

NBER WORKING PAPER SERIES

THE MACROECONOMIC IMPACT OF MICROECONOMIC SHOCKS:
BEYOND HULTEN'S THEOREM

David Rezza Baqaee
Emmanuel Farhi

Working Paper 23145
<http://www.nber.org/papers/w23145>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2017

We thank Daron Acemoglu, Pol Antras, Enghin Atalay, Costas Arkolakis, Natalie Bau, Francesco Caselli, Xavier Gabaix, Ben Golub, Dale Jorgenson, John Moore, Ricardo Reis, and Jean Tirole for useful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by David Rezza Baqaee and Emmanuel Farhi. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem
David Rezza Baqaee and Emmanuel Farhi
NBER Working Paper No. 23145
February 2017, Revised February 2017
JEL No. E01,E1,E23,E32,L16

ABSTRACT

We provide a nonlinear characterization of the macroeconomic impact of microeconomic TFP shocks in terms of reduced-form non-parametric elasticities for efficient economies. We also provide the mapping from structural parameters to these reduced-form elasticities, under general equilibrium. In this sense, the paper extends the foundational theorem of Hulten (1978) beyond first-order terms to capture nonlinearities. Key features ignored by first-order approximations that play a crucial role are: structural elasticities of substitution, network linkages, structural returns to scale, and the degree to which factors can be reallocated. Higher-order terms are large and economically interesting: they magnify negative shocks and attenuate positive shocks, resulting in an output distribution that is asymmetric (negative skewness), fat-tailed (excess kurtosis), and has a lower mean. They explain how small microeconomic shocks to critical sectors can have a large macroeconomic impact. To give a sense of magnitudes: in our benchmark calibration, output losses due to business cycle fluctuations are 0.6% of GDP, an order of magnitude larger than the cost of business cycles calculated by Lucas (1987), and are entirely due to a reduction in the mean of GDP because of nonlinearities in production; and accounting for second-order terms increases the estimated impact of the price shock to the critical sector of oil in the 1970s from 0.7% to 2.4% of world GDP.

David Rezza Baqaee
London School of Economics
UK
D.R.Baqaee@lse.ac.uk

Emmanuel Farhi
Harvard University
Department of Economics
Littauer Center
Cambridge, MA 02138
and NBER
emmanuel.farhi@gmail.com

1 Introduction

The foundational theorem of Hulten (1978) states that for efficient economies and under minimal assumptions, the first-order impact on output of a TFP shock to a firm or an industry is equal to that industry or firm's sales as a share of output. This surprising result, and related results in other fields like trade¹, have led macroeconomists to de-emphasize the role of microeconomic production structures in macroeconomics. After all, if the sales of a firm tell us the macroeconomic impact of a shock, and we directly observe these sales, then we need not concern ourselves with the details of the underlying microeconomic structure that gave rise to these sales.

A stronger form of this irrelevance argument suggests that we can also rule out microeconomic sources for aggregate fluctuations. If we can write aggregate quantities as a weighted average of individual-level data, Lucas (1977) and others, argue that idiosyncratic changes cannot explain changes in the aggregates. Since the economy consists of millions of workers and firms, the law of large numbers implies that idiosyncratic shocks to individual units should average out to zero with near-certainty, as long as the weights we use are sufficiently close to zero. However, in order to make this argument, we need to be able to write aggregates as a weighted average of individual quantities. Hulten's theorem gives a formal justification for this average as a first-order approximation and shows that the appropriate weights are observed expenditure shares. This is called Domar (1961) aggregation, and not only is it of theoretical interest, but it also underlies much of national accounting.

Recently, an active theoretical and empirical literature has brought this long-held conviction under renewed scrutiny. Broadly speaking, there have been three different ways in which the diversification argument has been challenged. The first branch questions the idea that the Domar weights are small in practice, the second points out that variance may not be the most interesting moment of GDP to focus on, and the last branch shows that in the presence of frictions, Hulten's theorem need not hold.

In this paper, we use a new line of attack on the diversification argument: we challenge the first-order approximation itself. We show that Hulten's theorem, powerful as it is, can in practice be very fragile due to significant nonlinearities in how shocks are mapped to output. We provide a reduced-form characterization of the second-order terms, and link these to deep parameters using a relatively general structural model. These second-order

¹See e.g. Burstein and Cravino (2015).

terms are shaped by structural elasticities of substitution, network linkages, structural returns to scale, and the degree to which factors can be reallocated, in a way that we precisely characterize. Although we maintain a focus on the impact of idiosyncratic shocks, we demonstrate that our results have important implications for the impact of correlated shocks, the average performance of the economy, and the shape of the distribution of output. These nonlinearities in production generate losses from business cycle fluctuations, and these losses are an order of magnitude larger than the ones owing to risk aversion identified by Lucas (1987).

Before describing our contribution in more details, we situate our work in the literature by briefly summarizing the other three branches. Seminal papers in the first branch by Gabaix (2011) and Acemoglu et al. (2012) challenge the idea that the expenditure shares are, in practice, close to zero. Gabaix (2011) points to the existence of very large, or in his language *granular* firms, as a possible source of aggregate volatility. If there exist very large firms, then shocks to those firms will not cancel out with shocks to much smaller firms, resulting in aggregate fluctuations. Acemoglu et al. (2012), working with a Cobb-Douglas model in the spirit of Long and Plosser (1983), observed that in an economy with input-output linkages, the equilibrium size of firms will depend on the shape of the input-output matrix. Central suppliers will be weighted more highly than peripheral firms, and therefore, shocks to those central players will not cancel out with shocks to small firms.² Carvalho and Gabaix (2013) show how Hulten’s theorem can be operationalized to decompose the sectoral sources of aggregate volatility.

The second type of objection to the diversification argument is due to Acemoglu et al. (2017) who argue that if the Domar weights are fat-tailed and if the underlying idiosyncratic shocks are fat-tailed, then GDP can exhibit non-normal behavior. Under these conditions, they argue that the variance of GDP is the wrong moment to focus on. Stated differently, GDP can inherit tail risk from idiosyncratic tail risk if the distribution of the Domar weights is fat-tailed. Our paper strengthens, but is distinct from, this point. We find that, for the empirically relevant range of parameters, the response of output to shocks is significantly asymmetric. Therefore, the nonlinearity inherent in the production structure can turn even symmetric thin-tailed sectoral shocks into rare disasters endogenously. This means that the economy could plausibly experience aggregate tail risk without either fat-tailed shocks or fat-tailed Domar weights.

²A related version of this argument was also advanced by Horvath (1998), who explored this issue quantitatively with a more general model in Horvath (2000). Separately, Carvalho (2010) also explores how the law of large numbers may fail under certain conditions on the input-output matrix.

The last line of objection to the diversification argument is typified by Baqaee (2016), Grassi (2017), and Bigio and La'O (2016) who show that the presence of frictions can cause Hulten's theorem to fail and that this failure may be extreme. Bigio and La'O (2016) work with a Cobb-Douglas model where financing constraints distort the equilibrium, and this distortion means that the Domar weights are no longer the correct weights.³ Baqaee (2016) works with a model with scale economies and imperfect competition. In his environment, Hulten's theorem fails, and the model's propagation and diffusion properties change. Grassi (2017) shows that the interaction of TFP shocks with the pricing power of firms can affect the volatility of GDP.

Stepping aside from the diversification argument, Hulten's theorem has, more generally, been something of a bugbear for the burgeoning literature on production networks, since it implies that, as long as we can observe the distribution of sales in the economy, to a first-order approximation, we do not need to concern ourselves with the underlying microeconomic details. In other words, from a macroeconomic perspective, it does not matter whether a firm is large because it produces a crucial intermediate input for other firms or because it sells a lot directly to the household. So, for example, up to a first order through the lens of Hulten's theorem, shocks to Walmart and shocks to the electricity production industry, the sales of which both currently stand at about 4% of U.S. GDP, have an equal impact on GDP.⁴ Furthermore, under Hulten's theorem, it does not matter whether a shock hits in a fragile, complicated ecosystem with high degrees of complementarity or in a robust, simple, and highly substitutable economy.

Because of this, in a recent survey article Gabaix (2016) writes "networks are a particular case of granularity rather than an alternative to it." This has meant that researchers studying the role of networks have either moved away from efficient models, or that they have retreated from aggregate volatility and turned their attention to the microeconomic implications of networks, namely the covariance of fluctuations between different industries and firms.⁵ However, models with the same sales distributions are only equivalent

³Altinoglu (2016) and Liu (2017) also investigate the impact of credit constraints in production economies with network structures.

⁴In Appendix D we show that this first-order approximation can be arbitrarily bad using a stylized model with an energy input.

⁵For instance, Foerster et al. (2011), Atalay (2016), Di Giovanni et al. (2014), and Stella (2015) investigate the importance of idiosyncratic shocks propagating through networks to generate cross-sectional covariances, but they refrain from analyzing the impact of these shocks on output. Atalay (2016) is particularly relevant in this context, since he finds that structural elasticities of substitution in production play a powerful role in generating covariance in sectoral output. Our paper complements this analysis by focusing instead on the way complementarities affect GDP.

up to the first order, and in this paper, we highlight the fragility of this first-order approximation. In particular, we argue that the agenda to trace the *network* origins of aggregate fluctuations (see Acemoglu et al., 2012) extends beyond the way networks affect steady-state Domar weights.

We proceed as follows. In Section 2, we derive a general formula describing the second-order impact on output of idiosyncratic shocks in terms of non-parametric sufficient statistics: reduced-form *macro* elasticities of substitution and *input-output* multipliers.⁶ We also explain the implications of this formula for the impact of correlated shocks and for the average performance of the economy. We then show how these sufficient statistics depend on deep structural parameters, taking into account general equilibrium forces.^{7,8} To do so, in Section 3, we set up a relatively general structural production model, which allows for any arbitrary network of nested CES production functions, heterogenous returns to scale in factors, and labor reallocation. In Sections 4, 5, and 6, using some stylized specifications of the model, we dig into the the roles of structural elasticities of substitution, returns to scale, factor reallocation, and intermediate inputs.⁹ In Section 7, we derive an industry-level network-centrality measure for the special case where every industry has constant returns to scale. In Section 8, we perform some illustrative exercises to investigate the quantitative implications of our results. First, using a calibrated structural multi-sector model, we find that the higher order terms can significantly degrade the average performance of output, reducing it by 0.6% in our benchmark calibration. Furthermore, output in this model is also negatively skewed, and has excess kurtosis even though our underlying technology shocks are lognormal. In addition, we find that negative shocks to crucial

⁶Studying the second-order terms is the first step in grappling with the nonlinearities inherent in multi-sector models with production networks. In this sense, our work illustrates the macroeconomic importance of *local* and strongly *nonlinear* interactions emphasized by Scheinkman and Woodford (1994). Other related work on nonlinear propagation of shocks in economic networks includes Durlauf (1993), Jovanovic (1987), Ballester et al. (2006) Acemoglu et al. (2015), Elliott et al. (2014), and especially Acemoglu et al. (2016).

⁷Our work is connected to the literature showing that macro and micro elasticities can, in principle, be very different. Houthakker (1955) is the archetypal example, though more recent work by Oberfield and Raval (2014) and Beraja et al. (2016) also fits into this category, albeit their focus is very different from ours. Indeed, the reduced-form sufficient statistics we define, the macro elasticity of substitution and the input-output multiplier, are general equilibrium objects and cannot be directly elicited using simple exogenous microeconomic variation, but can be estimated by combining such variation with a structural model.

⁸Some of our results are also related to the literature on the Le Chatelier principle in economics, like Samuelson (1960) and Milgrom and Roberts (1996), since we show that general equilibrium forces can increase effective elasticities of substitution by reallocating inputs and factors in response to shocks.

⁹Some of these results are closely related to the literature on growth and misallocation, especially Jones (2011), Jones (2013), and Kremer (1993), who have emphasized the importance of reallocation and complementarities in production for explaining the cross-sectional variation in aggregate GDP and TFP.

industries, like “oil and gas”, can have a significantly larger negative effect on output than negative shocks to larger but less crucial industries. Interestingly, the relative ranking of which industries are more important depends on both the sign and the size of the shock. Second, we derive and use a simple nonparametric formula, taking into account the observed change in the Domar weight for crude oil, to analyze the impact of the energy crisis of the 1970s up to the second order. We find that second-order terms amplified the impact of the oil price shocks from 0.7% to 2.4% of output.

Our results suggest that the Cobb-Douglas functional form, commonly used in the production network, growth, and multisector macro literatures, is a very special knife-edge case.¹⁰ For this special case, the second-order terms are identically equal to zero, and therefore, Hulten’s theorem is globally accurate. This means that the issues of complementarity, substitutability, returns to scale, factor reallocation, and network structure, which play a key role in our results, all disappear when one works with a Cobb-Douglas model. The empirical literature on production networks, like Atalay (2016), Boehm et al. (2015), and Barrot and Sauvagnat (2016) all find that structural elasticities of substitution in production are significantly below one, and sometimes very close to zero, across intermediate inputs, and between intermediate inputs and labor (at business cycle frequencies). Our results suggest we should be wary of Cobb-Douglas functional forms, or first-order approximations, under these scenarios.

This paper is focused on the implications of nonlinear production for business cycles. Hence, our quantitative exercises deal with within-country cyclical variations. However, our theoretical results can be applied just as easily to cross-country differences in output and TFP. The fact that we find lower TFP in crucial industries, like energy production, can have large effects on output may also help uncover the microeconomic origins of the large observed differences in cross-country output and aggregate TFP.

2 General Framework

First, we set up a nonparametric framework that demonstrates both Hulten’s theorem as well as our second-order approximation. Consider a perfectly competitive economy with

¹⁰A mixture of analytical tractability, as well as balanced-growth considerations, have made Cobb-Douglas *the* canonical production function for networks (Long and Plosser, 1983), multisector RBC models (Gomme and Rupert, 2007), and growth theory (Aghion and Howitt, 2008). Recent work by Grossman et al. (2016) shows how balanced growth can occur without Cobb-Douglas.

a representative consumer whose consumption-bundle metric is

$$C = \mathcal{C}(c_1, \dots, c_N),$$

where c_i is the household's consumption of good i . Aggregate consumption C is homogeneous of degree one and the household consumes a nonzero amount of every good. The budget constraint is

$$\sum_{i=1}^N p_i c_i = \sum_{i=1}^M w_i l_i + \sum_{i=1}^N \pi_i,$$

where p_i is the price and π_i is the profit of production unit i . For each labor type i , there is an endowment of labor l_i which is supplied inelastically and competitively on a spot market with a wage w_i . These labor markets may be common across producers or good-specific. Note that, in principle, c_i could represent consumption of different varieties of goods from the same industry, goods from different industries, or even goods in different periods of time, regions, or states of nature. Similarly, c_i could stand in for different types of leisure, thereby allowing for endogenous labor supply with a disutility of labor.

We interpret C as a cardinal measure of output and note that it is the correct measure of the household's "standard of living" in this model. We implicitly rely on the existence of complete financial markets, and ex ante symmetry of endowments, to ensure the existence of a representative consumer. Although the assumption of a representative consumer is not strictly necessary for the results in this section, it is a standard assumption in this literature since it allows us to unambiguously define and measure changes in real GDP without contending with the issue of the appropriate price index.¹¹

Each good i is produced by competitive firms using production function

$$y_i = A_i F_i(l_{i1}, \dots, l_{iM}, x_{i1}, \dots, x_{iN}),$$

where A_i is Hicks-neutral technology, x_{ij} are intermediate inputs of good j used in the production of good i , and l_{ij} is labor type j used by i . Once again, in principle, these production functions may be intertemporal or regional, however, for most of the analysis we interpret them as industries. Our assumption that labor markets may be segmented

¹¹In general, with heterogeneous households, without further assumptions on preferences, there is no uncontroversial way to boil down welfare into a single number. Although recent papers have relied on the existence of a representative consumer, Hulten (1978) instead derives his result by defining changes in real GDP as changes in final goods consumption holding prices fixed, which corresponds to the Laspeyres quantity index. Our results can be easily extended to cover this alternative set up.

is motivated by increasing evidence, like Acemoglu et al. (2016), Autor et al. (2016), and Notowidigdo (2011), that labor is not easily reallocated across industries or regions after shocks. Note that factor-augmenting technology shocks are a special case of this set-up since such shocks can always be rewritten as Hicks-neutral shocks simply by relabeling the industry's factor input as a separate industry. The profits earned by the producer of good i are

$$\pi_i = p_i y_i - \sum_{k=1}^M w_k l_{ik} - \sum_{j=1}^N p_j x_{ij}.$$

Competitive equilibrium is defined in the usual way, where all agents take prices as given, and markets for every good and every type of labor clears. Define $C(A_1, \dots, A_N)$ to be the equilibrium aggregate consumption as a function of the underlying technology shocks. Throughout the paper, unless otherwise specified, we refer to $C(1, \dots, 1)$ as the steady-state output of the model, and we derive results regarding the effects of shocks in the vicinity of this steady state. The relevant derivatives are all applied at $(A_1, \dots, A_N) = (1, \dots, 1)$.

Since this economy is efficient, an application of the envelope theorem has the following surprising but powerful implication.

Theorem 2.1 (Hulten). *Let λ_i denote industry i 's sales as a share of output. Then*

$$\frac{\partial \log C}{\partial \log A_i} = C_i \frac{A_i}{C} = \lambda_i,$$

where $C_i = \partial C / \partial A_i$.

In other words, to a first order, the underlying microeconomic details of the structural model are completely irrelevant as long as we observe the equilibrium sales distribution.¹²

Crucially, since Hulten's theorem is a consequence of the envelope theorem, as long as the steady state is efficient, it does not matter whether or not factors or inputs are reallocated in response to a shock. Since this is a first-order approximation, it captures all the relevant information for both idiosyncratic and correlated shocks — linearity implies that the impact of a common shock is simply the summation of the impact of idiosyncratic shocks.

Finally, in the special case where A_i is a labor-augmenting shock, the relevant λ_i corresponds to an industry's wage bill as a share of GDP. This is because if we relabel the

¹²This irrelevance result only holds for technology shocks. Even when Hulten's theorem is globally true, demand shocks will have different effects, as shown by, for example Baqaee (2015).

labor input of industry i as a new industry, we can represent a labor-augmenting shock to i 's labor as a Hicks-neutral shock to this new industry. Applying Hulten's theorem would then imply that the output elasticity of a shock to i 's labor are the sales of i 's labor as a share of GDP — in other words, the wage bill of i as a share of GDP.

Theorem 2.1 has long been a cornerstone of national accounting, starting with Domar (1961), since it justifies the creation of aggregate measures of inputs like capital and labor, as well as the construction of aggregate TFP from disaggregated data (see Hulten, 2001). More recently, it has also become prominent in theoretical work: it underlies the aggregation results used in the literature on the microeconomic origins of aggregate fluctuations (see Gabaix, 2011; Di Giovanni et al., 2014; Acemoglu et al., 2017; Carvalho and Gabaix, 2013) and the microeconomic origins of cross-country TFP differences (see Jones, 2011). It also emerges naturally as the correct measure of network centrality in perfectly competitive models such as Acemoglu et al. (2012).

In this paper, we highlight the fragility of this aggregation result and show that for quantitatively relevant cases, the first-order approximation can be misleading. In this section we provide a characterization of the second-order effects in terms of reduced-form elasticities that we define below. Later on, we show how these reduced-form elasticities arise from structural primitives using a structural model.

First, recall that for any homothetic function $f(A_1, \dots, A_N)$, the Morishima (1967) elasticity of substitution is

$$\frac{1}{\rho_{ij}} = -\frac{d \log(MRS_{ij})}{d \log(A_i/A_j)} = -\frac{d \log(f_i/f_j)}{d \log(A_i/A_j)}, \quad (1)$$

where MRS_{ij} is the ratio of partial derivatives with respect to A_i and A_j , and $f_i = \partial f / \partial A_i$. This is a generalization of the two-variable elasticity of substitution introduced by Hicks (1932) and analyzed in detail by Blackorby and Russell (1989).

When the homothetic function f corresponds to a CES utility function and A_i to quantities, ρ_{ij} is the associated elasticity of substitution parameter. However, we do not impose this interpretation, and instead treat it as a reduced-form measure of the curvature of isoquants. By analogy, and with a slight abuse of language, we define the reduced-form elasticity of substitution for non-homothetic functions in a similar fashion.

Definition 2.1. For an output function $C : \mathbb{R}^N \rightarrow \mathbb{R}$, define the *macro elasticity of substitution* as

$$\frac{1}{\rho_{ij}} \equiv -\frac{d \log(MRS_{ij})}{d \log(A_i)} = -\frac{d \log(C_i/C_j)}{d \log(A_i)}.$$

The macro elasticity of substitution ρ_{ij} is interesting because it measures changes in the relative sales shares of i and j when there is a shock to i . This follows from the fact that

$$\frac{d \log(\lambda_i/\lambda_j)}{d \log A_i} = \frac{d \log[(C_i A_i)/(C_j A_j)]}{d \log A_i} = \frac{d \log(C_i/C_j)}{d \log A_i} + 1 = 1 - \frac{1}{\rho_{ij}}, \quad (2)$$

where the first equality applies Hulten's theorem. Since the second-order impact of a shock to i can be measured in terms of the rate of change in the sales share of i , the macro elasticity of substitution will turn out to be an important sufficient statistic. A decrease in the productivity of i causes λ_i/λ_j to increase when $\rho_{ij} \in (0, 1)$, and to decrease otherwise. We say that a j is a *macro-complement* for i if $\rho_{ij} \in (0, 1)$, and a *macro-substitute* otherwise. When f is a CES aggregator, then this coincides with the standard definition of gross complements and substitutes. As usual, when f is Cobb-Douglas, i and j are neither substitutes nor complements. In general, macro-substitutability is not reflexive.

