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

Matthew L. Elliott and Matthew O. Jackson†

September 2023

Abstract

We develop a parsimonious multi-sector model of international production and use it to study how a disruption in the production of some 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 the final goods whose supply chains on which disrupted goods lie, while by contrast the long-run disruption depends on the cost of the disrupted goods. We use the model to show that increased complexity of supply chains leads to increased fragility. We also show that decreased transportation costs lead to increased specialization in production and a potential for greater fragility.

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 gratefully acknowledge support under NSF grant and 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.

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 completely stalled the production of downstream goods. More generally, goods vary widely in the positions and their potential to disrupt final goods production,[3] and so it is important to quantify that potential and how it depends on the production network.

We develop a parsimonious model of global trade and interlinked supply chains, and use it to examine disruptions to supply chains in detail. In particular, we characterize the impact of the disruption of goods and how they depend on structure of supply chains, the position of disruptions in the chains, and the value of the final goods in the impacted chains. We use the characterization to 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 supply chains and the potential for disruption change as transportation costs drop and supply chains involve more specialization and 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. If productivity drops by 5 percent then the circuits just become about 5 percent more expensive, and so the total resource loss is an extra 5 percent of what was being spent originally. This intuition applies at the margin, and might even overestimate the impact to the extent that production can be further substituted as one moves away from the margin. However, in the short run, the impact can be much larger. A 5 percent shortage of integrated circuits instead leads to a full delay of 5 percent of all final goods that use them as inputs. Those goods, valued at around $5 trillion, means lost production is around $250 billion. This comparison is even starker if a shock to basic computer chips is considered given the breadth of their use as inputs. While the short-run disruption might eventually be made up for, this stark contrast shows that

[1] Long 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, 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)).

[2] For 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 .

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

short-run impacts can be dramatic.

The contrast can be summarized as the short-run impact being proportional to the value of all downstream final products, while the long-run impact is simply the proportional change in cost of 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).[4]

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 firm(s) multiplied by the value of production of all final goods that used 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 iceberg costs are reduced sufficiently, countries specialize in what they produce and hence firm-specific shocks become industry-specific shocks, such that the bound is always obtained.

The more general calculation may differ from the bound depending on two issues. One is that some of the goods might source the same input from multiple sources, and if not all of those sources are shocked, then the reduction can be less than the percent of the shock. The other is that if there are cycles in the supply network, then these can feedback and lead to further impacts beyond the first iteration. 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 shock.

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 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 an increase in fragility as there are fewer sources of intermediate goods for production of the same final good. This is consistent with empirical evidence on globalization, specialization, and fragility (Giovanni and Levchenko, 2009; Kalemli-Özcan et al., 2022; Baldwin and Freeman, 2022).

The most closely related strands of literature to our work are the following three: 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

[4]See Carvalho and Tahbaz-Salehi (2019) for further discussion around this and ways in which Hulten's theorem can be relaxed. Of particular note is Baqaee and Farhi (2019).

et al. (2012), a series of papers examine the propagation of shocks through sectorial and firm inter-linkages in the economy.[5] Some of the recent work has incorporated cascading failures, production shut-downs and endogenized the network structure (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; 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-equilibriates. 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 realised. 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 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 of Pellet and Tahbaz-Salehi (2023) and their findings 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 short-run lack of adjustment 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.[6] 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. That work 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.[7] In contrast our analysis provides details on the how the size and scope of a disruption depend on

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

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

[7] The 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.

specific details of the supply chain, which are abstracted away from in their work.

Finally, there is related research on networks and international trade.[8] 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) and Bernard, Moxnes, and Saito (2019)). In terms of the macro modelling of international trade, the approach taken has tended to be very 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

Goods and countries:

We use the same notation for each set and its cardinality. There is a set N of countries indexed $n \in \{1, \ldots, N\}$. Goods consist of a set M of intermediate goods, including raw materials, indexed by $m \in \{1, \ldots, M\}$ that are used as inputs to production; and a set F of final goods indexed $f \in \{1, \ldots, F\}$ that are consumed. Labor is denoted by L. 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.[9]

Labor and Endowments:

Country n's endowment of labor is denoted $L_n > 0$, and it is supplied completely inelastically.

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 to that country to produce oil.

