

CAMBRIDGE WORKING PAPERS IN ECONOMICS

JANEWAY INSTITUTE WORKING PAPERS

Supply Chain Disruptions, the Structure of Production Networks, and the Impact of Globalization

Matthew
Elliott
University of
Cambridge

Matthew O.
Jackson
Stanford University
and Santa Fe
Institute

Abstract

We introduce a parsimonious multi-sector model of international production and use it to study how a disruption in the production of intermediate goods propagates through to final goods, and how that impact depends on the goods' positions in, and overall structure of, the production network. We show that the short-run disruption can be dramatically larger than the long-run disruption. The short-run disruption depends on the value of all of the final goods whose supply chains involve a disrupted good, while by contrast the long-run disruption depends only on the cost of the disrupted goods. We use the model to show how increased complexity of supply chains leads to increased fragility in terms of the probability and expected short-run size of a disruption. We also show how decreased transportation costs can lead to increased specialization in production, with lower chances for disruption but larger impacts conditional upon disruption.

Reference Details

2424 Cambridge Working Papers in Economics
2415 Janeway Institute Working Paper Series

Published 14 May 2024

Keywords Supply Chains, Globalization, Fragility, Production Networks, International Trade
JEL-codes D85, E23, E32, F44, F60, L14

Websites www.econ.cam.ac.uk/cwpe
www.janeway.econ.cam.ac.uk/working-papers

Supply Chain Disruptions, the Structure of Production Networks, and the Impact of Globalization

Matthew L. Elliott and Matthew O. Jackson*

May 2024

Abstract

We introduce a parsimonious multi-sector model of international production and use it to study how a disruption in the production of intermediate goods propagates through to final goods, and how that impact depends on the goods' positions in, and overall structure of, the production network. We show that the short-run disruption can be dramatically larger than the long-run disruption. The short-run disruption depends on the value of all of the final goods whose supply chains involve a disrupted good, while by contrast the long-run disruption depends only on the cost of the disrupted goods. We use the model to show how increased complexity of supply chains leads to increased fragility in terms of the probability and expected short-run size of a disruption. We also show how decreased transportation costs can lead to increased specialization in production, with lower chances for disruption but larger impacts conditional upon disruption.

Keywords: Supply Chains, Globalization, Fragility, Production Networks, International Trade

JEL Classification Numbers: D85, E23, E32, F44, F60, L14

*Elliott is at the Faculty of Economics at Cambridge University. Jackson is at the Department of Economics at Stanford University, and also an external faculty member of the Santa Fe Institute. Web sites: <https://sites.google.com/view/matthewlelliott/> and <http://www.stanford.edu/~jacksonm/>. We thank Vasco Carvalho, Glenn Magerman, Ankur Mani, and Chao Wei for helpful discussions. We gratefully acknowledge support under NSF grant SES-2018554, the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement number #757229) and the JM Keynes Fellowships Fund. Sarah Taylor and Alastair Langtry provided excellent research assistance.

1 Introduction

One among many global supply chain failures stemming from the labor and transport disruptions of the COVID pandemic was a worldwide-shortage of integrated circuits and, in particular, basic computer chips.¹ These appear in almost all consumer goods that involve electronics, and integrated circuits rank fourth among products traded internationally (Jeong and Strumpf, 2021). Although basic computer chips are a commodity good that in typical times can sell for a few cents each,² their shortage stalled the production of many downstream goods. More generally, goods vary widely in their positions in supply chains and their potential to disrupt final goods production.³ Further, supply chains are complex and yet the production of key inputs can be highly consolidated,⁴ and so it is important to quantify the potential for disruptions to propagate as a function of the production network.

We develop a parsimonious model of global trade and interlinked supply chains and use it to examine disruptions to supply chains. We characterize the impact of production disruptions and how this depends on structure of supply chains, the position of the disrupted goods in these chains, and the value of the final goods produced via the impacted chains. We contrast the short and long-run impact of a drop in some good’s productivity. We also show how the potential for short-term disruption increases linearly in the complexity of supply chains (for small disruption probabilities). Finally, we examine how the potential for disruption change as transportation costs drop and supply chains evolve to be more specialized and involve more trade across countries.

In the long run, when the change in productivity of an input can be completely compensated for by adjusting quantities and sources of all inputs, Hulten’s Theorem holds. That is, the marginal impact of a shock is proportional to the total amount spent on that input relative to total GDP. For instance, a shock that reduces the productivity of a \$500 billion market like that for integrated circuits by around 5 percent, would have a long-run impact of around \$25 billion. Roughly, at the margin, if productivity drops by 5 percent then the circuits become about 5 percent more expensive, and so the total resource loss is an extra 5 percent of what was being spent on that input originally.⁵ However, in the short run, the impact can be much larger. A 5 percent shortage of integrated circuits propagates to prevent the production of 5 percent of all

¹Long lead times in the necessary capital equipment, a drought in Taiwan that impacted on the production capacity of the industry, and high demand as people switched expenditure from experiences to products due to the Covid-19 pandemic, all played a role, along with several other factors including a fire at a production plant. See, for example, McLain (2021); Jeong and Strumpf (2021); Davidson and Farrer (2021).

²For example, an AC-DC Switching Power Supply Pulse-Width Modulation Integrated Circuit sells for around 6 cents per unit with a minimum order of 1000 .

³Several works show that disruptions propagate through supply chains Barrot and Sauvagnat (2016), Boehm, Flaaen, and Pandalai-Nayar (2019) and Carvalho et al. (2021).

⁴For example, studies suggest that the majority of complex computer chips are assembled in Taiwan (close to 90 percent for the most advanced), and yet the inputs cross borders more than 70 times before reaching the final stage of production (<https://www.gsaglobal.org/globality-and-complexity-of-the-semiconductor-ecosystem/>).

⁵This intuition applies at the margin, and can even overestimate the impact as it does not account for the possibility that the input mix can be changed to mitigate the impact of the shock as one moves away from the margin.

final goods that use them as inputs, directly or indirectly, at some stage of production. Those final goods, valued at around \$5 trillion, lose production valued at around \$250 billion, or ten times the long-run impact. This comparison is even starker if a shock to basic computer chips is considered given their low value and the breadth of their use as inputs.

The contrast can be summarized as the short-run impact being approximately proportional to the value of all downstream final products, while the long-run impact is approximately proportional to the change in the cost of producing the input itself. This is consistent with what Larry Summers wrote, “There would be a set of economists who would sit around explaining that electricity was only 4 percent of the economy, and so if you lost 80 percent of electricity, you couldn’t possibly have lost more than 3 percent of the economy... [However,] we would understand that [...] when there wasn’t any electricity, there wasn’t really going to be much economy.” (Summers, 2013).⁶

The details of the impact of a short-run disruption depend on the structure of the supply chains involved. An upper bound is that the disruption is equal to the percentage reduction in the output of the shocked technologies multiplied by the value of production of all final goods that use the shocked good as an input directly or indirectly. We identify several natural and intuitive situations in which this bound is tight. One such case is when there is conformity in the inputs used by producers within industries, and there is an industry shock—i.e., all producers of a certain good are shocked. Moreover, when all transportation costs are reduced sufficiently, countries specialize in what they produce and hence technology-specific shocks become industry-specific shocks, such that the bound is always obtained.

The more general calculation can differ from the bound depending on two issues. One is diversity of sourcing: some technologies might source the same input from multiple suppliers, and if not all of those suppliers are affected, then impact of the shock is reduced. The other is diversity of production technologies: different technologies that require different inputs might be used for producing the same good, and so be unaffected by a shock that propagates downstream to disrupt a competitor. However, if there are cycles in the supply network, then these can feed back and amplify disruptions allowing the bound to be obtained even in the presence of diversity in sourcing and technology. To calculate the short-run impact of a shock and how it propagates we provide a (convergent) algorithm, allowing for cycles in the production network. The algorithm converges to the maximum amount that can be produced of the final goods subject to the shocks.

We also examine the medium-run in which supply is limited, but in which rationing can be proportional to the value of the downstream goods that use a given input. For example, if production of chips are disrupted in the very short run, all manufacturers using those chips may face delayed deliveries and lost production. Over time, prices of chips can adjust and redirect the chips to the uses where they are most valued, and this can mitigate the overall loss in GDP. We show how this mitigation can range between the effect of the short versus long

⁶See [Carvalho and Tahbaz-Salehi \(2019\)](#) for further discussion around this and the impact of disruptions when Hulten’s theorem doesn’t hold. Of particular note is [Baqaee and Farhi \(2019\)](#). See also [Wei \(2003\)](#) for a discussion for the timing of adjustments.

run, depending on how diverse the downstream production using an input are.

With these results in hand, we then examine the expected disruption due to an independent probability of a shock to different inputs. We show that as the complexity of the supply chains, as measured by the number of inputs, increases, the expected loss in GDP holding all else fixed, increases linearly in the complexity. We go on to discuss how the long-run impact varies with the depth vs breadth of supply chains, while the short-run impact is more sensitive to how many final goods given inputs are upstream from.

We then use this to examine some comparative statics in trade costs. When all trade costs are reduced sufficiently, generically goods are sourced from the lowest cost technology and production becomes specialized. This leads to lower diversity in sourcing and diversity in production technologies is lost. There are countervailing effects: there are fewer sources to disrupt, which can lower the chance of a disruption occurring, but conditional upon occurrence, the impact is larger. This is consistent with empirical evidence on globalization, specialization, and fragility (Giovanni and Levchenko, 2009; Magerman et al., 2016; Di Giovanni et al., 2022; Bernard et al., 2022; Baldwin and Freeman, 2022).

Our results can be applied to a variety of situations in which supply chains become disrupted, including the imposition of sanctions and anticipating their effects. In this context, Bachmann et al. (2022) provide important intellectual foundations for our work. They estimate the impact of Russian energy sanctions for the German economy, varying assumptions about elasticities of substitution to represent the shorter run and longer run effects.

More generally, the three most closely related strands of literature to our work are: a macroeconomic literature on production networks, a more microeconomic literature on supply chain robustness, and the literature on international trade and global value chains.

Building on the seminal work of Leontief (1936), Long Jr and Plosser (1983) and Acemoglu et al. (2012), a series of papers examine the propagation of shocks through sectoral and firm inter-linkages in the economy.⁷ Some of the recent work has incorporated cascading failures, production shut-downs and endogenized the network structure (Costinot, Vogel, and Wang, 2013; Dhyne, Magerman, and Rubínová, 2015; Magerman et al., 2016; Brummitt et al., 2017; Baqaee, 2018; Oberfield, 2018; Acemoglu and Tahbaz-Salehi, 2020; Acemoglu and Azar, 2020; Baqaee and Farhi, 2021; Kopytov et al., 2021; König et al., 2022; Chen et al., 2023; Grossman, Helpman, and Lhuillier, forthcoming). However, within that literature, the focus has been on the long-run equilibrium impacts of shocks in which all factors are perfectly flexible and the economy re-equilibrates. We contribute to this literature by considering the short-run impact of a shock, with no adjustments.

Perhaps closest to us is Bui et al. (2022) and Pellet and Tahbaz-Salehi (2023) who incorporate rigidities into production networks. In their models some inputs are inflexible and must be committed to before shocks are realized. Firms imperfectly anticipate such shocks, and the flexibility of different inputs distorts the relative amounts that are used. While Bui et al. (2022) focus on rigidities in primary inputs, Pellet and Tahbaz-Salehi (2023) instead focus on

⁷See Carvalho and Tahbaz-Salehi (2019) and Baqaee and Rubbo (2022) for recent surveys.

rigidities related to the supply of some intermediate goods. This makes the network structure matter in a way that is a bit closer to our paper. However, the focus and findings of [Pellet and Tahbaz-Salehi \(2023\)](#) are different and complementary. They study how ex-ante adjustments by firms dampen the equilibrium impact of shocks, while we focus on estimating how the initial lack of adjustment in the short run amplifies shocks, and provide general results about how the network structure matters for this. Indeed, key results in their paper for which we have no counterpart include: (i) that the input mix used by firms is distorted, relative to the perfectly flexible benchmark, towards more flexible inputs; (ii) that the aggregate impact of a shocks is dampened by the rigidities in production choices; and (iii) how nominal rigidities impact inflation. Similarly, our main results on the propagation of the short-run impact of a shock with no adjustments has no counterpart in their paper.

There is also considerable work studying networks and fragility.⁸ Like ours, that work is mainly theoretical. The work closest to us in this area focuses on the fragility of supply chains and is complementary in so far as it focuses on better understanding the frictions that lead to the formation of inefficient supply networks. However, it abstracts from general equilibrium considerations and tends to focus on supply networks that are only a couple of layers deep. It includes, for example, [Bimpikis, Fearing, and Tahbaz-Salehi \(2018\)](#), [Bimpikis, Candogan, and Ehsani \(2019\)](#) and [Amelkin and Vohra \(2020\)](#). Perhaps closest is [Elliott, Golub, and Leduc \(2022\)](#). Like us the focus there is on the macroeconomic implications of shocks in the short-run and deep networks are accommodated. However, they investigate firms' strategic investments into the local robustness of their supply chains, which we do not consider, and study when equilibrium investments will yield fragile networks.⁹ In contrast our analysis provides details on the how the size and scope of a disruption depend on specific details of the supply chain, which are abstracted away from in their work. On the empirical side, [Diem et al. \(2022\)](#) develop and estimate using Hungarian value added tax data, a measure of the systemic risk different firms pose to the economy by considering the short-run impact on aggregate output of a randomly selected firm failing. Their approach differs from ours as they consider an influence measure that examines both upstream and downstream impacts. Our focus is on the GDP/consumption impact. Nonetheless, they show that the empirical impacts are large in magnitude, consistent with the contrast that we draw between the short and long runs.