The second object we need to define is the following.

Definition 2.2. Define the *input-output multiplier* to be

$$\xi \equiv \sum_{i=1}^N \frac{d \log C}{d \log A_i} = \sum_{i=1}^N \lambda_i.$$

When $\xi > 1$, total sales of the shocked factors exceed total income, a symptom of intermediate inputs. The impact of a uniform technology shock is correspondingly amplified due to the fact that goods are reproducible. Loosely speaking, ξ captures the percentage change in output in response to a uniform one-percent increase in technology. In this sense, it captures a notion of returns-to-scale at the aggregate level.

The input-output multiplier is called the intermediate input multiplier in a stylized model by Jones (2011), but it also appears under other names in many other contexts. It is also related to the network influence measure of Acemoglu et al. (2012), the granular multiplier of Gabaix (2011), international fragmentation measure of Feenstra and Hanson (1996), the production chain length multiplier in Kim et al. (2013), and even the capital multiplier in the neoclassical growth model.¹³ It also factors into how the introduction of intermediate inputs amplifies the gains from trade in Costinot and Rodriguez-Clare (2014). Although these papers feature multiplier effects due to the presence of round-about production (either via intermediate inputs or capital), they do not take into account the

¹³This follows from the fact that we can treat capital as an intertemporal intermediate input. For more details on how capital can be thought of in this framework, see Hulten (2001).

fact that this multiplier effect can respond to shocks. This is either because they assume constant factor shares or because they focus on first-order effects.

Having defined the macro elasticities of substitution and the input-output multiplier, we are in a position to characterize the second-order terms. We start by investigating the impact of an idiosyncratic shock.

Idiosyncratic Shocks

Theorem 2.2. *Suppose that C is homogenous of degree ξ , then*

$$\frac{d^2 \log C}{d \log A_i^2} = \frac{\lambda_i}{\xi} \sum_{j \neq i} \lambda_j \left(1 - \frac{1}{\rho_{ij}}\right).$$

When C is not homogeneous

$$\frac{d^2 \log C}{d(\log A_i)^2} = \frac{\lambda_i}{\xi} \sum_{j \neq i} \lambda_j \left(1 - \frac{1}{\rho_{ij}}\right) + \lambda_i \frac{\partial \log \xi}{\partial \log A_i}.$$

Proof. Differentiate $\sum_i \lambda_i = \xi$ to get

$$\begin{aligned} \lambda_i \frac{d \log \lambda_i}{d \log A_i} &= \xi \frac{d \log \xi}{d \log A_i} - \sum_{j \neq i} \lambda_j \frac{d \log \lambda_j}{d \log A_i}, \\ &= \xi \frac{d \log \xi}{d \log A_i} - \sum_{j \neq i} \lambda_j \frac{d \log \lambda_j / \lambda_i}{d \log A_i} - \sum_{j \neq i} \lambda_j \frac{d \log \lambda_i}{d \log A_i}, \end{aligned}$$

which using (2), we can rewrite as

$$\lambda_i \frac{d \log \lambda_i}{d \log A_i} = \xi \frac{d \log \xi}{d \log A_i} + \sum_{j \neq i} \lambda_j \left(1 - \frac{1}{\rho_{ij}}\right) - (\xi - \lambda_i) \frac{d \log \lambda_i}{d \log A_i},$$

Rearrange this to get

$$\xi \frac{d \log \lambda_i}{d \log A_i} = \xi \frac{d \log \xi}{d \log A_i} + \sum_{j \neq i} \lambda_j \left(1 - \frac{1}{\rho_{ij}}\right). \quad (3)$$

Finally, Theorem 2.1 implies that

$$\frac{d^2 \log C}{d \log (A_i)^2} = \lambda_i \frac{d \log \lambda_i}{d \log A_i}.$$

Substitute (3) into the expression above to get the desired result. Lastly, if C is homogeneous, Euler's theorem implies that

$$\sum_i \frac{dC}{dA_i} \frac{A_i}{C} = \sum_i \lambda_i = \xi,$$

hence, $d \log \xi / d \log A_i = 0$. ■

In words, the second-order impact of a shock to i is equal to the change in i 's sales share λ_i . The change in i 's share of sales is the change in the aggregate sales to GDP ratio, minus the change in the share of sales of all other industries. The former is measured by the elasticity of the input-output multiplier ξ , while the latter depends on the macro elasticities of substitution. Collectively, the sales shares λ_i , the reduced-form elasticities ρ_{ij} and the reduced-form elasticity of the input-output multiplier $d \log \xi / d \log A_i$ are sufficient statistics for how output responds to technology shocks up to a second order.

This result tells us that Hulten's approximation is globally accurate if reduced-form elasticities are unitary $\rho_{ij} = 1$ for every j and if the input-output multiplier ξ is independent of the shock A_i . We shall see that this amounts to assuming Cobb-Douglas production and consumption functions. At the opposite extreme, the output function is nearly singular if $\rho_{ij} \approx 0$ for any j . Hence, first-order approximations will perform more poorly as ρ_{ij} approaches zero, either from below or from above. These are the cases of extreme macro-complementarity or extreme macro-substitutability. In the limiting case $|\rho_{ij}| \rightarrow 0$, the first-order approximation is completely uninformative even for arbitrarily small shocks. Finally, the first-order approximation also behaves worse when $|\rho_{ij}| \rightarrow \infty$, although the size of the second-order terms remains bounded in this case.¹⁴

Therefore, although the Cobb-Douglas special case is very popular in the literature, it constitutes a very special case where the second-order terms are all identically zero.¹⁵ There is increasing evidence for strong complementarities in supply chains, and Theorem 2.2 indicates that Cobb-Douglas functional forms may be a poor guide to understanding

¹⁴In Appendix B, we derive a tight bound on the size of the error from the first-order approximation for a special case.

¹⁵See for example Acemoglu et al. (2012), Long and Plosser (1983), Bigio and La'O (2016), Acemoglu et al. (2017), Bartelme and Gorodnichenko (2015).

the behavior of the economy in the presence of these complementarities. Finally, Theorem 2.2 also shows that there is an interaction between the macro elasticity of substitution between i and j and the size of i and j . In the extreme case where either λ_i or λ_j is equal to zero, the macro elasticity of substitution between the two is irrelevant.

Theorem 2.2 also shows that deviations from Hulten's theorem need not be restricted to non-unitary macro elasticities of substitution, they can also arise from variations in the input-output multiplier ξ . Since the input-output multiplier is the ratio of sales to GDP, changes in the input-output multiplier can be interpreted as another kind of macro elasticity of substitution: namely the substitution between the underlying factors (whose payments are GDP) and the reproducible goods (whose payments are sales). If there is a strong tendency to substitute between labor and intermediate inputs in response to shocks, then this will hamper the accuracy of the first-order approximation.

The second-order approximation to the output function can then be written as

$$\log(C) \approx \log(\bar{C}) + \lambda_i \log(A_i) + \frac{\lambda_i}{\xi} \sum_{j \neq i} \lambda_j \left(1 - \frac{1}{\rho_{ij}}\right) (\log(A_i))^2 + \lambda_i \frac{\partial \log \xi}{\partial \log A_i} (\log(A_i))^2, \quad (4)$$

where \bar{C} is C evaluated at the steady-state technology values. When goods are macro-complements, the second-order terms amplify the effect of negative shocks and attenuate the effect of positive shocks relative to the first-order approximation. Instead when goods are macro-substitutes, the second-order approximation attenuates the negative shocks and amplifies the positive shocks instead. A similar intuition holds for the input-output multiplier: if the input-output multiplier is increasing, then the second-order approximation amplifies positive shocks and dampens negative shocks, and if this multiplier is decreasing, then the opposite is true.

Correlated Shocks

To consider shocks to several industries at once, we must extend these results to cover the off-diagonal terms in the Hessian.

Proposition 2.3. *Let ξ be the intermediate-input multiplier, or the ratio of total sales to output, then*

$$\frac{d^2 \log C}{d \log A_i d \log A_j} = \frac{\lambda_i}{\xi} \sum_{k \neq j} \lambda_k \left(1 - \frac{1}{\rho_{jk}}\right) + \lambda_i \frac{\partial \log \xi}{\partial \log A_j} - \lambda_i \left(1 - \frac{1}{\rho_{ji}}\right). \quad (i \neq j)$$

This result shows that the cross-partial derivatives are non-trivial, but are characterized by the same collection of sufficient statistics as the second-derivatives. Even in the simplest case, where $\rho_{ij} = \rho$ for all i and j and ξ is constant, the second-order effect of a common shock is not simply twice the second-order impact of an idiosyncratic shock.¹⁶ We revisit the issue of how common shocks to different industries may interact with one another in Appendix D.

Macro Moments

Finally, we can use the second-order terms to approximate an economy's macroeconomic moments. First, we begin by looking at average performance.

Proposition 2.4 (Average Performance). *Suppose that $(\log A_i)_{i=1}^N$ is a random vector with covariance matrix Σ and s_{ij} is the ij th element of Σ . Then*

$$E(\log(C/\bar{C})) \approx \frac{1}{2} \sum_{i=1, j \neq i}^N \left(\frac{\lambda_i}{\xi} \sum_{k \neq j} \lambda_k \left(1 - \frac{1}{\rho_{jk}} \right) + \lambda_i \frac{\partial \log \xi}{\partial \log A_j} - \lambda_i \left(1 - \frac{1}{\rho_{ji}} \right) \right) s_{ij} + N \sum_i \lambda_i s_{ii}.$$

In the special case where the shocks are i.i.d, $\Sigma = s^2 I$ and this simplifies to

$$E(\log(C/\bar{C})) \approx \frac{1}{\xi} \frac{s^2}{2} \sum_i \lambda_i \sum_j \lambda_j \left(1 - \frac{1}{\rho_{ij}} \right) + \frac{s^2}{2} \sum_i \lambda_i \frac{d \log \xi}{d \log A_i}.$$

The logic of proposition 2.4 is best seen by considering its absurd limit: we could have two economies with identical sales distributions and identical output evaluated at the steady-state technology. Up to a first order, these two economies are the same. However, if one of these economies has $\rho_{ij} = 1$ and the other has $\rho_{ij} > 0$ arbitrarily close to 0, then in the presence of any volatility $s > 0$, to a second order, the first economy will produce \bar{C} on average whereas the second economy will produce nothing.

Proposition 2.4 implies that $\frac{1}{2} d^2 \log C / d \log A_i^2$ represents both the second-order impact of a shock to i on GDP, and the log point difference between expected output and its certainty equivalent in units of variance. Loosely speaking, we can interpret this as the percent change in output relative to its certainty equivalent in units of variance, with the caveat that such a description is only approximately true.

¹⁶We can also use these ideas to capture the impact of an aggregate shock to the economy, since an aggregate shock is simply a common shock that affects all industries. If A is an aggregate TFP shock, then $\frac{d^2 \log C}{d \log A^2} = \xi \sum_i \frac{d \log \xi}{d \log A_i}$. So, for aggregate shocks, deviations from Hulten's theorem can only come from the input-output multiplier.

The second-order terms will also shape other moments of the distribution of GDP. To see this, suppose that the second-order terms are negative, corresponding to the case with high-degrees of complementarity. In this case, the distribution of GDP will endogenously be skewed to the left and fat-tailed, even if the technology shocks are symmetric and thin-tailed. This follows from the fact that the second-order terms magnify negative shocks and attenuate positive shocks, which makes the distribution skewed. Furthermore, since the negative shocks are magnified, this also fattens the left tail, giving rise to excess kurtosis.

To illustrate this intuition, let $\log A_i$ be a normal random variable with mean 0 and variance s^2 . Then the skewness of \log GDP is

$$E \left(\left(\frac{\log(C/\bar{C}) - \mu_c}{\sigma_C} \right)^3 \right) \approx \frac{1}{\sigma_C^3} \frac{d^2 \log C}{d \log A_i^2} \left[s^6 \left(\frac{d^2 \log C}{d \log A_i^2} \right)^2 + 3\lambda_i^2 s^4 \right],$$

where

$$\mu_c = E(\log(C/\bar{C})) = \frac{1}{2} \frac{d^2 \log C}{d \log A_i^2} s^2,$$

and

$$\sigma_c^2 = Var(\log(C/\bar{C})) = \left[\lambda_i^2 + 2 \left(\frac{\mu_c}{s} \right)^2 \right] s^2.$$

Hence \log output is negatively skewed if, and only if, the second-order term is negative. This asymmetry also helps explain why average \log GDP is lower than its deterministic steady state, since GDP is subject to larger recessions than booms. Next, we consider the thickness of tails, as measured by kurtosis. In this case, a second-order approximation gives

$$E \left(\left(\frac{\log(C/\bar{C}) - \mu_c}{\sigma_C} \right)^4 \right) \approx 3 \left(1 + \left(\frac{\mu_c}{s} \right)^2 \frac{[22(\mu_c/s)^2 + 7\lambda_i^2]}{(\lambda_i^2 + 2(\mu_c/s)^2)^2} \right) \geq 3.$$

So, output has excess kurtosis if, and only if, the second-order terms are nonzero.

To summarize, for a given variance of output, relatively more of the variance is due to negative, infrequent, extreme deviations, as opposed to symmetric, frequent, and modestly sized deviations (relative to a normal distribution). In Section 8 we revisit these issues with a calibrated model and show that they are quantitatively significant.

Welfare Costs of Business Cycles

For the majority of the paper, we focus on the performance of $\log GDP$, since this gives rise to unitless elasticities. This assumption is innocuous for welfare questions. One may

imagine that the losses from uncertainty that we identify depend on the concavity of the log function. In other words, a consumer with log utility in aggregate consumption prefers a mean-preserving reduction in uncertainty even when the GDP function is linear. However, as shown by Lucas (1987), such losses are extremely small in practice. The much larger effects we identify are nonlinearities in aggregate consumption itself, which are present even when the utility function is linear in aggregate consumption. The following proposition formalizes this intuition and shows that the Lucas welfare losses from risk-aversion, and the losses we identify from nonlinear production, do not interact with one-another up to a second-order approximation. In fact, it could easily be the case that a risk-averse household prefers the economy to be subject to stochastic shocks *if* the economy features macro-substitutability and the second-order terms are positive.¹⁷

Proposition 2.5. *Let $u : \mathbb{R} \rightarrow \mathbb{R}$ be a constant relative risk aversion utility function with parameter γ and let $C : \mathbb{R}^N \rightarrow \mathbb{R}$ be the GDP function. Suppose that TFP shocks have mean \bar{A} and a diagonal covariance matrix with k th diagonal element s_k^2 . Then*

$$E(u(C)) - u(C(\bar{A})) = -\frac{1}{2}\gamma \sum_k^N \lambda_k^2 s_k^2 + \frac{1}{2} \sum_k^N \frac{\partial^2 C}{\partial A_k^2} s_k^2.$$

The first term, which is quantitatively small, is the traditional Lucas cost arising from curvature in the utility function. The second term, which is quantitatively large, is due to the curvature inherent in production. Proposition 2.5 is stated for CRRA utility and idiosyncratic shocks for expositional clarity. In the appendix, we prove the result for more general utility functions and shocks.

3 Structural Model

Theorem 2.2 implies that the macro elasticities of substitution ρ_{ij} and the elasticity of the input-output multiplier $d \log \xi / d \log A_i$ are sufficient statistics for the second-order impact of shocks. However, these sufficient statistics are reduced-form elasticities, and unlike λ_i and ξ , they are not readily observable. Furthermore, since they are general equilibrium

¹⁷These ideas also relate to the concepts of fragility, resilience, and antifragility in Taleb (2013). In Section 3, we find that economies with immobile factors and structural complementarities are fragile, in the sense of having large negative second derivatives, whereas economies with mobile factors and structural complementarities are resilient, in the sense of having smaller negative second derivatives. However, we find that economies with mobile factors and structural substitutabilities are antifragile in the sense that their average performance improves with uncertainty.

objects, they cannot be identified through exogenous microeconomic variation. So, on the one hand, while careful empirical work can identify structural elasticities in production, the leap from micro-estimates to macro-effects is hazardous. On the other hand, while the macro elasticities could, in principle, be identified using exogenous macro-variation, such a reduced-form exercise will be susceptible to a form of the Lucas critique, because the estimated elasticities could shift in unpredictable ways. This is because the reduced-form macro elasticities will not necessarily be stable deep parameters. Worse still, plausibly exogenous macroeconomic variation is notoriously difficult to come by.

In this section, we make the mapping from structural micro parameters to the reduced-form macro elasticities explicit for a relatively general class of structural models. This helps to bridge the gap between estimates of microeconomic structural parameters and the macroeconomic elasticities of interest. More generally, the structural model shows that this mapping is highly nontrivial and clarifies some of the general equilibrium forces which must be accounted for.

Let

$$\frac{C}{\bar{C}} = \left(\sum_k b_k \left(\frac{c_k}{\bar{c}_k} \right)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}},$$

with

$$\sum_k p_k c_k = \sum_k w L_k + \sum_k w_k l_k + \sum_k \pi_k,$$

where we divide labor into an industry-specific component l_k , which cannot be reallocated across industries, and a common component L_k , which can be allocated to any industry. Here, w is the wage for the common labor and w_k is the wage for the industry-specific labor. We normalize $\sum_i b_i = 1$. Any variable with an overline \bar{x} is a normalizing constant denoted in the same units as x .

Let industry k 's production function be given by

$$\frac{y_k}{\bar{y}_k} = A_k \left(a_k \left(\left(\frac{L_k}{\bar{L}_k} \right)^{\beta_k} \left(\frac{l_k}{\bar{l}_k} \right)^{1-\beta_k} \right)^{\frac{\theta_k-1}{\theta_k}} + (1-a_k) \left(\frac{X_k}{\bar{X}_k} \right)^{\frac{\theta_k-1}{\theta_k}} \right)^{\frac{\theta_k-1}{\theta_k}},$$

where A_k is a Hicks-neutral shock, and X_k is a composite intermediate input which is combined with labor with elasticity of substitution θ_k . The labor input is a geometric average of industry-specific and common labor inputs. Therefore, β_k can be interpreted as a measure of the degree to which industry k 's labor can be reallocated. Since l_k is

industry-specific, in equilibrium, it is always equal to a constant, and so we can also treat β_k as technological returns-to-scale in labor. These two interpretations are equivalent for the purposes of this model. As before, this specification includes factor-augmenting shocks as a special case.

The composite intermediate input X_k is defined by

$$\frac{X_k}{\bar{X}_k} = \left(\sum_l \omega_{kl} \left(\frac{x_{lk}}{\bar{x}_{lk}} \right)^{\frac{\varepsilon_k - 1}{\varepsilon_k}} \right)^{\frac{\varepsilon_k - 1}{\varepsilon_k}},$$

and x_{lk} are intermediate inputs from industry l used by industry k . We normalize $\sum_k \omega_{ik} = 1$. Market clearing in each market requires that

$$y_k = c_k + \sum_l x_{lk}.$$

The production functions here allow the accommodation of any pattern of nested CES production functions and networks as a special case.

Proposition 3.1. *Let $D_i \log \lambda$ be the sales elasticities column vector whose j th element is $d \log \lambda_j / d \log A_i$. Then, at steady state, $D_i \log \lambda$ can be written explicitly as a function of observable expenditure shares a_i, b_i, ω_{ij} at steady state, structural elasticities $\sigma, \theta_i, \varepsilon_i$, and structural returns to scale β_i .*

The proof can be found in the appendix. Using this proposition, and equation (2), we can easily deduce the reduced-form macro elasticities:

$$\frac{d^2 \log C}{d \log A_k^2} = \frac{d \lambda_k}{d \log A_k} = \frac{\lambda_k}{\xi} \sum_{j \neq k} \lambda_j \left(1 - \frac{1}{\rho_{kj}} \right) + \lambda_k \frac{d \log \xi}{d \log A_k},$$

with

$$1 - \frac{1}{\rho_{ij}} = \frac{1}{\lambda_k} \frac{d \lambda_k}{d \log A_k} - \frac{1}{\lambda_j} \frac{d \lambda_j}{d \log A_k}, \quad (5)$$

and

$$\frac{d \log \xi}{d \log A_i} = \frac{1}{\xi} \sum_i \frac{d \lambda_k}{d \log A_i}. \quad (6)$$

In what follows, we work with increasingly complex special cases to show the various channels through which the reduced-form elasticities ρ_{ij} and ξ operate, both qualitatively and quantitatively.

4 Macro Elasticities of Substitution

In this section, we restrict ourselves to the case with no intermediate inputs. This means that the input-output multiplier is constant $\xi = 1$, and the deviations from Hulten's theorem only occur due to non-unitary macro elasticities of substitution. Since there are no intermediate inputs, $a_k = 1$ for every k . We emphasize how the following key structural parameters shape the macro elasticities of substitution: micro elasticities of substitution between sectors σ , the degree to which labor can be reallocated, and returns to scale in production.

