e_i) \tag{19}$$
$$\gamma_i^R = \lambda(e_i) \tag{20}$$

As in theorem 1, an immediate consequence of this result is that while tail centralities depend on whether each industry's intermediate inputs are substitutes or complements, they do not depend on the exact values of the corresponding elasticities of substitution. Similarly, except for the knife-edge case of $\sigma_i = 1$, input-output linkages only matter via the identities of each industry's suppliers, and the intensities of such relationships are immaterial for tail centralities. More importantly, the recursive characterizations in theorem 1 yield comparative static results on how various structural features of the economy shape tail centralities.

Proposition 2. *Let γ^L and γ^R denote the vectors of left and right tail centralities, respectively.*

1. *If $\sigma_i \leq 1$ for all i, then $\gamma^L \geq \lambda^{\text{ss}} \geq \gamma^R$.*
2. *If $\sigma_i \geq 1$ for all i, then $\gamma^L \leq \lambda^{\text{ss}} \leq \gamma^R$.*
3. *If $\sigma_i = 1$ for all i, then $\gamma^L = \lambda^{\text{ss}} = \gamma^R$.*

Statement (a) formalizes the intuitive idea that complementarities amplify the impact of large negative shocks, while attenuating that of large positive shocks. Put differently, when inputs are complements, a large negative shock to any industry results in lager contractions in aggregate output than what would have been predicted based on that industry's size, whereas a large positive shock would result in a smaller expansion compared to the Cobb-Douglas benchmark. Statement (b) shows that substitutability of inputs has the reverse effect: it amplifies the aggregate impact of positive shocks, while attenuating that of negative shocks. Trivially, and as expected, in the special case that all production functions and preferences are Cobb-Douglas, left and right centralities coincide with Domar weights.

To state the next result, define industry j to be *downstream* of industry i and i to be *upstream* of j if there is a directed path from industry i to industry j over the production network.[4]

[4]Note that if the production network has cycles, an industry can be simultaneously upstream and downstream to another industry. Furthermore, according to this definition, when the production network is strongly connected—in the sense that there exists a directed path from each industry to any other industry—all industries are upstream to one another.

Proposition 3. *The left and the right tail centralities of industry i are, respectively, weakly decreasing and weakly increasing in the elasticities of substitution of i's downstream industries.*

This result distills the essence of the recursive characterizations in theorem 1 part 2 by capturing how each industry's tail centralities are shaped by nonlinearities in the economy's disaggregated structure. In particular, it illustrates that replacing a production technology that exhibits input substitutability ($\sigma_j > 1$) with one that exhibits input complementarity ($\sigma_j < 1$) increases the left tail centrality of all its upstream industries, while simultaneously reducing their right tail centralities. The intuition for this result is once again straightforward: input complementarity in industry j makes j more vulnerable in the face of large negative shocks to any of its (direct or indirect) upstream suppliers i, thus amplifying the aggregate impact of a large negative shock to i.

A similar logic also implies that, all else equal, an expansion in the set of inputs of industry j would make the nonlinearity in j's production technology matter more:

Proposition 4. *Suppose industry i is upstream to industry j and let S_j denote the set of j's suppliers.*

1. *If $\sigma_j < 1$, the left and the right tail centralities of i are, respectively, weakly increasing and weakly decreasing in S_j.*

2. *If $\sigma_j > 1$, the left and the right tail centralities of i are, respectively, weakly decreasing and weakly increasing in S_j.*

A simple but important consequence of the above result is that the left tail centrality of industry i is larger if it is a direct or indirect input supplier to more industries whose inputs are complements. In particular, negative shocks to such an industry propagate to more downstream industries via intermediate inputs that are not easily substitutable, making the macroeconomic impact of large negative shocks to i more pronounced.

The result says that when the economy is more interconnected, there is more fragility if inputs are complements. If they are substitutes, more connections lead to more resilience

The final result in this section establishes that an industry's tail centralities are larger the more important intermediate inputs are for the production technologies of its downstream industries.

Proposition 5. *The left and right tail centralities of industry i are increasing in the intermediate input shares of all industries that are downstream to i.*

5 Examples

This section discusses two simple example economies. The examples demonstrate, among others, the following aspects of the model:

1. Tail centralities and steady-state Domar weights need not have any relationship: arbitrarily small units can have large tail centrality;

2. Tail centralities are unrelated to the number of units or the diversification of the economy, and can remain large even as the economy diversifies;

3. Under somewhat weak conditions, the tail centrality of a unit depends on the average downstream distance between the unit and final output

The first example is of a reasonably general network, which is highly similar to the parametric model studied in Baqaee and Farhi (2019). The second example is more specialized and features a universal input, which could be thought of as representing, for example, electricity.

5.1 Example 1: complementary production

Consider an economy in which all elasticities of substitution in production are less than 1: $\sigma_i < 1$ for $i \geq 1$, while the elasticity of substitution in consumption is 1, $\sigma_0 = 1$. Assume that $\alpha_i > 0$ and $A_{i,i} \in (0,1)$, which guarantees that every unit uses at least two inputs, one of which is its own output (which is true for 88 percent of units according to the BEA). Otherwise, the $A_{i,j}$ and β_i parameters are unconstrained.

Using the result from theorem 1 part 2, it is then immediate that $\gamma_i^R = \beta_i \;\forall\; i$, which also immediately implies that $\gamma_i^R \leq D_{ss,i}$. That is, in an economy with complementary production, each unit's right tail centrality is just equal to its share of the consumption basket – 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 ($f_i\left(x\right) = \min\left(x\right) \;\forall\; i$). That implies that

$$\gamma_i^L = \frac{1}{1-\alpha_i} \sum_{j=1}^{n} \beta_j \min_{\pi \in D(j,i)} \prod_{r \in \pi} \alpha_r \tag{21}$$

where $D(j,i)$ denotes the set of all directed paths from industry i to industry j over the production network. In the special case that all intermediate input shares are identical

($\alpha_i = \alpha$), (21) simplifies to

$$\gamma_i^L = \frac{1}{1-\alpha} \sum_{j=1}^{n} \beta_j \alpha^{d_{\min}(j,i)} \tag{22}$$

where $d_{\min}(j,i)$ is the length of the shortest directed path from i to j.[5] In the complementary economy, a unit's left tail centrality is measured by its average downstream closeness to final consumption. In equation (22), 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.

Equations (21) and (22) answer the question of what types of units have high tail centrality: those that are direct suppliers to producers of a large fraction of GDP (and that do not have substitutes). Those results are also consistent with the propositions in the previous section: γ_i^L is always greater than or equal to $D_{ss,i}$, is increasing in α, and can only increase when other units use more suppliers.

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
3. An increase in the share of expenditures on material inputs (α) by units downstream of i
4. An increase in the materials share of input expenditures, α

In some cases, one might think of the fundamental shock being to the relative price of a given good, or to the quantity available, as opposed to productivity. A shock to prices might correspond to a good in elastic supply whose price is determined by a global market, such as oil. A shock to supply could correspond to some fixed factor, such as capital or labor.

Proposition 6. *Denoting log output in unit i as a function of productivity by $Y_i(z)$, in the complementary economy,*

$$\lim_{t\to\infty} \frac{gdp(\zeta^i t)}{\log Y_i(\zeta^i t)} = (1-\alpha_i)\gamma_i^L \tag{23}$$

$$\lim_{t\to\infty} \frac{gdp(\zeta^i t)}{\log p_i(\zeta^i t)} = -(1-\alpha_i)\gamma_i^L \tag{24}$$

[5] $d_{\min}(j,i) = \min\{|\pi| : \pi \in D(j,i)\}$

That is, the same concept of tail centrality determines not just the relationship of GDP with unit productivity shocks, but also its relationship with unit output and prices.

Intuitively, the results in this section suggest that the out-degree of a unit – the number of units directly downstream of it – would be closely linked to tail centrality. Define the weighted out-degree of a unit to be

$$d_i \equiv \sum_{j:i\in S(j)} \beta_j \tag{25}$$

Proposition 7. *In the case where $\alpha_i = \alpha \; \forall \; i$, tail centrality satisfies*

$$\frac{1}{1-\alpha}\left(\beta_i + \alpha d_i\right) \leq T_i \leq \frac{1}{1-\alpha}\left(\beta_i + \alpha d_i + \alpha^2\left(1-d_i\right)\right) \tag{26}$$

5.2 Example 2: a universal input

Suppose there is a sector that produces a universal input, perhaps electricity, denoted by E, that is produced according to the simple function

$$E = \exp\left(z_e\right) \tag{27}$$

That is, it involves no inputs and depends just on an exogenous shock (one might also think of this as an exogenous supply of oil, for example).

Output in the N other units in the economy is produced according to the function

$$Y_i = \exp\left(z_i\right)\left(\left(1-b\right)^{1/\sigma} L_i^{(\sigma-1)/\sigma} + b^{1/\sigma} E_i^{(\sigma-1)/\sigma}\right)^{\sigma/(\sigma-1)} \tag{28}$$

where L_i is labor in each unit.[6] GDP is an equal-weighted Cobb–Douglas aggregate of the outputs of the N sectors. Appendix A.5 shows that in this case GDP is

$$gdp = \frac{1}{N}\sum_i z_i + \frac{1}{\sigma-1}\log\left(1-b+b\exp\left(\left(\sigma-1\right)z_e\right)\right) \tag{29}$$

[6]The production function here assumes an elasticity of substitution between labor and materials of σ instead of 1. Recall the discussion from section 2 that to account for such a situation, one can always think of a sector that uses the original labor input from the household to produce labor services with the trivial production function $L = \tilde{L}$, where L is the output of labor services and $\tilde{L}$ the input of actual labor from the household. Those labor services are then sold to sectors. If $\tilde{L}$ is normalized to 1 as before, then $\sum_i L_i = L = \tilde{L} = 1$. So the setup here fits within the general setup laid out in section 2.

while the steady-state Domar weight and tail centralities of the electricity unit are

$$D_{ss,E} = b \tag{30}$$

$$\gamma_E^R = 0 \tag{31}$$

$$\gamma_E^L = 1 \tag{32}$$

Local to steady-state, electricity is just a single sector, and if b is small, electricity will be relatively unimportant. For large positive shocks, the effect of electricity is eventually zero – how much electricity do we really need?