Finally, there is related research on networks and international trade.¹⁰ Some of the more closely related work in that literature seeks to better understand which importers match to which exporters, and how this is influenced by various frictions (see, for example, [Chaney \(2014\)](#); [Bernard, Moxnes, and Saito \(2019\)](#); [Grossmand, Helpman, and Redding \(2021\)](#)). In terms of the macro modelling of international trade, the approach taken has tended to be very

⁸See [Elliott and Golub \(2022\)](#) for a recent survey, while [Baldwin and Freeman \(2022\)](#) reviews the literature on risks in global supply chains.

⁹The set up of their model (simple networks with no cycles to facilitate percolation and small firms) is also very different, as is the amplification mechanism. In [Elliott, Golub, and Leduc \(2022\)](#) individual firms are small and in equilibrium production is robust with respect to the failure of any one firm. In contrast, we have representative firm/technologies and study how shocks to individual technologies propagate.

¹⁰For recent surveys see [Bernard \(2018\)](#) and [Antràs and Chor \(2022\)](#).

different from ours. For example, the workhorse models of Melitz (2003) and Caliendo and Parro (2015), while being well suited for answering a variety of questions and fitting various aspects of the trade data, are not so well suited to understanding how shocks’ propagate and amplify depending upon the position of disruption in the trade network.

2 A Model of (International) Supply Networks

2.1 The Model

We use the same notation for each set and its cardinality.

Goods and Countries:

There is a set N of countries indexed $n \in \{1, \dots, N\}$.

Goods consist of a set M of intermediate goods, including raw materials, indexed by $m \in \{1, \dots, M\}$ that are used as inputs to production; and a set F of final goods indexed $f \in \{1, \dots, F\}$ that are consumed. We assume that final goods are never used as inputs to production to simplify notation, but it is trivial to extend the model to permit this.¹¹

We use the term “goods” throughout the paper but emphasize that these include not only agricultural and manufactured goods, but also include services.

Labor and Endowments:

Country n ’s endowment of labor is denoted $L_n > 0$, and it is supplied completely inelastically so that in equilibrium L_n units are all used in production.

Access to raw materials (which are a special type of intermediate good) within a country is represented via the available production technologies: for instance, a country that has oil has a technology that outputs oil, while if there is no oil in a country then there is no technology available in that country to produce oil.

Technologies:

We work with Arrow-Debreu production economies. A technology is described, as in the classic model of Arrow and Debreu (1954) (a production plan in their parlance), by a vector $\tau \in \mathbb{R}^{1+M+F}$. A technology lists the combinations of labor and intermediate goods required as inputs to produce positive amounts of output; with the interpretation that $\tau_k < 0$ implies that good k is an input and $\tau_k > 0$ implies good k is an output.

We focus on constant returns to scale technologies. Although constant returns are not needed for our main results, we make it for simplicity as the model is already complicated by accounting for the supply network. Importantly, for short-term disruptions and measurements, constant returns are appropriate.

A technology τ satisfies three additional conditions.

¹¹Simply create a duplicate industry: If a country produces a good that is used as both an intermediate good and a final good, then let there be two industries producing the good, one of which sells it only as an intermediate good, and another that sells it only as a final good to consumers. Prices are dependent on costs of inputs and so the goods will naturally end up with the same prices in equilibrium.

First, $\{k : \tau_k > 0\}$ has exactly one element, interpreted as the output of the technology. Other entries are either 0 or negative, with the negative ones being the inputs. Let $O(\tau) = \{k : \tau_k > 0\}$ denote the output good associated with technology τ and $I(\tau) = \{k : \tau_k < 0\}$ denote the input goods.

Second, we normalize the output to 1 so that $\max_k \tau_k = 1$. Thus, a technology indicates the amounts of inputs needed to produce one unit of the output good, and this can be scaled to any level given the constant returns to scale.

Third, the leading entry in the technology vector, representing the amount of labor required, is strictly negative for all technologies, so that labor is needed in each production process. This ensures that there are no infinite production paths that are costless, and ensures existence.

An example of technologies appears in Figure 1.

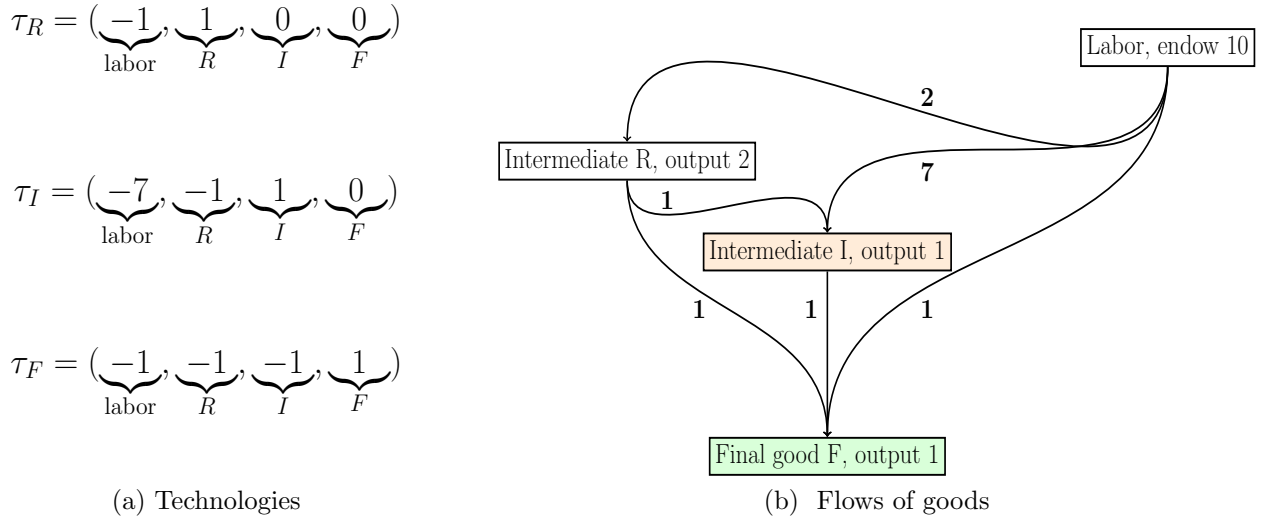


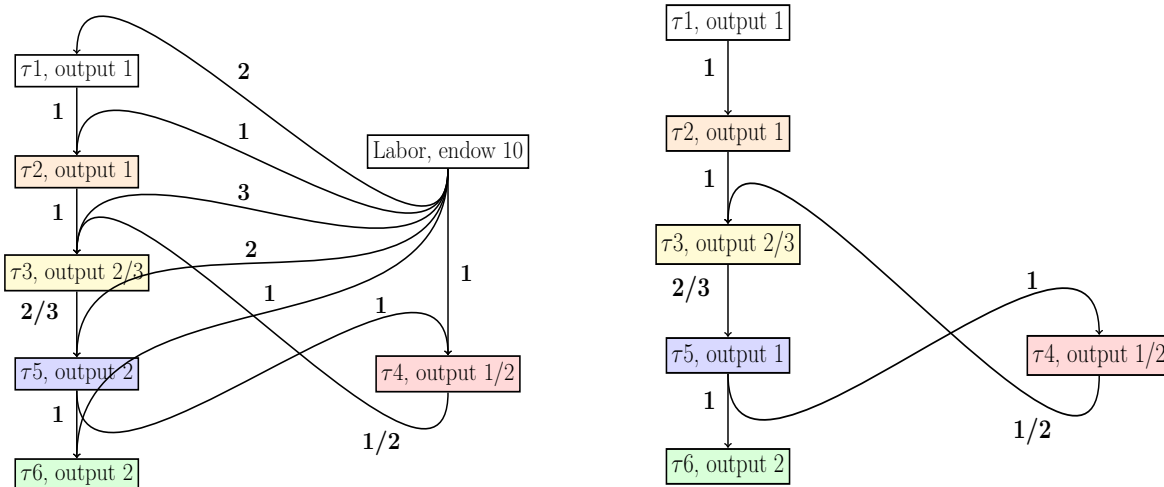
Figure 1: An example of an economy with three technologies and 10 units of labor, outputting 1 unit of a final good.

Figure 2 exhibits the possibility of cycles and more complex supply chains.

Let T_n be the set of technologies assigned to country n , and let $T = \cup_n T_n$ be the set of technologies. We take T to be finite.

This is for several reasons.

One is that any given Arrow-Debreu equilibrium in our setting can be rationalized with a finite set of production plans (our technologies) and this is without loss of generality for our results. For example, a Cobb Douglas production function of the form $y = \ell^\alpha m^{1-\alpha}$ is represented by following the set of triples as ℓ is varied: $(\tau_\ell)_\ell = (-\ell, -\ell^{\frac{\alpha}{1-\alpha}}, 1)_\ell$ where the first entry represents the quantity of labor input, the second entry represents the quantity of the intermediate good m input and a single unit of the final good is produced. This is pictured in Figure 3. Moreover, other productions (CES or some much more general function) can be



(a) Technological dependencies

(b) Technological dependencies without labor pictured

Figure 2: An example of technological interdependencies and good flows: The weight on a directed link from an input into a technology represents the number of units of that input shipped. It is convenient to occasionally omit the dependence on labor to simplify the figures in what follows. The final good technology is τ_6

approximated by our sets, and so by using general T^n we can avoid imposing restrictions on the nature of those functions and still derive our results.




Figure 3: Long-run production of good 2 according to a Cobb-Douglas function of labor and good 1: $y_2 = L^\alpha x_1^{1-\alpha}$, can be approximated by setting T^n to be the intersection of $\{(-l, -x_1, 1) : l^\alpha x_1^{1-\alpha} = 1\}$ with some grid.

Second, and importantly, is that this realistically matches short-run implications, and more

realistically than using some production function in which inputs can be readjusted. For instance, if a business loses electricity, or a car manufacturer loses steering wheels, or a phone producer loses chips, they cannot overcome this simply by shifting to use more labor. In the longer run, they can reoptimize their production recipe to adjust inputs and re-source or adopt different technologies, but in the short run such substitution is impossible. Thus, working with fixed sourcing and vectors (technologies) in the short run and then allowing for more potential variations in the longer run provides a natural contrast.

Third is that a finite set of available technologies that describe recipes for production is realistic when it comes to addressing things like innovations to technologies, which is something that we wish our model to address.

A final reason is that this makes having finite sets makes it easy to talk about generic variations in the technologies used. We can still prove existence of equilibrium despite the non-convexity.

When we say a *generic* T , we mean that there are no duplicate technologies: for any $\tau, \tau' \in T$, $\tau \neq \tau'$.¹²

Shipped units:

The matrix $\mathbf{x} \in \mathbb{R}_+^{N+T} \times \mathbb{R}_+^T$ denotes the number of units shipped. $x_{\tau'\tau}$ denotes the amount of the output, $O(\tau')$, produced by technology τ' that is shipped for use as an input by technology τ . We let $x_{n\tau}$ denotes the amount of labor endowed to country n that is used by technology τ , possibly in another country.

Transportation costs:

Transportation costs are captured via the matrix $\theta \in \mathbb{R}_+^{N+T} \times \mathbb{R}_+^T$. Entry $\theta_{n\tau} \geq 1$ denotes the amount of labor that must be supplied from country n to technology τ in order for technology τ to get one unit of labor input. This could reflect a cost of remote working, or a cost of migration, among other things. Entry $\theta_{\tau\tau'} \geq 1$ for $\tau, \tau' \in T$ denotes the number of units of good $O(\tau)$ that must be shipped for technology τ' to receive one unit of this good to use as an input. These costs can reflect many things, for instance international shipping costs or tariffs, among other things, but can also represent shipping costs internal to a country (for instance, shipping from the place where a raw material is produced to where a manufacturing plant is located). Thus, implicitly, technologies when coupled with this matrix encode transportation costs and locations.

Note also that differences in the effectiveness of labor and any intermediate goods across countries can be captured via transportation costs and available technologies.

Having final goods be costless to transport simplifies the consumption problem and enables us to concentrate on the production process. It is not necessary for our main results, but makes comparisons to key existing results (i.e., Hulten's Theorem) possible and so we maintain the assumption.

¹²Given that technologies can be seen as finite dimensional vectors and there are a finite set of them, one can use a Euclidean metric to define an open and dense set of collections T in the subspace of vectors satisfying our definition of a technology.

Prices:

Prices for each good and labor are described by a vector $p \in \mathbb{R}_+^{N+T}$, providing the cost of hiring labor in each country and technology-specific prices for all goods. These prices are prices at the point of sale per unit. Adjusting for transportation costs yields a price for use in any given technology.

As final goods are costless to ship, there is a world price for each final good. Abusing notation we let p_f denote the price of final good f . So, in equilibrium $p_\tau = p_{\tau'} = p_f$ for all (τ, τ') such that $O(\tau) = O(\tau') = f$.

Preferences:

Laborers are the consumers of the final goods. Consumers have preferences represented by

$$U(c_1, \dots, c_F)$$

that is increasing ($c \geq c', c \neq c'$ implies $U(c) > U(c')$) and strictly quasi-concave and homogeneous of degree 1.

As there are no transportation costs on final goods, equilibrium prices for final goods are the same across different countries. Further, as preferences are not country specific, and represented by a utility function that is homogeneous of degree 1, all agents consume scalings of the same bundle of final goods that are proportional to their wages. Thus, as we show in Lemma 1 in Appendix A, our formulation admits a representative consumer with preferences represented by $U(\cdot)$.

An Economy and Equilibrium

Competitive equilibrium and constant returns to scale imply that there are zero profits, and thus we ignore firm ownership and profits for the sake of eliminating unnecessary notation and definitions.

An *economy* is therefore a list specifying the set of countries, goods, technologies, labor endowments and transportation costs: $(N, M, F, \{T_n\}_n, \{L_n\}_n, \theta)$.

An *equilibrium* is defined in the usual way (following Arrow and Debreu (1954)) and is such that

- laborers supply their endowment of labor inelastically and choose final goods to consume to maximize their preferences,
- each technology is used to maximize profits of its output minus the costs of its inputs, and
- markets for all goods clear.