Technologies:

We work with Arrow-Debreu production economies, with one simplification. Instead of working with infinite production possibility sets, we work with finite sets. This is for three 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 so this is without loss of generality

[8] For recent surveys see Bernard (2018) and Antràs and Chor (2022).

[9] Simply 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.

for our results.[10] Second is that this is more realistic when it comes to addressing things like innovations to technologies, which is something that we wish our model to address. Another reason is that this makes it easy to talk about generic variations in the technologies used. We can still prove existence of equilibrium despite the non-convexity.

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. Moreover, any technology that has increasing returns can be approximated by a finite set of constant returns technologies.

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. The first entry in the vector, representing the amount of labor required, is assumed to be strictly negative for all technologies.

A technology τ satisfies two additional conditions. First, $\{k : \tau_k > 0\}$ has one element, interpreted as the output of the technology. 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. We 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. For notational reasons, we also associate each technology with a unique country. We let $n(\tau)$ be the country technology τ is assigned to. We allow there to be multiple identical technologies if these technologies are assigned to different countries.

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

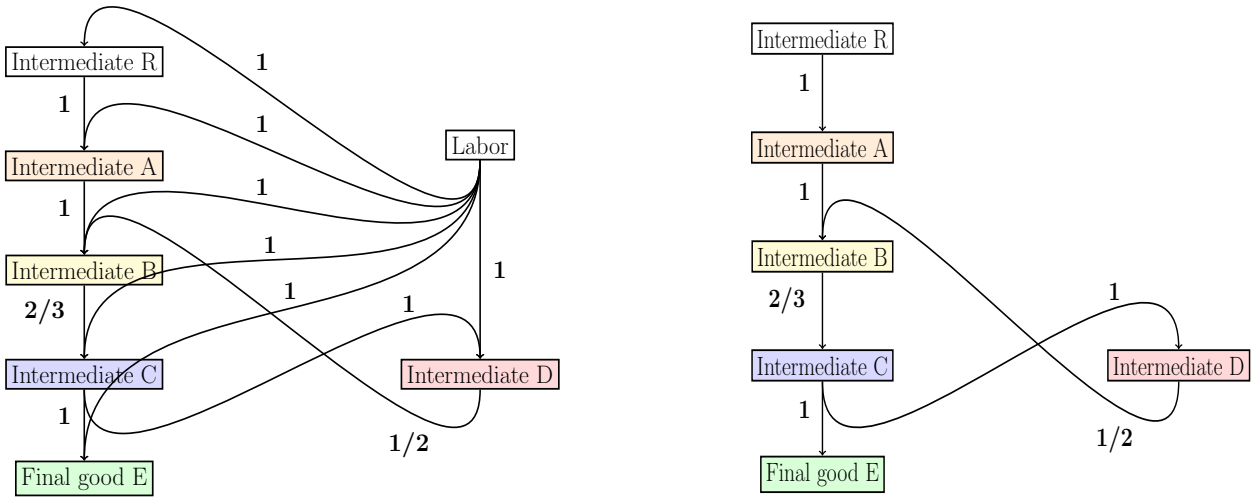
Shipped units:

The matrix $X \in \mathbb{R}_+^{N+T} \times \mathbb{R}_+^T$ denotes the number of units shipped, with an entry $x_{n\tau}$ denoting the amount of labor endowed to country n that is shipped for use by technology τ, and $x_{\tau'\tau}$ denoting the amount of good $O(\tau')$ that is shipped for use by technology τ.

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 by country n to technology τ in order for technology τ to get the benefits of 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

[10] 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.



(a) Technological dependencies

(b) Technological dependencies without labor pictured

Figure 1: An example of technological interdependencies: The weight on a directed link from an input into a technology represents the number of units of that input required to produce one unit of the output good associated with that technology. As all goods require labor, albeit in potentially different amounts, it is convenient to sometimes omit the dependence on labor.

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 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.

Prices:

Prices for the each good and labor are given by a vector $p \in \mathbb{R}_+^{N+T}$, providing the cost of hiring local labor and technology-specific prices for all goods that might be produced. These prices are the local prices at the point of sale per unit–including iceberg costs yield a price in another location.