Large negative shocks to the supply of electricity, on the other hand, have very large effects on GDP. Eventually, the supply of electricity becomes a binding constraint, so that a 1-percent change in the availability of electricity yields the same 1-percent change in final output.

This example demonstrates a few important features of the model. First, it formalizes the idea that a universal input can have small local effects on GDP but large effects when shocks become large. Second, it shows that there need not be any relationship between steady-state or average Domar weights and tail centralities. Third, tail centrality is separate from the number of units. The network structure can cause shocks to propagate such that tail centralities are, in fact, totally independent of the size of the economy, the number of units, and the local importance of any particular unit.

6 The risk of large deviations in GDP and their source

The results so far can be thought of as giving *ex post* results, in the sense that they describe what happens to GDP given some realization of productivity. This section presents *ex ante* results, which describe probabilities of various events conditional on distributional assumptions for productivity. The analysis here uses a reasonably general specification for shock distributions that captures the key forces in the model. Certain assumptions used here can be relaxed, but in general at the cost of both added complexity in the proofs and also typically having to put other restrictions on the distributions (e.g. limiting to particular families).

Two types of results are derived here: 1. characterizations of the probability of large deviations on GDP (and what about the economy determines that probability); and 2. characterizations about the probability of different combinations of shocks having been realized conditional on a large deviation in GDP. That is, we can describe what types of shocks cause "typical" large deviations. In particular, we can ask what units or combinations of units are most likely to be the sources of tail events.

6.1 Large deviations

For the purposes of studying large deviations, we only need to specify a distribution for z over its extreme values. There are various ways to do that (e.g. in terms of certain limits), but here I make the simplest possible assumption, which is simply to define the distribution of z only for large values of z, in a sense that will be formalized. That is, the distribution of z may take any form for relatively "small" values; we only need a specific characterization for the "large" values. This assumption is sufficient to illustrate the key determinants of large deviations in GDP.

I write z in the form

$$z = s\theta \tag{33}$$

where θ is an $N \times 1$ vector that is a member of some $(N-1)$-dimensional compact subset of $\mathbb{R}^N$, denoted Θ, with the property that any ray from the origin intersects Θ exactly once. For example, Θ could be a (hyper-)sphere, an ellipse, or a set such that some norm of θ is always equal to a fixed value.[7,8]

$s \geq 0$ is a scalar. The assumptions for θ and s mean that for any z there is a unique $s\theta$ decomposition. θ determines the angle and s the magnitude for the vector z.

I assume that there is an $\bar{s}$ such that for $s > \bar{s}$, the probability distribution for z can be described by a product measure in which s and θ are independent. That is, s and θ are asymptotically independent in the simple sense that they are independent for sufficiently large s.[9] Multiples of the set Θ then represent probability level sets for z relative to the origin.

I assume s has a cumulative distribution function F (with $\bar{F} \equiv 1 - F$) and θ a probability measure m_θ, so that

$$\Pr\left[s > \hat{s}, \theta \in \hat{\Theta}\right] = \bar{F}\left(\hat{s}\right) m_\theta\left(\hat{\Theta}\right) \text{ for } s > \bar{s} \tag{34}$$

[7]Spherical Θ corresponds to the family of spherical distributions, such as the normal. $\Theta \in \{\theta : \|\theta\|_1 = 1\}$, would arise in the case where the elements of z are independent and Laplace-distributed.

[8]Resnick (2007), for example, discusses this representation in the case of regularly varying distributions. Gupta and Song (1997) analyze this representation for random vectors that are spherically distributed under the L_p norm.

[9]A natural, though complicated, extension is to assume that s and θ are only ever approximately independent, converging toward being truly independent perhaps as $s \to \infty$. The details of that convergence in general depend on the distributional family of s, which is why I avoid it here. It also adds a second approximation term to handle in addition to that for GDP, making the analysis substantially more intricate.

and for some $\hat{\Theta} \subseteq \Theta$.

This representation is reasonably general and can accommodate many standard distributions studied in the literature such as multivariate normality, elliptical distributions more generally, transformations of Laplace distributed vectors, and the limiting form of regularly varying distributions (Resnick (2007)).

An alternative way of stating theorem 1 is

$$gdp = \mu(\theta) + s\lambda(\theta) + \varepsilon(gdp, \theta) \tag{35}$$

where the error term $\varepsilon(gdp, \theta) = o(1)$ in gdp and is bounded. In order for a given θ to be associated with asymptotically positive gdp is for $\lambda(\theta) > 0$. It is therefore useful to define the set

$$\Theta_+ \equiv \{\theta \in \Theta : \lambda(\theta) > 0\} \tag{36}$$

with Θ_- defined analogously and the same for any arbitrary $\Theta^* \subset \Theta$.

We then have the following general results

Theorem 2. *Under the conditions of theorem 1 and the distribution for z in (34), $\exists\, \bar{x}$ such that for $x > \bar{x}$*

$$\Pr[gdp < -x] = \int_{\Theta_-} \bar{F}\left(\frac{x - \mu(\theta) - \varepsilon(gdp, \theta)}{-\lambda(\theta)}\right) dm(\theta) \tag{37}$$

and by Bayes' theorem,

$$\Pr[\theta \in \Theta^* \mid gdp < -x] = \frac{\int_{\Theta_-^*} \bar{F}\left(\frac{x-\mu(\theta)-\varepsilon(x,\theta)}{-\lambda(\theta)}\right) dm(\theta)}{\int_{\Theta_-} \bar{F}\left(\frac{x-\mu(\theta)-\varepsilon(x,\theta)}{-\lambda(\theta)}\right) dm(\theta)} \tag{38}$$

Theorem 2 gives two key results for the tail distribution of GDP. First, the probability of large deviations in GDP depends on the probability of large deviations in productivity, scaled by the limiting slope, $\lambda(\theta)$. In other words, if one's goal is to understand the the risk of a large decline in GDP (or a large boom), the limiting representation of the economy in theorem 1 is the correct lens through which to analyze the economy. Other aspects of the economy – such as the steady-state Domar weights, or the precise values of the elasticities of substitution – are irrelevant.

Second, Bayes' theorem can be used to invert the probability distribution to find our what types of shocks – values of θ, specifically – typically cause large movements in GDP. The values of θ most likely to appear are, naturally, those for which $\lambda(\theta)$ is large, so that the magnitude of the shock, s, can be small. That is, on some level, not surprising – the units that GDP is most sensitive to are those that are most likely to create crashes. But part of the point of the result is that one must be careful to use the correct measure of sensitivity. It is not the steady-state Domar weight that matters, but rather the tail centrality. Large units need not be central, and central units need not be large.[10]

Finally, theorem 2 shows how nonlinearity in the economy generates increases in tail risk. If the economy were perfectly linear, the argument of $\bar{F}$ in (37) would be $\frac{x}{-D_{ss}'\theta}$. When $\lambda(\theta)$ is larger in magnitude than $D_{ss}'\theta$, there is a larger chance of a large movement in GDP. In the two examples studied in section 5, the left tail centralities are larger than the steady-state Domar weights. Equation (37) shows how that would increase left tail risk in the economy relative to what one would expect based on the Domar weights.

The next section specializes the above result to specific distributions, helping to clarify those two points and also to see precisely how s and θ interact.

6.2 Specific distributions

To see the use of those results more specifically, I now consider two major special cases, which cover most or all shock distributions typically analyzed in the literature.

6.2.1 Power law tails

Proposition 8. *Suppose s is distributed according to a power law for $s > \bar{s}$:*

$$\bar{F}(s) = c(s/\bar{s})^{-\kappa} \tag{39}$$

$$\text{where } c = \Pr(s \geq \bar{s}) \tag{40}$$

Then

1.

$$\lim_{x\to\infty} \Pr[gdp < -x] / \left(c\bar{s}^{\kappa}x^{-\kappa}\right) = \int_{\Theta_-} (-\lambda(\theta))^{\kappa}\, dm(\theta) \tag{41}$$

$$\lim_{x\to\infty} \Pr[gdp > x] / \left(c\bar{s}^{\kappa}x^{-\kappa}\right) = \int_{\Theta+} (\lambda(\theta))^{\kappa}\, dm(\theta) \tag{42}$$

[10]It is worth noting, and will be seen below, that the set Θ is also relevant. In some directions, the set Θ might be closer to the origin than others. The directions where Θ is close to the origin are ones where the shocks have a smaller scale, which will make them less likely to cause crashes, all else equal.

2.

$$\lim_{x\to\infty} \Pr\left[\theta\in\Theta^* \mid gdp < -x\right] = \frac{\int_{\Theta_-^*} \left(-\lambda\left(\theta\right)\right)^\kappa dm\left(\theta\right)}{\int_{\Theta_-} \left(-\lambda\left(\theta\right)\right)^\kappa dm\left(\theta\right)} \tag{43}$$

$$\lim_{x\to\infty} \Pr\left[\theta\in\Theta^* \mid gdp > x\right] = \frac{\int_{\Theta_+^*} \lambda\left(\theta\right)^\kappa dm\left(\theta\right)}{\int_{\Theta_+} \lambda\left(\theta\right)^\kappa dm\left(\theta\right)} \tag{44}$$

So when the shocks have power law tails, equation (41) gives two results for the tail of GDP. 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* depends on an average (with respect to the measure m) across all possible shocks of the tail slope, $\lambda\left(\theta\right)$. When the tail slopes tend to be larger, the probability of a large deviation in GDP is larger.