Given that details of the equilibrium and existence are standard, we present them in Supplementary Appendix I where we:

- (i) offer a fully formal definition of a general equilibrium of an economy,

- (ii) show that an equilibrium exists,
- (iii) show that in all equilibria the same amount of each final good is produced, and
- (iv) show that an equilibrium is fully specified by the flow of country-specific labor to technologies, the flow of goods between all technology pairs and (local) prices for the output of all technologies.

An equilibrium specifies the flow of goods between technologies. It is helpful to represent these flows as a directed, weighted network as in Figure 4.




Figure 4: An example of equilibrium flows, labor not pictured. Green technologies produce final goods, and colors indicate the output good.

3 Contrasting Short- and Long-Run Impacts of Supply Chain Shocks

So far, we have normalized the output of a technology τ producing good $k = O(\tau)$ to be $\tau_k = 1$. In what follows it is convenient to let τ_k vary at the margin to represent changes in productivity of that technology. To identify the impact of a shock, we consider a shock that changes the output of some technology τ , given by τ_k , from its initial value of 1.

We begin with the *long-run* impact of a shock, showing that Hulten’s (1978) Theorem holds in our setting.

3.1 Hulten’s Theorem: Long-Run Impacts of Changes in the Supply Chain

Recall that GDP denotes the total expenditures on final goods:

$$GDP = \sum_n \sum_f p_f c_{fn}.$$

Because the consumers in different countries have the same homothetic preferences, they demand (potentially) different quantities of the same bundle of goods. Thus, final total consumer demand equals the demand induced by a representative consumer with the same preferences and wealth equal to total labor income (Lemma 1, Appendix A). Therefore, the utility of the representative consumer, denoted by U , is a measure of overall welfare and, given the homogeneity of preferences, it is proportional to GDP.

PROPOSITION 1 (Hulten’s Theorem). *Consider an economy with a generic set of technologies and an equilibrium of that economy, and a technology τ used in positive amounts in equilibrium to produce good $k = O(\tau)$. The marginal impact on aggregate utility, and on GDP, of a change in the total factor productivity of τ is equal to the total expenditures on good k produced using technology τ , relative to overall GDP. That is,*

$$\frac{\partial \log(U)}{\partial \log(\tau_k)} = \frac{\partial \log(GDP)}{\partial \log(\tau_k)} = \frac{p_\tau y_\tau}{GDP}.$$

Hulten’s Theorem identifies the long-run marginal effect of a change in the productivity of a technology: it measures the full equilibrium adjustment of the economy to a new equilibrium. It shows that a sufficient statistic for long run marginal impact of a shock to an industry on GDP is the equilibrium value of the shocked industry. A simple intuition for the result is that, at the margin, the reduced productivity is compensated for by sourcing more inputs at their current prices. Genericity rules out that there are exact duplicates for the shocked technologies and so complete substitution away from the shocks is not costless.¹³

Hulten’s Theorem is illustrated in Figure 5, where a ten percent decrease in productivity of a technology that accounts for 1/5 of the overall expenditures in the economy ends up decreasing total output by 1/50. The figure shows how labor readjusts to increase the production of that technology, which ends up decreasing production overall. However, the ability for labor to be reallocated means that a ten-percent shock only leads to a two-percent decrease.

3.2 Bottlenecks and Disruptions: Short-Run Impacts of Shocks

The predictions of Hulten’s Theorem only apply in the long run when changes can be fully adapted to. In contrast, in the short run, the impact of changes in productivity can be much

¹³Bahal and Lenzo (2023) consider larger technology shocks that change a sector’s input mix, and aggregate the impact of these shocks into an alternative measure of a sector’s importance.

Labor endowment: 10

$$p = \left(\underbrace{\frac{1}{10}}_{\text{labor}}, \underbrace{\frac{1}{10}}_R, \underbrace{\frac{4}{5}}_I, \underbrace{1}_F \right)$$

$$GDP = \sum_f p_f c_f = 1; \quad p_R y_R = .1 * 2$$

$$\text{Marginal impact: } (p_R \cdot y_R) / GDP = 1/5$$

Extrapolating for a 10% shock (source more):

Long Run impact: $\approx 1/50$ th of GDP

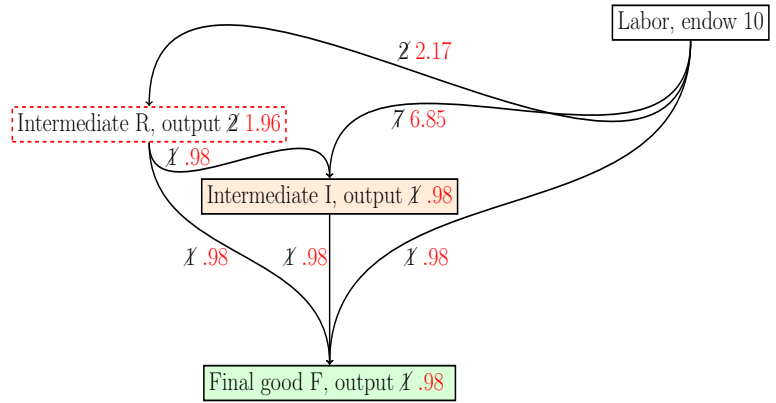


Figure 5: An example of the long-run impact of the shock to a technology. The expenditure on the shocked technology is 1/5 of GDP, and hence the long-run impact of a 10 percent shock to that technology is approximately 1/50th of GDP. The adjustments of equilibrium flows, accounting for the 10 percent lower productivity of Intermediate R, are pictured on the right. These flows are rounded to two decimal places. Equilibrium GDP, to three decimal places, is 0.978, so slightly more than 1/50th of GDP is lost in the long run.

more dramatic and depend on the structure of the supply network. This answers the point made by Larry Summers in the 2013 speech that was quoted in the introduction, highlighting the difference between the long and short run. Short-run supply disruptions can be substantial even for items whose costs are a small fraction of GDP, as we show in this section.

We define the short-run to be such that if there is a shortage of an input, it is rationed so that each customer suffers the same percentage shortfall in supply, and there are no other (compensating) adjustments to the inputs. (In Section 5, we return to the medium run in which there are shortages but adjustments in resourcing.) To measure the impact of a disruption of some technology, we examine the most that each technology can still produce given the shortage of some input(s) that it faces. Each technology that sources an input produced by a shocked technology is thus shocked, and so we also trace those impacts as they propagate downstream. Suppose, for example, that a single technology is used to produce a given final good and the supply chain for this final good technology involves a string of single-sourced inputs without any cycles. Then, the impact of an X percent drop in the output of an upstream good is to reduce the production of the good downstream of it in the supply chain by X percent, and this propagates directly down the chain and disrupts the output of the final good by X percent.

Although the short-run impact in Figure 6 is clear in how it works as the shock propagates directly downstream, there can be three complications in more intricate settings. First, some technologies downstream of the shock might source the same input from multiple suppliers, and if not all of those sources are shocked, then the reduction can be less than X percent. Second, the same good may be produced using different technologies that require different inputs, only

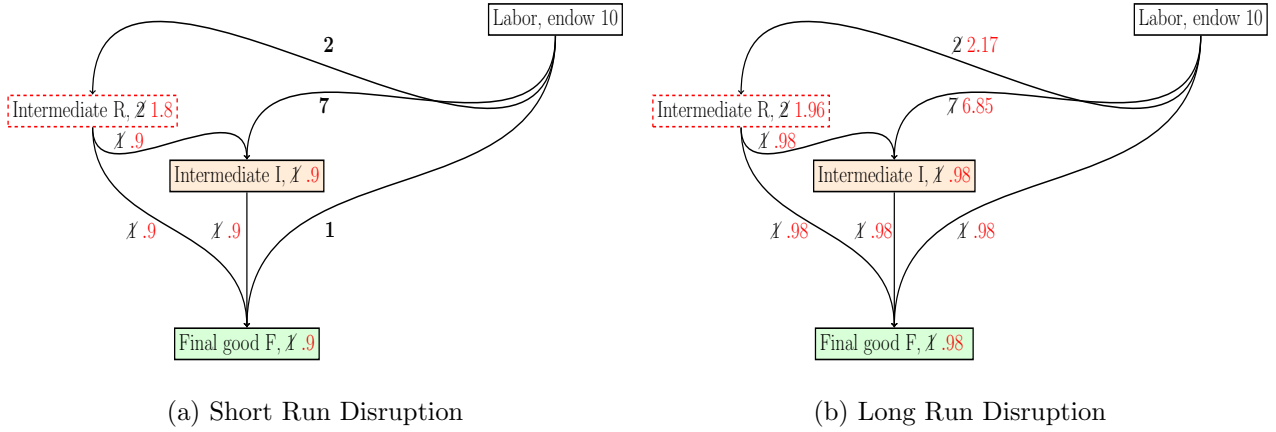


Figure 6: An example of the short-run impact of the shock to a technology, and the contrast to the long run. A 10 percent disruption of the production propagates through the network to the final good. Even though labor is not disrupted, it cannot produce the outputs without the corresponding inputs and so final good disruption is disrupted to the full extent of the input disruption. The disruption is 5 times larger than the corresponding long-run impact from Figure 5. In the long run, labor reallocates to even out the production needed as inputs downstream.

some of which are affected. Third, if there are cycles in the supply network, then these can feed back leading to repeated reductions that can amplify the effect of the shock in a way that helps negate the impact of sourcing and technological diversity.

The full impact of the disruption in an arbitrary supply network is the solution to an appropriate “minimum disruption problem.” We show below how the solution to this problem is found by an intuitive algorithm.

Starting from an equilibrium outputs and flows $(y_\tau)_\tau$ and $(x_{\tau\tau'})_{\tau\tau'}$, respectively, let $T^{Active} := \{\tau \in T : y_\tau > 0\}$ be the set of active technologies. Consider a shock that reduces the outputs of technologies $T^{Shocked} \subseteq T^{Active}$ to $\lambda < 1$ of their initial level. In the short run, outputs and flows adjust to $(\hat{y}_\tau)_\tau$ and $(\hat{x}_{\tau\tau'})_{\tau\tau'}$, which are the solution to the following minimum disruption problem:

$$\max_{(\hat{x}_{\tau\tau'})_{\tau\tau'}} \sum_{\tau: O(\tau) \in F} p_\tau \hat{y}_\tau$$

subject to

$$\widehat{y}_\tau \leq \lambda y_\tau \quad \text{for all } \tau \in T^{Shocked}, \quad (\text{shock constraints})$$

$$\widehat{y}_\tau \leq \left(\min_{k:\tau_k < 0} \frac{\sum_{\tau': O(\tau')=k} \widehat{x}_{\tau'\tau}}{\sum_{\tau': O(\tau')=k} x_{\tau'\tau}} \right) y_\tau \quad \text{for all } \tau \in T^{Active}, \quad (\text{technology constraints})$$

$$\widehat{y}_\tau = y_\tau = 0 \quad \text{for all } \tau \notin T^{Active}, \quad (\text{technology switching constraints})$$

$$\widehat{x}_{\tau'\tau} = x_{\tau'\tau} \left(\frac{\widehat{y}_{\tau'}}{y_{\tau'}} \right) \quad \text{for all } \tau', \tau \in T^{Active}. \quad (\text{proportional rationing})$$

The minimum disruption problem defines the maximum final production that can still be produced, subject to the reduced output of directly-shocked technologies, as well as technology constraints that do not allow new sources of inputs and are based on proportional rationing constraints that equally spread the impact of each technology's reduced production among its customers.

Proportional rationing applies in the short run in which firms are not able to seek out new suppliers, renegotiate contracts, or switch technologies to substitute disrupted inputs for non-disrupted ones. In the medium run prices, and thus the rationing, adjust. We explore this in Section 5, where we consider a time frame before new sources of production or alternative technologies emerge; but in which prices can redirect inputs to the most-valued downstream final good production chains.

The solution to the minimum disruption problem is found by the following (standard and intuitive) *Shock Propagation Algorithm*:

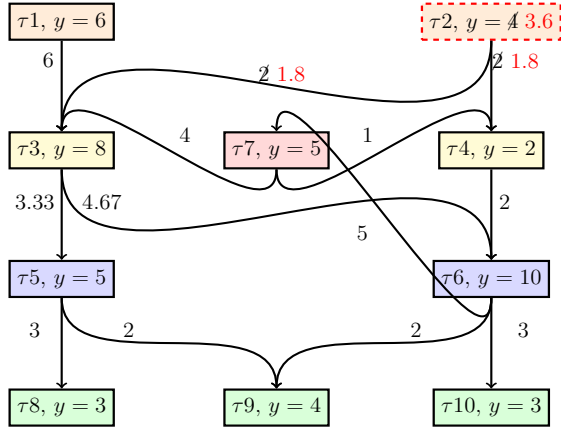
- Begin with the equilibrium flows and final good outputs $(x_{n\tau'})_{n\tau'}, (x_{\tau\tau'})_{\tau\tau'}, (y_\tau)_{O(\tau) \in F}$
- Reduce the outputs of the shocked technologies $\tau \in T^{Shocked}$ by $(1 - \lambda)$, so $y_\tau^1 = \lambda y_\tau$ and $x_{\tau\tau'}^1 = \lambda x_{\tau\tau'}$, and leave other flows and outputs are unchanged (so, for $\tau \notin T^{Shocked}$, $y_\tau^1 = y_\tau$ and $x_{\tau\tau'}^1 = x_{\tau\tau'}$ for all τ').