Different degrees of labor reallocations and returns to scale can be expected depending on the time scale of the response to shocks. At short horizons, labor and other factors such as capital are difficult to adjust, but such adjustments become easier at longer horizons. Some of these dynamic effects can be captured by comparative statics exercise in our model. We continuously move from no reallocation to full reallocation in labor markets by considering production functions with different $\beta_k \in [0, 1]$. As mentioned before, this parameter can either be interpreted as decreasing returns to scale in labor, or a geometric average of mobile and immobile workers, where the immobile workers are residual claimants of the firm's revenues net of other costs. Hence, the production function of good k is

$$c_k = A_k L_k^{\beta_k}.$$

Proposition 4.1. *[Limited Labor Reallocation] Consider the following special case of the structural model in Section 3. Aggregate consumption is CES with micro elasticity of substitution σ and expenditure share b_i at steady state. Each good is produced using labor. Assume uniform labor reallocation/returns to scale $\beta \in [0, 1]$ for every k . Then*

$$\rho_{ij} = \frac{\sigma(1 - \beta) + \beta}{\sigma(1 - \beta) + \beta + (1 - \sigma)}, \quad \lambda_i = b_i, \quad \xi = 1, \quad \frac{d \log \xi}{d \log A_i} = 0.$$

To build intuition, we first consider two polar cases with either $\beta_k = 0$ for every k or $\beta_k = 1$ for every k .

We start with the case in which $\beta_k = 0$ for every k . This is an endowment economy where the household simply consumes the output of every industry, and there are no intermediate inputs. Labor cannot be moved to increase the production of any good, either because production is simply an endowment or because labor cannot be reallocated

across industries in response to shocks. The latter case proxies for a situation where there are infinite adjustment costs in reallocating workers across industries. Then we can write

$$\frac{C}{\bar{C}} = \left(\sum_k b_k A_k^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}.$$

Unsurprisingly, for this special case, the macro and micro elasticity of substitution coincide

$$\rho_{ij} = \sigma, \quad \lambda_i = b_i, \quad \xi = 1, \quad \frac{d \log \xi}{d \log A_i} = 0.$$

Theorem 2.2 then implies that

$$\frac{d^2 \log C}{d \log A_i^2} = b_i(1 - b_i) \left(1 - \frac{1}{\sigma} \right).$$


It can immediately be seen that the second-order term changes sign depending on whether σ is greater or less than one. Hence, the second-order term amplifies negative shocks and attenuates positive shocks if $\sigma < 1$ relative to the first-order approximation, and the opposite is true if $\sigma > 1$. Since $\sigma \in [0, \infty)$ this means that $\rho \in [0, \infty)$ for this example.

In the Cobb-Douglas case $\sigma = 1$, the second-order term is identically equal to zero and the first-order approximation is globally accurate. The quality of the Hulten approximation deteriorates as we move away from $\sigma = 1$ in both directions. To understand why, it is useful to consider the extreme limits $\sigma \rightarrow 0$ and $\sigma \rightarrow \infty$.


We first consider the Leontief limit $\sigma \rightarrow 0$ where the first-order term becomes completely uninformative. To understand why this happens, consider how the sales share λ_i changes in response to a shock. We write

$$\frac{d \log(\lambda_i/\lambda_j)}{d \log A_i} = \frac{d \log(p_i/p_j)}{d \log A_i} + \frac{d \log(y_i/y_j)}{d \log A_i} = \frac{d \log(p_i/p_j)}{d \log A_i} + \frac{d \log(A_i/A_j)}{d \log A_i}.$$

When labor cannot be reallocated, the ratio of the quantities y_i/y_j is equal to the exogenously given A_i/A_j . However, close to the Leontief limit, the change in the sales share of i is very extreme, since the relative price of i to j goes to zero if the shock to i is positive, and to infinity if the shock is negative. Intuitively, this is because with extreme complementarity the scarcest good is the only one that has a positive marginal product. The extreme reaction of relative prices means that sales shares react very strongly to productivity shocks, and this means that the deviations from the first-order approxima-



(a) log GDP with no reallocation/extreme decreasing returns. Perfect substitutes and Hulten's approximation overlap almost perfectly.



(b) log GDP with full reallocation/constant returns. Leontief and Hulten's approximation overlap almost perfectly.

Figure 1: log GDP as a function of productivity $\log(A_i)$ in the economy with constant returns for different values of σ .

tion can be extreme. In this world, the second-order approximation amplifies the impact of negative shocks and attenuate the impact of positive shocks relative to the first-order approximation.

We then consider the perfect substitutes limit $\sigma \rightarrow \infty$. Then negative shocks are attenuated and positive shocks are amplified, but the effect is not nearly so dramatic. In this case, because goods are perfect substitutes, relative prices are always equal to 1. Therefore, ratio of the sales of i relative to j will move one-for-one with the ratio of the endowment of i relative to j . The situation is depicted graphically in Figure 1a.

Having analyzed the case with no labor reallocation, consider now the polar opposite case, where labor can be costlessly reallocated across industries and be used with constant returns to scale so that $\beta_k = 1$ for every k . The macro elasticity of substitution in this example is *not* necessarily equal to the structural micro elasticity of substitution in consumption:

$$\rho_{ij} = \frac{1}{2 - \sigma}, \quad \lambda_i = b_i, \quad \xi = 1, \quad \frac{d \log \xi}{d \log A_i} = 0.$$

Theorem 2.2 then implies that

$$\frac{d^2 \log C}{d \log A_i^2} = b_i(1 - b_i)(\sigma - 1).$$

As before, $\sigma > 1$ amplifies positive shocks and $\sigma < 1$ amplifies negative shocks, while the Cobb-Douglas case $\sigma = 1$ still ensures that the second-order terms are identically zero. However, this time, the second-order term becomes singular when the goods are highly substitutable rather than when they are highly complementary. Once again, we can unpack this result by noting that

$$\frac{d \log(\lambda_i/\lambda_j)}{d \log A_i} = \frac{d \log(p_i/p_j)}{d \log A_i} + \frac{d \log(y_i/y_j)}{d \log A_i} = \frac{d \log(A_j/A_i)}{d \log A_i} + \frac{d \log(C_i/C_j)}{d \log A_i}.$$

The ratio of relative prices is always equal to A_j/A_i , but the quantity of goods produced is endogenous since labor can be costlessly reallocated.

Contrary to what one may have assumed, a near-Leontief production function is *not* sufficient for generating large deviations from Hulten's theorem, as long as factors can be reallocated freely. With perfect reallocation of workers, the market always allocates workers to equate marginal products. This means that relative prices reflect relative productivities. So, near the Leontief limit, if i receives a negative productivity shock, workers are reallocated to that industry to reinforce the “weak link”, since otherwise the price of that good would soar. On the other hand, if i receives a positive productivity shock, workers are reallocated to other industries, to prevent the prices from collapsing. This means that near the Leontief limit, the relative sales shares responds one-for-one to changes in relative technology. So, while the second-order terms still amplify negative shocks and attenuate positive shocks, because they are negative, their magnitude is much smaller than in the case where labor could not be reallocated.

In the perfect substitutes limit, equating marginal products means a near-complete reallocation of workers to the most productive industry. This means that a positive shock to i will cause the sales share of i to increase dramatically relative to the rest, because while relative prices $p_i/p_j = A_j/A_i$, relative quantities change very rapidly. So, the second-order terms amplifies positive shocks and attenuates negative shocks relative to the first-order approximation, but their magnitude is much larger than in the case where labor could not be reallocated, because now the market can take advantage of the shocks to equate marginal products. Once again, the situation is depicted graphically in Figure 1b.

To recap, in the case where labor cannot be freely allocated, a negative shock can cause a large downturn due to complementarity but a positive shock does not make much of a difference. On the other hand, when labor can be allocated, a negative shock can be mitigated by a reallocation of workers to the affected industry, but a positive can be

amplified many times if goods are substitutable. These results are closely related to the findings in Jones (2011), who noted that in a model like this, the relevant CES parameter used in aggregating microeconomic TFP shocks depends on whether or not factors are allocated through the market or assigned exogenously.

Finally, having analyzed the cases with no labor reallocation ($\beta = 0$) and with full labor reallocation/constant returns to scale ($\beta = 1$), we revisit the general case with $\beta \in (0, 1)$. The situation is depicted graphically in Figure 2. The macro elasticity of substitution is now an intermediate value between the perfect reallocation and the no-reallocation cases.

As usual, in the Cobb Douglas case $\sigma = 1$, the macro elasticity is equal to 1, and Hulten's approximation is globally accurate. Therefore, regardless of $\beta \in (0, 1)$, the macro elasticity of substitution is always equal to 1 for a Cobb-Douglas model. This follows from the fact that in this case, the distribution of workers across sectors does not depend on the shock. Therefore, the degree to which workers can be reallocated is irrelevant, since they are not reallocated anyway.

When $\sigma < 1$, negative shocks are amplified and positive shocks are attenuated relative to the first-order approximation, and the opposite is true when $\sigma > 1$. However, we can show that $\rho(\sigma) \notin (-(1-\beta)/\beta, \beta/(1+\beta))$, so the macro elasticity of substitution is bounded away from 0 as long as $\beta > 0$. Furthermore, whether or not i and j are macro-complements or macro-substitutes does not depend on β . Therefore, partial reallocation preserves the general intuition of the last two sections, but dampens the size of the effects. Figure 2 illustrates these facts where we can see that the size of the second-order term looks like an average of the two polar cases and always has the same sign.

5 A Network Irrelevance Result

Now we extend the model to allow for intermediate inputs and arbitrary network interconnections. In this section, we provide a benchmark irrelevance result where the deviations from Hulten's approximation do not depend on the network structure. The key assumptions required for obtaining this irrelevance result are: (1) a constant input-output multiplier equal to unity, (2) uniformity of the micro structural elasticities. In Sections 6 and 7, we weaken each of these assumptions, in turn, and characterize the importance of the network structure in shaping the second-order terms.

To obtain our benchmark result, assume that structural micro elasticities are uniform across all agents and all inputs so that $\sigma = \theta_j = \varepsilon_j$ for every j . In this case, the model




Figure 2: The left panel plots the macro elasticity of substitution ρ against the structural elasticity of substitution σ for the model from Proposition 4.1. The second panel plots expected log GDP relative to its certainty equivalent, in units of variance, as a function of the structural elasticity of substitution for the same model.

becomes a generalization of the canonical Cobb-Douglas network model of Acemoglu et al. (2012) and Long and Plosser (1983), as well as the CES competitive network model of Baqaee (2016). Furthermore, assume that $\beta_k = 1$ or $\beta_k = 0$ for all k — the case with either full reallocation or no reallocation. Finally, technology shocks are labor-augmenting, which implies that $\xi \equiv 1$.

Proposition 5.1 (Network Irrelevance). *Consider the following special case of the structural model in Section 3. Network interconnections are arbitrary, micro elasticities are uniform across all agents and all inputs so that $\epsilon_k = \theta_k = \sigma$, and shocks are labor-augmenting. Then*

$$\rho_{ij} = \rho, \quad \xi = 1, \quad \frac{d \log \xi}{d \log A_i} = 0,$$

where

$$\rho = \begin{cases} \sigma & \text{if labor cannot be reallocated} \\ \frac{1}{2-\sigma} & \text{if labor can be reallocated} \end{cases}.$$

This implies that

$$\frac{d^2 \log C}{d \log A_i^2} = \lambda_i(1 - \lambda_i) \left(1 - \frac{1}{\rho}\right),$$

where $\lambda_i = w_i L_i / P_c C$.

First, note that, once again the Cobb-Douglas specification $\sigma = 1$ is the special case where labor reallocation becomes irrelevant. This is again a consequence of the fact that the distribution of workers does not depend on the shock in equilibrium, and therefore, reallocation of workers (or alternatively, returns to scale in labor) is irrelevant. However, this is only true in the Cobb-Douglas special case. Generically, returns to scale and factor reallocation will have a large effect on how the economy behaves.

Second and more importantly, when $\sigma \neq 1$, there is a deviation from Hulten's approximation, but the structure of the network remains irrelevant up to the second order, since the second-order approximation only depends on sales λ_i , the micro elasticity of substitution σ , and the extent of labor reallocation.

This network irrelevance result is driven by the fact that the micro elasticities of substitution and labor reallocation are the same in all industries, and $\xi \equiv 1$ (since technology shocks are labor augmenting). This makes the ρ_{ij} independent of the network structure, and ensures that the input-output multiplier ξ is inoperative. In Sections 6 and 7 we show that weakening either of these two assumptions breaks this irrelevance result. First, in Section 6, we show how input-output multipliers become variable in the presence of intermediate inputs and Hicks-neutral shocks. Then, in Section 7, we show how heterogeneity in the structural elasticities of substitution interact with the existence of general network linkages to accentuate output nonlinearities.

6 Input-Output Multiplier

So far, we have kept ξ constant and shown how macro elasticities of substitution can cause large deviations from Hulten's theorem. In this section, we instead focus on how variability in ξ can also generate large deviations from Hulten's theorem, even when ρ_{ij} are well-behaved.

We consider the simplest model with both intermediate inputs (a non-trivial network)

and Hicks-neutral technology shocks, weakening one of the assumptions (that shocks are labor augmenting) required for the network irrelevance result of Section 5. Indeed we find that this change upsets the result in the sense that the details of the network matter for at the second order in a way that we make precise.

To demonstrate the effect of a variable intermediate-input multiplier ξ , consider a special case where there are no deviations from Hulten's theorem arising from substitutability across sectors. In other words, consider the economy with a single good ($N = 1$) and with full labor reallocation/constant returns to scale (β_1), where total output is given by

$$\frac{Y}{\bar{Y}} = A \left(a \left(\frac{L}{\bar{L}} \right)^{\frac{\theta-1}{\theta}} + (1-a) \left(\frac{X}{\bar{X}} \right)^{\frac{\theta-1}{\theta}} \right)^{\frac{\theta}{\theta-1}}.$$

Suppose that labor is an endowment, and so $L/\bar{L} = 1$. GDP is given by

$$C = Y - X,$$

where X are intermediate inputs.

Proposition 6.1 (Variable IO multiplier). *Consider the following special case of the structural model in Section 3. There is a single good which is used both for consumption and as an intermediate input in production. Assume full labor reallocation/constant returns to scale. Then*

$$\xi = \frac{1}{a}, \quad \frac{d \log \xi}{d \log A} = (1-a)(\theta-1),$$

where θ is the micro structural elasticity of substitution in production between labor and the intermediate input and $1-a$ is the intermediate input share in steady-state.

This implies that

$$\frac{d^2 \log C}{d \log A^2} = \xi \frac{d \log \xi}{d \log A} = \frac{1-a}{a}(\theta-1).$$

Proposition 6.1 shows that although in partial equilibrium, the production function is homogenous in TFP, in general equilibrium, aggregate output is not homogeneous of degree 1. Furthermore, output is not homogenous of any degree in equilibrium, since ξ varies in response to the shock. Hence, Hulten's approximation is exact whenever there are no intermediate inputs ($a = 1$) or the economy is Cobb-Douglas $\theta = 1$. Otherwise, any non-unitary elasticity of substitution ($\theta - 1$) is increased by a factor $(1-a)/a$, which is

singular when the labor share goes to zero. The details of the network matter through the parameter $1 - a$ which indexes the steady-state intermediate input share.

Intuitively, this results from the fact that output is used as its own input, and if $\theta \neq 1$, then the intermediate input share of GDP changes with the shock, an effect which a first-order approximation neglects. The larger is the steady-state intermediate input share, the larger is the effect of this change in the intermediate input share. Figure 3 plots $\log C$ as a function of $\log A$ for the case where $\theta \approx 0$, $\theta = 1$, and $\theta = 2$. As expected, the Cobb-Douglas special case leads to a log-linear relationship where Hulten's approximation is globally true. However, for the case where $\theta < 1$, negative shocks are amplified and positive shocks are attenuated, and the reverse is true when $\theta > 1$.

In the case where $\theta = 1$, the intermediate input-multiplier is $1/a$, since a one percentage increase in TFP would increase output by $1 + (1 - a) + (1 - a)^2 + \dots$ percentages. This is because the increase in TFP makes intermediate inputs more productive, which makes output more productive, which makes intermediate inputs more productive, and so on. When $\theta \neq 1$, this effect is either attenuated or amplified depending on whether the economy can substitute between intermediate inputs and labor relatively more or less than the Cobb-Douglas benchmark. In the limit where $\theta = 0$, output is linear in TFP rather than log-linear with slope $1/a$, whereas when $\theta = 2$, output is hyperbolic in TFP.¹⁸

7 General Networks

In Section 5, we showed that as long as shocks are labor-augmenting and structural elasticities of substitution are homogeneous, departures from Hulten's theorem do not depend on the network structure. In this section, we allow for heterogeneity in the structural elasticities of substitution as well as Hicks-neutral shocks. We consider arbitrary network structures with full reallocation/constant-returns-to-scale, and we characterize the way in which the structure of the network matters at the second order.

Along the way, we uncover two important insights. First, we show that the macro elasticities of substitution are some weighted average of the underlying structural micro elasticities of substitution and this average depends on the network structure. Further-

¹⁸In this example, the economy with extreme complementarity $\theta = 0$ has $C = A/a$, where $1/a$ is the sales to output ratio in steady state. Therefore, although Hulten's approximation fails in log terms, Hulten's theorem is globally accurate in *linear* terms. This is an artefact of the fact that we have only one good. In Appendix C, we generalize this example to multiple goods, and show that output can be very strongly nonlinear even with full labor reallocation.




Figure 3: Output as a function of productivity shocks $\log(A)$ with variable input-output multiplier effect with steady-state intermediate input share $1 - a = 0.9$.

more, we show that the structural elasticities of substitution in j 's production matters only so far as j is exposed to the shock in a heterogenous fashion. Second, we show that when there are constant-returns-to-scale at the micro level, *only* an industry's role as a supplier of inputs matters in terms of how output responds to shocks, up to the second order.

To do this, we need another definition.

Definition 7.1. The $N \times N$ *input-output* matrix Ω is the the matrix whose ij th element is equal to the steady-state value of

$$\Omega_{ij} = \frac{p_j x_{ij}}{p_i y_i}.$$

The *Leontief inverse* is

$$\Psi = (I - \Omega)^{-1}.$$

Intuitively, the ij th element ψ_{ij} of the Leontief inverse is a measure of i 's total reliance on j as a supplier. It captures both the direct and indirect ways through which i uses j in its production.

Proposition 7.1. *Consider the structural model in Section 3 with full labor reallocation/constant returns to scale ($\beta_i = 1$ for every i). Define*

$$\begin{aligned} c_i = & b_i(1 - \sigma)[\lambda_k - \psi_{ik}] + \sum_j \omega_{ji}(1 - a_j)\lambda_j(\theta_j - 1)[\sum_l \omega_{lj}\psi_{lk} - \psi_{jk}] \\ & + \sum_j \omega_{ji}(1 - a_j)\lambda_j(1 - \varepsilon_j)[\sum_l \omega_{lj}\psi_{lk} - \psi_{ik}] + (\theta_k - 1)\omega_{ki}(1 - a_k)\lambda_k. \end{aligned}$$

Then the sales elasticities vectors are given by

$$D_i \log \lambda' = \text{diag}(\lambda) c' \Psi, \quad (7)$$

where $\text{diag}(\lambda) = \sum_i \lambda_i e_i e_i'$ with e_i being the i th standard basis vector.

Recall, from (5) and (6), that we can easily use this proposition to recover the macro elasticities of substitution and the elasticity of the input-output multiplier.

We decompose Proposition 7.1 into something more interpretable by relabeling the industries such that we can assume $\theta_i = \varepsilon_i$ for every i . To do this, for every industry i , relabel its composite intermediate input X_i to be a new industry which uses no labor. Then, without loss of generality, we can impose the assumption that for every industry $\varepsilon_i = \theta_i$, as long as we keep in mind that we need to adjust the input-output matrix under this relabeling scheme.

Then define

$$\text{Cov}_{\Omega_{(j)}}(\Psi_{(m)}, \Psi_{(i)}) = \sum_k \Omega_{jk} \psi_{ki} \psi_{km} - \left(\sum_k \Omega_{jk} \psi_{ki} \right) \left(\sum_k \Omega_{jk} \psi_{km} \right).$$

In words, this is the covariance between the m th and i th column of the Leontief inverse using the j th row of the input-output matrix as the distribution. What makes this different to a standard covariance is that the weighting distribution, $\Omega_{(j)}$, do not necessarily have to sum to one. More generally, we could write $\text{Cov}_b(c, d)$, which takes the covariance of vectors c and d according to a weighting vector b . Now we can reinterpret (7) in the following way.