The integral on the right-hand side of (41) thus gives a formal measure of the fragility of the economy. Recall that $\lambda\left(\theta\right)$ comes from theorem 1 and depends only on the structure of the economy, not the shock distribution. So the integral shows how the structure of the economy determines the probability of a large decline in GDP. Economies in which that integral are larger have greater risk of large declines – in a sense are more fragile. The parameter κ affects the relative weighting of the integral. When κ is small, it is essentially an average of $\lambda\left(\theta\right)$ (with respect to the measure $m\left(\theta\right)$). When κ is larger, on the other hand, the integral is determined more and more by the largest values of $\lambda\left(\theta\right)$.

Equation (42) gives the probability of a large increase in GDP. The tail shape is again κ. Differences in the probability of large increases in large declines are determined by the average value of $\lambda\left(\theta\right)$ when $\lambda\left(\theta\right)$ is positive versus negative. Asymmetry in the distribution of GDP comes from asymmetry in the tail slopes.

Applying Bayes' rule as above yields the second part of the result. The probability of any a large deviation in GDP being caused by any particular combination of shocks, θ, depends on the value of $\lambda\left(\theta\right)$ relative to the average value, again scaling by κ and with respect to the measure m. When κ is relatively small, shocks are very heavy tailed, so that in some sense any combination becomes close to equally likely (that is the $\kappa \to 0$ limit). On the other hand, as $\kappa \to \infty$, only the θ that has the largest $\lambda\left(\theta\right)$ has any likelihood of causing a large deviation in GDP.

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

$$\lim_{x\to\infty} \Pr\left[gdp < -x\right] / \left(c\bar{s}^{\kappa} x^{-\kappa}\right) = N^{-1} \sum_i \left(-\gamma_i^L\right)^{\kappa} \tag{45}$$

$$\lim_{x\to\infty} \Pr\left[\theta = -e_i \mid gdp < -x\right] = \frac{\left(-\gamma_i^L\right)^{\kappa}}{N^{-1} \sum_i \left(-\gamma_i^L\right)^{\kappa}} \tag{46}$$

With independent Pareto tails, then, the left tail distribution for GDP depends on the left tail centralities, and the right tail of GDP on the right tail centralities. That immediately yields a prediction for asymmetry in GDP: any time the left tail centralities are uniformly larger than the right tail centralities – e.g. if all elasticities of substitution are less than 1 – the distribution of *gdp* will be skewed left in the sense that the $\Pr\left[gdp < -x\right] / \Pr\left[gdp < x\right] > 1$.

6.2.2 Weibull tails

Proposition 9. *Suppose the shocks have a Weibull-type tail, such that for $s > \bar{s}$:*

$$\bar{F}(s) = c \exp\left(-\eta \left(s - \bar{s}\right)^{\kappa}\right) \tag{47}$$

$$\textit{where } c = \Pr\left(s \leq \bar{s}\right) \tag{48}$$

for parameters κ and η. Then

$$\lim_{x\to\infty} \Pr\left[gdp < -x\right]^{1/(x^{\kappa})} = \left\| \exp\left(-\eta \left(\frac{-1}{\lambda(\theta)}\right)^{\kappa}\right) \right\|_{\infty;m} \tag{49}$$

where $\|\cdot\|_{\infty;m}$ denotes the essential supremum (the L_∞ norm) with respect to the measure m.[11] *Furthermore, for any set Θ^* such that $\|-\lambda(\theta)\|_{\infty;m;\Theta^*} < \|-\lambda(\theta)\|_{\infty;m;\Theta}$,*

$$\lim_{x\to\infty} \Pr\left[\theta \in \Theta^* \mid gdp < -x\right] = 0 \tag{50}$$

In the class of Weibull-type distributions, $\kappa = 1$ corresponds to an exponential tail and $\kappa = 2$ to a Gaussian-type tail. For $\kappa < 1$, the tail is heavy in the sense that the it is not exponentially bounded. The probability that GDP has an extreme event in the Weibull case is determined only by the shock that has the largest asymptotic slope. As in the Pareto case, GDP again inherits the tail shape of the shocks – now determined by κ – with

$$\Pr\left[gdp < -x\right] \sim \exp\left(-\psi x^{\kappa}\right) \tag{51}$$

[11] In particular, $\|f(\theta)\|_{\infty;m} = \inf\left\{a \in \mathbb{R} : m\left(\left\{\theta \in \Theta : f(\theta) > a\right\}\right) = 0\right\}$

for some constant ψ.

The stronger decay means that asymmetry in the economy is more pronounced. In particular,

$$\|-\lambda(\theta)\|_{\infty;m;\lambda(\theta)<0} > \|\lambda(\theta)\|_{\infty;m;\lambda(\theta)>0} \Rightarrow \lim_{x\to\infty} \frac{\Pr[gdp < -x]}{\Pr[gdp > x]} = 0 \tag{52}$$

$$\|-\lambda(\theta)\|_{\infty;m;\lambda(\theta)<0} < \|\lambda(\theta)\|_{\infty;m;\lambda(\theta)>0} \Rightarrow \lim_{x\to\infty} \frac{\Pr[gdp < -x]}{\Pr[gdp > x]} = \infty \tag{53}$$

In the power law case the probability of an extreme event depended on an integral across all values of $\lambda(\theta)$, so that any combination of shocks could produce a crash. Here, though, the probability that a crash comes from any direction except those attaining the maximum for $-\lambda(\theta)$ is zero in the limit. So except for the knife-edge case where the θ associated with the maximum is not unique (e.g. in a perfectly symmetrical economy), essentially only a *single* combination of shocks will be associated with the largest crashes.

Fragility in the economy with Weibull-distributed shocks is thus very different from in the Pareto case. The fragility here depends only on the unit or combination of units that has the largest asymptotic effect on GDP.

7 What do crashes look like?

The analysis so far has discussed only the behavior of aggregate output. This section studies unit-level output and relative prices. The results help understand the patterns produced in unit-level data that can help in testing the model. In certain cases above, the paper analyzed the effect of a large technology shock that affects just one unit. One might think that that would imply that such shocks would be easily identified in the data, since the output of the shocked unit might drastically fall and the price rise. While that will happen, shocks to systemically important units will also be associated with large declines in output and increases in prices in all downstream units – they can even look like aggregate shocks under the right conditions.

7.1 General result

Proposition 10. *Under the conditions of theorem 1, unit output and prices follow*

$$\lim_{t\to\infty} t^{-1} y_i = \phi_i(\theta) \tag{54}$$

$$\lim_{t\to\infty} t^{-1} p_i = -\phi_i(\theta) \tag{55}$$

where ϕ_i is the same function as in theorem 1 part 2, solving the equation

$$\phi_i = \theta_i + \alpha_i f_i(\phi) \text{ for } i \in \{0, 1, ..., N\} \tag{56}$$

Proposition 11. *The vector of Domar weights, D, satisfies, as a function of θt*

$$\lim_{t\to\infty} D(\theta t) = \nabla f_0(\phi(\theta)) \mathbb{J}\phi(\theta) \tag{57}$$

where ∇f_0 denotes the gradient of f_0 and $\mathbb{J}\phi$ the jacobian of ϕ.

Proposition 10 shows that unit output and prices are determined by the same recursion that determines aggregate output. That result is reminiscent of what one obtains in the usual Cobb–Douglas case (since they must match when $\sigma_i = 0 \ \forall \ i$). For large shocks, unit output and prices move inversely with each other. As usual, they depend on productivity shocks upstream. That means that a shock to productivity in a single unit will propagate downstream, affecting prices and output in all downstream units. So when there is an isolated shock to one unit, it is not in general the case that output will fall and prices rise in just that unit – they will do the same in all downstream units, though by less.

Note also that since the prices and outputs of the units move exactly inversely proportionally with each other in the tail, it is not the case that nominal output of any unit will diverge, so that Domar weights will not in general converge to zero or 1.

Proposition 11 gives the Domar weights, which are obtained simply by differentiating *gdp* with respect to productivity. There are two terms. The first is the Jacobian of ϕ with respect to θ. Shifts in productivity in a given unit will in general affect both that unit and all of the units downstream. However, what is "downstream" will depend only on the links that are operating. That is, there is a sense of a tail network. On the margin, the vector ϕ does not depend on all units that are upstream, but only those affecting f_i. If f_i is either a max or a min, it will be just that unit that achieves the extreme. In the Cobb–Douglas case, all inputs always contribute.

The second term is the gradient of *gdp* with respect to the vector $\phi(\theta)$. Again, this is evaluated at the point $\phi(\theta)$. So what determines the response of GDP to a given movement in productivity on the margin is effectively the tail network – what is the binding input for GDP. If final production is Cobb–Douglas, then all inputs matter on the margin regardless of θ, while if final production has an elasticity different from 1, then in general just a single good will matter. In that case, $\nabla f_0 (\phi (\theta))$ selects just the associated row of $J\phi^{(i)} (\theta)$.

8 Measuring tail centrality

This section measures tail centrality empirically and compares it to the Domar weights or sales shares that have been studied in past work. We use the unit detail input-output tables from 2012 reported by the BEA. The tables have 379 private units.[12] For our purposes, it is important to use a detailed version of the input-output tables because at higher levels of aggregation, the units become very strongly connected. The disaggregated table has much more sparse links.

We define an $a_{i,j}$ coefficient to be positive, so that there is an upstream link, if unit i spends at least 0.5 percent of its expenditures on materials for the output of unit j. We set $\alpha_i = 1/2 \forall i$ in calculating T_i. Finally, the β_i parameters are calculated from the fraction of nominal final expenditure going to each unit.