- Iteratively in k :

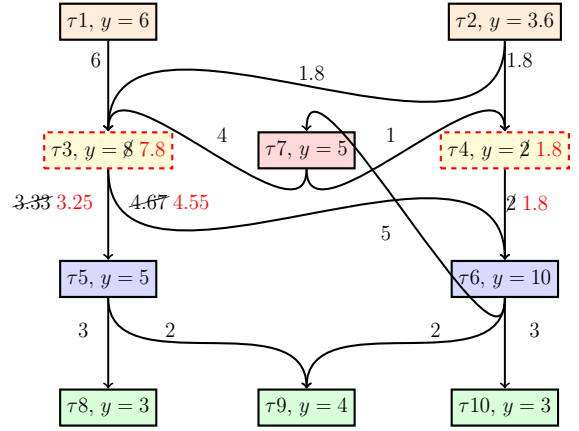
$$- \text{ let } y_{\tau'}^k = y_{\tau'} \left(\min_{\tau'': x_{\tau\tau''} > 0} \frac{x_{\tau\tau''}^{k-1}}{x_{\tau\tau''}} \right) \text{ and } x_{\tau'\tau''}^k = \frac{y_{\tau'}^k}{y_{\tau'}} x_{\tau'\tau''} \text{ for all } \tau''.$$

- Iterate until $y^{k-1} = y^k$ (or $\max_\tau (y_\tau^k - y_\tau^{k-1}) \leq \Delta$ for some threshold $\Delta \geq 0$).


At each stage, we examine nodes downstream from any shocked node and calculate what it can produce given the available inputs. Any nodes that have outputs change are then shocked, and we repeat the calculation. In particular, we calculate, for each affected input, the proportion of the original supplied level of that input type which can still be successfully sourced from the shocked flow network. From this, we calculate how much the output of each node declines, and continue to trace affected technologies. The overall impact on the economy is then given by the value of lost final good production resulting from the shock(s).




(a) 10 percent shock to τ_1




(b) Shock propagates to τ_3, τ_4




(c) Shock continues to τ_5, τ_6



(d) Shock hits $\tau_7, \tau_8, \tau_9, \tau_{10}$



(e) Shock further disrupts τ_3 but not τ_4



(f) Algorithm converges to these flows

Figure 7: An illustration of the Shock Propagation Algorithm on the example from Figure 4.

The Shock Propagation Algorithm is illustrated in Figure 7 in which there is a 10 percent shock to technology τ_2 . It traces the impact of the shock and updates the output of a node each time the supply of one of its inputs is reduced. In this example there are several cycles, and note that τ_6 and τ_7 are on cycles involving τ_3 and τ_4 . The feedback via τ_4 stops after a few steps in the algorithm since the diminished flow directly from τ_2 ends up being the binding one. However, the feedback via τ_3 continues infinitely, and eventually converges to the final levels pictured in panel (f).

If there are no cycles, the shock propagation algorithm terminates in a finite number of steps at a fixed point which is the unique solution to the minimum disruption problem. If there are cycles, the algorithm may not terminate in finite time. However, in that case, it still converges to the fixed point which is a solution to the minimum disruption problem. This is formalized in Proposition 2.

PROPOSITION 2. *The Shock Propagation Algorithm (with $\Delta = 0$, which may continue ad infinitum) converges to a flow of goods that is weakly lower for all links, and strictly lower for all links on a directed path from any shocked node. The limit output vector solves the minimum disruption problem.*

Proposition 2 shows that the Shock Propagation Algorithm converges to a solution of the minimum disruption problem. The uniqueness of the limit, and that it solves the minimum disruption, follow from several facts. First, the algorithm iterates on the flows between technology pairs, and these flows can be represented by vectors. Bounding each entry of the vector from above by the initial equilibrium flow, and from below by 0, the vectors can then be partially ordered such that one vector is weakly ordered above another when all flows are weakly higher. Moreover, this partially ordered set is a complete lattice. Second, reductions in inputs in the short run are complementary to each other, so one iteration of the algorithm outputs weakly higher flows when the initial flows are weakly higher. As such, an iteration of the algorithm is an isotone (monotonic) function, and Tarski's fixed point theorem tells us that the fixed points of this mapping form a complete lattice (given the same partial order). Third, the algorithm terminates at/converges to a fixed point in which no more reductions are necessary, and every reduction implemented by the algorithm is necessary. Thus the highest possible production levels that satisfy all the constraints is the limit, and the algorithm finds a solution to the minimum disruption problem. The full proof, including the infinite case appears in the appendix. It is worth noting that flows that solve the minimum disruption problem solution are not uniquely tied down: there could be several flows that are changing in combination and only some of them end up binding the production, and so nonbinding flows can be lowered and not change overall production. One can check directly that convex combinations of flows that solve the problem also solve the problem, and hence given that the flows that can solve the problem are bounded above and below, the set of flows that solve the problem form a complete lattice, of which the Shock Propagation Algorithm finds the maximum.

We have specified the minimum disruption problem to minimize the lost value of final goods at the initial equilibrium prices. One might wonder whether the new long-run equilibrium prices

should be used instead, or perhaps some other prices. It turns out that the flows that solve the minimum disruption problem are invariant to such choices. This is because proportional rationing is imposed in the problem, and this determines the final good levels. In Section 5 we show that this can change when that proportionality constraint is removed.

OBSERVATION 1. *Flows that solve the minimum disruption problem remain a solution as prices in the objective function are changed, and the total production of each final good is the same.*

Given the monotonicity of the algorithm, from it we can also deduce an upper bound on the impact of a shock that is to easy-to-calculate, applies in many cases of interest, and contrasts starkly with Hulten’s Theorem.

Consider a shock to some subset of the technologies $T^{Shocked}$. Let $F(T^{Shocked})$ denote the set of final good technologies that lie downstream (on a directed path) from a shocked technology, so either (i) are in $T^{Shocked}$ directly or (ii) use an input good from a technology (directly or indirectly) that is shocked. These are the final goods that ever change at some stage k of the algorithm. Thus $F(T^{Shocked})$ gives the set of all final good technologies that are impacted by the shock.

PROPOSITION 3. *Consider a shock that reduces the output of all technologies $k \in T^{Shocked}$ to $\lambda < 1$ of their original levels. Then the proportion of GDP that is lost to this shock is bounded above by*

$$(1 - \lambda) \left(\frac{\sum_{\tau \in F(T^{Shocked})} \sum_{\tau} p_{\tau} y_{\tau}}{GDP} \right).$$

Proposition 3 provides an upper bound on the impact of a shock that is tied to the value of all final goods related to the shocked technologies, rather than the cost of production of the shocked technologies, which is what matters in the long-run by Hulten’s Theorem. We turn now to identifying sufficient conditions under which the upper bound is achieved.

Consider the sub-network which describes the supply chains of all final goods that are impacted by the shock. Specifically, the *disrupted industries sub-network*, $\mathcal{G}(T^{Shocked})$, is the sub-network induced on \mathcal{G} by all technologies that are on a directed path that terminates at a final good technology in $F(T^{Shocked})$.

When a *disrupted industries sub-network*, $\mathcal{G}(T^{Shocked})$, is acyclic, there exist a set of nodes (technologies) that have no in-links. We denote these technologies $R(T^{Shocked})$. In this case, it is helpful to introduce the notion of a (s,t) -cut set. For a directed network with disjoint sets of nodes s and t , an (s,t) -cut set is a set of edges that when removed from the network there are no remaining paths between s and t .

PROPOSITION 4. *If the disrupted industries sub-network $\mathcal{G}(T^{Shocked})$ is acyclic and the set of edges adjacent to the set of shocked technologies $T^{Shocked}$ forms an $(R(T^{Shocked}), F(T^{Shocked}))$ -cut set of the disrupted industries sub-network $\mathcal{G}(T^{Shocked})$,¹⁴ then the bound from Proposition 3 binds.*

¹⁴Note that this permits the set of shocked technologies $T^{Shocked}$ to intersect both $R(T^{Shocked})$ and $F(T^{Shocked})$.

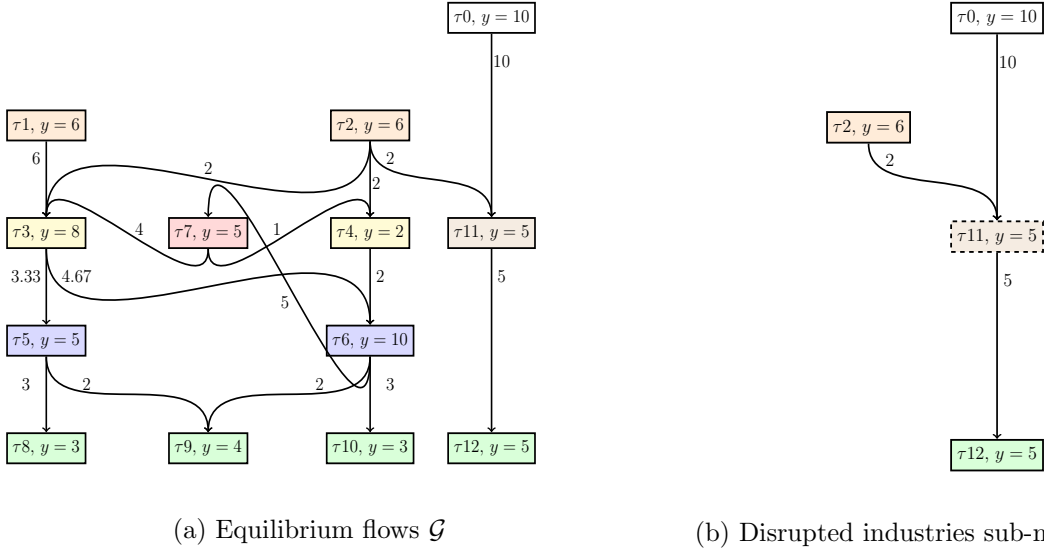


Figure 8: Disrupted industries sub-networks: For the equilibrium flows in panel (a), panel (b) shows the disrupted industries subnetwork for $T^{Shocked} = \{\tau_{11}\}$. The disrupted industries sub-network for $T^{Shocked} = \{\tau_2\}$ is the network shown in panel (a), while the disrupted industries sub-network for $T^{Shocked} = \{\tau_1\}$ is the same as the equilibrium flows network shown in Figure 4.

Proposition 4 shows that the upper bound for the impact of a shock from Proposition 3 is tight when the technology network is acyclic and removing the edges adjacent to the shocked technologies $T^{Shocked}$ constitutes a cut on the network that restricts attention to related goods (i.e., disconnects raw materials from their related final goods). The example shown in Figure 8b illustrates a shock which, by Proposition 3, must obtain the bound.

Figure 7 shows that the bound can be obtained when the shocked technologies do not constitute a cut set but the supply network contains cycles. Further, a variation on Figure 6(a) shows that even when the technology network is acyclic, the bound can be obtained when the edges adjacent to the shocked technologies do not constitute a cut set. To see this simply relabel the Labor node to be a technology, and the shocked technologies do not constitute a cut set.

We say there is *no technological diversity* if all technologies for producing any given good use the same set of inputs (albeit possibly in different ratios or with different efficiencies). We say there are *industry-wide shocks* (as opposed to technology-specific shocks) if for any shocked technology $\tau \in T^{Shocked}$ all producers of good $O(\tau)$ are also in $T^{Shocked}$.

PROPOSITION 5. *If there is no technological diversity and there are industry-wide shocks, then the bound from Proposition 3 binds.*

Proposition 5 does not follow from Proposition 4 as it allows for cyclic production and also for cases in which the shocked technologies do not constitute a cut, as there could be other technologies that are not affected that are on separate paths to the affected final goods.

The proof is straightforward, and so we just sketch it. Consider a final good technology

$\tau \in F(T^{Shocked})$. By definition of $F(T^{Shocked})$, there is a path from a shocked technology $\tau' \in T^{Shocked}$ to τ . Consider that path, starting with a shocked technology τ' that is at maximal distance from a final good. The output of τ' is reduced to λ of its initial level as it is shocked. Consider the next technology τ'' on the path, which sources good $O(\tau')$ from τ' . As shocks are industry wide, all producers of good $O(\tau')$ are shocked, and so τ'' is only able to source λ of the initial amount of good $O(\tau')$ that it sourced, and thus its output will be reduced to λ its initial level. As all shocked inputs are shocked to λ of their initial levels, it does not matter if just one input or multiple inputs are shocked. Moreover, as there is no technological diversity, all producers of good $O(\tau'')$ also use good $O(\tau')$ as an input and are only able to source λ units per unit they initially source of this input. Thus all producers of good $O(\tau'')$ have their output reduced to λ their initial levels. This implies that the next technology on path, τ''' , also has its output reduced to λ of its initial level and so on. Thus the shock propagates down the path, reducing output to λ of its initial level at each step, until the output of final good $\tau \in F(T^{Shocked})$ has its output reduced to λ of its initial level. This implies that all final goods in $F(T^{Shocked})$ have their output reduced to λ of their initial level, and so bound from Proposition 3 is obtained.

Proposition 5 highlights the value of diversity in production technologies and the location of industries. For example, if all production of a given good is located in the same country, then country-specific shocks become industry shocks, and by Proposition 5 the upper bound is obtained.

3.3 Contrasting the Short- and Long-Run Impacts of Shocks

Comparing Figures 6a and 6b gave us some idea of the differences that can occur between the long and the short run. This message is further reinforced when the bound from Proposition 3 is tight. For example, for an industry-wide shock, the long-run (marginal) impact on GDP is, by Hulten's theorem, proportional to the value of the output of the affected industry. By contrast, if there is no technological diversity, the short-run impact on GDP is proportional to the value of the output of all final goods industries that use the output of the affected industry directly or indirectly by Proposition 5. If, for example, production of several final good depends on a basic type of computer chip, and there is a 20 percent disruption in the supply of these computer chips, then 20 percent of final good production would be lost for all affected final goods. This can constitute a substantial short run impact, while in the long run the impact on GDP would only be 20 percent of the amount spent on these basic computer chips, which could be tiny in comparison.