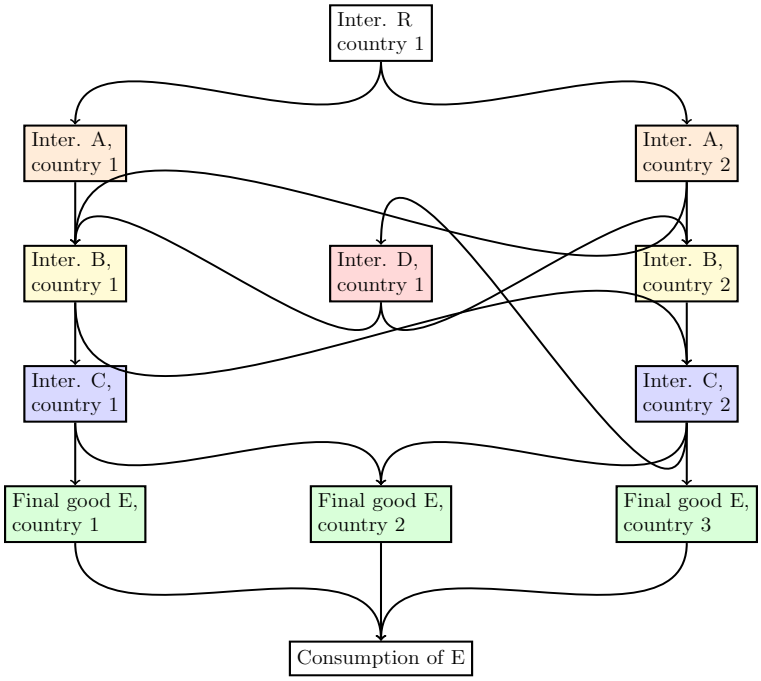


Figure 2: An example of the flow of goods: A directed link from one technology to another represents that a positive amount of the output of the first technology is used as an input for the second technology. Here country 1 has technologies for producing goods R (representing a raw material), A, B, C, D and E, country 2 has technologies for producing goods A, B, C and E while country 3 only has a technology for producing good E. The technologies for producing a given good are similar in this example insofar as they all require the same combination of inputs (but in possibly different ratios as these are not specified).

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 $O(\tau) = O(\tau') = f$.

It is helpful to define the iceberg-cost adjusted prices that technology τ faces. We let $\hat{p}_\tau \in \mathbb{R}_+^{1+M+F}$ denote these prices. The first entry of this vector records the adjusted cost of labor for technology τ. This is given by $\hat{p}_{\tau L} = \min_{n' \in N} \theta_{n'\tau} p_{n'}$. The next M entries record the adjusted sourcing costs of inputs, with entry $\hat{p}_{\tau m} = \min_{\tau': O(\tau')=m} \theta_{\tau'\tau} p_{\tau'}$. The final F entries are redundant, as by assumption final goods are never used as inputs, but for consistency we set $\hat{p}_{\tau f} = \min_{\tau': O(\tau')=f} p_{\tau'} = p_f$.

Preferences:

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

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

that is increasing and strictly quasi-concave and homogeneous of degree 1.

As there are no iceberg 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, our formulation admits a representative consumer with preferences represented by $U(\cdot)$.

An Economy:

Given competitive equilibrium and constant returns to scale, there are zero profits, so 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: $(N, M, F, \{T_n\}_n, \{L_n\}_n, \theta)$.

2.2 Equilibrium

An *equilibrium* of an economy is $(N, M, F, \{T_n\}_n, \{L_n\}_n, \theta)$ is a specification of

- prices $p \in \mathbb{R}_+^{N+T}$ for labor and technologies (whether in use or not), and
- for all countries n and technologies $\tau \in T^n$:
  - a corresponding amount of the output $O(\tau)$ for each $\tau \in T$ denoted by $y^\tau \geq 0$,
  - Amounts (in units shipped) of inputs for us by the technology τ $x_{k\tau} \in \mathbb{R}_+$ for all $k \in N \cup T$, and
- the total amount of each final good f consumed in each country n, c_{fn};

that satisfy the following conditions:

- **Consumers optimize:** Labor is fully supplied at the highest available wage (accounting for transportation costs) and wages are spent on final goods to maximize utility. That is, laborers in country n choose consumption $(c_{1n}, \ldots, c_{Fn})$ to maximize

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

- **Producers optimize (and earn zero profits):** For each active technology τ (such that $y^\tau > 0$)

$$\hat{p}_\tau \cdot \tau = 0,$$

and for each inactive technology τ (such that $y^\tau = 0$)

$$\hat{p}_\tau \cdot \tau \leq 0.$$

- **Production Plans are Feasible** For each active technology τ

$$\sum_{\tau' : O(\tau') = k} \frac{x_{\tau'\tau}}{\theta_{\tau'\tau}} = -\tau_k y_\tau, \tag{1}$$

which says that the total inputs sourced, adjusted for shipping losses, equal the amount of inputs required under the technology to produce the desired level of output.