The top panel of figure 3 plots Domar weights (nominal output divided by nominal GDP) against 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 units, such as Petroleum Refineries, that have sales shares noticeably higher than most others. But there are numerous units with tail centralities close to 0.5. 13 units have $T_i > 0.8 \max (T_i)$, while only two have $S_i > 0.8 \max (S_i)$.

One can also see that the top units by Domar weight have very different tail centralities – Petroleum Refineries at 0.41, Oil and Gas Extraction at 0.25, and Hospitals at 0.06. Oil and Gas Extraction is lower because it is one more step up the supply chain from refineries. Hospitals are low because they produce almost entirely final output – they are not an important input for any unit.

Table 1 further examines the top units sorted by sales and tail centrality. The top panel lists the top by tail centrality. Their most common feature is that they are universal inputs. The first is electricity, which is why we have discussed it as our prime example. The second highest tail centrality is for legal services – again, simply because every unit purchases legal services. Does it make sense to claim that a large negative shock to the legal services unit

[12] We 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.

could cause a crash in GDP? To us, this result is surprising but, on reflection, reasonable. There is ample evidence that legal institutions are necessary for the growth of the economy. All aspects of business rely on property rights and contract enforcement. If, for some reason, the legal system literally shut down and legal services were actually no longer available to firms, it is entirely plausible that there would be massive 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 lawyers by firms. That actually illustrates a key advantage of T_i: measuring it does not require measuring *all* of each unit's expenditures on each input. All that we need to know is that a unit uses some input. And the input-output tables are certainly correct that all units directly use legal services.

In addition to utilities (electricity, communications) and professional services like lawyers and accountants, the last major category of units 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 unit 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 unit, financial services, and legal and accounting institutions. The analysis here shows how those shocks are linked: they all represent shocks to inputs that are used nearly universally, which is why they might have effects larger than would be implied just from looking at their average sales shares. In table 1, we show that the tail centralities are at least an order of magnitude larger than the Domar weights, if not more.

The bottom section of table 1 reports the top units sorted by sales share. As discussed above, not all have particularly high tail centralities for two reasons: some are too far upstream, like Oil and Gas Extraction, while others produce just final outputs, like Hospitals, Offices of Physicians, Pharmaceuticals, and Scientific R&D.

We have claimed so far that the units with the highest tail centralities are universal inputs. To show that formally, we now examine out-degrees. The asterisks in the bottom panel of figure 3 represent a scatter plot of T_i against weighted out-degree, d_i. There is clearly a positive relationship, and the units with the highest tail centrality have the highest out-degree. In fact, they have the same ordering for the top eight and 14 of the top 16 units. The $\bigtriangledown$ and $\bigtriangleup$ series represent the upper and lower bounds for T_i, respectively, from equation (26). Again, the lower bound is based on the fraction of GDP immediately downstream of unit i, while the upper bound holds if all of GDP is no more than two steps downstream.

For the top units, the upper bound is highly accurate. That is, for those units, it is approximately true that all of GDP is either one or two steps downstream. That approximation for T_i – which just requires knowing out-degree – is highly accurate for all the units with d_i above about 0.2, which is the top 10 percent of units.

The empirical analysis overall shows that tail centralities are very different from Domar weights. They are much larger, and among the top units closely related to out-degree. The units with the highest tail centrality are not necessarily those with the highest sales, but those that sell to the most units downstream. They represent utilities, professional and financial services, and petroleum products.

9 Conclusion

This paper studies large deviations in GDP in the context of a general nonlinear network production model. It shows that, generically, log GDP approaches a linear asymptote in terms of productivity. That result then leads to a host of results, including a novel measure of tail centrality, giving a way to measure the systemic importance of different units and elucidate the determinants of systemic riskiness. The paper's second key result is to combine the tail approximation for GDP with assumptions about the tail distribution of the productivity shocks to characterize the tail distribution of GDP.

The results show that sales shares in normal times are not a good indication of which units are systemically important in the sense of being critical to GDP when shocks become large. The most important units, regardless of their size, are those that are inputs to producers of a large fraction of GDP and that do not have substitutes.

References

Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, "Microeconomic origins of macroeconomic tail risks," *American Economic Review*, 2017, *107* (1), 54–108.

——, **Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi**, "The network origins of aggregate fluctuations," *Econometrica*, 2012, *80* (5), 1977–2016.

Alfaro, Laura, 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.

Baqaee, David Rezza and Emmanuel Farhi, "The macroeconomic impact of microeconomic shocks: beyond Hulten's Theorem," *Econometrica*, 2019, *87* (4), 1155–1203.

Barro, Robert J., "Rare Disasters and Asset Markets in the Twentieth Century," *Quarterly Journal of Economics*, 2006, *121(3)*, 823–866.

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.

Bigio, Saki and Jennifer Lao, "Distortions in production networks," *The Quarterly Journal of Economics*, 2020, *135* (4), 2187–2253.

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.

Carvalho, Vasco M, Makoto Nirei, Yukiko Saito, and Alireza Tahbaz-Salehi, "Supply chain disruptions: Evidence from the great east japan earthquake," 2020. 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.

Dew-Becker, Ian, Alireza Tahbaz-Salehi, and Andrea Vedolin, "Macro skewness and conditional second moments: evidence and theories." Working paper.

Dupraz, Stephane, Jon Steinsson, and Emi Nakamura, "A Plucking Model of Business Cycles." Working paper.

Gabaix, Xavier, "The granular origins of aggregate fluctuations," *Econometrica*, 2011, *79* (3), 733–772.

Gofman, Michael, Gill Segal, and Youchang Wu, "Production networks and stock returns: The role of vertical creative destruction," *The Review of Financial Studies*, 2020, *33* (12), 5856–5905.

Gupta, AK and D Song, "Lp-norm spherical distribution," *Journal of Statistical Planning and Inference*, 1997, *60* (2), 241–260.

Jr, John B Long and Charles I Plosser, "Real business cycles," *Journal of political Economy*, 1983, *91* (1), 39–69.

Liu, Ernest, "Industrial policies in production networks," *The Quarterly Journal of Economics*, 2019, *134* (4), 1883–1948.

Resnick, Sidney I, *Heavy-tail phenomena: probabilistic and statistical modeling*, Springer Science & Business Media, 2007.

Rietz, Thomas A., "The Equity Risk Premium: A Solution," *Journal of Monetary Economics*, 1988, *22(1)*, 117–131.

A Proofs

A.1 Theorem 1

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

$$GDP = W/P_0 \tag{58}$$

where W is the wage and P_0 is the price of the consumption bundle. Setting labor to be the numeraire, so that W is normalized to 1, the CES preferences for the consumer immediately imply

$$p_0 = -\frac{1}{1-\sigma_0}\log\left(\sum_{j=1}^{n}\beta_i \exp\left(\left(1-\sigma_0\right)p_i\right)\right) \tag{59}$$

$$gdp = -p_0 \tag{60}$$

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

$$p_i = -z_i + \frac{\alpha_i}{1-\sigma_i}\log\left(\sum_{j=1}^{n}a_{ij}\exp\left(\left(1-\sigma_i\right)p_i\right)\right) \tag{61}$$

Now define $\phi_i = -\lim_{t\to\infty}(1/t)\log p_i$ and set the vector $\phi \equiv [\phi_1, ..., \phi_N]$. If that limit exists and is finite (a claim established below), then diving by t and taking limits of both sides of equations (59) and (61) gives

$$\lim_{t\to\infty} t^{-1} gdp = f_0\left(\phi\right) \tag{62}$$

$$\phi_i = \zeta + \alpha_i f_i\left(\phi\right) \tag{63}$$

where the mapping $f_i : \mathbb{R}^n \to \mathbb{R}$ is defined in equation (11). So then as long as the above system for ϕ has a unique and finite solution, there is a unique and finite λ in theorem 1.

To show that the system has a unique solution (guaranteeing that ϕ is also finite), define a mapping $g : \mathbb{R}^N \to \mathbb{R}^N$ as

$$g_i\left(\phi\right) = \theta_i + \alpha_i f_i\left(\phi\right) \tag{64}$$

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. This completes the proof of **part 2** of the theorem.

To get the constant $\mu(\theta)$, again consider a series expansion, $p_i = \mu_i + \phi_i t + o(1)$. Inserting that into (61) taking limits, and using (63) yields

$$\mu_i = \frac{\alpha_i}{1-\sigma_i} \log \left(\sum_{j \in j^*(i)} A_{i,j} \exp\left(\left(1-\sigma_i\right) \mu_j\right) \right) \tag{65}$$

where

$$j^*(i) \equiv \begin{cases} \{j : \phi_j = \min_{k \in S_i} \phi_k\} & \text{if } \sigma_i < 1 \\ \{j : \phi_j = \max \phi_k\} & \text{if } \sigma_i > 1 \end{cases} \tag{66}$$

with 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. For GDP, the same analysis gives

$$\mu_0 = \frac{1}{1-\sigma_0} \log \left(\sum_{j \in j^*(0)} \beta_0 \exp\left(\left(1-\sigma_0\right) \mu_j\right) \right) \tag{67}$$

That proves **part 4**.