An important point to make is that the long-run impact of a shock is dependent only on the expenditures on the shocked technology, and not on the details of the network beyond that.¹⁵ To see this, consider an economy with two variations of the sets of technologies from

¹⁵Again, those expenditures are dependent on the network, so this is not to say that the network is irrelevant. It simply says that the information needed to determine the impact is captured by a very simple sufficient statistic.

the example in Figures 1, 5, and 6, but, for simplicity, let us omit the labor inputs. As we see in Figure 9, the long-run impact of the 10 percent shock to the technology in the upper left is the same regardless of whether which combination of technologies is used to produce which final good. The new long-run equilibrium differs across the two networks, but the ultimate impact of the shock on GDP does not.




Figure 9: Long run: In both cases have supply networks that have two copies of technologies similar to those in the example from Figure 1. Each final good needs one resource and one intermediate good, but which combination of inputs are needed downstream differs between the networks. In the long run the details of the network structure do not matter if the amount spent on the shocked technology is the same. In both cases the labor endowment is 20, the initial prices are $p = (\frac{1}{10}, \frac{1}{10}, \frac{4}{5}, 1)$ and $GDP = \sum_f p_f c_f = 2$. Thus, from Hulten’s Theorem, the marginal impact is $\frac{\partial R1YR1}{GDP} = \frac{1}{10}$ and then extrapolating for a 10% shock, the *long-run* impact is 1/100th of GDP. We do see, however, that the new long-run equilibrium flows differ across the two variations, but the GDP impact is similar.

In contrast, the short-run impact in the same economies differs substantially depending on the details of the network as we see in Figure 10.




Figure 10: Short run: the details of the network structure matter even with identical initial prices and technological structures. Here the impact is either 5 or 10 times more than the long-run impact (which was 1/100th), and here it depends on the network structure.

Although our short-run calculation is accurate for non-marginal shocks, unlike Hulten’s theorem which only applies at the margin, it is also instructive to compare the short-run

marginal impact of a shock to the long-run marginal impact. For example, if the bound from Proposition 3 is obtained, the short-run impact of a disruption to some technology τ with output good k is

$$\frac{\partial \log(U)}{\partial \log(\tau_k)} = \frac{\partial \log(GDP)}{\partial \log(\tau_k)} = \frac{\sum_{f \in F(\tau)} p_f y_f}{GDP},$$

while by contrast the long-run (marginal) impact is, by Hulten’s Theorem,

$$\frac{\partial \log(U)}{\partial \log(\tau_k)} = \frac{\partial \log(GDP)}{\partial \log(\tau_k)} = \frac{p_\tau y_\tau}{GDP},$$

illustrating the large difference that is possible, given the huge potential difference between the value of all affected final goods $\sum_{f \in F(\tau)} p_f y_f$ compared to the cost of the affected input $p_\tau y_\tau$.

We remark on a couple of implications of these comparisons. In the short run, the impact of a disruption is not dependent upon how expensive an input is, but instead by the value of all final goods that lie downstream. In the long run, it is the reverse. That does not mean that the long run impact is independent of how upstream or downstream a good is. Goods that are nearer to final goods and incorporate more inputs from upstream will be more expensive, all else held equal, and so the long run disruption of them is more costly. So, how upstream or downstream a good is makes a difference in the short run since that might affect how many final goods it reaches, while how upstream or downstream a good is makes a difference in the long run since that affects how expensive it is. Roughly, shocks to goods that are more upstream are more disruptive in the short run since they affect more final goods, while shocks to goods that are more downstream tend to be more disruptive in the long run since they are more costly. The details are given by the formulas, but these rough intuitions are still useful to note.

Although we have so far discussed the implications of negative shocks, the Shock Propagation Algorithm can also be applied to consider positive shocks. Consider a positive shock to technology τ that increases output to a proportion $\lambda > 1$ of its initial equilibrium level. If $O(\tau)$ is a final good, the impact on final good production is an immediately apparent proportional increase in production of technology τ . On the other hand, if $O(\tau)$ is an intermediate good and all producers that source from τ use intermediate goods other than just $O(\tau)$, then there is no impact on final good production. This again contrasts with the long-run impact of such changes as captured by Hulten’s Theorem, where the (marginal) impact of positive and negative shocks is symmetric.

3.4 Disruption Centrality

Beyond cases in which the intuitive bound is tight, we can calculate the exact impact of a disruption in a much wider set of cases, even when the bound is not tight. In this subsection, we use our model to estimate the impact of the disruption of any given technology under a condition that there is a (weak) ordering of upstream to downstream goods. Such an ordering rules out cycles. We call the impact “disruption centrality,” which we now define.

Let $S(f, \tau')$ denote the percentage of final good f that is produced by technology τ' . Let $S(\tau', \tau'')$ denote the fraction of $O(\tau'')$ that τ' sources as an input from τ'' . For most pairs this will be 0, but for a technology that sources from another it will be positive.

Let us say that a set of technologies T' is *well-ordered* if there is an order \succ over the technologies such that, for any $\tau, \tau' \in T'$, if $\tau \succ \tau'$ then output good $O(\tau)$ is not used as an input by τ' . Thus technologies higher in the \succ order are further downstream in a clearly defined sense. This rules out cycles, but still allows upstream goods to be used in the production of multiple downstream goods and so for rich structures. In general when a set of technologies is well-ordered, then there will be many such orders, as some goods are never used up or down stream from each other.¹⁶

Consider a technology $\tau \in T$ that has equilibrium paths to some set of final goods $F(\tau)$, and consider some $f \in F(\tau)$. Let $T_{\tau,f}$ be the set of all technologies that lie on any equilibrium path between τ and some τ_f producing f (inclusive). Suppose that $T_{\tau,f}$ is well-ordered with corresponding order $\succ_{\tau,f}$. Let $G_{\tau,f}$ be the set of goods that are produced by any $\tau' \in T_{\tau,f}$, and number them 1 to K_f , ordered according to $\succ_{\tau,f}$ with good 1 being $O(\tau)$ and good K_f being f .

Let $d_{\tau,f}(\tau) = 1$ and let $d_{\tau,f}(\tau') = 0$ for all $\tau' \notin G_{\tau,f}$. Inductively, in all goods in $G_{\tau,f}$, $k \in \{1, \dots, K_f\}$, consider $\tau' \in T_{\tau,f}$ producing good k . Let

$$d_{\tau,f}(\tau') = \max_{i \in I(\tau')} \left[\sum_{\tau'': O(\tau'')=i} d_{\tau,f}(\tau'') S(\tau', \tau'') \right].$$

Then

$$D(\tau) \equiv \frac{\sum_{f \in F(\tau)} p_f y_f \left(\sum_{\tau': O(\tau')=f} d_{\tau,f}(\tau') S(f, \tau') \right)}{GDP}$$

is the disruption centrality of τ .

In particular, if the output of τ is reduced by some amount γ , then the total short run fraction of GDP that is lost is exactly $\gamma D(\tau)$.

If all technologies are single-sourced (i.e., $S(\cdot, \cdot)$ is always 0 or 1) then $D(\tau) = \sum_{f \in F(\tau)} p_f y_f$, which corresponds to the upper bound we identified before. However, more generally $D(\tau)$ captures all the downstream disruptions accounting for all the fractions and multiple paths that lie between some technology and all the final goods that are produced downstream from it.

Figure 11 provides an example of how the $d_{\tau,f}(\tau')$ terms are calculated for the equilibrium flows shown in Figure 8(a) with the disrupted technology being $\tau = \tau_2$. If the equilibrium prices of the final goods are all 1 (which will depend on the labor inputs of the different technologies),

¹⁶For example, if the output of technology τ is good 1, and it is used by technologies τ' and τ'' to produce goods 2 and 3 respectively, which are then used by technology τ''' to produce the final good, then the orderings $\tau''' \succ \tau' \succ \tau'' \succ \tau$ and $\tau''' \succ \tau'' \succ \tau' \succ \tau$ are both valid.

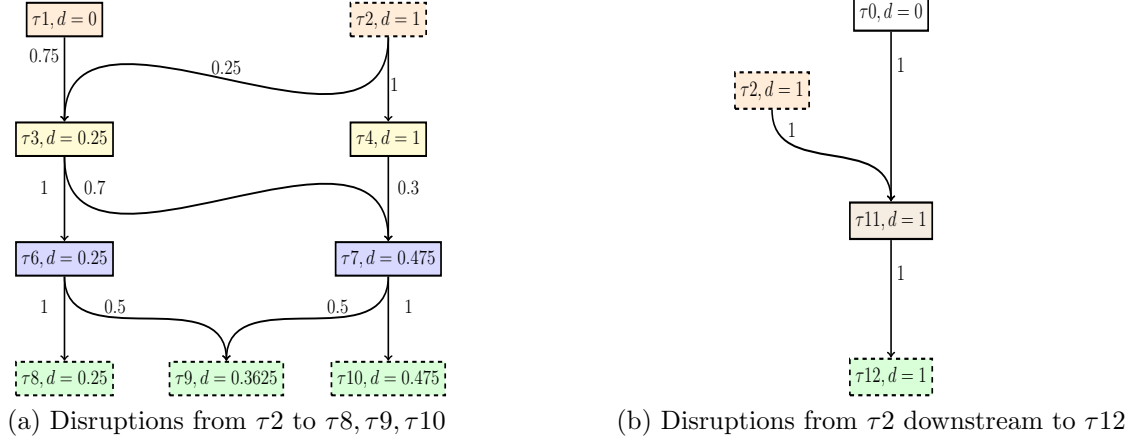


Figure 11: Consider the equilibrium flows shown in Figure 8(a). Technology τ_2 affects parts of the network, that leading to τ_8, τ_9 and τ_{10} , as well as that produced by τ_{12} . Panel (a) shows the disruptions for goods between τ_2 and τ_8, τ_9 and τ_{10} , as well as for good τ_1 as this is needed for the recursive calculations. Panel (b) shows the disruptions for goods between τ_2 and τ_{12} as well as for technology τ_0 . In both panels the weight on an edge from τ' to τ'' represents $S(\tau', \tau'')$.

the disruption centrality of technology τ_2 for these equilibrium flows is

$$D(\tau_2) = \frac{10}{15} \left(0.25(0.3) + 0.3625(0.4) + 0.475(0.3) \right) + \frac{5}{15} = 0.575$$

4 Complexity, Fragility, and Globalization

4.1 Supply Chain Complexity and Increased Fragility

Next we examine how the impact of a disruption depends on the complexity of a supply chain. We show that shocks have more impact as production becomes more complex and relies on more intermediate goods, all else held equal.

Consider a generic setting in which the bound from Proposition 3 is obtained (e.g., the setting of Proposition 5), and let us consider a random shock. Let S denote the average number of different (non-final) technologies used directly or indirectly to produce a final good—averaged across final goods. Thus, S is a measure of the complexity of supply chains. Let ΔGDP denote the initial equilibrium GDP minus the GDP after any shocks. Finally, let q denote the ratio of the average expenditures spent on a randomly picked (non-final) technology to the average expenditures on a randomly picked final good, and m denote the expected number of final goods that lie downstream from a randomly picked (non-final) technology.

PROPOSITION 6. [Supply Chain Disruption as a Function of Complexity] *Consider a generic economy in which the bound from Proposition 3 binds, final goods prices are independent of the*

complexity of their supply chains. Let disruptions to intermediate good technologies—in which a proportion $(1 - \lambda)$ of output is lost—occur with probability π , independently of other disruptions. Then for small π (so that the probability of two or more shocks on the same chain is vanishingly small relative to the probability of a single shock) and small $1 - \lambda$ (so that Hulten’s Theorem applies):

$$\begin{aligned} \text{Short-Run:} \quad \mathbb{E} \left[\frac{\Delta GDP}{GDP} \right] &\approx -(1 - \lambda)\pi S, \\ \text{Long-Run:} \quad \mathbb{E} \left[\frac{\Delta GDP}{GDP} \right] &\approx -(1 - \lambda)\pi S \frac{q}{m}. \end{aligned}$$

Some intuition behind Proposition 6 is as follows. As supply chains become more complex, more intermediate goods are used (directly or indirectly) in the production of final goods, and so the probability of short-run of disruptions increases. In the long-run, complexity still matters because it still increases the probability of disruption. However, this is now modulated by the size of the impact which is proportional to the average cost of a shocked intermediate good to the value of final goods it affects, as we would expect by Hulten’s Theorem.

The proof of Proposition 6 is straightforward and so we simply outline it here. First, note that $m = SF/M$. Next, given that $M\pi$ is expected number of disruptions, and each hits a fraction of m/F of the total set of goods, it follows that $\pi Mm/F = \pi S$ is the expected fraction of final goods shocked and the expected probability of disrupting a typical final good (appealing to the fact that when π is small the probability that more than one technology is shocked at a time is vanishingly small compared to a single technology being shocked, and so we do not worry about multiple shocks to any supply chain). In the short run, a disrupted chain reduces consumption of the final good by a fraction $(1 - \lambda)$, and so disrupting a typical final good by $(1 - \lambda)$ with a probability πS gives the short-run result. For the long-run result, the expected number of disruptions is approximately $M\pi$ and the expenditures on the technology compared to GDP is q/F . By Hulten’s Theorem the expected marginal value is approximately $M\pi q/F = \pi S q/m$.

The contrast of short and long-run disruptions is again stark. If supply chains are completely horizontal, so that inputs go directly into final goods, then $q \approx 1/S$ (ignoring labor cost in assembling final goods, which are otherwise added to the denominator) and $m = 1$, and so the long-run effect is approximately $(1 - \lambda)\pi$. This is $1/S$ of the short-run effect. If supply chains are completely vertical and $m = 1$, and inputs use similar amounts of labor, then $q \approx 1/2$, and so the overall effect is approximately $(1 - \lambda)\pi S/2$. If there are parallel supply chains so that each input reaches a different final good, then the impact of the long run is like the horizontal case, while the short run impact is more compartmentalized.