Proposition 7.2 (Second-Order Network Centrality). *Consider the structural model in Section 3 with full labor reallocation/constant returns to scale ($\beta_i = 1$ for every i).*

Without loss of generality, assume $\epsilon_i = \theta_i$ for every i . Then

$$\frac{d^2 \log C}{d \log A_i^2} = (\sigma - 1)Var_b(\Psi_{(i)}) + \sum_j (\theta_j - 1)\lambda_j Var_{\Omega_j}(\Psi_{(i)}), \quad (8)$$

Recall that

$$\frac{d^2 \log C}{d \log A_k^2} = \frac{d \lambda_k}{d \log A_k} = \frac{\lambda_k}{\xi} \sum_{j \neq k} \lambda_j \left(1 - \frac{1}{\rho_{kj}}\right) + \lambda_k \frac{d \log \xi}{d \log A_k},$$

with

$$1 - \frac{1}{\rho_{ij}} = \frac{1}{\lambda_k} \frac{d \lambda_k}{d \log A_k} - \frac{1}{\lambda_j} \frac{d \lambda_j}{d \log A_k},$$

and

$$\frac{d \log \xi}{d \log A_i} = \frac{1}{\xi} \sum_i \frac{d \lambda_k}{d \log A_i}.$$

Finally, using these relationships and the following fact

$$\lambda_m \frac{d \log \lambda_m}{d \log A_i} = (\sigma - 1)Cov_b(\Psi_{(i)}, \Psi_{(m)}) + \sum_j (\theta_j - 1)\lambda_j Cov_{\Omega_j}(\Psi_{(i)}, \Psi_{(m)}),$$

we can deduce the full set of reduced-form elasticities. This means that the macro elasticities are effectively a weighted average of the underlying structural elasticities of substitution, where the weighting does not need to sum to 1 and the weights depend on the network structure. In the special case where all industries have the same size λ , the macro elasticities take an especially simple form:

$$1 - \frac{1}{\rho_{mi}} = (\sigma - 1)Cov_b(\Psi_{(i)}, \Psi_{(m)} - \Psi_{(i)}) + \sum_j (\theta_j - 1)\lambda_j Cov_{\Omega_j}(\Psi_{(i)}, \Psi_{(m)} - \Psi_{(i)}),$$

$$\frac{d \xi}{d \log A_i} = (\sigma - 1)Cov_b(\Psi_{(i)}, \sum_m \Psi_{(m)}) + \sum_j (\theta_j - 1)\lambda_j Cov_{\Omega_j} \left(\Psi_{(i)}, \sum_m \Psi_{(m)} \right).$$

Equation (8) gives, as a special case, the constant returns irrelevance result of Section 5, since in the case where all structural elasticities of substitution are the same, the formula collapses to the one in Proposition 5.1. Equation 8 is also related to the *concentration* centrality defined by Acemoglu et al. (2016), but generalizes their result by allowing for heterogeneity in the interaction functions, non-symmetric network structures, and micro-founds its use for production networks. It has a simple intuition: the second-order impact

of a shock to i depends on how the sales share λ_i changes. This, in turn, depends on how demand for i changes— which is composed of demand from the household and demand from other industries, indexed by j . The extent to which the structural elasticity of substitution θ_j , for industry j , matters depends on how unequally j is exposed to i through its different inputs, and on how big j is. If j is small, or is exposed in the same way to i through all of its inputs, then the extent to which it can substitute amongst its inputs is irrelevant. The same holds for the household, which can substitute across consumption goods with elasticity σ . Therefore, non-unitary elasticities can be amplified by concentrated linkages.

A simple example, motivated by a universal intermediate input like electricity, helps explain some of the intuition of Proposition 8. Consider an example where

$$C = \left(\sum_i b_i c_i^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}},$$

and

$$c_i = \left(a_i l_i^{\frac{\epsilon-1}{\epsilon}} + (1 - a_i) E_i^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}}.$$

Assume that the universal intermediate input E is produced using labor with constant returns to scale. For this example, industry i 's steady-state sales share is $\lambda_i = b_i$, the intermediate input share of industry i is $1 - a_i$, and the sales share of electricity is $\lambda_E = \sum_i \lambda_i (1 - a_i)$. This economy is depicted in figure 4.




Figure 4: An illustration of the economy with a universal intermediate input which we treat as energy. Each industry has gross labor share a_i and substitutes across labor and energy with elasticity ϵ . The household can substitute across goods with elasticity of substitution σ

For this example, (8) implies that

$$\frac{d^2 \log C}{d \log A_E^2} = \lambda_E(1 - \lambda_E)(\sigma - 1) + \sum_i \lambda_i(1 - a_i)a_i(\varepsilon - \sigma). \quad (9)$$

We know from Section 5 that the first term is the second-order term if there was no network structure (or if the structural elasticities of substitution were homogeneous), the second term is a correction that takes into account the fact that $\varepsilon \neq \sigma$. We simplify this example further by supposing that all final sectors are equally sized $\lambda_i = 1/N$, and that $M \leq N$ sectors use electricity with steady-state intermediate input share $1 - a_i = 1 - a$, while $N - M$ use no electricity at all ($1 - a_i = 0$). We set a to ensure λ_E stays constant. Then (9) implies

$$\frac{d^2 \log C}{d \log A_E^2} = \lambda_E(1 - \lambda_E)(\sigma - 1) + (\varepsilon - \sigma)\lambda_E \left(1 - \frac{N}{M}\lambda_E\right).$$

For concreteness, take $\varepsilon < \sigma < 1$. Then the second-order term is negative and decreasing in M , since negative shocks to electricity have a smaller impact on output if electricity is not an input into everything, and therefore, the household can substitute to consumption goods that do not rely on electricity.

When every sector uses electricity $M = N$, this simplifies further to

$$\frac{d^2 \log C}{d \log A_E^2} = \lambda_E(1 - \lambda_E)(\varepsilon - 1),$$

and the elasticity of substitution in consumption σ drops out completely. Hence, even if σ is much greater than ε , this makes no difference to the second-order term. However, severing just one link between electricity and final goods can significantly increase the second-order term and even flip its sign, as long as consumption is sufficiently more substitutable than production. The fact that σ is irrelevant when $M = N$ is a manifestation of the general principle stated in proposition 7.2. Since the household is symmetrically exposed to shocks from the electricity industry, it does not matter how well the household can substitute amongst its own inputs. In Appendix D, we show that this result holds much more generally. We also show how heterogeneity in intermediate input shares, as well as decreasing returns to scale in energy production, affects these results.

Another noteworthy special case of Proposition 7.2 is for nested CES production functions which appear frequently in various literatures, even those not concerned with the role of input-output relationships (e.g. see Oberfield and Raval, 2014). In this case the

macro elasticities take an especially simple form where they are a weighted average of the micro elasticities of substitution within and across the different nests, where the weights depend on the relative expenditure shares. The details for this special case can be found in Appendix F.

More intuition about (8) can be gleaned by focusing on regular economies. We call an economy *regular* if its input-output matrix satisfies $\Omega \mathbf{1} = c_1 \mathbf{1}$ and $\Omega' \mathbf{1} = c_2 \mathbf{1}$ for some scalars c_1 and c_2 , and the household's consumption expenditures are $\frac{1}{N} \mathbf{1}$. Intuitively, a regular economy is symmetric. If an economy is regular, as a matter of accounting, it must be the case that $c_1 = c_2 = 1 - a$, where a is the value-added share of gross output. Three leading examples of regular economies are the ring, the completely connected, and the completely disconnected examples in Figure 5.




Figure 5: Three examples of regular economies: a ring, a complete network, islands.

Consider the “island” economy where each industry only buys inputs from itself $\Omega = (1 - a)I$, the complete economy where each industry buys from every other industry in equal proportions $\Omega = (1 - a)/N \mathbf{1} \mathbf{1}'$, and a ring economy. Suppose that all structural elasticities of substitution in production and consumption are the same $\sigma = \theta_i = \epsilon_i$, and labor is fully mobile $\beta_i = 1$. Then, for the ring economy

$$\frac{d^2 \log C}{d \log A_i^2} = (\sigma - 1) \frac{1}{Na} \left(\frac{2}{1 - (1 - a)^N} - 1 - \frac{1}{Na} \right),$$

for the island economy

$$\frac{d^2 \log C}{d \log A_i^2} = (\sigma - 1) \frac{1}{Na} \left(\frac{2}{a} - 1 - \frac{1}{Na} \right),$$

and for the complete network

$$\frac{d^2 \log C}{d \log A_i^2} = (\sigma - 1) \frac{1}{Na} \left(\frac{1}{aN} + 1 - \frac{2}{N} \right).$$

These economies are identical up to a first-order approximation, but their second-order properties can be very different. Especially when $a \approx 0$ and $N \neq 1$, the island economy can be vastly more nonlinear than the complete economy when judged by the size of the second-order term.

Figure 6 plots the degradation in the average performance of these economies as function of the labor share a . For a fixed intermediate input share, the impact of a microeconomic shock is much greater in the island economy than in the complete and ring economies. This difference is especially great for low values of a as long as $N \neq 1$. This is driven by the strength of *diversification* in each of these economies, and equation (8) can provide intuition for this fact: since $\sigma < 1$, for a fixed intermediate input share a , more concentrated linkages (higher variance $\Psi_{(i)}$'s) degrade the average performance of output. The island economy, with the least diversified links, is the worst performing, whilst the complete economy, with complete diversification, performs the best. When a negative shock hits a sector, the impact of the shock diffuses over all sectors in the complete economy but is concentrated on the affected sector in the island economy. Hence, output drops by more in the latter case due complementarities in consumption and production. The ring constitutes an intermediate case with less diversification than the complete economy but more diversification than the island economy.¹⁹

On the other hand, as we reduce a , we see that all of these economies perform worse. In the language of Elliott et al. (2014), we can think of this as varying *integration* rather than diversification, since the input-output matrix Ω scales with $1 - a$ increasing the intensity of network connections uniformly. As a approaches zero, the intermediate input share approaches unity, and the elements of Leontief inverse $\Psi_{(i)}$ become larger and larger, as does their variance. This magnifies the size of the terms in (8).²⁰ Proposition 7.2 allows

¹⁹This analysis allows us to connect with one of the most surprising, and well-known, results in the literature on production networks due to Dupor (1999) who showed that, for a Cobb-Douglas economy, all regular economies with the same a have the same output function $C(A_1, \dots, A_n)$. In this sense, there is nothing more unstable about rings versus complete graphs. Since $\lambda_i = \frac{1}{Na}$ in steady state, Hulten's theorem can be used to extend this intuition to non Cobb-Douglas economies up to a first order. However, the examples in this section show that this result is only true up to a first-order approximation, and this first-order approximation can, in practice, be very poor.

²⁰Our results show that the issues of integration and diversification, whose importance for volatility has been recognized by Koren and Tenreyro (2013) and Kurz and Senses (2016), also affect average performance

us to easily characterize the best-performing and worst-performing networks as measured by their certainty-equivalents when we vary the gross labor share.




Figure 6: $E(\log(C)/\bar{C})$ as a function of the gross labor share a . For this example, the shocks are idiosyncratic with mean zero and standard deviation one (in logs), the structural elasticity of substitution $\sigma = 1/2$, and the number of industries $N = 3$. If $\sigma > 1$, the sign of this plot would flip, but the relative magnitudes would be preserved.

Proposition 7.3 (Ranking Networks). *Consider different versions of the structural model in Section 3 with full labor reallocation/constant returns to scale ($\beta_i = 1$ for every i). Without loss of generality, assume $\epsilon_i = \theta_i$ for every i . We rank networks by the second-order approximation to $E(\log(C/\bar{C}))$ for i.i.d shocks.*

1. *When all micro elasticities are below one ($\sigma < 1$ and $\theta_j < 1$ for every j), the degenerate network ($a_i = 1$ for every i) performs the best. The performance of a network can be arbitrarily bad in the limit of a zero gross labor share ($a_k \rightarrow 0$ for every k).*
2. *When all micro elasticities are above one ($\sigma > 1$ and $\theta_j > 1$ for every j), the degenerate network ($a_i = 1$ for every i) performs the worst. The performance of a network can be arbitrarily good in the limit of a zero gross labor share ($a_k \rightarrow 0$ for every k).*

when the output function is not loglinear.

Intuitively, when the structural elasticities are less than one, we are in a world with complementarities, and so performance is improved if there are no intermediate inputs at all, minimizing the magnitude of the elements of the Leontief inverse $\Psi_{(i)}$'s in (8). On the other hand, performance can be made arbitrarily bad by increasing the intermediate input share, driving the elements of the Leontief inverse to infinity.

A final implication of equation (8) is that even with variable elasticities, in this context, only the industry's role as a supplier matters, not its role as a consumer. This generalizes Proposition 4.1 in Baqaee (2016), showing that even with different CES parameters, the output elasticity of a TFP shock to k only depends on k 's role as a supplier.

Proposition 7.4 (Direction of Diffusion). *Consider the structural model in Section 3 with full labor reallocation/constant returns to scale ($\beta_i = 1$ for every i). Consider two industries k and l that sell the same share to all other industries and the household ($\omega_{ik} = \omega_{il}$ for each i and $b_k = b_l$). Then*

$$\frac{d \log C}{d \log A_k} = \frac{d \log C}{d \log A_l},$$

and

$$\frac{d^2 \log C}{d \log A_k^2} = \frac{d^2 \log C}{d \log A_l^2}.$$

Proof. Denote the i th standard basis vector by e_i . Then, by assumption, $\Omega e_k = \Omega e_l$. Repeated multiplication implies that $\Omega^n e_k = \Omega^n e_l$. This then implies that $\Psi e_k = \Psi e_l$. In steady state, $\lambda_k = b' \Psi e_k = b' \Psi e_l = \lambda_l$. So the first-order impact of a shock is the same. Furthermore, substitution into (8) shows that the second-order impact of a shock is also the same. ■

This is an implication of the full labor reallocation/constant-returns-to-scale assumption at the micro level, and this result would break down otherwise. The intuition is that output then depends only on the prices of consumption goods. Furthermore a change in the size of the i th industry will not affect its price. Hence, a productivity shock will travel downstream from suppliers to their consumers, by lowering their marginal costs, but it will not travel upstream from consumers to their suppliers, since the supplier's price does not depend on its size. The general model of Section 3 does not satisfy this property since it allows for imperfect labor reallocation/decreasing returns to scale. In Appendix E, we work through an explicit example to show how decreasing returns to scale will break this

result. Other ways to break this result are explored by Baqaee (2016), who studies a CES network with scale economies and fixed costs of operating.

8 Quantitative Illustration

In this section, we perform some illustrative quantitative simulations to gauge whether or not the nonlinearities we have identified are likely to be important in the data. We perform two exercises. First, we calibrate a multi-sector model to match input-output data and use the best available information to calibrate the structural elasticities of substitution. We shock this model and compare its performance relative to the first-order approximation. This first exercise necessarily imposes an unrealistic degree of homogeneity across the structural elasticities of substitution due to a lack of information.²¹ In the second exercise, we study the macroeconomic impact of the energy crisis of the 1970s using a non-parametric generalization of Hulten (1978) that takes the second-order terms into account. Both exercises suggest that production is highly nonlinear.

8.1 A Quantitative Structural Model

In this section, we quantitatively explore the importance of the nonlinearities in production that we have emphasized. To do this, we calibrate a simplified version of the structural model in Section 3. To calibrate the model, we need estimates for the industry-specific structural elasticities of substitution. Unfortunately, disaggregated estimates of these elasticities do not exist. We rely instead on estimates from Atalay (2016) and Comin et al. (2015) who estimate a small number of structural elasticities. These are the elasticities of substitution between value-added and intermediate inputs, amongst intermediate inputs, and among consumption goods. For our model, this imposes $\theta_i = \theta$, and $\varepsilon_i = \varepsilon$. We set $\theta = 0.3$, $\sigma = 0.4$, and $\varepsilon = 0.0001$.

Our values of θ and σ are on the lower end of the values estimated by Atalay (2016) and Comin et al. (2015). We justify this by the fact that we use more disaggregated data, and for more disaggregated data, the elasticities of substitution are smaller for reasons emphasized by Oberfield and Raval (2014). Indeed, firm-level estimates like Boehm et al.

²¹In Appendix D we use a more stylized example but zoom in on how heterogeneous elasticities of substitution can give an outsized importance to an industry even though it is small – a situation that is impossible for a first-order approximation.

(2015) suggest that the firm-level elasticity of substitution between intermediate inputs and value-added θ is zero.

We verify that the benchmark model with these values of $(\sigma, \varepsilon, \theta)$ matches the volatility of observed industry-level sales shares. We target $\sum_i \bar{\lambda}_i \sigma_{\lambda_i} = 0.0197$, where $\bar{\lambda}_i$ is the time-series average and σ_{λ_i} is the time-series standard deviation of industry i 's Domar weight. A Cobb-Douglas model would imply that this should always be zero, since the Domar weights would be constant. We also consider two robustness cases, one with more substitutability $(\sigma, \theta) = (0.6, 0.8)$ and one with less substitutability $(\sigma, \theta) = (0.3, 0.1)$ than the benchmark model. As expected, the volatility of the Domar weights approaches zero as the model approaches the Cobb-Douglas limit. Our benchmark model matches $\sum_i \bar{\lambda}_i \sigma_{\lambda_i} = 0.0197$, but the more substitutable economy undershoots with 0.0110 and the less substitutable economy overshoots with 0.0413.

We work with the 88 sector US KLEMS annual input-output data from Dale Jorgenson and his collaborators, dropping the government sectors. The dataset contains sectoral output and inputs from 1960 to 2005. The advantage of this dataset over the more detailed input-output table from the BEA is that it contains both price and quantity data, which allows for the construction of sectoral gross TFP at annual frequency. We use the sector-level TFP series computed by Carvalho and Gabaix (2013) using the methodology of Jorgenson et al. (1987).²²

We calibrate the expenditure share parameters to match the input-output table, using 1982 (the middle of the sample) as the base year for the calibration. For our benchmark results, we set sectoral TFP $\log A_i \sim \mathcal{N}(-\Sigma_{ii}/2, \Sigma_{ii})$, where Σ_{ii} is the sample variance of $\Delta \log TFP$ for the industry i . We work with uncorrelated sectoral shocks since the average correlation between sectoral growth rates is extremely small (less than 5%). Our results are not significantly affected if we matched the whole covariance matrix of sectoral TFP instead.

Table 1 displays the mean, standard deviation, and skewness of log GDP for various specifications. For comparison, the table also shows these moments for GDP growth and aggregate TFP growth.²³

Our benchmark model, without reallocation, assumes that the labor market for each

²²As shown by Diewert (1976), TFP accounting with the Tornqvist index is equivalent to assuming that the production function is a translog function of inputs, which is consistent with our emphasis on accounting for the second-order impact of shocks. For a review of these issues, see also Diewert and Nakamura (1993).

²³Since our model has inelastic factor supply, its output is more comparable to aggregate TFP than GDP. As shown by Gabaix (2011) and Carvalho and Gabaix (2013), elastic capital and labor supply would further amplify TFP shocks.

sector is completely segmented so that no workers can be reallocated across industries in response to shocks. The benchmark model has three important ingredients that we have emphasized in the paper: (1) non-unitary structural elasticities of substitution, (2) labor market segmentation, (3) network structures. In table 1, we not only show the macro moments for our benchmark model, we also show these moments for every combination of these three ingredients. Since the model is nonlinear, these ingredients interact with one another, and therefore, this is as far as we can go in providing a decomposition of the importance of the various channels. The model with full reallocation has (1) and (3) but not (2). The log-linear approximation has (3) and (2) but not (1).²⁴ The “no network, no reallocation” model has (1) and (2) but not (3). The “no network, full reallocation” model has (1) but not (2) or (3). To emphasize the importance of these nonlinearities in production, we also display the results for a linear (as opposed to log-linear) approximation. Finally, we also include results for more volatile shocks, which we discuss in detail later. For now, let’s consider each moment of in turn.

We start with the mean. In the benchmark model, the mean is -0.0057 log points. This means that the welfare cost of business cycles are 0.57% of output. These costs are entirely due to nonlinearities in production. They are an order of magnitude larger than the welfare gains of around 0.05% of output arising from risk aversion in consumption estimated by Lucas (1987).²⁵ In the model with full reallocation the mean is -0.0026 , while for the model with no network but labor market segmentation it is -0.0014 . These numbers are significantly smaller than the benchmark model but they are still non-trivial. Hence, all three ingredients are important. The model with “no network and full reallocation” yields an almost zero mean. This confirms our finding from Section 4 that low structural elasticities of substitution, without either labor market segmentation or intermediate inputs, cannot generate substantial deviations from linearity. Lastly, under the log-linear approximation the mean is -0.0010 , and under the linear approximation it is exactly 0 – this is due to the fact that the TFP shocks have a negative mean in logs but not in levels.

Now, consider the standard deviation. The benchmark has a standard deviation of 0.0117 , which is a slight amplification relative to the log-linear model. Overall, the standard deviation is fairly constant across specifications, which is intuitive since larger second-

²⁴For this case, labor market segmentation is irrelevant.

²⁵As shown by proposition 2.5, looking at log consumption rather than consumption does not alter these results by much. We manually computed the same numbers for GDP instead of log GDP and the results are very similar.