To show **parts 1 and 3**, consider the next term in a series expansion,[13] $p_i = b_i t^{-1} + \mu_i + \phi_i t + o\left(t^{-1}\right)$, and take limits as $t \to \infty$,

$$p_i t - \left(\mu_i + \phi_i t\right) t = \left(-\zeta_i t + \frac{\alpha_i}{1-\sigma_i} \log \left(\sum_{j=1}^{n} A_{i,j} \exp\left(\left(1-\sigma_i\right) p_j\right) \right) - \mu_i - \phi_i t \right) t \tag{68}$$

$$b_i = \lim_{t\to\infty} \left(\begin{array}{c} -\zeta_i t + \frac{\alpha_i}{1-\sigma_i} \log \left(\sum_{j=1}^{n} A_{i,j} \exp\left(\left(1-\sigma_i\right)\left(b_j t^{-1} + \mu_j + \phi_j t + o\left(t^{-1}\right)\right)\right) \right) \\ -\mu_i - \phi_i t \end{array} \right) t \tag{69}$$

$$b_i = \lim_{t\to\infty} \left[\left\{ \begin{array}{c} \frac{\alpha_i}{1-\sigma_i} \log \left(\sum_{j=1}^{n} A_{i,j} \exp\left(\left(1-\sigma_i\right)\left(b_j t^{-1} + \mu_j + \left(\phi_j - f_i(\phi)\right) t + o\left(t^{-1}\right)\right)\right) \right) \\ -\mu_i \end{array} \right\} t \right] \tag{70}$$

The recursion from above for μ_i immediately implies that the limit of the term in braces is

[13]Formally 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 $\tau = t^{-1}$ aroud $\tau = 0$. There is a pole at $\tau = 0$ since GDP and all prices go to $\pm\infty$. The τ^{-1} terms remove the pole, at which point we just have a standard Taylor series in τ. Part 4 of theorem 1 says that all terms in that Taylor series of order higher than 0 (i.e. eveyrthing but the constant) is equal to zero).

zero. Applying L'Hopital's rule then yields the result that the whole limit is equal to zero. The same analysis goes through for terms of any order, so that we have the statement from **part 4** of theorem 1. **Part 1** is a special case. ■

A.2 Corollary 1.1

Based on the results above, given δ there exists a t^* such that $|gdp\,(\theta t) - (\mu\,(\theta) + \lambda\,(\theta)\,t)| < \delta$ for $t > t^*$. For $t < t^*$,

$$\begin{aligned} |gdp\,(\theta t) - (\mu\,(\theta) + \lambda\,(\theta)\,t)| &\leq |gdp\,(\theta t)| + |\mu\,(\theta) + \lambda\,(\theta)\,t| & (71) \\ &\leq \max_{t \leq t^*} |gdp\,(\theta t)| + |\mu\,(\theta)| + |\lambda\,(\theta)|\,t^* & (72) \end{aligned}$$

So the fact that gdp is finite for any finite t gives the result. ■

An Auxiliary Lemma

Lemma 1. *Let $g : \mathbb{R}^n \times \Theta \to \mathbb{R}^n$ be increasing and a contraction in its first argument for all $\theta \in \Theta$. If $g(\phi;\theta) \geq g(\phi;\hat{\theta})$ for all ϕ, then $\phi^*(\theta) \geq \phi^*(\hat{\theta})$, where $\phi^*(\theta) = g(\phi^*(\theta);\theta)$ is the unique fixed point of $g(\cdot,\theta)$.*

Proof: Using an inductive argument, we first show that

$$g^{(k)}(\phi;\theta) \geq g^{(k)}(\phi;\hat{\theta}) \tag{73}$$

for all $\phi \in \mathbb{R}^n$ and all $k \in \mathbb{N}$, where $g^{(k)}(\phi;\theta) = g(g^{(k-1)}(\phi;\theta);\theta)$. The base of the induction for $k = 1$ holds by assumption. As the induction hypothesis, suppose that (73) holds for some k. Since g is increasing in its first argument, (73) implies that

$$g^{(k+1)}(\phi;\theta) \geq g(g^{(k)}(\phi;\hat{\theta});\theta) \geq g(g^{(k)}(\phi;\hat{\theta});\hat{\theta}) = g^{(k+1)}(\phi;\hat{\theta}),$$

where the second inequality follows from the assumption that $g(\phi;\theta) \geq g(\phi;\hat{\theta})$. The above inequality therefore establishes that (73) is also satisfied for $k+1$, which completes the inductive argument.

Having established (73) for all k, taking a limit from both sides of this inequality as $k \to \infty$ and using the contraction mapping theorem implies that

$$\phi^*(\theta) = \lim_{k\to\infty} g^{(k)}(\phi;\theta) \geq \lim_{k\to\infty} g^{(k)}(\phi;\hat{\theta}) = \phi^*(\hat{\theta}),$$

thus completing the proof. ■

Proposition 3

We prove this result for left tail centralities. The proof for right tail centralities is identical. We also restrict the proof to monotonicity in $(\sigma_1, \ldots, \sigma_n)$ as the proof of monotonicity in σ_0 is straightforward.

Let Σ denote the set of vectors of industry-level elasticities of substitution with a typical element $\sigma = (\sigma_1, \ldots, \sigma_n)$ and define the mapping $g : \mathbb{R}^n \times \Sigma \to \mathbb{R}^n$ as in equation (64), where the index σ signifies the fact that function f depends on the underlying elasticities. In the proof of theorem 1, we established that $g(\cdot, \sigma)$ is a contraction irrespective of the value of σ. It is also immediate that $g(\cdot, \cdot)$ is weakly increasing in both arguments. Therefore, Lemma 1 guarantees that the unique fixed point of (64), $\phi^*(\theta; \sigma)$, is weakly increasing in σ.

On the other hand, recall from Corollary 1.2 that the left tail centrality of industry i is equal to $\gamma_i^L(\theta; \sigma) = -f_0(\phi^*(\theta; \sigma))$, where $\phi^*(\theta; \sigma)$ is the unique fixed point of mapping (64). As a result, γ_i^L is weakly decreasing in the vector of elasticities of substitution σ.

Next, we show that γ_i^L is invariant to σ_j for any industry j that is not downstream to i. To see this, recall from Corollary 1.2 that the left tail centrality of industry i is given by $\gamma_i^L = f_0(\phi^*)$, where ϕ^* is the unique solution to the system of equations

$$\phi_j = -1\{i = j\} + \alpha_j f_j(\phi) \qquad \text{for all } j. \tag{74}$$

A simple inductive argument then implies that if j is not downstream to i, then $\phi_j = 0$ irrespective of the value of σ_j. This means that the unique solution to (74) and hence γ_i^L are invariant to the σ_j. ■

Proposition 2

Once again, we prove this result for left tail centralities. The proof for right tail centralities is identical.

The (weak) monotonicity of left tail centralities in elasticities of substitution (proved in Proposition 3) implies that if $\sigma_i \leq 1$ for all i, then $\gamma^L \geq D_{ss}$, using the fact that, in a Cobb-Douglas economy, left tail centralities coincide with Domar weights. Similarly, when $\sigma_i \geq 1$ for all i, the monotonicity of left tail centralities guarantees that $\gamma^L \leq D_{ss}$. ■

Proposition 4

part (a) We present the proof for the left tail centrality as the proof for the right tail centrality follows from an identical argument.

Suppose $\sigma_j < 1$ and let $\Sigma = \{S_j : S_j \subseteq N\}$ denote the set of all possible set of suppliers of industry j. Define the mapping $g : \mathbb{R}^n \times \Sigma \to \mathbb{R}^n$ as in equation (64), where indexing by elements of Σ signifies the fact that f_j depends on industry j's set of input suppliers. In the proof of Theorem 1, we established that $g(\cdot, S_j)$ is a contraction for all $S_j \in \Sigma$. It is also immediate that $g(\cdot, \cdot)$ is increasing in its first argument. Furthermore, since $\sigma_j < 1$, it is immediate that $g(\phi, S_j) \geq g(\phi, \hat{S}_j)$, whenever $S_j \subseteq \hat{S}_j$. Therefore, Lemma 1 guarantees that $\phi^*(\theta; S_j) \geq \phi^*(\theta; \hat{S}_j)$. Thus, by Corollary 1.2, $\gamma_i^L(S_j) \leq \gamma_i^L(\hat{S}_j)$ when $S_j \subseteq \hat{S}_j$, that is, the left tail centrality of i is weakly increasing in j's set of suppliers.

The proof is therefore complete once we show that γ_i^L is invariant to the set of j's suppliers if j is not downstream to i. This follows from a simple inductive argument similar to the proof of Proposition 3. In particular, it is immediate that $\phi_k^*(\theta; S_j) = 0$ irrespective of S_j for all k that are not downstream to i. Therefore, if j is not downstream to i, $\phi^*(\theta; S_j) = \phi^*(\theta; \hat{S}_j)$, which in turn implies that γ_i^L is also independent of S_j. ■

part (b) The proof of this part parallels that of part (a). Since $\sigma_j > 1$, it is immediate that $g(\phi, S_j) \leq g(\phi, \hat{S}_j)$, whenever $S_j \subseteq \hat{S}_j$. Therefore, Lemma 1 guarantees that $\phi^*(S_j) \leq \phi^*(\hat{S}_j)$. Thus, by Corollary 1.2, $\gamma_i^L(S_j) \geq \gamma_i^L(\hat{S}_j)$ when $S_j \subseteq \hat{S}_j$, that is, the left tail centrality of i is weakly decreasing in j's set of suppliers. When j is not downstream to i, the invariance of γ_i^L to the set of suppliers of j then follows from an inductive argument similar to part (a). ■

Proposition 5

Once again, we present the proof for left tail centralities, as the result for right tail centralities follows from an identical argument.

Let Σ denote the set of vectors of industry-level intermediate input shares, with a typical element $\alpha = (\alpha_1, \dots, \alpha_n)$ and define the mapping $g : \mathbb{R}^n_- \times \Sigma \to \mathbb{R}^n_-$ as in equation (64), where index α signifies the fact that function f depends on the underlying intermediate input shares. In the proof of Theorem 1, we established that $g(\cdot, \alpha)$ is a contraction irrespective of the value of α. It is also immediate that $g(\cdot, \alpha)$ is weakly increasing in its first argument and that $g(\phi, \alpha) \geq g(\phi, \alpha)$ for all ϕ whenever $\alpha \leq \hat{\alpha}$. Therefore, Lemma 1 guarantees that $\phi^*(\theta; \alpha) \geq \phi^*(\theta; \hat{\alpha})$ whenever $\alpha \leq \hat{\alpha}$. That is, $\phi^*(\theta; \alpha)$ is weakly decreasing in all industry-level intermediate input shares.