This shows that there are systematic ways in which network shape affects both the short- and long-run impacts, and that they depend on different features of the network. The long-run impact is more dependent upon depth vs breadth, while the short-run impact is more dependent upon how many final goods given inputs are upstream from.

The contrast as a function of the supply network shape is illustrated in Figure 12. We see




Figure 12: Three different supply chain configurations and the corresponding impacts. In each case the endowment of labor is 5 units and each input technology uses one unit of labor, with the remaining unit of labor used in the production of the final good(s). In the vertical and horizontal supply chains the complexity is $S = 4$ while it is $S = 1$ in the parallel case. The number of final goods downstream of any input is $m = 1$ in each case, and the corresponding q_s are .5, .2, .8. The corresponding price vectors are $p_{vertical} = (\frac{1}{5}, \frac{1}{5}, \frac{2}{5}, \frac{3}{5}, \frac{4}{5}, 1)$, $p_{horizontal} = (\frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5}, 1)$, $p_{parallel} = (\frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$.

that the short run is the same regardless of whether the supply chain is horizontal or vertical, since each input still disrupts the final good. By contrast, when the supply chains are parallel with the same number of inputs but used for four different final goods, then the short-run disruption is lowered. The long-run disruption shows very different patterns. Not only are those disruptions much smaller, but here they differ between the horizontal and vertical, but not between the horizontal and parallel, as vertical supply chains build up costs, while the input costs are the same across the horizontal and parallel.

4.2 Globalization

We next show that globalization, modeled as reduced transportation costs, can lead to the conditions hypothesized in Proposition 5. The idea is simple. As transportation costs become sufficiently low, technological diversity decreases and each good becomes single-sourced, and technology shocks become industry shocks.

To formalize these ideas we first show that how specialization increases with reduced costs, which we then trace through to resulting changes in the equilibrium production network and how they affect overall fragility.

PROPOSITION 7 (Specialization). *If T is generic and transportation costs are sufficiently low, then there is full specialization: there exists a threshold on transportation cost $\bar{\theta} > 1$ such that if $\max_{\tau, \tau'} \theta_{\tau\tau'} < \bar{\theta}$ and $\max_{n, \tau'} \theta_{n\tau'} < \bar{\theta}$, then $y^\tau > 0$ and $y^{\tau'} > 0$ implies that $O(\tau) \neq O(\tau')$.*

With low enough costs to shipping, only the cheapest technology for any good is used. Generically, this is a unique technology. As, generically, no two countries have access to exactly equivalent technologies the proposition says that there is a unique technology used to produce

each good and each good is produced in just one country. Although outside of our model (which has a given set of final goods), it could be that more final goods are produced as transportation costs decrease, as some technologies might be too expensive to be part of an equilibrium in the face of higher transportation costs.

If two countries have identical technologies, then they could both survive in equilibrium for any small transportation costs, so the genericity condition is important for the full specialization equilibrium to emerge. Although this is clearly a benchmark result, it captures the idea that as shipping costs drop, the most efficient production technologies can displace other technologies and specialization in technologies can emerge. If one extends Proposition 7 to allow multiple countries have access to *exactly* the same technology, it could be that the same goods are produced in more than one place, but still all producers would use the same technology for that production. This still leads to fragility in the short term with respect to technology-specific shocks.

COROLLARY 1 (Fragility). *If T is generic and transportation costs are sufficiently low, then there is full specialization: there exists a threshold on transportation cost $\bar{\theta} > 1$ such that if $\max_{\tau, \tau'} \theta_{\tau\tau'} < \bar{\theta}$ and $\max_{n, \tau'} \theta_{n\tau'} < \bar{\theta}$, then each final good has a supply chain in which all goods are single-sourced and the bound from Proposition 3 binds.*

The fact that the bound from Proposition 3 binds can be seen as follows. Every affected final good technology has a path from some shocked technology to it. Since goods are all single-sourced, each technology on the path between an affected final good technology and a shocked technology is fully reduced by $(1 - \lambda)$. Intuitively, as all transportation costs decrease sufficiently we get specialization (Proposition 7), which eliminates technological or sourcing diversity and turns technology-specific shocks into industry-wide shocks. Thus, by Proposition 5, the bound from Proposition 3 binds.

Another implication of Propositions 3 and 7 is that moving to a frictionless economy changes fragility. In fact, more generally consolidation of supply chains changes fragility.

PROPOSITION 8. *Consider two economies, indexed 1, 2, that share the same technology sets T^1, \dots, T^n and have equilibria that satisfy the bound from Proposition 3 from the shock to any active technology. Consider some final good f that is produced in equilibrium by both economies by some technology τ_f , and suppose that the equilibrium output of technology τ_f is higher in economy 2 than 1. If the set of technologies that lie on a directed path to τ_f is smaller in economy 2 (i.e., $\mathcal{G}^2(\tau_f) \subsetneq \mathcal{G}^1(\tau_f)$), and shocks are independent across technologies with the same proportional disruption, then the probability of a disruption to τ_f is lower in economy 2 than economy 1, but the expected short-run size of that disruption conditional upon occurrence is higher in economy 2 than 1.*

Proposition 8 is straightforward to prove, and so we omit its proof.¹⁷ Nonetheless, Proposition 8 is important to state, since it has as a corollary that a lowering of trade costs that leads

¹⁷The only nuance is that there is a higher chance in economy 1 that multiple upstream technologies are shocked simultaneously, but given that the bound holds, then whether one or more technologies are shocked by the same amount, the consequences are identical.

to consolidation of supply chains and larger outputs of final goods from some single technology, has this implication. There are fewer technologies upstream from the final good that might fail, but greater consequences when any of them are shocked.

The trade off between the probability of a disruption and the size of the disruption is a fundamental point that applies to increasingly specialized supply chains. Of course, this holds all else constant. As new technologies emerge in production due to globalization and supply chains cross more borders there are increased probabilities of political disruptions and any issues that disrupt shipping. This changes the overall distribution of shocks and so opens the possibility that one could see both an increase in the probability of disruption *and* the size of the disruption.

5 Flexible Prices: Shock Impacts in the Medium Run

While in the short run existing contracts may prevent prices from adjusting, as embedded within our proportional rationing assumption, once prices can adjust existing production should be rerouted to minimize the overall impact of the shock. We now consider this possibility, while at the same time fixing the production technologies of different firms and preventing firms from increasing their output beyond pre-shock levels.

Flexible prices can correlate downstream disruptions in a way that minimizes the overall disruption. Consider, for example, the short run impact of a shock to producer τ_1 (without flexible prices) shown in panel (a) of Figure 13.

Here a disruption to τ_1 affects the production of the final goods τ_4 through two potential channels. First τ_1 is used directly as an input, and secondly it is used to produce τ_3 which is an input for τ_4 . In this example τ_1 and τ_2 produce the same good, and so by redirecting some of the output of τ_1 to τ_3 there is less disruption of that good which is critical to the production of τ_4 , while there is a substitute for τ_1 in the direct input for τ_4 .




Figure 13: How price adjustments can partly mitigate shocks. Here τ_1 and τ_2 are producing the same goods, and so can be substituted for each other.

A second example is illustrated in Figure 14. Suppose labor is priced at 1, one unit of labor is needed to make one unit of the intermediate good R_1 , 1 unit of R_1 and no units of labor are

needed to make final good $F1$, while $1/9$ th of a unit of $R1$ and $8/9$ ths of a unit of labor are needed to make a unit of final good $F2$. Thus both $F1$ and $F2$ are priced at 1. Suppose there are 10 units of labor, and at these prices, labor demands 9 units of good $F2$ and 1 unit of $F1$. Thus the overall production by $F2$ is much more valuable than the overall production by $F1$: $F2$'s output contributes 9 to a GDP of 10, while $F1$'s output contributes 1.

Consider now a 10% total factor productivity shock to $R1$. In the short run, proportional rationing leads the output of both final goods to be reduced by 10%, and GDP is correspondingly reduced by 10% (achieving the short run upper bound). However, if we relax the proportional rationing constraint and allow intermediate good to flow to its most valuable use, then production of good $F2$ will not be affected at all, while output of $F1$ will decrease by 20%. However, because production of good $F2$ is substantially more valuable, the reduction in GDP is now just 2%.

For a real world example where such a reallocation would have been valuable, consider the semiconductor / computer chip crisis. In the short run the production of very valuable downstream goods like cars was disrupted, along with many other consumer goods. With flexible prices, in the medium run, the disruptions should become concentrated on less valuable consumer goods (like, for example, cheap toys). The extent to and speed with which this happened is an empirical question of interest.

Another illustrative example along these lines is the rolling electricity blackouts in California in the summer of 2020. This was applied without discrimination (save some emergency services and legal restrictions) and so fits our proportional rationing assumption. However, in anticipation of further future blackouts, more complicated contracts have emerged that result in priorities in rationing and corresponding differences in prices, with electricity being allocated to its highest value uses when there are shortages. Again, it is an empirical question regarding the extent to which efficient contracts have been put in place.




Figure 14: The benefits of pricing flexibility. If the downstream goods have equal values the short run and medium run outcomes are the same and 10% of output is lost in the medium run (Panel (a)). However, the downstream goods have different values, then the adjustment in the medium run is helpful as it can direct the disruption down the least valuable path. Now in the medium run only a 2% of output is lost, versus 10% in the short run that instead propagates equally down both paths (Panel (b)).

It is immediate that the output loss in the medium run (with flexible prices) is weakly less than the output loss in the short run, because we are relaxing the proportional rationing con-

straint. However, there are occasions on which it will make no difference. If the conditions for Proposition 5 hold, such that there is no technological diversity and there are industry specific shocks, then there will be no scope for reducing the impact of the shock by correlating the downstream instances of the disruption—all downstream producers will be equally impacted under proportional rationing, with the same shortages of the same inputs inhibiting their production. On the other hand, it is possible to construct examples in which price flexibility makes a very big difference. Specifically, the ratio of lost output in the short run to lost output in the medium run with price flexibility is unbounded (see Supplementary Appendix II).

In this section we have considered a minimal relaxation of our formulation for the short run impact of a shock by allowing prices to adjust. A next step would be to allow further flexibility in the economies response to a shock to consider what would happen over a longer time horizon. A natural way to do this would be to start to relax the no increased output constraints for used technologies. In a somewhat different setting, this is problem studied in [Carvalho, Elliott, and Spray \(2024\)](#). They define exogenous capacity constraints on outputs and good flows, and allow production of used technologies to be scaled up subject to these constraints. This can result in some firms, both upstream and downstream of a disruption, increasing their outputs and it is no longer possible to trace through the impact of a shock with a simple and intuitive algorithm like the shock propagation algorithm we consider here.

6 Concluding Remarks

We comment on the flexibility of the model and identify further explorations for which the model can serve as a foundation.

6.1 Sanctions

One potential further application of the model is to estimate the impact of sanctions. Our approach would provide an estimate of the impact of targeted sanctions, and can be used as a foundation for selectively targeting goods and services that would most impact parts of the world economy and not others. As is clear from our analysis, the short-run and long-run impacts of sanctions can differ dramatically, depending on the ability of target countries or industries to reallocate over time.¹⁸

6.2 Edge shocks

We specified shocks as being incident on nodes (country-specific industries or technologies), but our analysis extends directly to analyze edge shocks. Suppose for example that a good is produced and output is shipped through the Red Sea and other output is shipped elsewhere, and we want to consider the possibility of a shock that disrupts shipping through the Red Sea.

¹⁸See [Bachmann et al. \(2022\)](#) for an analysis of the impact of Russian energy sanctions for the German economy.

Then we can simply divide the original technology into two different nodes, one that ships via the Red Sea and another that ships elsewhere. Shocking the first node is equivalent to shocking the Red Sea edge of the original technology, and we have seen that such shocks can cause relatively severe short term disruptions.¹⁹

6.3 Inventories and Endogenous Robustness

In our benchmark of a short-run disruption, disrupted goods are fully missing from production. Firms can maintain (costly) inventories of inputs to avoid issues with supply chain disruptions. An existing inventory can buffer some of the impact depending on a disruption’s magnitude and how long it lasts. Firms might even maintain alternative production technologies, some which are inefficient, but can serve as backups in times of disruption. This differs across industries and the perceived dangers faced from disruptions.

Although producers prefer to avoid disruptions, excess inventory costs (beyond minimum ones to run production processes) may not be compensated in the face of competitive pressures, and thus excess inventories tend to be low. To the extent that fully contingent contracts are not written in advance for final consumption goods (and, indeed, most goods are sold on spot markets), market incompleteness favors firms with lower costs and thus pressure them to avoid excess inventory costs. Firms that are more robust, gain by sales when others are disrupted, but to the extent that they cannot capitalize on those profits *ex ante* (due to incomplete contracts), the market may be inefficient. This is an interesting topic for further consideration and, again, our results provide a foundation on which to build.

6.4 Services

While it might be natural to think of the model in terms of physical goods rather than services, the model includes services (as both intermediate and final ‘goods’). Inputs, including labor, are used to create outputs in the same way, and disruptions propagate through the supply network equivalently. Given that labor is often less mobile than other (intermediate) goods, it might be that structure of networks around services vary depending on what types of service it is: e.g., coding which can be done remotely, versus healthcare which is done more locally. The model provides the structure with which to estimate shock effects and how those differ across different supply network structures, and this application presents an interesting future agenda.

References

Acemoglu, Daron and Pablo D Azar. 2020. “Endogenous Production Networks.” *Econometrica* 88 (1):33–82.

¹⁹For example, in an article title “Tesla to Halt Production in Germany as Red Sea Conflict Hits Supply Chains,” The Wall Street Journal reports that Tesla had to shut down its main European factory for two weeks following the Houthi attacks in the Red Sea (January 12, 2024).