	Mean	Standard Deviation	Skewness
GDP Data	–	0.0238	-0.6190
TFP Data	–	0.0147	-0.2888
Benchmark	-0.0057	0.0117	-0.5229
Full reallocation	-0.0026	0.0110	-0.0745
Log Linear Hulten	-0.0010	0.0110	0.0000
Linear Hulten	0.0000	0.0110	0.0432
No Network, no reallocation	-0.0014	0.0053	-0.0420
No Network, full reallocation	0.0000	0.0053	0.0301
$(\theta, \sigma) = (0.1, 0.3)$	-0.0102	0.0138	-1.2864
$(\theta, \sigma) = (0.6, 0.8)$	-0.0035	0.0112	-0.1648
High Volatility Benchmark	-0.0117	0.0180	-0.8821
High Volatility Hulten	-0.0015	0.0155	0.0000

Table 1: Simulated and estimated moments. For GDP and aggregate TFP, we use the demeaned growth rates. For the model, we use the sample moments of log GDP. The simulated moments are calculated from 50,000 draws.

order terms will magnify some shocks but attenuate other shocks, leaving the variance relatively stable. The only case where the standard deviation is substantially different is the case with no intermediate inputs. In this case, the input-output multiplier ξ is counterfactually equal to one. The lack of an input-output multiplier means that the model generates less variance than the log-linear model, which has $\xi \approx 2$.

Skewness fits the same pattern as the results on the mean. The benchmark model generates strong skewness in output, which is substantially mitigated if we remove either the network or allow reallocation of workers. Reducing substitutability significantly decreases the mean and increases the negative skewness of output. As expected, the log-linear, linear, and almost log-linear models generate no skewness or a slight positive skew (since the lognormal distribution is positively skewed).

For comparison, we also compute the same moments for the second-order approximation of the model in logs. The second-order approximation performs well in approximating the mean and standard deviation of the underlying structural model, but does less well on higher moments like skewness and kurtosis (though these are in the right direction). For the benchmark model without reallocation, the second-order approximation to log GDP has mean of -0.0056 , with a standard deviation of 0.0113 , skewness -0.3679 . The moments for other parameterizations are shown in Appendix G.

In figure 7 we plot the histograms for the model with no reallocation for different val-

ues of ε and θ . The model exhibits significant negative skewness when the elasticity of substitution is low, but it also has excess kurtosis or fat tails. For instance, the benchmark model has an excess kurtosis of 0.96, meaning that it is subject to occasional and endogenous large negative fluctuations (rare disasters). These features become significantly more pronounced as θ and σ are lowered, to the point where the model with $\theta = 0.1$ and $\sigma = 0.3$ has excess kurtosis of 3.47. We do not formally tabulate our results for kurtosis since our sample size of 45 observations is far too small to be able to estimate kurtosis with any degree of confidence in the data.




Figure 7: The distribution of GDP for the benchmark calibration, as well as two robustness cases.

Unlike Acemoglu et al. (2017) or Barro (2006), to achieve rare disasters, we do not need to assume fat-tailed exogenous shocks nor rule out “rare bonanzas” a priori, instead these features are endogenously generated by the nonlinearities in the model. This can be seen in figure 8, where we plot the histograms for the benchmark model and for a log-linear approximation subject to the same shocks. To drive this point home, we also simulate and plot the benchmark model with more volatile shocks. We double the variance of the sectoral shocks, which roughly corresponds to the average variance of sectoral TFP shocks during the 1970s and before the Great Moderation. For the log-linear model, increasing the variance of lognormal shocks has no effect on the skewness or kurtosis of GDP, since

these standardized moments are scale invariant. However, for the benchmark model, more volatile shocks significantly increase the deviations from normality and dramatically decrease the economy's average performance to -0.012 . Excess kurtosis increases to 2.3 for the benchmark model when shocks are more volatile. These results suggest that complementarities are likely to be significantly more costly in eras or countries where volatility is high.

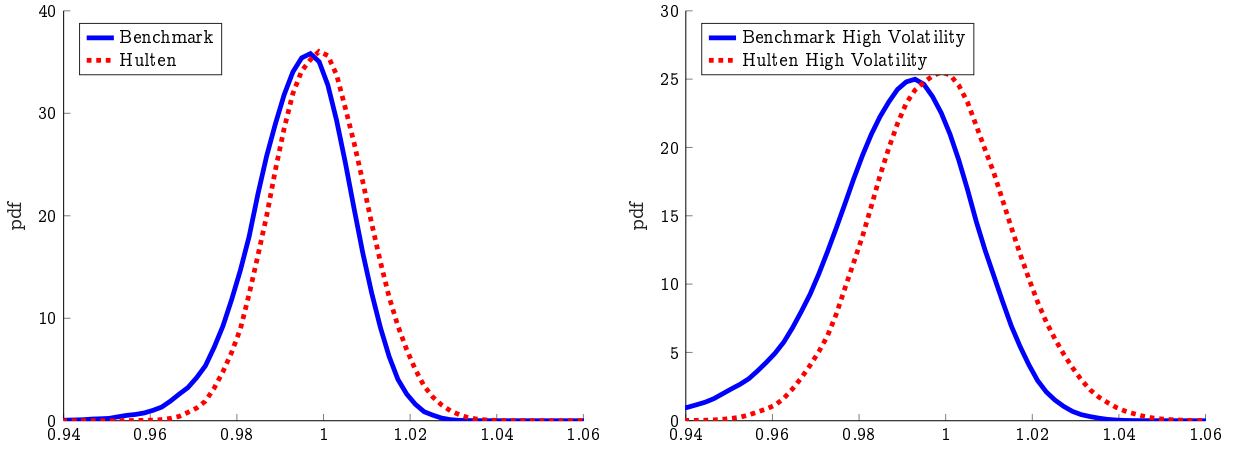


Figure 8: The left panel shows the distribution of GDP for the benchmark model and log-linearized model. The right panel shows these for shocks whose variance is twice as high.

Finally, we consider the response of GDP to shocks to specific industries. It turns out that for a large negative shock, the “oil and gas” industry produces the largest negative response in GDP – this despite the fact that the oil and gas industry is not the largest industry in the economy. Figure 9 plots the response of GDP for shocks to the oil and gas industry as well as for the “retail trade (excluding automobiles)” industry. The retail trade industry has a similar sales share, and therefore, to a first order, both industries are equally important. As expected, the nonlinear model is significantly more fragile to both kinds of shocks (negative shocks are amplified and positive ones are attenuated). However, output is more sensitive to oil and gas for large negative shocks. On the other hand, output is more sensitive to the retail industry for positive shocks. The strong asymmetry is consistent with the empirical findings of Hamilton (2003) that oil price increases are much more important than oil price decreases.

Figure 9 gives an impression that the relative ranking of industries is stable as a function of the size of the shock. The oil industry is always more important than the retail trade industry for negative shocks, and always less important for positive shocks. However,




Figure 9: The effect of TFP shocks to the oil and gas industry and the retail trade industry. Both industries have roughly the same sales share, and so they are equally important up to a first-order approximation (dotted line). The nonlinear model is more fragile to both shocks than the log-linear approximation. The oil and gas industry is significantly more important than retail trade for large negative shocks. The histogram is the empirical distribution of sectoral TFP shocks pooled over the whole sample.

this need to be the case. In figure 10 we plot GDP as a function of shocks to the oil industry and the construction industry. Unlike retail trade, the construction industry is significantly larger than the oil industry. Therefore, the first-order approximation implies that it should be more important. The nonlinear model also behaves the same way for positive shocks, and small negative shocks. However, for very large negative shocks, the oil and gas industry once again becomes more important. In fact, for a negative 20% TFP shock, the oil and gas industry is *the* most important industry in the economy.

8.2 The Effect of Oil Shocks

The oil shocks of the 1970s serve as a useful example of the way industry-level shocks can have macroeconomic consequences. To recap the history, in October 1973, the Arab members of OPEC proclaimed an oil embargo limiting shipments of oil to the United States and

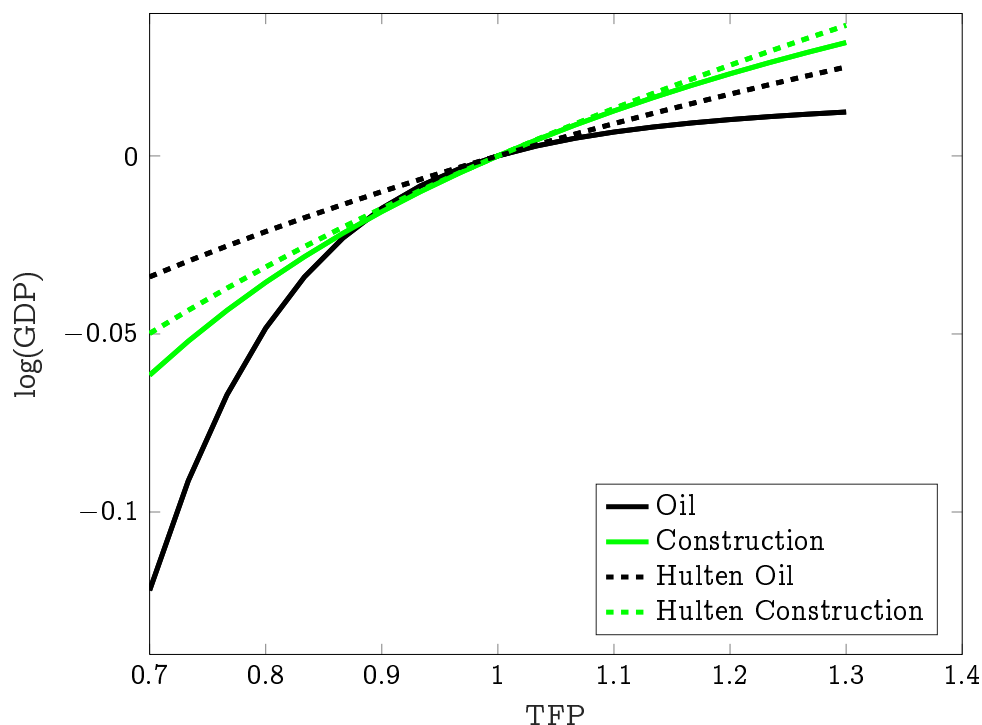


Figure 10: The effect of TFP shocks to the oil and gas industry and the construction industry. Construction has a bigger sales share, but oil and gas is more important for large negative shocks. This graph shows that the ranking of which industry is more important is not monotonic in the size of the shock.

some of its allies (including Japan and the United Kingdom). Although the oil embargo was lifted in 1974, coordinated action by OPEC kept prices elevated throughout the mid-1970s. The price of crude oil increased from \$3.5 a barrel in 1972 to \$11 a barrel in 1974. In 1979, OPEC implemented a second round of price increases which caused the price of crude to soar to \$31 a barrel. At the same time, the Iranian revolution in 1979, as well as the ensuing Iraqi invasion of Iran in 1980, caused further disruptions to global crude oil supply. The price peaked at \$37 in 1980. Starting in the early 1980s, with the departure of the Shah, OPEC's pricing structure collapsed as Saudi Arabia flooded the market with inexpensive oil. In real terms, the price of crude oil declined back to its pre-crisis levels by 1986. According to the NBER's business cycle dating committee, both oil price shocks coincided with recessions in the US. Although our structural model suggests that the "oil and natural gas" extraction industry is important, it abstracts away from trade, by assuming all intermediate inputs are sourced domestically, with net imports showing up only in final demand. Hence, the Domar weight of the oil and natural gas industry measures

domestic production, rather than domestic consumption. Since the oil price shocks did not directly affect the productivity of domestic oil production, this means that they are not measured in our sectoral TFP data (which is for domestic production). Furthermore, our industry classification is too coarse to isolate crude oil separately from other petrochemicals. Therefore, in this section, we pursue a non-parametric calibration of the impact of the oil price shocks up to second-order terms for the world as a whole, avoiding strong structural assumptions as well as the complications arising from international trade.

Proposition 8.1. *Up to the second order in the vector Δ , we have*

$$\log(C(A + \Delta)/C(A)) = \frac{1}{2} [\lambda(A + \Delta) + \lambda(A)]' (\log(A + \Delta) - \log(A)).$$

Proof. Let $C^{2nd}(A + \Delta)$ be the quadratic function of Δ that provides a second-order approximation to output at point A as a function of Δ . The quadratic approximation lemma from Diewert (1976) implies that

$$\begin{aligned} \log(C^{2nd}(A + \Delta)/C^{2nd}(A)) \\ = \frac{1}{2} [\nabla \log C^{2nd}(A + \Delta) + \nabla \log C^{2nd}(A)]' [\log(A + \Delta) - \log(A)]. \end{aligned} \quad (10)$$

Hulten (1978) then implies that $\nabla \log C(A) = \lambda(A)$ and $\nabla \log C(A + \Delta) = \lambda(A + \Delta)$. The result follows since $\nabla \log C^{2nd}(A) = \nabla \log C(A)$, $\nabla \log C^{2nd}(A + \Delta) = \nabla \log C(A + \Delta)$ up to the first order in Δ , and $\log(C^{2nd}(A + \Delta)/C^{2nd}(A)) = \log(C(A + \Delta)/C(A))$ up to the second order in Δ . ■

The idea of averaging weights across two periods is due to Leo Törnqvist (1936). Proposition 8.1 relates the impact of the oil shocks on GDP to the size of the shock and the corresponding Domar weights before and after the shock.²⁶

²⁶One can always compute the full nonlinear impact of a shock on output by computing $\int_A^{A+\Delta} \lambda(\tilde{A}) d \log \tilde{A}$, and our formula approximates this integral by performing a first-order (log) approximation of the Domar weight $\lambda(\tilde{A})$ or equivalently a second-order (translog) approximation of GDP. In theory, if TFP is a continuous diffusion then one can disaggregate time-periods and compute the impact of shocks over a time period $[t, t+\delta]$ as $\int_t^{t+\delta} \lambda(A_s) d \log(A_s)$ which can be seen as a repeated application of Hulten's theorem at every point in time over infinitesimal intervals of time. However, when TFP has jumps, then this decomposition no longer applies. In any case, even when it does apply, and when the required high-frequency data regarding TFP shocks and Domar weights is available, it can only be useful ex post to assess the changes in GDP over an elapsed period of time due to the TFP shocks $d \log(A_s)$ to a given sector given the observed path of Domar weight $\lambda(A_s)$. It is of no use ex ante to predict how these future shocks will affect GDP because one would need to know how the Domar weight will change over time as a result of the shocks, and hence of no use to

We measure the price of oil using the West Texas Intermediate Spot Crude Oil price from the Federal Reserve Database. Global crude oil production, measured in thousand tonne of oil equivalents, is from the OECD. World GDP, in current USD, is from the World Bank national accounts data. The choice of the pre and post Domar weight is not especially controversial. Crude oil, as a fraction of world GDP, increased from 5% in 1972 to 31% in 1980. Reassuringly, the Domar weight is back down to its pre-crisis level by 1986 (see figure 11). This means that, taking the second-order terms into account, we need to weight the shock to the oil industry by $1/2(5 + 31) = 18\%$. Hence, the second-order terms amplify the shock by a factor of $18/5 \approx 3.6$.




Figure 11: Global expenditures on crude oil as a fraction of world GDP.

Calibrating the size of the shock to the oil industry is more tricky, since it's not directly observed. If we assume that oil is an endowment, then we can measure the shock simply via changes in the physical quantity of production. To do this, we demean the log growth rate in global crude oil production, and take the shock to be the cumulative change in demeaned growth rates from 1973 to 1981, which gives us a shock of -13% .

Putting this altogether, the first-order impact on GDP is therefore

$$0.05 \times -0.13 = -0.0078.$$

run counterfactuals. This latter part is precisely what the second-order approximation at the heart of our paper accomplishes.

On the other hand, the second-order impact on GDP is

$$\frac{1}{2}(0.05 + 0.31) \times -0.13 = -0.0244.$$

Hence, accounting for the second-order terms amplifies the impact of the oil shocks significantly, so that oil shocks can be macroeconomically significant even without any financial or demand side frictions.²⁷

9 Conclusion

The paper points to many unanswered questions. For instance, it shows that the macroeconomic impact of a microeconomic shock depends greatly on how quickly factors can be reallocated across production units. Since our structural model is static, we are forced to proxy for the temporal dimension of reallocation by resorting to successive comparative statics. In ongoing work, we investigate the dynamic adjustment process more rigorously and find that although we can think of the no-reallocation and perfect-reallocation cases as the beginning and end of the adjustment, the speed of adjustment also greatly depends on the microeconomic details. This means that the dynamic response of output to different shocks is greatly affected by issues like geographic or sectoral mobility of labor, even with perfect and complete markets that allow us to abstract from distributional issues. Our model also lacks capital accumulation and reallocation, and incorporating these into the present analysis is an interesting area for future work.

Similarly, although our structural model assumes away geography, we speculate that our results are likely to have implications for work on the aggregate gains from trade, as well as on how, even with complete markets, the speed and extent of readjustment can play an important role in how the economy as a whole responds to changes in international and regional trade policy.

Finally, although our results assume away market frictions beyond impediments to reallocation, the forces we identify are unlikely to disappear in richer models with inefficient

²⁷As noted by Hamilton (2013), first-order approximations of efficient models assign a relatively small impact to oil price shocks. Hence, the literature has tended to focus on various frictions that may account for the strong statistical relationship between oil shocks and output. Our calculations suggests that nonlinearities in production, even in an efficient model, may help to explain the outsized effect of oil shocks. Furthermore, our calculation also makes no allowance for amplification of shocks through endogenous labor supply and capital accumulation, which are the standard channels for amplification of shocks in the business cycle literature. Hence, coupled with the standard amplification mechanisms of those models, we would expect the reduction in aggregate output to be even larger.

equilibria. As suggested by the results of Jones (2011) and Baqaee (2016), networks and macro-substitutability can amplify or attenuate the underlying frictions. A systematic characterization of these effects seem to us to be valuable areas for further work.

References

- Acemoglu, D., D. Autor, D. Dorn, G. H. Hanson, and B. Price (2016). Import competition and the great US employment sag of the 2000s. *Journal of Labor Economics* 34 (S1), S141–S198.
- Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012). The network origins of aggregate fluctuations. *Econometrica* 80(5), 1977–2016.
- Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi (2015). Systemic risk and stability in financial networks. *The American Economic Review* 105(2), 564–608.
- Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi (2016). Networks, shocks, and systemic risk. In *The Oxford Handbook of the Economics of Networks*.
- Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi (2017). Microeconomic origins of macroeconomic tail risks. *The American Economic Review* 107(1), 54–108.
- Aghion, P. and P. W. Howitt (2008). *The economics of growth*. MIT press.
- Allcott, H., A. Collard-Wexler, and S. D. O’Connell (2016). How do electricity shortages affect industry? evidence from india. *The American Economic Review* 106(3), 587–624.
- Altinoglu, L. (2016). The origins of aggregate fluctuations in a credit network economy.
- Atalay, E. (2016). How important are sectoral shocks? Technical report.
- Autor, D. H., D. Dorn, and G. H. Hanson (2016). The china shock: Learning from labor-market adjustment to large changes in trade. *Annual Review of Economics* 8, 205–240.
- Ballester, C., A. Calvó-Armengol, and Y. Zenou (2006). Who’s who in networks. wanted: the key player. *Econometrica* 74(5), 1403–1417.
- Baqaee, D. R. (2015). Targeted fiscal policy.

- Baqae, D. R. (2016). Cascading failures in production networks.
- Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, 823–866.
- Barrot, J.-N. and J. Sauvagnat (2016). Input specificity and the propagation of idiosyncratic shocks in production networks. *The Quarterly Journal of Economics*.
- Bartelme, D. and Y. Gorodnichenko (2015). Linkages and economic development. Technical report, National Bureau of Economic Research.
- Beraja, M., E. Hurst, and J. Ospina (2016). The aggregate implications of regional business cycles. Technical report, National Bureau of Economic Research.
- Bigio, S. and J. La'O (2016). Financial frictions in production networks. Technical report.
- Blackorby, C. and R. R. Russell (1989). Will the real elasticity of substitution please stand up?(a comparison of the Allen/Uzawa and Morishima elasticities). *The American economic review* 79(4), 882–888.
- Boehm, C., A. Flaaen, and N. Pandala-Nayar (2015). Input linkages and the transmission of shocks: Firm-level evidence from the 2011 tōhoku earthquake. *US Census Bureau Center for Economic Studies Paper No. CES-WP-15-28*.
- Burstein, A. and J. Cravino (2015). Measured aggregate gains from international trade. *American Economic Journal: Macroeconomics* 7(2), 181–218.
- Carvalho, V. and X. Gabaix (2013). The great diversification and its undoing. *The American Economic Review* 103(5), 1697–1727.
- Carvalho, V. M. (2010). Aggregate fluctuations and the network structure of intersectoral trade.
- Comin, D. A., D. Lashkari, and M. Mestieri (2015). Structural change with long-run income and price effects. Technical report, National Bureau of Economic Research.
- Costinot, A. and A. Rodriguez-Clare (2014). Trade theory with numbers: quantifying the consequences of globalization. *Handbook of International Economics* 4, 197.
- Di Giovanni, J., A. A. Levchenko, and I. Méjean (2014). Firms, destinations, and aggregate fluctuations. *Econometrica* 82(4), 1303–1340.