- **Markets clear:**
  - Labor markets clear: the total amount of labor from each country n being used in production throughout the world is equal to the its endowment:

$$L_n = \sum_{\tau} x_{n\tau}. \tag{2}$$

  - Intermediate goods markets clear: For each technology τ such that $O(\tau) \in M$, the amount being used in production across all countries is equal to the amount produced. So:

$$\sum_{\tau'} x_{\tau\tau'} = y_{\tau}. \tag{3}$$

  - Final goods markets clear: For each technology τ such that $O(\tau) \in F$, the amount being consumed across all countries is equal to the amount produced. So:

$$\sum_{n} c_{fn} = \sum_{\tau \in T : O(\tau) = f} y_{\tau}. \tag{4}$$

Given constant returns to scale, in equilibrium firms make zero profits and hence all goods that are produced in positive amounts are priced at their respective unit costs. In particular, for all countries n and technologies $\tau \in T^n$, if $y_\tau > 0$ and $O(\tau) = k$, then

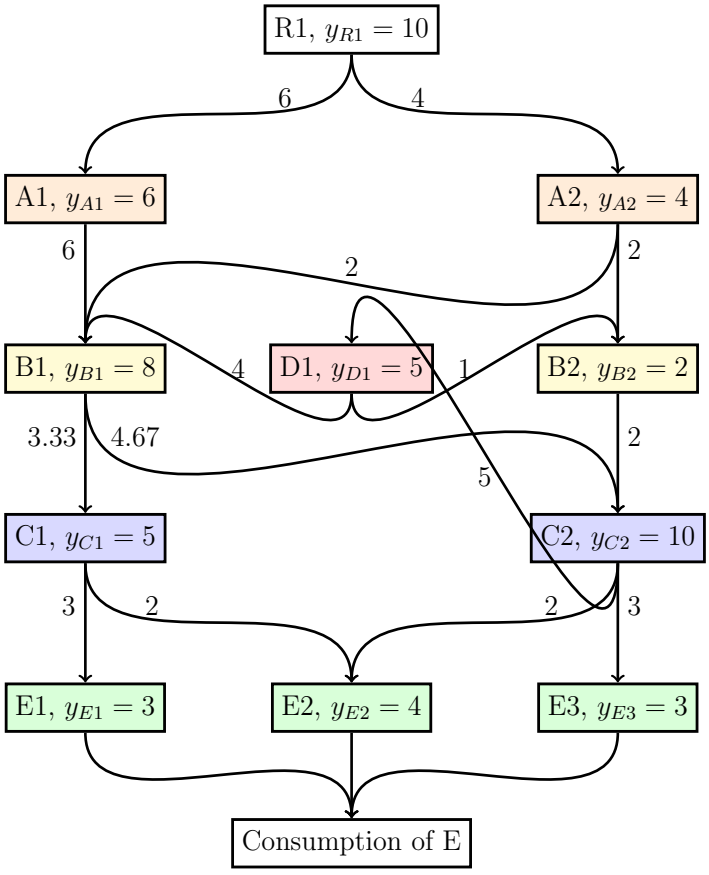
$$p_\tau = \sum_{k' \neq k} -\tau_{k'} \hat{p}_{\tau k'}.$$

Note that the left hand side is p and the right hand side is $\hat{p}$, capturing that this is the cost of output in n, while inputs are sourced from their cheapest source.