On the other hand, recall from Corollary 1.2 that the left tail centrality of industry i is equal to $\gamma_i^L(\alpha) = -f_0(\phi^*(e_i;\alpha))$, where $\phi^*(\theta;\alpha)$ is the unique fixed point of mapping (64). As a result, γ_i^L is weakly increasing in the vector of intermediate input shares $\alpha = (\alpha_1, \ldots, \alpha_n)$.

The proof is therefore complete once we show that γ_i^L is invariant to α_j if industry j is not downstream to i. But this follows from an inductive argument identical to that of Proposition 3. In particular, it is immediate that $\phi_k^*(e_i;\alpha) = 0$ irrespective of α for all k that are not downstream to i. Therefore, if j is not downstream to i, any change in α_j has no impact on ϕ^* and hence not on γ_i^L. ■

A.3 Proposition 6

This result follows from the expression for the sector output tail and the price tail. ■

A.4 Proposition 7

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

A.5 Solution for example 2

Electricity is produced according to the trivial production function $e = \exp(z_e)$, where z_e is productivity in the electricity sector, with the resource constraint

$$\sum_i e_i = \exp(z_e) \tag{75}$$

The price of electricity is then simply $p_e = -z_e$.

The price of good i is

$$p_i = -z_i + \frac{1}{1-\sigma}\log\left(1 - b + b\exp\left((1-\sigma)\,p_e\right)\right) \tag{76}$$

while the price of the final consumption good is

$$p_0 = \sum_i N^{-1} p_i \tag{77}$$

with the wage normalized to 1 so that, again, $gdp = -p_0$. Combining the prices gives

$$p_0 = -N^{-1}\sum_i z_i + \frac{1}{1-\sigma}\log\left(1-b+b\exp\left(-\left(1-\sigma\right)z_e\right)\right) \tag{78}$$

$$gdp = N^{-1}\sum_i z_i + \frac{1}{\sigma-1}\log\left(1-b+b\exp\left(\left(\sigma-1\right)z_e\right)\right) \tag{79}$$

To get the Domar weights and tail centralities, simply differentiate gdp with respect to z_e,

$$\lambda_e = \frac{d}{dz_e}gdp = \frac{b\exp\left(\left(\sigma-1\right)z_e\right)}{1-b+b\exp\left(\left(\sigma-1\right)z_e\right)} \tag{80}$$

$$\lambda_e^{ss} = b \tag{81}$$

$$\lim_{z_e\to\infty}\lambda_e = \lambda_e^R = 0 \tag{82}$$

$$\lim_{z_e\to-\infty}\lambda_e = \lambda_e^L = 1 \tag{83}$$

■

A.6 Theorem 2

$$gdp\left(z\right) = \mu\left(\theta\right) + \lambda\left(\theta\right)s + o\left(s\right) \tag{84}$$

So to have $gdp\left(z\right) < -x$, we have

$$\mu\left(\theta\right) + \lambda\left(\theta\right)s + o\left(x\right) < -x \tag{85}$$

$$s > \frac{x + o\left(x\right) + \mu\left(\theta\right)}{-\lambda\left(\theta\right)} \text{ for } \lambda\left(\theta\right) < 0 \tag{86}$$

As long as $\lambda\left(\theta\right)$ is bounded for $\theta \in \Theta$, and because the $o\left(x\right)$ term is bounded (from theorem 1), there exists an $\bar{x}$ such that $s > \bar{s}$ (where $\bar{s}$ is the cutoff for the tail distribution from the text) when $gdp < -\bar{x}$. We then have, for $x > \bar{x}$,

$$\Pr\left[gdp < -x\right] = \int_{\theta\in\Theta:\lambda(\theta)<0} \bar{F}\left(\frac{x + o\left(x\right) + \mu\left(\theta\right)}{-\lambda\left(\theta\right)}\right) dm\left(\theta\right) \tag{87}$$

The second part of the result is just an application of Bayes' rule. ■

A.7 Proposition 8

We have

$$\bar{F}(s) = c(s/\bar{s})^{-\alpha} \tag{88}$$

$$\text{where } c = \Pr(s \geq \bar{s}) \tag{89}$$

Inserting those into the formula from theorem 2,

$$\Pr[gdp < -x] = c\bar{s}^{\alpha} \int_{\theta\in\Theta:\lambda(\theta)<0} \left(\frac{x + o(x) + \mu(\theta)}{-\lambda(\theta)}\right)^{-\alpha} dm(\theta) \tag{90}$$

$$\lim_{x\to\infty} \Pr[gdp < -x] // \left(c\bar{s}^{\alpha}x^{-\alpha}\right) = \lim_{x\to\infty} \int_{\theta\in\Theta:\lambda(\theta)<0} \left(-\lambda(\theta)^{-1} + x^{-1}\frac{o(x) + \mu(\theta)}{-\lambda(\theta)}\right)^{-\alpha} dm(\theta) \tag{91}$$

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

$$\left\|\left(-\lambda(\theta)^{-1} + x^{-1}\frac{o(x) + \mu(\theta)}{-\lambda(\theta)}\right)^{-\alpha} - (-\lambda(\theta)^{\alpha})\right\|_{\infty} \leq \left\|\left(-\lambda(\theta)^{-1} + x^{-1}\frac{o(x) + \mu(\theta)}{-\lambda(\theta)}\right)^{-1}\right\|_{\infty}^{\alpha} + \|-\lambda(\theta)\|_{\infty}^{\alpha}$$

$$\leq \left\|-\lambda(\theta)\left(1 + \frac{o(x) + \mu(\theta)}{x}\right)^{-1}\right\|_{\infty}^{\alpha} + \|-\lambda(\theta)\|_{\infty}^{\alpha} \tag{92}$$

$$\leq \|-\lambda(\theta)\|_{\infty}^{\alpha} \left\|\frac{x}{x + \inf_{\theta\in\Theta}\{\mu(\theta)\}}\right\|_{\infty}^{\alpha} + \|-\lambda(\theta)\|_{\infty}^{\alpha} \tag{93}$$

Passing the limit through the integral yields the result from the text. The second claim is again an application of Bayes' rule. ■

A.8 Proposition 9

$$\bar{F}(s) = c\exp\left(-\beta(s - \bar{s})^{\alpha}\right) \tag{94}$$

$$\text{where } c = \Pr(s \leq \bar{s}) \tag{95}$$

$$\Pr\left[gdp < x\right] = \int_{\theta\in\Theta} \exp\left(-\left(\frac{x - o\left(x\right) - \mu\left(\theta\right)}{-\lambda\left(\theta\right)} - \bar{s}\right)^{\alpha}\right) dm\left(\theta\right) \tag{96}$$

$$\Pr\left[gdp < x\right]^{1/x^{\alpha}} = \left[\int_{\theta\in\Theta} \exp\left(-\left(\frac{1}{-\lambda\left(\theta\right)} - \frac{o\left(x\right) + \mu\left(\theta\right)}{x}\frac{1}{\lambda\left(\theta\right)} - \frac{\bar{s}}{x}\right)^{\alpha}\right)^{x^{\alpha}} dm\left(\theta\right)\right]^{1/x^{\alpha}} \tag{97}$$

Now consider the limit as $x \to \infty$. I claim that the limit of the right-hand side is the essential supremum of $\exp\left(-\left(\frac{1}{-\lambda(\theta)}\right)^{\alpha}\right)$ with respect to the measure $m\left(\theta\right)$ (i.e. the measure of the set of θ such that $\exp\left(-\left(\frac{1}{\lambda(\theta)}\right)^{\alpha}\right)$ is above the essential supremum is zero). Denote that by $\left\|\exp\left(-\left(\frac{1}{\lambda(\theta)}\right)^{\alpha}\right)\right\|_{\infty;m}$.

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.

From here on, the "m" is dropped from the norm notation, so that $\left\|\cdot\right\|_{\infty}$ denotes the essential supremum with respect to the measure m.

Define, for simplicity

$$f\left(\theta\right) = \exp\left(-\left(\frac{1}{\lambda\left(\theta\right)}\right)^{\alpha}\right) \tag{98}$$

$$f\left(\theta;x\right) = \exp\left(-\left(\frac{1}{\lambda\left(\theta\right)} - \frac{o\left(x\right) + \mu\left(\theta\right)}{x}\frac{1}{\lambda\left(\theta\right)} - \frac{\bar{s}}{x}\right)^{\alpha}\right) \tag{99}$$

Lemma 2. $\lim_{x\to\infty}\left\|f\left(\theta;x\right)\right\|_{\infty} = \left\|f\left(\theta\right)\right\|_{\infty}$.