- Acemoglu, Daron, Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2012. “The Network Origins of Aggregate Fluctuations.” *Econometrica* 80 (5):1977–2016.
- Acemoglu, Daron and Alireza Tahbaz-Salehi. 2020. “Firms, Failures, and Fluctuations: The Macroeconomics of Supply Chain Disruptions.” Tech. rep., National Bureau of Economic Research.
- Amelkin, Victor and Rakesh Vohra. 2020. “Strategic Formation and Reliability of Supply Chain Networks.” In *Proceedings of the 21st ACM Conference on Economics and Computation*. 77–78.
- Antràs, Pol and Davin Chor. 2022. “Global value chains.” *Handbook of international economics* 5:297–376.
- Arrow, Kenneth J. and Gerard Debreu. 1954. “Existence of an equilibrium for a competitive economy.” *Econometrica* :265–290.
- Bachmann, Rüdiger, David Baqaee, Christian Bayer, Moritz Kuhn, Andreas Löschel, Benjamin Moll, Andreas Peichl, Karen Pittel, and Moritz Schularick. 2022. “What if? The economic effects for Germany of a stop of energy imports from Russia.” Tech. rep., ECONtribute Policy Brief. No. 028.
- Bahal, Girish and Damian Lenzo. 2023. “Beyond Domar Weights: A New Measure of Systemic Importance in Production Networks.” .
- Baldwin, Richard and Rebecca Freeman. 2022. “Risks and global supply chains: What we know and what we need to know.” *Annual Review of Economics* 14:153–180.
- Baqaee, David and Emmanuel Farhi. 2021. “Entry vs. rents: Aggregation with economies of scale.” Tech. rep., National Bureau of Economic Research. Working Paper No. 27140.
- Baqaee, David and Elisa Rubbo. 2022. “Micro propagation and macro aggregation.” Tech. rep., National Bureau of Economic Research. Working Paper No. 30538.
- Baqaee, David Rezza. 2018. “Cascading failures in production networks.” *Econometrica* 86 (5):1819–1838.
- Baqaee, David Rezza and Emmanuel Farhi. 2019. “The macroeconomic impact of microeconomic shocks: beyond Hulten’s Theorem.” *Econometrica* 87 (4):1155–1203.
- Barrot, Jean-Noël and Julien Sauvagnat. 2016. “Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks.” *The Quarterly Journal of Economics* 131 (3):1543–1592.
- Bernard, Andreas, Andrew B. and Moxnes. 2018. “Networks and trade.” *Annual Review of Economics* 10:65–85.

- Bernard, Andrew B, Emmanuel Dhyne, Glenn Magerman, Kalina Manova, and Andreas Moxnes. 2022. “The origins of firm heterogeneity: A production network approach.” *Journal of Political Economy* 130 (7):1765–1804.
- Bernard, Andrew B, Andreas Moxnes, and Yukiko U Saito. 2019. “Production networks, geography, and firm performance.” *Journal of Political Economy* 127 (2):639–688.
- Bimpikis, Kostas, Ozan Candogan, and Shayan Ehsani. 2019. “Supply Disruptions and Optimal Network Structures.” *Management Science* 65 (12):5504–5517.
- Bimpikis, Kostas, Douglas Fearing, and Alireza Tahbaz-Salehi. 2018. “Multisourcing and miscoordination in supply chain networks.” *Operations Research* 66 (4):1023–1039.
- Boehm, Christoph E, Aaron Flaaen, and Nitya Pandalai-Nayar. 2019. “Input linkages and the transmission of shocks: firm-level evidence from the 2011 tōhoku earthquake.” *Review of Economics and Statistics* 101 (1):60–75.
- Brummitt, Charles D, Kenan Huremović, Paolo Pin, Matthew H Bonds, and Fernando Vega-Redondo. 2017. “Contagious Disruptions and Complexity Traps in Economic Development.” *Nature Human Behaviour* 1 (9):665.
- Bui, Ha, Zhen Huo, Andrei A Levchenko, and Nitya Pandalai-Nayar. 2022. “Information frictions and news media in global value chains.” Tech. rep., National Bureau of Economic Research. Working Paper No. 30033.
- Caliendo, Lorenzo and Fernando Parro. 2015. “Estimates of the Trade and Welfare Effects of NAFTA.” *The Review of Economic Studies* 82 (1):1–44.
- Carvalho, Vasco M., Matthew Elliott, and John Spray. 2024. “Network bottlenecks and market power.” *Unpublished manuscript* .
- Carvalho, Vasco M, Makoto Nirei, Yukiko U Saito, and Alireza Tahbaz-Salehi. 2021. “Supply chain disruptions: Evidence from the great east japan earthquake.” *The Quarterly Journal of Economics* 136 (2):1255–1321.
- Carvalho, Vasco M and Alireza Tahbaz-Salehi. 2019. “Production networks: A primer.” *Annual Review of Economics* 11:635–663.
- Chaney, Thomas. 2014. “The network structure of international trade.” *American Economic Review* 104 (11):3600–3634.
- Chen, Kedong, Ankur Mani, Kevin Wayne Linderman, and Bing Wang. 2023. “Where to Invest in Resilience in a Facility Network?” *Available at SSRN 4012928* .
- Costinot, Arnaud, Jonathan Vogel, and Su Wang. 2013. “An elementary theory of global supply chains.” *Review of Economic studies* 80 (1):109–144.

- Davidson, Helen and Martin Farrer. 2021. “Chips with everything: How one Taiwanese Company Drives the World Economy.” URL <https://www.theguardian.com/world/2021/may/19/chips-with-everything-how-one-taiwanese-company-drives-the-world-economy>.
- Dhyne, Emmanuel, Glenn Magerman, and Stela Rubínová. 2015. “The Belgian production network 2002-2012.” Tech. rep., NBB Working Paper.
- Di Giovanni, Julian, ?ebnem Kalemli-Özcan, Alvaro Silva, and Muhammed A Yildirim. 2022. “Global supply chain pressures, international trade, and inflation.” Tech. rep., National Bureau of Economic Research.
- Diem, Christian, András Borsos, Tobias Reisch, János Kertész, and Stefan Thurner. 2022. “Quantifying firm-level economic systemic risk from nation-wide supply networks.” *Scientific reports* 12 (1):7719.
- Elliott, Matthew and Benjamin Golub. 2022. “Networks and economic fragility.” *Annual Review of Economics* 14:665–696.
- Elliott, Matthew, Benjamin Golub, and Matthew V Leduc. 2022. “Supply network formation and fragility.” *American Economic Review* 112 (8):2701–47.
- Giovanni, Julian di and Andrei A Levchenko. 2009. “Trade openness and volatility.” *The Review of Economics and Statistics* 91 (3):558–585.
- Grossman, Gene M, Elhanan Helpman, and Hugo Lhuillier. forthcoming. “Supply Chain Resilience: Should Policy Promote International Diversification or Reshoring?” *Journal of Political Economy* .
- Grossmand, Gene, Elhanan Helpman, and Stephen Redding. 2021. “When tariffs disrupt global supply chains.” *American Economic Review* .
- Hulten, Charles R. 1978. “Growth accounting with intermediate inputs.” *The Review of Economic Studies* 45 (3):511–518.
- Jeong, Eun-Young and Dan Strumpf. 2021. “Why the Chip Shortage Is So Hard to Overcome.” *Wall Street Journal* URL <https://www.wsj.com/articles/why-the-chip-shortage-is-so-hard-to-overcome-11618844905>.
- König, Michael D, Andrei Levchenko, Tim Rogers, and Fabrizio Zilibotti. 2022. “Aggregate fluctuations in adaptive production networks.” *Proceedings of the National Academy of Sciences* 119 (38):e2203730119.
- Kopytov, Alexandr, Bineet Mishra, Kristoffer Nimark, and Mathieu Taschereau-Dumouchel. 2021. “Endogenous production networks under supply chain uncertainty.” Available at SSRN 3936969 .

- Leontief, Wassily W. 1936. “Quantitative Input and Output Relations in the Economic Systems of the United States.” *The Review of Economic Statistics* :105–125.
- Long Jr, John B and Charles I Plosser. 1983. “Real business cycles.” *Journal of political Economy* 91 (1):39–69.
- Magerman, Glenn, Karolien De Bruyne, Emmanuel Dhyne, and Jan Van Hove. 2016. “Heterogeneous firms and the micro origins of aggregate fluctuations.” Tech. rep., NBB Working Paper.
- Mas-Colell, Andreu, Michael Dennis Whinston, Jerry R Green et al. 1995. *Microeconomic theory*, vol. 1. Oxford university press New York.
- McLain, Sean. 2021. “Auto Makers Retreat From 50 Years of ‘Just in Time’ Manufacturing.” *Wall Street Journal* URL <https://www.wsj.com/articles/auto-makers-retreat-from-50-years-of-just-in-time-manufacturing-11620051251>.
- Melitz, Marc J. 2003. “The impact of trade on intra-industry reallocations and aggregate industry productivity.” *econometrica* 71 (6):1695–1725.
- Oberfield, Ezra. 2018. “A Theory of Input-Output Architecture.” *Econometrica* 86.
- Pellet, Thomas and Alireza Tahbaz-Salehi. 2023. “Rigid production networks.” *Journal of Monetary Economics* .
- Summers, Lawrence. 2013. “Crises yesterday and today.” URL <http://larrysummers.com/imf-fourteenth-annual-research-conference-in-honor-of-stanley-fischer>. Speech at 14th Jacques Pollack Annual IMF Research Conference in Honor of Stanley Fischer.
- Wei, Chao. 2003. “Energy, the stock market, and the putty-clay investment model.” *American Economic Review* 93 (1):311–323.

A Omitted proofs

LEMMA 1. *In any equilibrium aggregate consumption is equivalent to the consumption choice of a single representative consumer with preferences represented by the utility function $U(c_1, \dots, c_F)$ and with wealth $\sum_n L_n p L_n$.*

Proof of Lemma 1:

Country n 's consumer solves

$$\max_{c_{1n}, \dots, c_{Fn}} U(c_{1n}, \dots, c_{Fn}) \text{ subject to } \sum_{f \in F} p_f c_{fn} = L_n p_n.$$

Given finite prices for all final goods and labor (which follows in equilibrium, as shown below), and given that the utility function is $U(\cdot)$ is increasing and strictly quasi-concave there is a unique solution to this problem. Further, as consumers' utility functions are identical and all consumers in all countries face the same prices for final goods (because there are no transportation costs on final goods), each country's consumer solves the same problem except for differences in their labor endowments inducing differences in their wealth levels. Thus, as $U(\cdot)$ is homogeneous of degree 1, the solution to each country's consumer problem is a re-scaling of the same bundle of goods. Moreover, it follows that a representative consumer with utility function $U(c_1, \dots, c_F)$ and with wealth $\sum_n L_n p_{L_n}$ chooses a re-scaling of the same bundle of goods, and that re-scaling is the aggregate consumption. ■

A.1 Proof of Lemma 2

Proof. We prove existence of an equilibrium by mapping our economy into one without transportation costs, and then applying standard results. In order to do this we use the transportation costs to map each technology $\tau \in \mathbb{R}^{1+M+F}$ into source-specific technologies $t \in \mathbb{R}^{N+T}$. Specifically, for each technology τ that required k inputs, we create N^k source-specific technologies allowing for each possible combination of sourcing choices from different technologies across the different inputs. Further, we adjust the number of units of each input required to represent the number of units that need to be sourced from that technology, including those units that will be lost to transportation costs. So, if technology τ requires k units of good g , and the transportation costs associated with sourcing good g from τ' are 2, then corresponding technology source-specific technologies t that source good g from technology τ' , requires $2k$ units of good g ; and similarly for the N sources of labor.

Replacing transportation costs and technologies with technology source-specific technologies, existence of equilibrium in our environment can be proven using standard techniques (noting that each technology can then be represented as a production possibility set by adding free disposal), such as that used to prove Theorem 17BB2 of [Mas-Colell et al. \(1995\)](#). This establishes existence.

An equilibrium of our economy must be Pareto efficient by the first Welfare Theorem. Further, as by Lemma 1 our economy admits a representative consumer with utility $U(c_1, \dots, c_F)$, all Pareto efficient allocations must maximize this utility function subject to feasibility constraints. We show that the set of feasible consumption bundles is convex and compact, and hence, as the representative consumer's utility function is increasing and quasi-concave, that there is a unique bundle of final goods that solves the Pareto problem.

A consumption bundle $c \in \mathbb{R}_+^F$ is feasible if

1. $c_f \leq \sum_{\tau: O(\tau)=f} y_\tau$ for all f ,
2. $y_\tau \leq \min_{k \in I(\tau)} \sum_{\tau': O(\tau')=k} \left(\frac{x_{\tau'\tau}}{-\tau_k \theta_{\tau'\tau}} \right)$ for all τ
3. $\sum_\tau x_{n\tau} \leq L_n$ for all n .

The first condition requires that enough final goods are produced. The second condition requires that sufficient inputs are sourced to support the required output for each technology. The final condition requires that the use of labor satisfies the labor endowments.

Consider two feasible consumption bundles c and c' . We first show that this implies the consumption bundle $c'' = \lambda c + (1 - \lambda)c'$ is also feasible for all $\lambda \in [0, 1]$. First suppose that we reduced the labor endowments of all countries to λ their initial levels. As all technologies are constant returns to scale, and it was feasible before the reduction to produce the bundle c , it must be feasible after the reduction to produce the bundle λc . This can be obtained by reducing all inputs (and hence all outputs) of all used technologies to λ their initial levels. Equivalently, if all labor endowments are reduced to $(1 - \lambda)$ their initial levels, then the bundle $(1 - \lambda)c'$ is feasible. Hence, the bundle c'' is feasible with the initial labor endowments. This shows that set of feasible consumption bundles is convex.