2.3 Example

Consider an economy with three countries 1, 2 and 3, intermediate goods R, A, B, C and D and a unique final good E. The used recipes for producing these goods are shown in figure 1, and the countries that produce different goods in equilibrium are shown in Figure 2. The labor endowments across the three countries are $\{37, 20, 3\}$ respectively, as we assume that all goods are perfectly mobile (with iceberg costs of 1), except for labor which incurs an iceberg cost strictly above 1 whenever labor from a given country is used in a technology located in a different country. When labor is used within a country it has an iceberg cost of 1. For this economy equilibrium flows of goods are shown in Figure 3, while the price of a given type of

Figure 3: Equilibrium flow network



goods is constant across countries and equal to

$$\left(\underbrace{\frac{1}{6}}_{\text{labor}}, \underbrace{\frac{1}{6}}_{R}, \underbrace{\frac{1}{3}}_{A}, \underbrace{1}_{B}, \underbrace{\frac{5}{6}}_{C}, \underbrace{1}_{D}, \underbrace{1}_{E} \right).$$

GDP is given by the value of the final goods produced, so in this example

$$GDP = \sum_n \sum_f p_f c_{fn} = 10.$$

2.4 Existence and Generic Uniqueness

To ensure equilibrium existence, it is necessary that the available technologies avoid money pumps. We have assumed that each technology uses a positive amount of labor. This assumption is sufficient for avoiding money pumps and is quite natural, although weaker conditions

could be used.[11]

LEMMA **1.** *There exists an equilibrium. Moreover, the same bundle of final goods is produced in all equilibria. Generically in T, the active set of technologies is the same in all equilibria.*

The proof of Lemma 1 appears in Appendix A. We show that production of final goods can be represented by direct technologies that directly convert (a vector of country-specific) labor into final goods. There is then an induced economy in terms of direct technologies, corresponding to sourcing choices that could be made to produce the final goods in the initial economy, such that there is an equilibrium of the induced economy if and only if there is an equilibrium of the original economy. It is easier to work in the space of the induced economy, and prove the existence of an equilibrium for this economy (by solving the Pareto problem and showing that these allocations can be supported by prices in an equilibrium). Moreover, this equilibrium is unique up to wages and consumptions across countries.

Although there is unique total consumption in equilibrium, this can sometimes be supported by multiple wage profiles, and the existence of multiple wage profiles is generic. To see this, suppose there are three countries, A, B and C and each have a unit of labor endowment. There is one final good f and two intermediate goods 1 and 2. Suppose that labor iceberg costs are sufficiently large that labor never moves, but in contrast intermediate goods have iceberg costs of 1, so that they ship with no loss. Suppose country A can only produce intermediate good 1, and it takes 1 unit of its labor to produce one unit of it. Country 2 can only produce intermediate good 2, and it takes 1 unit of its labor to do so. Country 3 can only produce the final good, and to produce one unit it has to combine 1 unit of intermediate good 1, 1 unit of intermediate good 2, and 1 unit of its labor. Thus it takes 1 unit of labor from each country to produce 1 unit of the consumption good, and in equilibrium exactly 1 unit of the final good is produced. However, any vector of wages that sums to 1 can be supported as part of an equilibrium. These wages just pin down the relative amounts of consumption that the three different countries can afford, and the consumption good market clears (as well as the labor good markets).

Note that this example is "robust" in the following sense: the multiplicity is robust to the ratios at which intermediate goods are combined to produce the final good, and how much labor is needed by each country in its production technology.[12] If instead, labor transportation costs are sufficiently low that labor in one country can substitute for that in another, then that begins to bracket relative wages, until in the extreme wages are equal across countries.

[11] We could instead just assume that each final good supply chain uses a positive amount of labor, but at the expense of some technicalities (including the possibility of zero prices for some intermediate goods) that lead to more complex and obscure proofs.

[12] Suppose we added in a technology for each country that converted its labor into the final good f. For country i suppose that $k_i > 3$ units of labor are needed to produce one unit of good f. Then the equilibrium wages that could be supported would be any vector summing to 1 such that the wage in country i is weakly greater than $1/k_i$.

3 Contrasting Short- and Long-Run Supply Chain Changes

3.1 Hulten's Theorem: Long-Run Adjustments to Changes in the Supply Chain

So far, we have normalized the output of each 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 good's technology.

Recall that GDP denotes the total expenditures on final goods:

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

Further, 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. Thus the utility of the representative consumer, denoted by U, is a measure of overall welfare.

Note that if we normalize the price of the final consumption bundle to 1, then given the homogeneity of preferences, GDP is proportional to the overall total utility.