Proof. $f\left(\theta;x\right) \to f\left(\theta\right)$ pointwise trivially. The difference $\left|f\left(\theta;x\right) - f\left(\theta\right)\right|$ is also bounded using the same argument as in the proof of proposition 8. $f\left(\theta;x\right)$ then converges uniformly to $f\left(\theta\right)$, from which $\left\|f\left(\theta;x\right)\right\|_{\infty} \to \left\|f\left(\theta\right)\right\|_{\infty}$ follows, since, using the reverse triangle inequality,

$$\left|\left\|f\left(\theta;x\right)\right\|_{\infty} - \left\|f\left(\theta\right)\right\|_{\infty}\right| \leq \left\|f\left(\theta\right) - f\left(\theta;x\right)\right\|_{\infty} \tag{100}$$

■

Lemma 3. $\limsup_{x\to\infty}\left[\int_{\theta\in\Theta} f\left(\theta;x\right)^{x^{\alpha}}\right]^{1/x^{\alpha}} dm\left(\theta\right) \leq \left\|f\left(\theta\right)\right\|_{\infty}$

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

$$\left\|f\left(\theta;x\right)\right\|_{x^{\alpha}} \leq \left\|\left\|f\left(\theta;x\right)\right\|_{\infty}\right\|_{x^{\alpha}}$$

Taking limits of both sides

$$\begin{aligned}
\lim_{x\to\infty} \|f(\theta;x)\|_{x^\alpha} &\leq \lim_{x\to\infty} \|\|f(\theta;x)\|_\infty\|_{x^\alpha} &(101)\\
&= \lim_{x\to\infty} \|f(\theta;x)\|_\infty &(102)\\
&= \|f(\theta)\|_\infty &(103)
\end{aligned}$$

where the second line follows from the fact that $\|f(\theta;x)\|_\infty$ is constant and the third line uses lemma 2. ■

Lemma 4. $\liminf_{x\to\infty} \left[\int_{\theta\in\Theta} f(\theta;x)^{x^\alpha}\right]^{1/x^\alpha} dm(\theta) \geq \|f(\theta)\|_\infty$

Proof. Consider some $\varepsilon > 0$, and set $A = \left\{\theta \in \Theta : \exp\left(-\left(\frac{1}{-\lambda(\theta)}\right)^\alpha\right) \geq \left\|\exp\left(-\left(\frac{1}{-\lambda(\theta)}\right)^\alpha\right)\right\|_\infty - \varepsilon\right\}$. Consider also the set $A' = \left\{\theta \in \Theta : \exp\left(-\left(\frac{1}{\lambda(\theta)} - \frac{o(x)+\mu(\theta)}{x}\frac{1}{\lambda(\theta)} - \frac{\bar{s}}{x}\right)^\alpha\right) \geq \left\|\exp\left(-\left(\frac{1}{\lambda(\theta)}\right)^\alpha\right)\right\|_{\infty;m} - \varepsilon\right\}$. For any ε sufficiently small that A has positive measure, there exists an $\bar{x}(\varepsilon)$ sufficiently large that A' has positive measure for all $x > \bar{x}(\varepsilon)$ due to the continuity of $\exp\left(-\left(\frac{1}{\lambda(\theta)} - \frac{o(x)+\mu(\theta)}{x}\frac{1}{\lambda(\theta)} - \frac{\bar{s}}{x}\right)^\alpha\right)$ and the fact that $\exp\left(-\left(\frac{1}{\lambda(\theta)} - \frac{o(x)+\mu(\theta)}{x}\frac{1}{\lambda(\theta)}\right)^\alpha\right) \to \exp\left(-\left(\frac{1}{\lambda(\theta)}\right)^\alpha\right)$ as $x \to \infty$.

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

$$\begin{aligned}
&\int_{\theta\in\Theta} \exp\left(-\left(\frac{1}{\lambda(\theta)} - \frac{o(x)+\mu(\theta)}{x}\frac{1}{\lambda(\theta)} - \frac{\bar{s}}{x}\right)^\alpha\right)^{x^\alpha} dm(\theta) &(104)\\
\geq\ &\int_{A'} \exp\left(-\left(\frac{1}{\lambda(\theta)} - \frac{o(x)+\mu(\theta)}{x}\frac{1}{\lambda(\theta)} - \frac{\bar{s}}{x}\right)^\alpha\right)^{x^\alpha} dm(\theta) &(105)\\
\geq\ &\left(\left\|\exp\left(-\left(\frac{1}{\lambda(\theta)}\right)^\alpha\right)\right\|_{\infty;S(\theta)} - \varepsilon\right)^{x^\alpha} \mu(A') &(106)
\end{aligned}$$

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

$$\liminf_{x\to\infty} \left[\int_{\theta\in\Theta} \exp\left(-\left(\frac{1}{\lambda(\theta)} - \frac{o(x)+\mu(\theta)}{x}\frac{1}{\lambda(\theta)} - \frac{\bar{s}}{x}\right)^\alpha\right)^{x^\alpha}\right]^{1/x^\alpha} dm(\theta) \geq \left\|\exp\left(-\left(\frac{1}{\lambda(\theta)}\right)^\alpha\right)\right\|_{\infty;m} \tag{107}$$

■

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

For the second part, in the set Θ^*, there exists an ε such that $|\lambda(\theta)| < \|\lambda(\theta)\|_{\infty;\Theta} - \varepsilon$. Therefore

$$\Pr\left[\theta \in \Theta^* \mid gdp < -x\right] = \frac{\int_{\theta\in\Theta^*} \exp\left(-\left(\frac{x-o(x)-\mu(\theta)}{-\lambda(\theta)} - \bar{s}\right)^\alpha\right) dm(\theta)}{\int_{\theta\in\Theta} \exp\left(-\left(\frac{x-o(x)-\mu(\theta)}{-\lambda(\theta)} - \bar{s}\right)^\alpha\right) dm(\theta)}$$

$$\leq \frac{\int_{\theta\in\Theta^*} \exp\left(-\left(\frac{x-o(x)-\mu(\theta)}{-\left(\|\lambda(\theta)\|_{\infty;\Theta}-\varepsilon\right)} - \bar{s}\right)^\alpha\right) dm(\theta)}{\int_{\theta\in\Theta:|\lambda(\theta)|>|\lambda(\theta)|-\varepsilon/2} \exp\left(-\left(\frac{x-o(x)-\mu(\theta)}{-\lambda(\theta)} - \bar{s}\right)^\alpha\right) dm(\theta)}$$

$$\leq \frac{\exp\left(-\left(\frac{x-o(x)-\mu(\theta)}{-\left(\|\lambda(\theta)\|_{\infty;\Theta}-\varepsilon\right)} - \bar{s}\right)^\alpha\right)}{\exp\left(-\left(\frac{x-o(x)-\mu(\theta)}{-\left(\|\lambda(\theta)\|_{\infty;\Theta}-\varepsilon/2\right)} - \bar{s}\right)^\alpha\right)} \frac{1}{m\left(\left\{\theta \in \Theta : |\lambda(\theta)| > \|\lambda(\theta)\|_{\infty;\Theta} - \varepsilon/2\right\}\right)} \quad (108)$$

$$\to 0 \quad (109)$$

■

A.9 Proposition 10

To prove this, we will show that the above is consistent with the model's equilibrium conditions,

$$Y_i = \exp(z_i) L_i^{(1-\alpha_i)} \left(\sum_j a_{i,j} x_{i,j}^{\gamma}\right)^{\alpha_i/\gamma_i} \quad (110)$$

$$Y_j = C_j + \sum_i X_{i,j} \quad (111)$$

$$P_j = \beta_j C_j^{-1} \quad (112)$$

$$P_j = \alpha_i P_i \exp(z_i) (Y_i / \exp(z_i))^{(\alpha_i-\gamma_i)/\alpha_i} A_{i,j} X_{i,j}^{\gamma_i-1} \quad (113)$$

$$1 = (1-\alpha_i) P_i Y_i / L_i \quad (114)$$

We first prove some small lemmas.

$$\gamma - 1 = \frac{-1}{\sigma} \tag{115}$$

$$\frac{1}{1-\gamma} = \sigma \tag{116}$$

$$\frac{\gamma}{1-\gamma} = \sigma - 1 \tag{117}$$

Lemma 5. *$\phi_{j^*(i)} + \sigma_i \left(\phi_j - \phi_{j^*(i)}\right) \leq \phi_j$ for all $j \in S(i)$*

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

$$\phi_{j^*(i)} + \sigma_i \left(\phi_j - \phi_{j^*(i)}\right) = \phi_j + (1 - \sigma_i) \left(\phi_j - \phi_{j^*(i)}\right) \tag{118}$$

Since $1 - \sigma_i > 0$ and $\phi_j - \phi_{j^*(i)} \leq 0$ in thise case, the result again follows. It holds trivially for $\sigma_i = 1$. ■

Lemma 6. *$f_i \left(\left[\phi_{j^*(i)} + \sigma_i \left(\phi_j - \phi_{j^*(i)}\right)\right]\right) = \phi_{j^*(i)}$*

Proof. If $\sigma_i > 1$, then $f_i = \max_{j \in S(i)}$. Note that $\phi_{j^*(i)} + \sigma_i \left(\phi_j - \phi_{j^*(i)}\right) \leq \phi_j$, and $\phi_{j^*(i)} = \max_{j \in S(i)} \phi_j$. Then the result immediately follows. Suppose $\sigma_i < 1$. Then $f_i = \min_{j \in S(i)}$ and $\phi_{j^*(i)} = \min_{j \in S(i)} \phi_j$. In this case, $\phi_{j^*(i)} + \sigma_i \left(\phi_j - \phi_{j^*(i)}\right) \leq \phi_{j^*(i)}$, with equality if $j = j^*(i)$, again giving the result. ■

As in the main text, define

$$j^*(i) = \begin{cases} \arg\min_{j \in S(i)} \phi_j \text{ if } \sigma_i < 1 \\ \arg\max_{j \in S(i)} \phi_j \text{ if } \sigma_i > 1 \end{cases} \tag{119}$$

Proposition 10. *We have the following limits,*

$$\lim_{t \to \infty} \frac{y_j}{t} = \lim_{t \to \infty} \frac{c_j}{t} = \lim_{t \to \infty} \frac{-p_j}{t} = \phi_j \tag{120}$$

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

$$\lim_{t \to \infty} \frac{x_{i,j}}{t} = \phi_{j^*(i)} + \sigma_i \left[\phi_j - \phi_{j^*(i)}\right] \tag{121}$$

We need to verify that the above, along with the solution in the proposition, satisfies, in the limit, the equilibrium conditions (110)-(114).