- Diewert, W. E. (1976). Exact and superlative index numbers. *Journal of econometrics* 4(2), 115–145.
- Diewert, W. E. and A. Nakamura (1993). *Essays in index number theory*, Volume 217. Emerald Group Publishing.
- Domar, E. D. (1961). On the measurement of technological change. *The Economic Journal* 71(284), 709–729.
- Dupor, B. (1999). Aggregation and irrelevance in multi-sector models. *Journal of Monetary Economics* 43(2), 391–409.
- Durlauf, S. N. (1993). Nonergodic economic growth. *The Review of Economic Studies* 60(2), 349–366.
- Elliott, M., B. Golub, M. O. Jackson, et al. (2014). Financial networks and contagion. *American Economic Review* 104(10), 3115–53.
- Feenstra, R. C. and G. H. Hanson (1996). Globalization, outsourcing, and wage inequality. *The American Economic Review* 86(2), 240.
- Foerster, A. T., P.-D. G. Sarte, and M. W. Watson (2011). Sectoral versus aggregate shocks: A structural factor analysis of industrial production. *Journal of Political Economy* 119(1), 1–38.
- Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica* 79(3), 733–772.
- Gabaix, X. (2016). Power laws in economics: An introduction. *Journal of Economic Perspectives* 30(1), 185–206.
- Gomme, P. and P. Rupert (2007). Theory, measurement and calibration of macroeconomic models. *Journal of Monetary Economics* 54(2), 460–497.
- Grassi, B. (2017). IO in I-O: Competition and volatility in input-output networks. Technical report.
- Grossman, G. M., E. Helpman, E. Oberfield, and T. Sampson (2016). Balanced growth despite Uzawa. Technical report, National Bureau of Economic Research.
- Hamilton, J. D. (2003). What is an oil shock? *Journal of econometrics* 113(2), 363–398.

- Hamilton, J. D. (2013). Historical oil shocks. In R. E. Parker and R. M. Whaples (Eds.), *Routledge Handbook of Major Events in Economic History*, Chapter 21, pp. 239–265. Routledge.
- Hicks, J. (1932). *Theory of wages*. Springer.
- Horvath, M. (1998). Cyclicalities and sectoral linkages: Aggregate fluctuations from independent sectoral shocks. *Review of Economic Dynamics* 1(4), 781–808.
- Horvath, M. (2000). Sectoral shocks and aggregate fluctuations. *Journal of Monetary Economics* 45(1), 69–106.
- Houthakker, H. S. (1955). The pareto distribution and the Cobb-Douglas production function in activity analysis. *The Review of Economic Studies* 23(1), 27–31.
- Hulten, C. R. (1978). Growth accounting with intermediate inputs. *The Review of Economic Studies*, 511–518.
- Hulten, C. R. (2001). Total factor productivity: a short biography. In *New developments in productivity analysis*, pp. 1–54. University of Chicago Press.
- Jones, C. I. (2011). Intermediate goods and weak links in the theory of economic development. *American Economic Journal: Macroeconomics*, 1–28.
- Jones, C. I. (2013). Input-Output economics. In *Advances in Economics and Econometrics: Tenth World Congress*, Volume 2, pp. 419. Cambridge University Press.
- Jorgenson, D. W., F. Gollop, and B. M. Fraumeni (1987). Productivity and US economic growth.
- Jovanovic, B. (1987). Micro shocks and aggregate risk. *The Quarterly Journal of Economics*, 395–409.
- Kim, S.-J., H. S. Shin, et al. (2013). Working capital, trade and macro fluctuations. Technical report.
- Koren, M. and S. Tenreyro (2013). Technological diversification. *The American Economic Review* 103(1), 378–414.
- Kremer, M. (1993). The O-ring theory of economic development. *The Quarterly Journal of Economics*, 551–575.

- Kurz, C. and M. Z. Senses (2016). Importing, exporting, and firm-level employment volatility. *Journal of International Economics* 98, 160–175.
- Liu, E. (2017). Industrial policies and economic development. Technical report.
- Long, J. B. and C. I. Plosser (1983). Real business cycles. *The Journal of Political Economy*, 39–69.
- Lucas, R. E. (1977). Understanding business cycles. In *Carnegie-Rochester conference series on public policy*, Volume 5, pp. 7–29. North-Holland.
- Lucas, R. E. (1987). *Models of business cycles*, Volume 26. Basil Blackwell Oxford.
- Milgrom, P. and J. Roberts (1996). The LeChatelier principle. *The American Economic Review*, 173–179.
- Morishima, M. (1967). A few suggestions on the theory of elasticity. *Keizai Hyoron (Economic Review)* 16, 144–150.
- Notowidigdo, M. J. (2011). The incidence of local labor demand shocks. Technical report, National Bureau of Economic Research.
- Oberfeld, E. and D. Raval (2014). Micro data and macro technology. Technical report, National Bureau of Economic Research.
- Samuelson, P. A. (1960). An extension of the LeChatelier principle. *Econometrica*, 368–379.
- Scheinkman, J. A. and M. Woodford (1994). Self-organized criticality and economic fluctuations. *The American Economic Review*, 417–421.
- Stella, A. (2015). Firm dynamics and the origins of aggregate fluctuations. *Journal of Economic Dynamics and Control* 55, 71–88.
- Taleb, N. N. (2013). ‘Antifragility’ as a mathematical idea. *Nature* 494 (7438), 430–430.
- Törnqvist, L. (1936). The Bank of Finland’s consumption price index. Bank of Finland Monthly Bulletin 10, Bank of Finland.

A Proofs

Proof of theorem 2.1. Since the first welfare theorem holds, the equilibrium allocation solves

$$C(A_1, \dots, A_N) = \max_{c_i, x_{ij}, l_{ij}} \mathcal{C}(c_1, \dots, c_n) + \sum_i \mu_i \left(A_i F_i((l_{ij})_j, (x_{ij})_j) - \sum_j x_{ij} - c_i \right) + \sum_i \lambda_i \left(\bar{L}_i - \sum_j l_{ji} \right),$$

where \bar{L}_i is the endowment of each labor type, and μ_i and λ_i are Lagrange multipliers. The envelope theorem then implies that

$$\frac{\partial C}{\partial A_i} = \mu_i F_i((l_{ij})_j, (x_{ij})_j) = \mu_i y_i.$$

If we can show that μ_i is equal to the price of i in the competitive equilibrium, then we are done.

To prove this note that competitive equilibrium requires that

$$\frac{\partial \mathcal{C} / \partial c_i}{\partial \mathcal{C} / \partial c_1} = \frac{\mu_i}{\mu_1} = \frac{p_i}{p_1},$$

for every i , since the household consumes a nonzero amount of every good. In other words, we have

$$\frac{\partial \mathcal{C}}{\partial c_i} = \frac{p_i}{p_1} \frac{\partial \mathcal{C}}{\partial c_1}. \quad (11)$$

Hence, using Euler's theorem on homogenous functions we can write

$$C = \sum_i \frac{\partial \mathcal{C}}{\partial c_i} c_i = \frac{\partial \mathcal{C} / \partial c_1}{p_1} \sum_i p_i c_i. \quad (12)$$

Define the expenditure function for the household to be $e(p, C)$. Since \mathcal{C} is homogenous of degree one, we can write $e(p, C) = e(p)C$. In other words, we must have that

$$\sum_i p_i c_i = e(p)C.$$

Normalize the unit cost of consumption $e(p) = 1$, so that

$$\sum_i p_i c_i = C.$$

Combine this with (12) to get

$$\frac{\partial \mathcal{C}}{\partial c_1} = p_1,$$

which can be substituted into (11) to yield $p_i = \mu_i$ for every i . ■

Proof of proposition 2.3.

$$\frac{d^2 \log C}{d \log A_j d \log A_i} = \frac{d \lambda_i}{d \log A_j}. \quad (13)$$

By definition

$$\frac{d \log \lambda_i}{d \log A_j} = \left(\frac{1}{\rho_{ji}} - 1 \right) + \frac{d \log \lambda_j}{d \log A_j}, \quad (14)$$

which simplifies to

$$\frac{d \lambda_i}{d \log A_j} = \lambda_i \left(\frac{1}{\rho_{ji}} - 1 \right) + \frac{\lambda_i}{\lambda_j} \frac{d \lambda_j}{d \log A_j}. \quad (15)$$

Now apply theorem 2.2 to the second summand to obtain the desired result. ■

Proof of proposition 2.4. The second-order approximation to $\log C$ is given by

$$\log(C/\bar{C}) = \sum_i \lambda_i d \log A_i + \sum_{ij} \frac{\partial^2 \log C}{\partial \log A_i \partial \log A_j} d \log A_i d \log A_j.$$

Take expectations of both side and substitute the result from proposition 2.3 to get the desired result. ■

Proof of proposition 2.5. We prove a slightly more general formulation with arbitrary variance covariance matrix and an arbitrary twice-differentiable utility function.

$$\begin{aligned} E(f(g(A))) &\approx E \left(f(g(\bar{A})) + f'(g(\bar{A})) \nabla g(\bar{A})(A - \bar{A}) + \frac{1}{2} f''(g(\bar{A}))(A - \bar{A})' (\nabla g(\bar{A}) \circ \nabla g(\bar{A})') (A - \bar{A}) + \right. \\ &\quad \left. \frac{1}{2} f'(g(\bar{A}))(A - \bar{A})' \nabla^2 g(\bar{A})(A - \bar{A}) \right), \\ &= f(g(\bar{A})) + \frac{1}{2} f''(g(\bar{A})) \text{tr} \left((\nabla g(\bar{A}) \circ \nabla g(\bar{A})') \Sigma \right) + \frac{1}{2} f'(g(\bar{A})) \text{tr} \left(\nabla^2 g(\bar{A}) \Sigma \right). \end{aligned}$$

Now apply Hulten's theorem to get

$$= f(g(\bar{A})) + \frac{1}{2} f''(g(\bar{A})) \sum_{k,j}^N \lambda_k \lambda_j \sigma_{jk} + \frac{1}{2} f'(g(\bar{A})) \sum_{j,k}^N \frac{\partial^2 g}{\partial A_k \partial A_j} \sigma_{jk},$$

with idiosyncratic shocks, this simplifies to

$$= f(g(\bar{A})) + \frac{1}{2} f''(g(\bar{A})) \sum_k \lambda_k^2 \sigma_k^2 + \frac{1}{2} f'(g(\bar{A})) \sum_{j,k} \frac{\partial^2 g}{\partial A_k^2} \sigma_k^2.$$

Finally, observe that $f'(g(\bar{A})) = 1$ and $f''(g(\bar{A})) = -\gamma$ for CRRA utility with risk aversion parameter γ . The second summand is the Lucas term (which equals zero when f is linear), and the third summand is our term. \blacksquare

Proof of proposition 3.1. Define q_i to be the ideal price index for intermediate inputs used by industry i . The second-order approximation is given by

$$\begin{aligned} \frac{d \lambda_i}{d \log A_k} = & b_i(1 - \sigma) \frac{d \log p_i}{d \log A_k} + \\ & \sum_j \left[\omega_{ji}(1 - a_j) \lambda_j \left\{ (\theta_j - 1) \left(\mathbf{1}(j = k) - \frac{d \log q_j}{d \log A_k} + \frac{d \log p_j}{d \log A_k} \right) \right. \right. \\ & \left. \left. + (1 - \varepsilon_j) \left(\frac{d \log p_i}{d \log A_k} - \frac{d \log q_j}{d \log A_k} \right) + \frac{d \lambda_j}{d \log A_k} \frac{1}{\lambda_j} \right\} \right], \end{aligned} \quad (16)$$

$$\frac{d \log q_i}{d \log A_k} = \sum_j \omega_{ij} \frac{d \log p_j}{d \log A_k}, \quad (17)$$

$$\frac{d \log y_i}{d \log A_k} = \frac{d \log \lambda_i}{d \log A_k} + \lambda_k - \frac{d \log p_i}{d \log A_k}, \quad (18)$$

$$0 = \sum_i \frac{a_i \beta_i \lambda_i}{\theta_i(1 - \beta_i) + \beta_i} \left[\frac{d \log y_i}{d \log A_k} + \theta_i \left[\frac{d \log p_i}{d \log A_k} - \frac{d \log w}{d \log A_k} \right] - \mathbf{1}(i = k) \right], \quad (19)$$

$$\begin{aligned} \frac{d \log y_i}{d \log A_k} = & \theta_i \left[1 - \frac{a_i}{\theta_i(1 - \beta_i) + \beta_i} - (1 - a_i) \right]^{-1} \left\{ \left[\frac{a_i}{\theta_i(1 - \beta_i) + \beta_i} + (1 - a_i) \right] \frac{d \log p_i}{d \log A_k} \right. \\ & \left. - \frac{a_i}{\theta_i(1 - \beta_i) + \beta_i} \frac{d \log w}{d \log A_k} - (1 - a_i) \frac{d \log q_i}{d \log A_k} \right\} + \mathbf{1}(i = k), \end{aligned} \quad (20)$$

To see this, note that first-order conditions for industry i are given by

$$q_i = p_i \left(\frac{y_i}{A_i \bar{y}_i} \right)^{\frac{1}{\theta_i}} (1 - a_i) \left(\frac{1}{\bar{X}_i} \right)^{\frac{\theta_i - 1}{\theta_i}} X_i^{\frac{-1}{\theta_i}}.$$

We can rearrange the intermediate input demand function as

$$X_i = \left(\frac{q_i}{p_i}\right)^{-\theta_i} (1 - a_i)^{\theta_i} \left(\frac{\bar{y}_i A_i}{\bar{X}_i}\right)^{\theta_i-1} y_i,$$

and

$$x_{ij} = \left(\frac{p_j}{q_i}\right)^{-\varepsilon_i} \omega_{ij}^{\varepsilon_i} \left(\frac{\bar{X}_i}{\bar{x}_{ij}}\right)^{\varepsilon_i-1} X_i.$$

Using the input demand functions, the household demand functions, and market clearing, we can deduce that

$$y_i = b_i^\sigma \left(\frac{\bar{C}}{\bar{c}_i}\right)^{\sigma-1} \left(\frac{p_i}{p_c}\right)^{-\sigma} C + \sum_j \left(\frac{p_i}{q_j}\right)^{-\varepsilon_j} \omega_{ji}^{\varepsilon_j} \left(\frac{\bar{X}_j}{\bar{x}_{ji}}\right)^{\varepsilon_j-1} \left(\frac{q_j}{p_j}\right)^{-\theta_j} (1 - a_j)^{\theta_j} \left(\frac{\bar{y}_j A_j}{\bar{X}_j}\right)^{\theta_j-1} y_j.$$

Multiplying both sides of this equation by $p_i/p_c C$ gives a recursive characterization of expenditure shares:

$$\lambda_i = b_i^\sigma \left(\frac{\bar{C}}{\bar{c}_i}\right)^{\sigma-1} \left(\frac{p_i}{p_c}\right)^{1-\sigma} + \sum_j \left(\frac{p_i}{q_j}\right)^{1-\varepsilon_j} \omega_{ji}^{\varepsilon_j} \left(\frac{\bar{X}_j}{\bar{x}_{ji}}\right)^{\varepsilon_j-1} \left(\frac{q_j}{p_j}\right)^{1-\theta_j} (1 - a_j)^{\theta_j} \left(\frac{\bar{y}_j A_j}{\bar{X}_j}\right)^{\theta_j-1} \lambda_j. \quad (21)$$

On the other hand, we know from cost minimization that

$$q_i = \left(\sum_j \omega_{ij}^{\varepsilon_i} \left(\frac{\bar{X}_i}{\bar{x}_{ij}}\right)^{\varepsilon_i-1} p_j^{1-\varepsilon_j} \right)^{\frac{1}{1-\varepsilon_j}}. \quad (22)$$

The first-order condition for labor for industry i is given by

$$w = p_i \left(\frac{y_i}{A_i \bar{y}_i}\right)^{\frac{1}{\theta_i}} a_i \beta_i \left(\frac{z_i}{\bar{l}_i}\right)^{\beta_i \frac{\theta_i-1}{\theta_i}} l_i^{\frac{\theta_i-1}{\theta_i} \beta_i - 1}, \quad (23)$$

We can rearrange the labor demand function as

$$l_i = \left(\frac{w}{p_i}\right)^{\frac{-\theta_i}{\theta_i(1-\beta_i)+\beta_i}} (a_i \beta_i)^{\frac{\theta_i}{\theta_i(1-\beta_i)+\beta_i}} \left(\frac{y_i}{A_i \bar{y}_i}\right)^{\frac{1}{\theta_i(1-\beta_i)+\beta_i}} \left(\frac{z_i}{\bar{l}_i}\right)^{\frac{(\theta_i-1)\beta_i}{\theta_i(1-\beta_i)+\beta_i}}.$$

Summing this across i and equating it to labor supply \bar{l} gives

$$\bar{l} = \sum_i \left(\frac{w}{p_i} \right)^{\frac{-\theta_i}{\theta_i(1-\beta_i)+\beta_i}} (a_i \beta_i)^{\frac{\theta_i}{\theta_i(1-\beta_i)+\beta_i}} \left(\frac{y_i}{A_i \bar{y}_i} \right)^{\frac{1}{\theta_i(1-\beta_i)+\beta_i}} \left(\frac{z_i}{\bar{l}_i} \right)^{\frac{(\theta_i-1)\beta_i}{\theta_i(1-\beta_i)+\beta_i}}. \quad (24)$$

Finally, substituting the input demand back into the production function of i yields a relationship between i 's price (marginal cost) and the price of its inputs as well as output, since the industry potentially exhibits diminishing returns to scale

$$\begin{aligned} \left(\frac{y_i}{A_i \bar{y}_i} \right) &= a_i \left[\left(\frac{p_i}{w} \right)^{\frac{\theta_i}{\theta_i(1-\beta_i)+\beta_i}} \left(\frac{y_i}{A_i \bar{y}_i} \right)^{\frac{1}{\theta_i(1-\beta_i)+\beta_i}} (a_i \beta_i)^{\frac{\theta_i}{\theta_i(1-\beta_i)+\beta_i}} \left(\frac{z_i}{\bar{l}_i} \right)^{\frac{\theta_i}{\theta_i(1-\beta_i)+\beta_i}} \right]^{\frac{\theta_i-1}{\theta_i}} \\ &\quad + (1 - a_i)^{\theta_i} \left(\frac{p_i}{q_i} \right)^{\theta_i-1} \left(\frac{y_i}{A_i \bar{y}_i} \right)^{\frac{\theta_i-1}{\theta_i}} \left(\frac{1}{\bar{X}_i} \right)^{\theta_i-1}. \end{aligned} \quad (25)$$

Last, we set consumption to be numeraire so that $p_c = 1$ and observe that by definition

$$y_i = \frac{\lambda_i C}{p_i}. \quad (26)$$

The set of equations (21), (22), (24), (25), and (26) collectively define the equilibrium. To derive the second-order approximation, observe that

$$\frac{d^2 \log C}{d \log A_k^2} = \frac{d \lambda_k}{d \log A_k}.$$

To find this, we must find the first-order approximation of the changes to λ_i . To this end,

linearize the system of equations that collectively define the equilibrium. This gives

$$\begin{aligned} \frac{d \lambda_i}{d \log A_k} = & b_i^\sigma \left(\frac{\bar{C}}{\bar{c}_i} \right)^{\sigma-1} (1-\sigma) \left(\frac{p_i}{p_c} \right)^{1-\sigma} \frac{d \log p_i}{d \log A_k} \\ & + \sum_j \omega_{ij}^{\varepsilon_j} (1-a_j)^{\theta_j} \left(\frac{\bar{y}_j A_j}{\bar{X}_j} \right)^{\theta_j-1} \left(\frac{p_i}{q_j} \right)^{1-\varepsilon_j} \left(\frac{\bar{X}_j}{\bar{x}_{ji}} \right)^{\varepsilon_j-1} \left(\frac{q_j}{p_j} \right)^{1-\theta_j} \lambda_j \\ & \left[(\theta_j - 1) \left(1(j=k) + \frac{d \log p_j}{d \log A_k} - \frac{d \log q_j}{d \log A_k} \right) + (1-\varepsilon_j) \left(\frac{d \log p_i}{d \log A_k} - \frac{d \log q_j}{d \log A_k} \right) + \frac{d \log \lambda_j}{d \log A_k} \right], \end{aligned} \quad (27)$$

$$\frac{d \log q_i}{d \log A_k} = \sum_j \omega_{ij}^{\varepsilon_i} \left(\frac{\bar{X}_i}{\bar{x}_{ij}} \right)^{\varepsilon_i-1} \left(\frac{p_j}{q_i} \right)^{1-\varepsilon_i} \frac{d \log p_j}{d \log A_k}, \quad (28)$$

$$\frac{d \log y_i}{d \log A_k} = \frac{d \log \lambda_i}{d \log A_k} + \lambda_k - \frac{d \log p_i}{d \log A_k}, \quad (29)$$

$$\begin{aligned} 0 = & \sum_i \frac{1}{\theta_i(1-\beta_i) + \beta_i} \left(\frac{w}{p_i} \right)^{-\frac{\theta_i}{\theta_i(1-\beta_i) + \beta_i}} \left(\frac{y_i}{A_i \bar{y}_i} \right)^{\frac{1}{\theta_i(1-\beta_i) + \beta_i}} (a_i \beta_i)^{\frac{\theta_i}{\theta_i(1-\beta_i) + \beta_i}} \left(\frac{z_i}{\bar{l}_i} \right)^{\frac{\theta_i-1}{\theta_i(1-\beta_i) + \beta_i}} \\ & \left[\frac{d \log y_i}{d \log A_k} - 1(i=k) + \theta_i \left(\frac{d \log p_i}{d \log A_k} - \frac{d \log w}{d \log A_k} \right) \right] \end{aligned} \quad (30)$$

$$\begin{aligned} \left(\frac{\theta_i - 1}{\theta_i} \right) \left(\frac{\bar{y}_i A_i}{y_i} \right)^{\frac{1-\theta_i}{\theta_i}} \left[\frac{d \log y_i}{d \log A_k} - 1(i=k) \right] = & \frac{\theta_i - 1}{\theta_i^2(1-\beta_i) + \theta_i \beta_i} a_i \left[\left(\frac{p_i}{w} \right)^{\theta_i} \left(\frac{y_i}{A_i \bar{y}_i} \right) (a_i \beta_i)^{\theta_i} \right]^{\frac{\theta_i-1}{\theta_i} \frac{1}{\theta_i(1-\beta_i) + \beta_i}} \\ & \left[\theta_i \left(\frac{d \log p_i}{d \log A_k} - \frac{d \log w}{d \log A_k} \right) + \left(\frac{d \log y_i}{d \log A_k} - 1(i=k) \right) \right] \\ & + (\theta_i - 1)(1-a_i)^{\theta_i} \left(\frac{p_i}{q_i} \right)^{\theta_i-1} \left(\frac{y_i}{A_i \bar{y}_i} \right)^{\frac{\theta_i-1}{\theta_i}} \left(\frac{1}{\bar{X}_i} \right)^{\theta_i-1} \\ & \left[\frac{d \log p_i}{d \log A_k} - \frac{d \log q_i}{d \log A_k} + \frac{1}{\theta_i} \left(\frac{d \log y_i}{d \log A_k} - 1(i=k) \right) \right], \end{aligned} \quad (31)$$