To consider the long run impact of a shock we let $\tilde{p}_\tau$ be the weighted average price obtained for technology τ:

$$\tilde{p}_\tau = p_\tau \sum_{\tau'} \frac{\theta_{\tau\tau'} x_{\tau\tau'}}{y_\tau},$$

and then a total factor productivity shock that changes the output of a technology τ to λ of its initial value.

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,*

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

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 network impact of a shock 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.

3.2 Bottlenecks and Disruptions: Short-Run Adjustments to Changes in the Supply Chain

The predictions of Hulten's Theorem only apply in the extreme long run when changes can be fully adapted to. In contrast, in the short run, the impact of changes in productivity can be much 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. 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 the shocked technology is thus shocked, and so we also trace those impacts as they propagate downstream. For example if a supply chain involved a string of single-sourced inputs, without any cycles, then the impact of an X percent drop in some output would simply reduce the production of each downstream good by X percent, and so would reduce each downstream final good by X percent.

There are two complications to this calculation. One is that some of the goods might source the same input from multiple sources, and if not all of those sources are shocked, then the reduction can be less than X percent. The other is that if there are cycles in the supply network, then these can feedback and leading to repeated reductions.

The total impact of the disruption is captured via an appropriate minimum disruption problem. The solution to the problem is found by an algorithm that we subsequently define.

Starting from an equilibrium outputs and flows $(y_\tau)_\tau$ and $(x_{\tau\tau'})_{\tau\tau'}$, respectively, let $\Phi := \{\tau : y_\tau > 0\}$ be the set of active technologies. Consider a shock that reduces the outputs of technologies $\Psi \subseteq \Phi$ to $\lambda < 1$ of their initial level. In the short run, outputs and flows adjust to $(\hat{y}_\tau)_\tau$ and $(\hat{x}_{\tau\tau'})_{\tau\tau'}$ that 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

$$\begin{aligned}
&\hat{y}_\tau \leq \lambda y_\tau \quad \text{for all } \tau \in \Psi, && \text{(shock constraints)}\\
&\hat{y}_\tau \leq \left(\min_{k:\tau_k<0} \frac{\sum_{\tau':O(\tau')=k} \hat{x}_{\tau'\tau}}{\sum_{\tau':O(\tau')=k} x_{\tau'\tau}} \right) y_\tau \quad \text{for all } \tau \in \Phi, && \text{(technology constraints)}\\
&\hat{x}_{\tau\tau'} = x_{\tau\tau'} \left(\frac{\hat{y}_\tau}{y_\tau} \right) \quad \text{for all } \tau, \tau' \in \Phi. && \text{(proportional rationing)}
\end{aligned}$$

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 makes sense in the extreme short run when firms are not able to seek out new suppliers or switch technologies to substitute disrupted inputs for non-disrupted ones.. As prices begin to adjust, the rationing could adjust in a medium run, before new sources of production or alternative technologies could emerge.

The solution to the minimum disruption problem is found by the following algorithm. As illustrated in Figure 3, an equilibrium induces a directed and weighted network representing the flow of goods between technologies, which we denote by Ω. The node-set of Ω, $\mathcal{N}$, comprises the set of used technologies Φ and terminal nodes. The terminal nodes (denoted by $\mathcal{N}_f \subset \mathcal{N}$) represent the final goods. We index the nodes by i. The edges of Ω, denoted by $\mathcal{E}$, represent the flow of goods to the next stage of production. Each edge corresponds to both a quantity and the type of good. Abusing notation, we let Ω also denote the adjacency matrix for magnitudes of the flows with i, jth entry $\omega_{ij} = x_{ij} \geq 0$ for all $i, j \in \Psi$ and $\omega_{ij} = y_i \geq 0$ for all $i \in \Psi$, $j \in \mathcal{N}_f$.

We calculate, for each affected input of each affected node, 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 affected node declines. The overall impact on the economy is then given by the value of lost final good production resulting from the shock(s). This process is described in the Shock Propagation Algorithm[13] below.

The Shock Propagation Algorithm traces the impact of the shock and updating the output of a node each time the supply of one of its inputs is reduced. This continues until a step at which the outputs of all nodes has changed by less than ϵ. The algorithm terminates in finite time for any $\epsilon > 0$. One can also run it infinitely (which is only required if there are cycles in the affected part of the network), in which case it converges to a limit that is the solution to the problem.