We now show that the set of feasible consumption bundles is compact. Take any feasible use of technologies that produces a non-zero consumption bundle c . As all technologies use a strictly positive amount of labor, producing this bundle uses a strictly positive amount of labor. Let $L(c) \in \mathbb{R}_+^n$ denote the vector of labor inputs used across countries. Fixing the use of technologies, as all technologies are constant returns to scale, we can increase all inputs to $\lambda \geq 1$ their initial level to produce the consumption bundle λc . Thus there exists a unique $\bar{\lambda} \geq 1$ that maximizes λc subject to $\lambda L(c)_n \leq L_n$ for all n . As the consumptions must be non-negative and the 0 bundle is feasible, the set of feasible consumption bundle is compact.

The Pareto problem therefore involves maximizing a strictly quasi-concave function subject to a convex and compact constraint set, and thus has a unique solution. Hence, in all equilibria, the same aggregate consumption must occur.

Finally, we argue that, generically, the use of technologies (and hence flow of goods) is unique in equilibrium. As utility is increasing, and labor is not perfectly immobile, all labor must be employed in equilibrium. Suppose then that there exist two different equilibrium uses of technologies. These must both fully employ labor and produce the same bundle of final goods. However, this cannot occur generically, because a slight perturbation to the total factor productivity of the technologies would, with probability 1, allow a strictly more preferred final bundle of goods to be produced by one usage of technologies than the other. \square

A.2 Proof of Hulten's Theorem: Proposition 1

Proof of Proposition 1:

We first prove the result for final goods, and then extend it to intermediate goods.

Euler's Homogeneous Function Theorem implies that for any $(c_1, \dots, c_F) \in \mathbb{R}_+^F$:

$$U(c_1, \dots, c_F) = \sum_f c_f \frac{\partial U}{\partial c_f}. \tag{1}$$

By Lemma 1 total world consumption choices are as if there is a representative consumer

with preferences represented by the homogeneous of degree 1 utility function $U(c_1, \dots, c_F)$ and wealth $I = \sum_n L_n p_{Ln}$. The representative consumer's problem is choosing non-negative amounts of the final goods to consume to

$$\text{maximize } U(c_1, \dots, c_F) \text{ subject to } \sum_f p_f c_f \leq I.$$

Thus, in equilibrium,

$$\frac{\partial U}{\partial c_f} = \lambda p_f$$

for all final goods consumed in positive amounts, where $\lambda > 0$ is the Lagrange multiplier on the wealth constraint. Thus, from (1) it follows that in equilibrium

$$U(c_1, \dots, c_F) = \lambda \sum_f c_f p_f. \quad (2)$$

Moreover, in equilibrium, constraint (11) holds. Allowing τ_f to change from one this implies that:²⁰

$$c_f = \sum_{\tau \in T: O(\tau)=f} \tau_f y_\tau \quad (3)$$

Substituting (3) into (2) yields

$$U(c_1, \dots, c_F) = \lambda \sum_f p_f \sum_{\tau \in T: O(\tau)=f} \tau_f y_\tau. \quad (4)$$

As this is an equilibrium expression for utility, the envelope theorem can be applied and hence, for a given production technology τ with $O(\tau) = f$,

$$\frac{\partial U}{\partial \tau_f} = \lambda p_f y_\tau \quad (5)$$

Then from (5) and (2),

$$\frac{\partial \log(U)}{\partial \log(\tau_f)} = \left(\frac{\partial U}{\partial \tau_f} \right) \left(\frac{\tau_f}{U} \right) = \frac{\lambda p_f y_\tau \tau_f}{U} \Big|_{\tau_f=1} = \frac{p_f y_\tau}{\sum_f c_f p_f}.$$

Note that as $GDP = \sum_f p_f c_f$, by (2) $U = \lambda GDP$, and so we also have

$$\frac{\partial \log(U)}{\partial \log(\tau_f)} = \frac{\partial [\log(GDP) + \log(\lambda)]}{\partial \log(\tau_f)} = \frac{\partial \log(GDP)}{\partial \log(\tau_f)} = \frac{p_f y_\tau}{GDP}.$$

We now extend this to intermediate goods.

²⁰One can interpret y_τ as a number of units of operation of the technology, and then τ_f different from one scales the amount produced.

The zero profit conditions for each firm allow the revenues of a firm to be equated to its costs, which can be expressed in terms of its suppliers' revenues, which are also equated to costs, and so on. Repeating this process, if a technology $\hat{\tau}$ produces an intermediate good that is used directly or indirectly in the production of a final good, the revenues generated by sales of this final good can be expressed in terms of the revenues generated by technology $\hat{\tau}$, and the remaining direct and indirect labor costs associated with the production of the final good.

Consider a technology τ with $O(\tau) = f \in F$. By the zero profit condition for technology τ

$$p_f y_\tau = \sum_{k \in I(\tau)} \sum_{\tau': O(\tau')=k} p_{\tau'} x_{\tau'\tau} + \sum_n p_n x_{n\tau}.$$

But, for an input technology τ' such that $O(\tau') = k \in I(\tau)$ and $x_{\tau'\tau} > 0$, we also have, by the zero profit condition

$$p_{\tau'} x_{\tau'\tau} = \left(\frac{x_{\tau'\tau}}{y_{\tau'}} \right) \sum_{k' \in I(\tau')} \sum_{\tau'': O(\tau'')=k'} p_{\tau''} x_{\tau''\tau'} + \sum_n p_n x_{n\tau'}.$$

Iteratively substituting in these expressions for all intermediate good technologies except the shocked one, $\hat{\tau}$, the obtained expression will converge to one with the following form:

$$p_f y_\tau = p_{\hat{\tau}} \hat{x}_{\hat{\tau}\tau} + \sum_n p_n \hat{x}_{n\tau}, \quad (6)$$

where $\hat{x}_{\hat{\tau}\tau}$ is the amount of good $O(\hat{\tau})$ produced by technology $\hat{\tau}$ that ultimately ends up being used (directly or indirectly) by final good technology τ and $\hat{x}_{n\tau}$ is the amount of labor, other than that used by technology $\hat{\tau}$, from country n that ultimately ends up being used (directly or indirectly) by final good technology τ .

Note that as there is no waste in equilibrium all production of an intermediate goods can be assigned to final goods and hence

$$\sum_{f \in F} \sum_{\tau: O(\tau)=f} \hat{x}_{\hat{\tau}\tau} = y_{\hat{\tau}}. \quad (7)$$

Consider an intermediate good technology $\hat{\tau}$ with $O(\hat{\tau}) = k$, and substitute 6 into 4, setting $\tau_f = 1$ (as we are not varying it) and adding in $\hat{\tau}_k = 1$ (as we will be varying it). This gives

$$U(c_1, \dots, c_F) = \lambda \sum_f \sum_{\tau \in T: O(\tau)=f} \left(\hat{\tau}_k p_{\hat{\tau}} \hat{x}_{\hat{\tau}\tau} + \sum_n p_n \hat{x}_{n\tau} \right).$$

As this is an equilibrium expression for utility, the envelope theorem can again be applied and hence

$$\frac{\partial U}{\partial \hat{\tau}_k} = \lambda p_{\hat{\tau}} \sum_f \sum_{\tau \in T: O(\tau)=f} \hat{x}_{\hat{\tau}\tau}.$$

Substituting in 7 we get

$$\frac{\partial U}{\partial \hat{\tau}_k} = \lambda p_{\hat{\tau}} y_{\hat{\tau}}.$$

Thus,

$$\frac{\partial \log(U)}{\partial \log(\tau_f)} = \frac{\partial [\log(GDP) + \log(\lambda)]}{\partial \log(\hat{\tau}_k)} = \frac{\partial \log(GDP)}{\partial \log(\hat{\tau}_k)} = \frac{p_{\hat{\tau}} y_{\hat{\tau}}}{GDP},$$

which is the claimed expression. ■

A.3 Proof of Proposition 2

Proof. We begin by showing the first half of the proposition: that the output of the algorithm, ω^* , is unique and that $\omega_{ij}^* \leq \omega_{ij} \forall i, j \in \mathcal{N}$. Moreover, we show that for any path \mathcal{P} starting at an affected node i we have that $\omega_{jk}^* < \omega_{jk} \forall j, k \in \mathcal{P}$.

Start with the original equilibrium flow network ω , and reduce the value of the outlinks for the shocked nodes to λ their initial level. Let Φ be the vector consisting of all non-zero link-weights in this network. Now, consider the space $S = \prod_{j=1, \dots, |\Phi|} [0, \phi_j]$. Define the partial ordering on this space such that $s \succeq \hat{s}$ for $s, \hat{s} \in S$ if it is weakly greater entry by entry (i.e., $s \succeq \hat{s}$ if and only if $s_i \geq \hat{s}_i$ for all i). Note that (S, \succeq) is a complete lattice.

We represent the flow along each of the links in each iteration of the algorithm by some $\omega \in S$. It is clear that each iteration of the algorithm provides a continuous mapping $\Gamma : S \rightarrow S$ with $\Gamma(\omega) \preceq \omega$. This implies that $\Gamma(\cdot)$ is an isotone function with respect to (S, \succeq) and hence, by the Knaster-Tarski theorem, the set of fixed points of $\Gamma(\cdot)$ is a complete lattice under \succeq . There is thus a largest fixed point, with respect to the partial ordering \succeq , which we denote by ω^* .

As the space S is compact, and in each iteration $\sum_i \omega_i$ weakly decreases, the algorithm converges by the monotone convergence theorem. Thus, the limit of the algorithm, ω^* , is well-defined and unique and $\omega_{ij}^* \leq \omega_{ij} \forall i, j \in \mathcal{N}$.

The fact that ω^* is a solution to the minimum disruption problem is argued as follows. We have shown that the shock propagation algorithm converges to a fixed point of Γ . By the construction of Γ any fixed point of it must satisfy the constraints in the minimum disruption problem (if there is any violation then the algorithm keeps iterating). Moreover, any flows ω that are not a fixed point of Γ must violate one of the constraints of the minimum disruption problem—otherwise the algorithm would terminate. This implies that a solution to the minimum disruption problem must be a fixed point of Γ , and specifically the fixed point that maximizes the value of final goods produced. So we just need to show that the algorithm converges to this fixed point of Γ . To show maximality with respect to network structure notice that each reduction in the flow imposed by the algorithm is inevitable, so the flows ω^* represent the maximum that can be produced final good by final good after the disruption.

Finally, note that for each country-industry node that ends up with a strictly lower value of a flow on an in-link relative to the initial equilibrium (i.e., $\omega_i^* < \omega_i$), its output, and hence the flow on each of its out-links, must be strictly lower than in \mathcal{G} . This proves that for any path \mathcal{P}

starting at an affected node i we have that $\omega_{jk}^* < \omega_{jk} \forall j, k \in \mathcal{P}$. \square

A.4 Proof of Proposition 3

Proof. The upper bound follows from the algorithm directly, noting that if none of a firm's suppliers have production levels below λ of their initial levels, then, given the structure of the production functions, the firm's output is ever below λ of its initial level. Thus, given that all goods start with at least λ of their initial levels, any limit point of the algorithm has the production of each good at λ of its initial level or higher. \square

A.5 Proof of Proposition 4

Proof. Suppose that the edges adjacent to $T^{Shocked}$ constitute a $(R(T^{Shocked}), F(T^{Shocked}))$ -cut. One way in which we can assign the equilibrium flows (before the shock) to supply chains is by assigning them to supply chains in which all input goods are single-sourced and where, as there are no cycles in the disrupted industries subnetwork, there are no cycles in each such supply chain. We then observe that (i) for a set of nodes $T^{Shocked}$ to be an $(R(T^{Shocked}), F(T^{Shocked}))$ -cut every path in the network leading to every final good $k \in F(T^{Shocked})$ must contain at least one node in $T^{Shocked}$; and (ii), that for any path that passes through a node in $T^{Shocked}$, output is reduced by $(1 - \lambda)$. Part (ii) holds because along each path all goods are single-sourced—thus the output of each node downstream of the shocked node must have its output reduced by exactly $(1 - \lambda)$ of its initial level. Parts (i) and (ii) together imply that the output of paths in the network leading to all final goods in $F(T^{Shocked})$ are reduced by a proportion $1 - \lambda$ of their initial level, and the upper bound is obtained. \square

A.6 Proof of Proposition 7

Proof. Suppose $\max_{\tau, \tau'} \theta_{\tau\tau'} = \max_{n, n'} \theta_{nn'} = 1$ (so there are no transportation cost). By Lemma 2 there is a unique equilibrium bundle of final goods produced. By the zero profit conditions, in equilibrium the revenues generated by the sale of a unit of any good must equal its unit cost. However, as there are no transportation costs, there will be unique world price at which any given good is sold—if there were two different prices only the lower priced product would be purchased. Similarly, there will be a unique world price for labor.

As argued in the proof of Proposition 1, the unit cost of any technology with positive sales can be represented as the bundle of labor costs incurred in corresponding supply chains per unit produced. Moreover, as there is a world price of labor each technology's labor costs is just the world price of labor times their total use of labor (through their supply chains) per unit produced. As sales only go to a producer of a given good if it is priced weakly below all other producers of the same good, and as in equilibrium all firms must make zero profits, only supply chains that use the least total amount of labor possible to produce each unit can be active in

equilibrium. Thus, generically in the space of technologies, as each country is endowed with a finite set of technologies, there will be a unique technology used to produce any given good.

Now suppose the maximum transportation cost is greater than 1. Note that for all transportation cost any combination of supply chains must produce less than was produced in the frictionless economy. By the above argument, in the frictionless economy generically there will be a unique produced of each good in equilibrium, and by the first welfare theorem, this must maximize GDP while, generically, the use of any other technologies to produce a strictly positive amount of a good would reduce GDP. Hence, by the continuity of the production technologies, there exists a $\bar{\theta} > 1$ such that if $\max_{\tau, \tau'} \theta_{\tau\tau'} < \bar{\theta}$ and $\max_{n, \tau'} \theta_{n\tau'} < \bar{\theta}$ such that the same technologies must be used as in the frictionless economy. \square