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

$$\phi_{g,j} \equiv \lim_{t\to\infty} \frac{g_j}{t} \tag{122}$$

Inpsection of equation (114) 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 (110)-(114), respectively) by t and taking limits as $t \to \infty$ yields

$$\phi_{y,i} = \theta_i + (1-\alpha_i)\,\phi_{l,i} + \alpha_i f_i\,(\phi_{x,i,j}) \tag{123}$$

$$\phi_{y,j} = \max\left\{\phi_{c,j}, \max_i \phi_{x,i,j}\right\} \tag{124}$$

$$\phi_{p,j} = -\phi_{c,j} \tag{125}$$

$$\phi_{p,j} = \phi_{p,i} + \theta_i + \frac{\alpha_i - \gamma_i}{\alpha_i}\,(\phi_{y,i} - \theta_i) - \sigma_i^{-1}\phi_{x,i,j} \tag{126}$$

$$0 = \phi_{p,i} + \phi_{y,i} - \phi_{l,i} \tag{127}$$

Equation (123) holds using lemma 6 and the recursion for ϕ_i. Equation (124) holds trivially using the guesses and lemma 5. Equation2 (125)-(127) hold trivially after inserting the various guesses. ■

Proposition 11. *The vector of Domar weights, D, satisfies, as a function of θt*

$$\lim_{t\to\infty} D\,(\theta t) = \nabla f_0\,(\phi\,(\theta))\,\mathbb{J}\phi\,(\theta) \tag{128}$$

Proof. We get this result simply by differentiating the limiting expression for gdp with respect to z. Specifically, consider the derivative in some direction e_i, where e_i is a unit vector equal to 1 in element i and zero elsewhere.

$$\frac{dgdp}{dz_i} = \lim_{k\to 0} \frac{f_0\,(\phi\,(z+ke_i)) - f_0\,(\phi\,(z)) + \mu_0\,(z+ke_i) - \mu_0\,(z)}{k} \tag{129}$$

$$= \lim_{k\to 0} \frac{f_0\,(\phi\,(z+ke_i)) - f_0\,(\phi\,(z))}{k} + \frac{\mu_0\,(\theta + ke_i/t) - \mu_0\,(\theta)}{k} \tag{130}$$

$$= \nabla f_0\,(\phi\,(\theta t))\,\frac{\partial}{\partial z_i}\phi\,(\theta t) + \frac{\partial}{\partial z_i}\mu_0\,(z)\,\frac{1}{t} \tag{131}$$

where the second line uses the fact that μ_0 is homogeneous of degree 0. We then have

$$\lim_{t\to\infty} \frac{dgdp}{dz_i} = \nabla f_0\,(\phi\,(\theta))\,\frac{\partial}{\partial z_i}\phi\,(\theta) \tag{132}$$

using the facts that ϕ and f_0 are homogeneous of degree 1 and that their derivatives are (necessarily, then) homogeneous of degree 0. ■

B Relaxing the CES assumption

This section extends the baseline result to a broader class of production functions and shows that the main result 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 it each sector has constant returns to scale. Again, without loss of generality, assume that labor and materials are combined with a unit elasticity of subsitution. Those assumptions assumptions imply that, in competitive equilibrium, the price of good i is given by

$$P_i = \frac{1}{Z_i} W^{1-\alpha_i} (C_i(P_1, \ldots, P_n))^{\alpha_i} \tag{133}$$

where Z_i is the productivity shock to industry i, C_i is a homogenous function of degree one, and $\alpha_i < 1$. 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 (133), with the convention that $\alpha_0 = 1$ and $Z_0 = 1$.

The only additional assumption imposed on C_i is that

$$\lim_{t\to\infty} \frac{1}{t} \log C_i \left(\exp\left(\phi_l t\right), \exp\left(\phi_1 t\right), ..., \exp\left(\phi_n t\right)\right) = \tilde{f}_i \left(\phi_l, \phi_1, ..., \phi_n\right) \tag{134}$$

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

$$\lim_{t\to\infty} \frac{d}{dt} C_i \left(\exp\left(\phi_l t\right), \exp\left(\phi_1 t\right), ..., \exp\left(\phi_n t\right)\right) \tag{135}$$

exists (that is, 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.

For concision, this section just presents the generalization of the main result from theorem 1.

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

$$\lim_{t\to\infty} \left(gdp\left(z\right) - \lambda\left(\theta\right) t\right) t^{-1} = 0 \tag{136}$$

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*

$$\phi_i = \theta_i + \alpha_i \tilde{f}_i(\phi) \;\; \text{for } i \in \{0, 1, ..., N\} \tag{137}$$

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

$$p_i = -\log z_i + \alpha_i \log C_i(\exp(p)) \tag{138}$$

Let

$$\phi_i = -\lim_{t\to\infty} t^{-1} p_i \tag{139}$$

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

$$t^{-1} p_i = -\theta + \alpha_i t^{-1} \log C_i(\exp(p)) \tag{140}$$

$$\phi_i = -\theta + \alpha_i \lim_{t\to\infty} t^{-1} \log C_i(\exp(p)) \tag{141}$$

$$= -\zeta + \alpha_i f_i(\phi) \tag{142}$$

where the second line takes the limit as $t \to \infty$ 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

$$\log GDP = -\log P = f_0(\phi)\, t + o(t) \tag{143}$$

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

$$\hat{g}_i(\phi) = \zeta_i + \alpha_i f_i(\phi) \tag{144}$$

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. ■

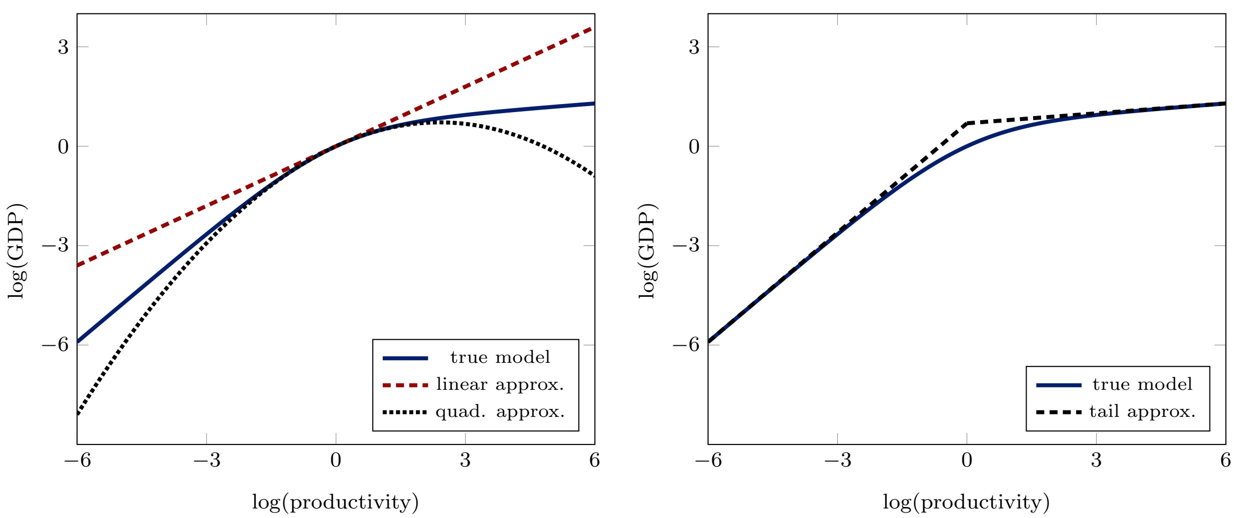

(a) Small-shock approximations

(b) Large-shock approximation

Figure 1: Linear, quadratic, and tail approximations

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

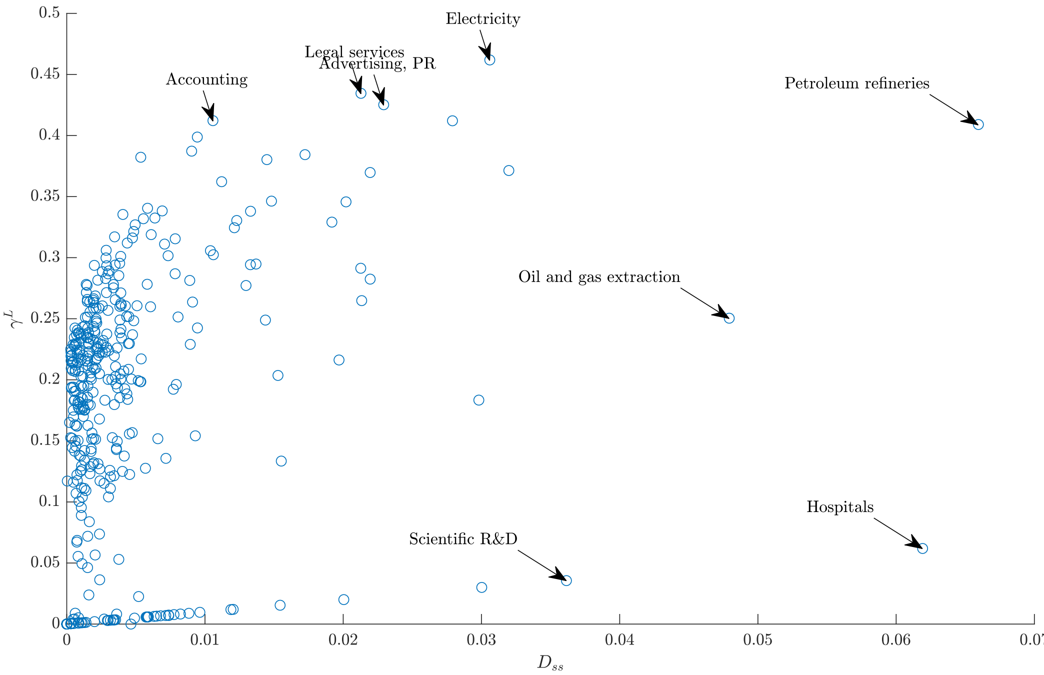

Figure 2: Domar weights and tail centralities

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.