Evaluating these derivatives at the steady-state, where $A_i = 1$, $p_i = q_i = 1$, $\bar{x}_{ij} = (1-a_i)\omega_{ij}$, $c_i/C = b_i$, $y_i/Y = \lambda_i$, $l_i/y_i = \beta_i a_i$ gives the desired result. \blacksquare

Proof of proposition 4.1. The first-order equation for the allocation of labor is

$$\beta \beta_i A_i^{\frac{\theta-1}{\theta}} L_i^{\beta \frac{\theta-1}{\theta}} = \lambda L_i C^{-\frac{1}{\theta}}.$$

where λ is the Lagrange multiplier on labor. Using the first-order conditions and the labor

market clearing condition implies that

$$L_i = \left(\frac{\beta_i A_i^{\frac{\theta-1}{\theta}}}{\sum \beta_j A_j^{\frac{\theta-1}{\theta}}} \right)^{\frac{1}{1-\beta} \frac{\theta-1}{\theta}}.$$

Substituting this into the utility function gives

$$C = \left(\sum_i \beta_i^{\frac{\theta}{\theta(1-\beta)+\beta}} A_i^{\frac{\theta-1}{\theta(1-\beta)+\beta}} \right)^{\frac{\theta(1-\beta)+\beta}{\theta-1}}.$$

Then for this economy

$$\rho_{ij} = \rho = \frac{\theta(1-\beta) + \beta}{\theta(1-\beta) + \beta + (1-\theta)},$$

where $\beta = 0$ corresponds to

$$\rho = \theta,$$

which is the same as no reallocation case. On the other hand, for $\beta = 1$,

$$\rho = \frac{1}{2-\theta},$$

which is the same as the fully reallocative case. Note that this explodes when $\theta \geq 2$, since in that case, the reallocative solution is more substitutable than perfectly substitutes! For $\rho \in (0, 1)$ we get something in between the perfectly reallocative and no reallocation special cases. ■

Proof of proposition 5.1. First consider the case with reallocation, and note that

$$\tilde{\beta}_k = \frac{\bar{p}_k^\theta \bar{y}_k}{\bar{p}_c^\theta \bar{C}} = \bar{\lambda}_i \frac{\bar{p}_k^{\theta-1}}{\bar{p}_c^{\theta-1}}.$$

Substitute this into the expression for C in proposition E.1 to get

$$C = \left[\sum_k \bar{\lambda}_k \frac{\bar{p}_k^{\theta-1}}{\bar{p}_c^{\theta-1}} A_k^{\theta-1} \alpha_k z_k^{\theta-1} \right]^{\frac{1}{\theta-1}} L.$$

On the other hand, we have that

$$\alpha_k = \frac{\bar{w}^\theta \bar{L}_k}{\bar{p}_k^\theta \bar{y}_k} \left(\bar{A}_k \bar{z}_k \right)^{1-\theta}.$$

Substitute this into the previous expression to get

$$C = \left[\sum_k \bar{a}_k \frac{\bar{w}_k^{\theta-1}}{\bar{p}_c^{\theta-1}} \left(\frac{A_k z_k}{\bar{A}_k \bar{z}_k} \right)^{\theta-1} \right]^{\frac{1}{\theta-1}} L.$$

Finally, substitute in $\bar{p}_c = \bar{w}L/\bar{C}$ and rearrange to get the result.

Next, consider the case without reallocation and note that

$$\tilde{\beta}_k = \frac{\bar{p}_k^\theta \bar{y}_k}{\bar{p}_c^\theta \bar{C}},$$

and

$$\alpha_k = \frac{\bar{w}_k^\theta \bar{L}_k}{\bar{p}_k^\theta \bar{y}_k},$$

which we can substitute into the expression for GDP in proposition E.1 to get our desired expression. ■

Proof of proposition 6.1. Consumption is given by

$$C = A\bar{Y} \left(\bar{a} \left(\frac{L}{\bar{L}} \right)^{\frac{\theta-1}{\theta}} + (1-\bar{a}) \left(\frac{X}{\bar{X}} \right)^{\frac{\theta-1}{\theta}} \right)^{\frac{\theta}{\theta-1}} - X.$$

The first-order condition gives

$$\frac{X}{\bar{X}} = (\bar{Y}A)^{\theta-1} (1-\bar{a})^\theta \bar{X}^{-\theta} Y.$$

Substituting this into the production function gives

$$Y = \frac{A\bar{Y}\bar{a}^{\frac{\theta}{\theta-1}}}{\left(1 - (1-\bar{a})^\theta (\bar{Y}A/\bar{X})^{\theta-1} \right)^{\frac{\theta}{\theta-1}}}.$$

This means that

$$C = \frac{A\bar{Y}\bar{a}^{\frac{\theta}{\theta-1}}}{\left(1 - (1 - \bar{a})^\theta \left(\bar{Y}A/\bar{X}\right)^{\theta-1}\right)^{\frac{1}{\theta-1}}}.$$

Finally, note that

$$\frac{d \log C}{d \log A} = \xi.$$

■

Proof of proposition 7.1. We can write

$$\frac{d \log p_i}{d \log A_k} = a_i \lambda_k + \sum_j (1 - a_i) \omega_{ij} \frac{d \log p_j}{d \log A_k} - \mathbf{1}(i = k).$$

We can solve this to get

$$\frac{d \log p_i}{d \log A_k} = \lambda_k - \psi_{ik},$$

where ψ_{ik} is the ik th element of the Leontief inverse. This says that the change in the price of i depends on the change in the real wage λ_k to the intensity with which i consumes from k . We know that $\psi_{kk} \geq \psi_{ik}$. This expression, which greatly simplifies matters, stems from constant-returns-to-scale, which means that relative prices don't depend on quantities.

Next, we know that

$$\frac{d \lambda_i}{d \log A_k} = b_i (1 - \sigma) \frac{d \log p_i}{d \log A_k} + \tag{32}$$

$$\begin{aligned} & \sum_j \left[\omega_{ji} (1 - a_j) \lambda_j \left\{ (\theta_j - 1) \left(\mathbf{1}(j = k) - \frac{d \log q_j}{d \log A_k} + \frac{d \log p_j}{d \log A_k} \right) \right. \right. \\ & \left. \left. + (1 - \varepsilon_j) \left(\frac{d \log p_i}{d \log A_k} - \frac{d \log q_j}{d \log A_k} \right) + \frac{d \lambda_j}{d \log A_k} \frac{1}{\lambda_j} \right\} \right], \end{aligned} \tag{33}$$

where

$$\frac{d \log q_j}{d \log A_k} = \lambda_k - \sum_l \omega_{jl} \psi_{lk}.$$

Combining these facts, we can write

$$D\lambda' = c'\Psi,$$

where

$$\begin{aligned} c_i = & b_i(1 - \sigma)[\lambda_k - \psi_{ik}] + \sum_j \omega_{ji}(1 - a_j)\lambda_j(\theta_j - 1)[\sum_l \omega_{lj}\psi_{lk} - \psi_{jk}] \\ & + \sum_j \omega_{ji}(1 - a_j)\lambda_j(1 - \varepsilon_j)[\sum_l \omega_{lj}\psi_{lk} - \psi_{ik}] + (\theta_k - 1)\omega_{ki}(1 - a_k)\lambda_k. \end{aligned}$$

We can interpret c_i as how λ_i changes conditional on the change in prices holding fixed other λ_{-i} . ■

Proof of proposition 7.2. In the special case where $\theta_j = \varepsilon_j$,

$$c_i = b_i(1 - \sigma)(\lambda_k - \psi_{ik}) + \sum_j \lambda_j \omega_{ji}(1 - a_j)(\theta_j - 1)[\psi_{ik} - \psi_{jk}] + (\theta_k - 1)\omega_{ki}(1 - a_k)\lambda_k.$$

Therefore,

$$\begin{aligned} \frac{d\lambda_m}{d\log A_k} &= c'_m \Psi e_m, \\ &= \sum_i c_i \psi_i m, \\ &= (1 - \sigma) \left(\lambda_k \lambda_i - \sum_i b_i \psi_{ik} \psi_{im} \right) \\ &\quad + \sum_i \sum_j \omega_{ji} \lambda_j (1 - a_j) (\theta_j - 1) [\psi_{im} \psi_{ik} - \psi_{jk} \psi_{im}] \\ &\quad + \lambda_k (\theta_k - 1) \sum_i \omega_{ki} (1 - a_k) \psi_{im}, \\ &= (\sigma - 1) \left(\sum_i b_i \psi_{ik} \psi_{im} - \sum_i b_i \psi_{ik} \sum_i b_i \psi_{im} \right) \\ &\quad + \sum_j (\theta_j - 1) \lambda_j \left(\sum_i \omega_{ji} (1 - a_j) \psi_{im} \psi_{ik} - \left[\sum_l \omega_{jl} (1 - a_j) \psi_{lk} + \mathbf{1}(j = k) \right] \left[\sum_i \omega_{ji} (1 - a_j) \psi_{im} \right] \right) \\ &\quad + \lambda_k (\theta_k - 1) (\psi_{km} + \mathbf{1}(k = m)), \\ &= (\sigma - 1) \left(\sum_i b_i \psi_{ik} \psi_{im} - \sum_i b_i \psi_{ik} \sum_i b_i \psi_{im} \right) \\ &\quad + \sum_j (\theta_j - 1) \lambda_j \left(\sum_i \omega_{ji} (1 - a_j) \psi_{im} \psi_{ik} - \left[\sum_l \omega_{jl} (1 - a_j) \psi_{lk} \right] \left[\sum_i \omega_{ji} (1 - a_j) \psi_{im} \right] \right) \\ &\quad + \lambda_k (\theta_k - 1) (\psi_{km} + \mathbf{1}(k = m)) + (1 - \theta_k) \lambda_k \sum_i \omega_{ki} (1 - a_k) \psi_{im}, \end{aligned}$$

where the third and fourth line make repeated use of the fact that

$$\Psi - I = (I - \Omega)^{-1} \Psi$$

and

$$\lambda' = b' \Psi.$$

To complete the proof, observe that $d^2 \log C / d \log A_k^2 = d \lambda_k / d \log A_k$, and substitute $m = k$. ■

B A Bound on the Approximation Error

When GDP is homogeneous of degree 1 and the macro elasticities of substitution are all constant, we can bound the size of the first-order approximation error in terms of observable expenditure shares before the shock.

Proposition B.1. *Assume that ξ is constant and equal to 1 and that ρ_{ij} is constant and equal to ρ for every $i \neq j$ (at and away from steady state). Denote λ_i at steady state by $\bar{\lambda}_i$. Then*

$$C(A_i) = \left(\bar{\lambda}_i A_i^{\frac{\rho-1}{\rho}} + (1 - \bar{\lambda}_i) \right)^{\frac{\rho}{\rho-1}},$$

where with some abuse of notation, we denote output as a function of A_i by $C(A_i)$. Furthermore, if $\lambda_i < 1/2$, then

$$\left| \log \left(\frac{C}{\bar{C}} \right) - \bar{\lambda}_i \log(A_i) \right| \geq \frac{1}{2} \left| \frac{\rho-1}{\rho} \right| \bar{\lambda}_i (1 - \bar{\lambda}_i) |\log(A_i)|^2. \quad (34)$$

whenever $(1 - A_i)(\rho - 1) < 0$.

Proposition B.1 shows that when ρ_{ij} are uniform and constant, we can globally characterize output with a CES aggregator, where the relevant weights are the steady-state λ_i . Furthermore, if $\lambda_i < 1/2$, then we can put a lower bound on the size of the approximation error from Hulten's theorem.

Proof of proposition B.1. Using theorem 2.2, we know that

$$\frac{d^2 \log C}{d \log z_i^2} = \left(\frac{\rho-1}{\rho} \right) \lambda_i(z_i) (1 - \lambda_i(z_i)).$$

We wish to provide a lower bound for $\frac{d^2 \log C}{d \log z_i^2}$ over some interval I with $\bar{z}_i \in I$. Note that if $\bar{\lambda}_i < 1/2$ then $\lambda_i(z_i)(1 - \lambda_i(z_i))$ is minimized at $\min_{z \in I} \lambda(z)$.

First, consider $\rho \leq 1$. In this case, $\min_{z \in I} \lambda(z) = \lambda(\max_{z \in I} z)$. If $I = [z, \bar{z}_i]$, then $\min_{z \in I} \lambda_i(z_i)(1 - \lambda_i(z_i)) = \bar{\lambda}_i(1 - \bar{\lambda}_i)$.

Next, consider $\rho > 1$. In this case, $\min_{z \in I} \lambda(z) = \lambda(\min_{z \in I} z)$. If $I = [\bar{z}_i, z]$, then $\min_{z \in I} \lambda_i(z_i)(1 - \lambda_i(z_i)) = \bar{\lambda}_i(1 - \bar{\lambda}_i)$.

Therefore, as long as $(\rho - 1)(z_i - \bar{z}_i) < 0$, then

$$\min_{z \in I} \frac{d^2 \log C}{d \log z_i^2} = \frac{\rho - 1}{\rho} \bar{\lambda}_i(1 - \bar{\lambda}_i).$$

Finally, Taylor's theorem states that

$$\log(C) = \log(\bar{C}) + \bar{\lambda}_i \log(z_i/\bar{z}_i) + \frac{1}{2} \frac{d^2 \log C}{d \log z_i^2} \Big|_{\bar{z} \in I} (\log(z_i/\bar{z}_i))^2, \quad (35)$$

$$\geq \log(\bar{C}) + \bar{\lambda}_i \log(z_i/\bar{z}_i) + \min_{\bar{z} \in I} \frac{1}{2} \frac{d^2 \log C}{d \log z_i^2} \Big|_{\bar{z} \in I} (\log(z_i/\bar{z}_i))^2, \quad (36)$$

$$\geq \log(\bar{C}) + \bar{\lambda}_i \log(z_i/\bar{z}_i) + \frac{1}{2} \left(\frac{\rho - 1}{\rho} \right) \bar{\lambda}_i(1 - \bar{\lambda}_i) (\log(z_i/\bar{z}_i))^2. \quad (37)$$

Rearrange this to get the desired result. ■

C Generalization of Section 6 to Multiple Goods

For the example in Section 6, the economy with extreme complementarity $\theta = 0$ has $C = A/a$, where $1/a$ is the sales to output ratio in steady-state. Therefore, in this example, although Hulten's approximation fails in log terms, Hulten's theorem is globally accurate in *linear* terms. In other words, our examples so far may suggest that extreme complementarities can only have outsized effects, in linear terms, if we restrict the movement of labor across industries.

However, this impression is false. To see this, consider a slightly more complex example where we generalize the example above by allowing multiple industries. Aggregate consumption is Cobb-Douglas across goods ($\sigma = 1$) with equal weights ($b_i = 1/N$). Each good is produced using labor and the good itself as an intermediate input. We assume full

labor reallocation/constant returns to scale ($\beta_i = 1$). We have

$$C = \prod_i c_i^{1/N},$$

and

$$Y_i = \bar{Y}_i A_i \left(a_i \left(\frac{L_i}{\bar{L}_i} \right)^{\frac{\theta_i - 1}{\theta_i}} + (1 - a_i) \left(\frac{X_i}{\bar{X}_i} \right)^{\frac{\theta_i - 1}{\theta_i}} \right)^{\frac{\theta_i}{\theta_i - 1}},$$

with

$$Y_i = c_i + X_i,$$

and perfect reallocation of labor. Then we have the following.

Proposition C.1. *Consider the following special case of the structural model in Section 3. Aggregate consumption is Cobb-Douglas across goods with equal weights, where each good is produced using labor and the good itself as an intermediate input, with expenditure share on labor a_i in steady state and micro elasticity of substitution θ_i . Assume full labor reallocation/constant returns to scale. Then*

$$1 - \frac{1}{\rho_{ij}} = (\theta_i - 1) \left(\frac{1}{a_i} - 1 \right),$$

and

$$\frac{\partial \log \xi}{\partial \log A_i} = \frac{1}{N} (\theta_i - 1) \left(\frac{1}{a_i} - 1 \right).$$

In Figure 12 we plot output as a function of TFP shocks in linear terms. As promised, this economy features strong aggregate complementarities in the sense that a negative TFP shock can cause a drastic reduction in output even in linear terms, despite the fact that labor can be costlessly reallocated across sectors. This happens because, in equilibrium, a negative shock to industry i does not result in more labor being allocated to production in industry i . This follows from the fact that consumption has a Cobb-Douglas form, and so the income and substitution effects from a shock to i offset each other. Since no new labor is allocated to i , if i faces a low structural elasticity of substitution $\theta_i \approx 0$, its output falls dramatically in response to a negative shock. This can then have a large effect on aggregate consumption. Of course, Cobb-Douglas consumption is simply a clean way to illustrate this intuition. If the structural elasticity of substitution in consumption $\sigma < 1$, then these effects would be even further amplified.




Figure 12: GDP for the Leontief case $\theta_i \approx 0$ with two industries.

Proof of proposition C.1. First, consider

$$\max_{X_i} Y_i - X_i,$$

which has the first-order condition

$$X_i = Y_i(1 - \bar{a}_i)^{\theta_i} \left(\frac{A_i \bar{Y}_i}{\bar{X}_i} \right)^{\theta_i - 1} = Y_i(1 - \bar{a}_i) A_i^{\theta_i - 1},$$

where we use the fact that $\bar{X}_i = \bar{Y}_i(1 - \bar{a}_i)$. Substitute this into the production function for Y_i to get

$$Y_i = \frac{A_i \bar{Y}_i \bar{a}^{\theta_i / (\theta_i - 1)} L_i / \bar{L}_i}{\left(1 - (1 - \bar{a}) A_i^{\theta_i - 1} \right)^{\frac{\theta_i}{\theta_i - 1}}}.$$

Substitute this into $c_i = Y_i - X_i$ to get

$$c_i = \frac{A_i \bar{Y}_i \bar{a}^{\theta_i / (\theta_i - 1)} L_i / \bar{L}_i}{\left(1 - (1 - \bar{a}) A_i^{\theta_i - 1} \right)^{\frac{1}{\theta_i - 1}}}.$$

Substitute these into the utility function to get aggregate consumption when labor cannot be reallocated. To get aggregate consumption when labor is reallocated, maximize

aggregate the non-reallocative solution with respect to L_i .

$$\frac{C^r}{\bar{C}} = \left(\sum_i^N \bar{b}_i^\sigma \left(\frac{A_i \bar{Y}_i \bar{a}_i^{\frac{\theta_i}{\theta_i-1}} / \bar{L}_i}{(1 - (1 - \bar{a}_i) A_i^{\theta_i-1})^{\frac{1}{\theta_i-1}}} \right)^{\sigma-1} \right)^{\frac{1}{\sigma-1}} \bar{L}.$$

■

D A Quantitative Illustration in the Case of Electricity

Finally, we end with a simple illustration of some of the limitations of our quantitative exercise. Due to data limitations, we assumed uniform symmetry of ε_i . In this section, we instead work with a stylized representation of the economy, still calibrated to match the relevant expenditure shares, but highlight the unique role an industry like power generation can have when the elasticities of substitution can be heterogeneous. This exercise helps shed some light on how we can reconcile the fact that an industry like electricity production is a small part of the economy, but a priori, we think that it is a critical industry. This tension between size and importance is exemplified by a 2013 speech from Lawrence Summers: “Electricity is only four percent of the economy, and so if you lost eighty percent of electricity you couldn’t possibly have lost more than three percent of the economy, ... we understand that somehow, even if we didn’t exactly understand in the model, that when there wasn’t any electricity there wasn’t really going to be much economy.”

Consider an economy where

$$C = \left(\sum_i b_i c_i^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}},$$

and

$$c_i = \left(a_i l_i^{\frac{\varepsilon_i-1}{\varepsilon_i}} + (1 - a_i) E_i^{\frac{\varepsilon_i-1}{\varepsilon_i}} \right)^{\frac{\varepsilon_i}{\varepsilon_i-1}},$$

with

$$E = \sum_i E_i = A l_E^\beta,$$

where $\beta \in [0, 1]$.

Proposition D.1. *Consider the special case of the structural model in Section 3 outlined above where each good is produced with a structural elasticity of substitution ϵ_i from labor and energy with steady-state shares a_i and $1 - a_i$. All industries, except energy, have constant returns to scale and hire their workers from a common labor market. However, energy is produced with decreasing returns to scale $\beta < 1$. Let A denote TFP shocks to the energy industry, then at steady state,*

$$\frac{d^2 \log C}{d \log A^2} = \frac{\sum_i (1 - a_i) \lambda_i [(\sigma - 1)(1 - \lambda_E) + a_i(\epsilon_i - \sigma)]}{1 + \frac{1 - \beta}{\lambda_i \lambda_E} \sum_i (1 - a_i) \lambda_i [a_i(\epsilon_i - \sigma) + (1 - \beta)(1 - \lambda_E) \lambda_E (\sigma - 1)]}, \quad (38)$$

where λ_E is the Domar weight of energy and λ_i is the Domar weight of labor.




Figure 13: A calibration of the first-order and second-order terms as a function of σ with $\epsilon_i = 0.01$ and $A_E = 0.2$.

We calibrate (38) for different values of consumption elasticity of substitution σ assuming $\beta_E = 0$. To calibrate the model, we set $(1 - a_i)$ to match the E th element of the Leontief inverse. We then set b_i equal to the sales share of the i th industry. This calibration preserves the equilibrium steady-state size of all industries, as well as the total dependence of each industry on electricity. Finally, we set $\epsilon_i = 0.01$ to reflect the low structural elasticity of substitution energy inputs have with labor. Figure 13 plots the first and second-order terms of the Taylor approximation to $\log(GDP)$, following (4), for

an 80% reduction in the productivity of the energy industry $A_E = 0.2$. In our calibration, $\lambda_E \approx 0.04$, so the first-order loss from an 80% decline in electricity generation is just over 3%. However, the second-order losses are roughly $-0.5 \times \log(0.2)^2 = -1.318$, or 73% when $\sigma = 1$, and they are even larger if $\sigma < 1$.

We can also use the second-order terms to approximate the expected value of log GDP in the presence of shocks to the energy industry. If we take Cobb-Douglas consumption as a benchmark, then if the energy industry experiences volatility with variance v , then average GDP will be around $v/2$ times lower than its certainty equivalent. So despite the electricity industry being smaller 4% of GDP, it can transmit idiosyncratic shocks to expected GDP 1-for-2. This may help explain why unreliable electricity production in developing countries can have very harmful effects on output (see Allcott et al., 2016).

We can simplify proposition D.1 by considering a special case where all industries use electricity by the same amount.

Corollary. *For the model of proposition D.1, in the special case where $a_i = a$ for all i , denote the expenditure-weighted average $\sum_i \lambda_i \varepsilon_i$ by $\bar{\varepsilon}$. Then we have*

$$\frac{d^2 \log C}{d \log A^2} = \frac{(1-a)a(\beta + a(1-\beta))(\bar{\varepsilon} - 1)}{a(1-\beta)\bar{\varepsilon} + \beta},$$

this simplifies to

$$\frac{d^2 \log C}{d \log A^2} = \begin{cases} \lambda_E(1 - \lambda_E)^{\frac{\bar{\varepsilon}-1}{\bar{\varepsilon}}}, & \beta = 0 \\ \lambda_E(1 - \lambda_E)(\bar{\varepsilon} - 1), & \beta = 1 \end{cases}$$

regardless of the value of σ . Our decomposition is $\xi = 1 + \lambda_E$, and

$$\begin{aligned} \frac{d^2 \log C}{d \log A^2} &= \lambda_E \frac{\partial \log \xi}{\partial \log A} + \frac{\lambda_E}{1 + \lambda_E} \sum_i \lambda_i \left(\frac{\rho_{Ei} - 1}{\rho_{Ei}} \right), \\ &= \frac{\lambda_E}{1 + \lambda_E} \frac{(1-a)a(\beta + a(1-\beta))(\bar{\varepsilon} - 1)}{a(1-\beta)\bar{\varepsilon} + \beta} + \frac{1}{1 + \lambda_E} \frac{(1-a)a(\beta + a(1-\beta))(\bar{\varepsilon} - 1)}{a(1-\beta)\bar{\varepsilon} + \beta}. \end{aligned}$$

The fact that σ disappears when a_i are symmetric is a manifestation of general principle stated in Section 7. Since the household is symmetrically exposed to shocks from the electricity industry, it does not matter how well the household can substitute amongst its inputs! This example suggests that there are three attributes of the energy sector that make it important to the economy. First, energy appears with a low elasticity of substitution in

production, or ε_i is low for every i , second, energy appears in every production function, or $a_i \gg 0$ for every i , and finally, energy is produced inelastically, or $\beta \ll 0$. All three of these attributes must be present in order for a negative shock to industry E to have large second-order effects on output.

The result below formalizes the idea that although the energy industry E and another industry $i \neq E$ can have the exact same share of sales and, therefore, first-order effects on the economy. Their second-order impact can be arbitrarily different from one another. Let $b_i = 1/N$, $a_i = 1 - 1/N$. If $\beta = 1$, then

$$\frac{d^2 \log C}{d \log A_E^2} = \frac{N-1}{N} (\bar{\varepsilon} - 1),$$

and if $\beta = 0$, then

$$\frac{d^2 \log C}{d \log A_E^2} = \frac{N-1}{N} \left(\frac{\bar{\varepsilon} - 1}{\bar{\varepsilon}} \right),$$

while

$$\frac{d^2 \log C}{d \log A_i^2} = \frac{N-1}{N} (\sigma - 1) \quad i \in \{1, \dots, N\}.$$

However, $\lambda_i = \lambda_E = 1/N$.

Finally, we turn our attention to the non-trivial way in which common shocks will affect this economy. Suppose that we split energy production into hydroelectric power and wind, and we calibrate the expenditure shares so that $\lambda_h = \lambda_w = 1/N$, where h and w denote hydro and wind. For simplicity, suppose that there is only one final goods producer λ_3 . Suppose that $\rho_{hw} = \rho_{wh} = \bar{\rho}$ and $\rho_{h3} = \rho_{w3} = \underline{\rho}$. Finally, assume ξ is constant. Then the second-order impact of a shock to hydro is

$$\frac{1}{2} \frac{1}{N^2} \left(1 - \frac{1}{\bar{\rho}} \right) + \frac{1}{2} \frac{1}{N} \frac{N-2}{N} \left(1 - \frac{1}{\underline{\rho}} \right),$$

whereas the second-order impact of a common shock to hydro and wind is

$$\frac{1}{N} \left(\frac{2}{N} - 1 \right) \left(1 - \frac{1}{\bar{\rho}} \right) + \frac{2}{N} \frac{N-2}{N} \left(1 - \frac{1}{\underline{\rho}} \right).$$

We see that the effect of the common shock is not simply twice that of the idiosyncratic shock (which would be the case if the Hessian was diagonal). Furthermore, the sign on $(1 - 1/\bar{\rho})$ flips when comparing the idiosyncratic and the common shock. The difference between the second-order approximation of a common shock and the sum of two

idiosyncratic shocks is

$$-\frac{1}{N} \frac{N-1}{N} \left(1 - \frac{1}{\underline{\rho}}\right) + \frac{1}{N} \frac{N-2}{N} \left(1 - \frac{1}{\underline{\rho}}\right).$$

We take $1 - 1/\underline{\rho} < 0$, since hydro and wind are macro-complements with the rest of the economy. Then the difference between the idiosyncratic shock and the common shock is greater if the two forms of power generation are macro-substitutes than if they are macro-complements. Intuitively, the closer a substitute that wind is for hydro, the bigger the difference between the impact of an idiosyncratic shock and the impact of a common shock.

Proof of proposition D.1. Denote electricity's sales as a share of GDP by λ_E . Then

$$\lambda_E = \sum_i (1 - a_i) \left(\frac{p_E}{p_i}\right)^{1-\varepsilon_i} \lambda_i,$$

while

$$\lambda_i = b_i \left(\frac{p_i}{p_c}\right)^{1-\sigma},$$

and

$$p_i^{1-\varepsilon_i} = a_i w^{1-\varepsilon_i} + (1 - a_i) p_E^{1-\varepsilon_i},$$

and

$$p_E = \frac{w l_E^{1-\beta}}{A \beta},$$

and

$$l_E = \left(\frac{p_E}{w}\right) \beta E.$$

Finally, we know that

$$C = w l + (1 - \beta) p_E E,$$

where consumption is the numeraire. These equations can be differentiated to give

$$\frac{d^2 \log C}{d \log A^2} = \frac{d \lambda_E}{d \log A} = \sum_i (1 - a_i) \left(\frac{p_E}{p_i}\right)^{1-\varepsilon_i} \left[(1 - \varepsilon_i) \left[\frac{d \log p_E}{d \log A} - \frac{d \log p_i}{d \log A} \right] \lambda_i + \frac{d \lambda_i}{d \log A} \right],$$

we also have that

$$\frac{d \lambda_i}{d \log A} = b_i \left(\frac{p_i}{p_c}\right)^{1-\sigma} (1 - \sigma) \left[\frac{d \log p_i}{d \log A} - \frac{d \log p_c}{d \log A} \right],$$

and

$$\lambda_E = \frac{d \log w}{d \log A} \lambda_l + (1 - \beta) \frac{d \log p_E}{d \log A} \lambda_E = \frac{d \log w}{d \log A} \lambda_l + (1 - \beta) \left(\frac{d \lambda_E}{d \log A} + \lambda_E^2 \right). \quad (39)$$

Next, note that

$$\frac{d \log p_i}{d \log A} = a_i \frac{d \log w}{d \log A} + (1 - a_i) \frac{d \log p_E}{d \log A}, \quad (40)$$

and

$$\begin{aligned} \frac{d \log p_E}{d \log A} &= \frac{d \log w}{d \log A} + (1 - \beta) \frac{d \log l_E}{d \log A} - 1, \\ \frac{d \log l_E}{d \log A} &= \frac{1}{\lambda_E} \frac{d \lambda_E}{d \log A} + \lambda_E - \frac{d \log w}{d \log A}. \end{aligned} \quad (41)$$

Combining the last two equations gives

$$\frac{d \log p_E}{d \log A} = \beta \frac{d \log w}{d \log A} + (1 - \beta) \left[\frac{1}{\lambda_E} \frac{d \lambda_E}{d \log A} + \lambda_E \right] - 1, \quad (42)$$

Combine this with (39) to get

$$\frac{d \log w}{d \log A} = \frac{\lambda_E}{\lambda_l} - \frac{1 - \beta}{\lambda_l} \left(\frac{d \lambda_E}{d \log A} + \lambda_E^2 \right).$$

Substitute this equation into (42) to get

$$\frac{d \log p_E}{d \log A} = \beta \frac{\lambda_E}{\lambda_l} + \left[\frac{1 - \beta}{\lambda_E} - \frac{\beta(1 - \beta)}{\lambda_l} \right] \left(\frac{d \lambda_E}{d \log A} + \lambda_E^2 \right) - 1,$$

Substitute these into (40) to get

$$\frac{d \log p_i}{d \log A} = (a_i + (1 - a_i)\beta) \left(\frac{\lambda_E}{\lambda_l} - \frac{1}{\lambda_l} \left(\frac{d \lambda_l}{d \log A} + \lambda_l^2 \right) \right) + (1 - a_i) \frac{(1 - \beta)}{\lambda_E} \left(\frac{d \lambda_E}{d \log A} + \lambda_E^2 \right) - (1 - a_i)$$

Some tedious algebraic manipulation, as well as the fact that

$$\lambda_l = 1 - (1 - \beta)\lambda_E$$

gives the desired result. ■

E Tractable Special Case with Uniform Elasticities

GDP as a function of the underlying productivity shocks can be written in closed form when $\theta = \varepsilon = \sigma$. In this case, the network model collapses into the same setup as the one in Section 1, we basically recover the same result except the weight parameters have to now be “weighted” by the network. Let \mathbf{A} and \mathbf{z} denote the vector of TFP and labor-augmenting technology shocks.

In this special case, we have

$$\begin{aligned} p^\theta y &= \beta' (I - \text{diag}(1 - \alpha) \text{diag}(A)^{\theta-1} \Omega)^{-1} p_c^\theta C, \\ p^{1-\theta} &= (I - \text{diag}(1 - \alpha) \text{diag}(A)^{\theta-1} \Omega)^{-1} \text{diag}(A)^{\theta-1} \text{diag}(\alpha \circ z^{\theta-1}) w^{1-\theta}, \\ w^\theta &= \text{diag}(A)^{\theta-1} \text{diag}(L)^{-1} \text{diag}(\alpha \circ z^{\theta-1}) p^\theta y, \end{aligned}$$

and

$$C = w' L.$$

Define the Leontief inverse to be

$$\Psi = (I - \text{diag}(1 - \alpha) \text{diag}(A)^{\theta-1} \Omega)^{-1},$$

and define the network-adjusted consumption share to be

$$\tilde{\beta}' = \beta' \Psi,$$

noting that in the case where $\alpha = \mathbf{1}$, we have $\beta = \tilde{\beta}$.

Proposition E.1. *In the case where labor cannot be reallocated,*

$$C = \left[\sum_k \tilde{\beta}_k^{\frac{1}{\theta}} A_k^{\frac{\theta-1}{\theta}} \alpha_k^{\frac{1}{\theta}} z_k^{\frac{\theta-1}{\theta}} L_k^{\frac{\theta-1}{\theta}} \right]^{\frac{\theta}{\theta-1}}.$$

In the case where labor can be perfectly reallocated,

$$C = \left[\sum_k \tilde{\beta}_k A_k^{\theta-1} \alpha_k z_k^{\theta-1} \right]^{\frac{1}{\theta-1}} L.$$

Proof. For perfect reallocation, note that the vector w is now a constant which can be set

to 1 without loss of generality. Then

$$p^{1-\theta} = \Psi \text{diag}(A)^{\theta-1} \alpha.$$

We know that

$$C = \frac{wL}{p_c},$$

which by definition is

$$C = \left(\beta' p^{1-\theta}\right)^{\frac{1}{\theta-1}} L.$$

Substitute the formula for $p^{1-\theta}$ to get the desired result. ■

Proposition E.2. *At the steady-state, when labor cannot be reallocated,*

$$\frac{d^2 \log C}{d \log A_i^2} = \left(\frac{\theta-1}{\theta}\right) \lambda_i a_i \psi_{ii} (\psi_{ii} - \lambda_i) - \frac{(\theta-1)^2}{\theta} \sum_{k \neq i} a_k \frac{\lambda_i^2 \psi_{ik}^2}{\lambda_k} \quad (43)$$

$$+ (\theta-1) \lambda_i (2\psi_{ii} - 1) - (\theta-1) \lambda_i a_i \psi_{ii}^2. \quad (44)$$

If i buys from no one else, then this simplifies to

$$\frac{d^2 \log C}{d \log A_i^2} = \left(\frac{\theta-1}{\theta}\right) \lambda_i (1 - \lambda_i). \quad (45)$$

When labor can be reallocated

$$\frac{d^2 \log C}{d \log A_i^2} = (\theta-1) \lambda_i (\psi_{ii} - 1 + \psi_{ii} - \lambda_i),$$

when i buys from no one else, then this simplifies to

$$\frac{d^2 \log C}{d \log A_i^2} = (\theta-1) \lambda_i (1 - \lambda_i).$$

F Nested Networks with Heterogeneous Micro Elasticities

Now we turn our attention to examples where the macro elasticities of substitution can vary in response to shocks. These examples show that having heterogeneous (but constant) micro elasticities of substitution can lead to variable elasticities at the aggregate level.

Consider a simple nested CES production network with a single final goods producer

$$Y = \left(\sum_k \beta_k x_k^{\frac{\theta-1}{\theta}} \right)^{\frac{\theta}{\theta-1}},$$

where

$$x_k = \left(\sum_i \frac{1}{N} x_{ik}^{\frac{\varepsilon_k-1}{\varepsilon_k}} \right)^{\frac{\varepsilon_k}{\varepsilon_k-1}},$$

and

$$x_{ik} = A_{ik} L_{ik}.$$

This is the simplest example of a network with non-uniform micro-elasticities of substitution. Since we only consider labor-augmenting shocks, we still have that $\xi = 1$ and $d \log \xi / dA_i = 0$.

We consider three cases: (1) labor cannot be reallocated at all, (2) labor can be reallocated within the inner nest but not the outer nest, (3) labor can be fully reallocated. Using these examples, we see how extreme complementarity or substitutability can propagate through the network. Define

$$\rho_r = \begin{cases} \varepsilon_r & \text{if labor cannot be reallocated within industry } r, \\ \frac{1}{2-\varepsilon_r} & \text{if labor can be reallocated within industry } r, \end{cases}$$

and

$$\rho = \begin{cases} \theta & \text{if labor can be reallocated across industries,} \\ \frac{1}{2-\theta} & \text{if labor can be reallocated across industries.} \end{cases}$$

Then we can show the following.

Proposition F.1. *Consider the following special case of the structural model in Section 3. There is a nested CES network structure with heterogeneous micro elasticities of substitution and labor-augmenting shocks. The macro elasticity of substitution is given by*

$$1 - \frac{1}{\rho_{ij}} = \begin{cases} 1 - \frac{1}{\rho_k} & i, j \in N_k \\ \frac{a(1-\frac{1}{\rho_r}) + b(1-\frac{1}{\rho_k}) + c(1-\frac{1}{\rho})}{a+b+c}, & i \in N_k, j \in N_r \end{cases}, \quad (46)$$

where a, b, c are positive and $a = \frac{1}{\lambda_r^{(i)}} - \frac{1}{\lambda_r}$, $b = \frac{1}{\lambda_k^{(j)}} - \frac{1}{\lambda_k}$, and $c = \frac{1}{\lambda_r} + \frac{1}{\lambda_k}$ with $\lambda_r^{(i)}$ being good i 's expenditure share of total expenditure, and λ_r is industry r 's expenditure share of GDP.

This example shows how networks can mix elasticities of substitution, such that the macro elasticities of substitution are not constant even though all the micro elasticities of substitution are constant. In the special case where the network has a simple nested structure, the macro elasticities of substitution take an especially simple form: a harmonic average of the *effective* elasticities of substitution within and across the different nests, where the weights depend on the relative expenditure shares. The effective elasticity of substitution here is the micro elasticity of substitution adjusted for the general equilibrium effect of reallocation. This example also shows that if *any* of the effective elasticities of substitutions are close to zero, from below or from above, then this will strongly affect the macro elasticity of substitution. In other words, near-singularities will be propagated to macro elasticities regardless of where they show up in the nests. This gives a simple way of deducing the ultimate impact of complementarity or substitutability within nests, up to a second order, by simply averaging across them with appropriate weights.

For these special cases, we can go one step further and even provide the expression for output in closed form.

Proposition F.2. *For the case with no-reallocation, GDP as a function of productivity shocks is given by*

$$C = \left(\sum_k \beta_k \left(\sum_i^{N_k} \frac{1}{N} (A_{ik} \bar{L}_{ik})^{\frac{\varepsilon_k - 1}{\varepsilon_k}} \right)^{\frac{\varepsilon_k}{\varepsilon_k - 1} \frac{\theta - 1}{\theta}} \right)^{\frac{\theta}{\theta - 1}}.$$

Proposition F.3. *For the case with sectoral reallocation, GDP as a function of productivity shocks is given by*

$$C = \left(\sum_k \beta_k \left(\left(\sum_i^{N_k} \left(\frac{1}{N} \right)^{\varepsilon_k} A_{ik}^{\varepsilon_k - 1} \right)^{\frac{1}{\varepsilon_k - 1}} \bar{L}_k \right)^{\frac{\theta - 1}{\theta}} \right)^{\frac{\theta}{\theta - 1}}.$$

Proposition F.4. *For the case with full reallocation, GDP as a function of productivity shocks is given by*

$$C = \left(\sum_k \beta_k^\theta \left(\sum_i^{N_k} \left(\frac{1}{N} \right)^{\varepsilon_k} A_{ik}^{\varepsilon_k - 1} \right)^{\frac{1}{\varepsilon_k - 1} (\theta - 1)} \right)^{\frac{1}{\theta - 1}} \bar{L}.$$

G Macro Moments of Second-Order Approximation

	Mean	Standard Deviation	Skewness
Model without reallocation	-0.0056	0.0113	-0.3679
Model with full reallocation	-0.0036	0.0110	-0.1081
$(\theta, \sigma) = (0.1, 0.3)$	-0.0089	0.0117	-0.6496
$(\theta, \sigma) = (0.6, 0.8)$	-0.0035	0.0112	-0.1648

Table 2: For each model, we compute the second-order approximation of $\log GDP$ as a translog function (quadratic in log of TFP shocks). We then simulate the quadratic model with 50,000 draws of TFP shocks.