Consider a shock to node $A2$ in Figure 3 which halves its output to 2. At each iteration of the algorithm the flows adjust as shown in Table 1.

Note that in the last iteration of the Shock Propagation Algorithm shown in Table 1 (iteration 5), the flow from $B1$ to $C1$ and $C2$ is updated but the flow from $B2$ to $C2$ remains fixed. This is because the supply of input D becomes the binding constraint for $C1$, while the supply of input B remains the binding constraint for $C2$. As the algorithm continues to iterate the output of $B1$ is reduced, which in turn reduces the output of $C2$, which subsequently reduces the output of $D1$ further, and hence the output of $B1$ once more. This cycle of diminishing outputs does not end in a finite number of updates to the network. Nevertheless, the output on each link in the network converges, and thus the algorithm provides a limit result. The limit outputs and flows are shown in Figure 4.

For this example the short run impact of the shock is to reduce GDP by 1/2 of its initial

[13]The notation $x \leftarrow y$ is used to denote the instruction to assign variable x the value of variable y: ie. set x equal to y.

Algorithm 1: Shock Propagation Algorithm

input : the flow network Ω, with edge weights ω_{ij}, edge types $\theta_{ij} \in \Theta$, and nodes $\mathcal{N}$, including terminal nodes $\mathcal{N}_c$
input : the nodes that are shocked Ψ
input : the shock amount $(1-\lambda) \in (0,1)$
input : the convergence threshold $\epsilon > 0$
output: the flow network after the shock Ω^* with edge weights ω^*_{ij}

$\Omega^* \leftarrow \Omega$;
$V \leftarrow \{j \in \mathcal{N} \setminus \mathcal{N}_c \ : \ \omega^*_{ij} > 0, \ i \in \Psi\}$;
forall j **do**
  $\omega^*_{ij} \leftarrow \lambda \omega^*_{ij}$;
$\delta \leftarrow \max_{j \in V}(1-\lambda)\omega^*_{ij}$;
repeat
  $V' \leftarrow \emptyset$;
  forall $j \in V$ **do**
    $\Theta(j) \leftarrow \{O(h) \text{ such that } \omega_{hj} > 0\}$;
    $u_j \leftarrow \min_{\theta \in \Theta(j)} \frac{\sum_{h|O(h)=\theta} \omega^*_{hj}}{\sum_{h|O(h)=\theta} \omega_{hj}}$;
    forall k **do**
      $\hat{\omega}_{jk} \leftarrow \omega^*_{jk}$;
      $\omega^*_{jk} \leftarrow u_j \omega_{jk}$;
    if $\max_k(\hat{\omega}_{jk} - \omega^*_{jk}) > 0$ **then**
      $V' \leftarrow V' \cup j$
  $\delta \leftarrow \max_{j \in V'} \max_k \left(\frac{\hat{\omega}_{jk} - \omega^*_{jk}}{\omega_{jk}} \right)$;
  $V \leftarrow \{k \in \mathcal{N} \setminus \mathcal{N}_c \ : \ \omega^*_{jk} > 0 \text{ for some } j \in V'\}$;
until $\delta < \epsilon$;

level. By contrast, Hulten's Theorem estimates the long run impact of the shock to be to reduce GDP by 1/15th of its initial level.

PROPOSITION **2.** *The Shock Propagation Algorithm 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 is the unique solution to the minimum disruption problem.*

The uniqueness of the limit and the fact that it solves the minimum disruption follows from several facts. The vector of inputs forms a complete lattice, and reductions in inputs in the short run are complementary to each other. Thus, the reductions at any step yield a new vector of input levels weakly above the solution to the problem, and any reductions at any step are necessary reductions. Once the algorithm terminates, there are no longer any necessary reductions and so the highest possible production levels that satisfy all the constraints is reached, and hence is the unique solution to the problem. The full proof, including the infinite case appears in the appendix.

	0	1	2	3	4	5	...	∞
R1-A1	6.000	6.000	6.000	6.000	6.000	6.000	...	**6.000**
R1-A2	4.000	4.000	4.000	4.000	4.000	4.000	...	**4.000**
A1-B1	6.000	6.000	6.000	6.000	6.000	6.000	...	**6.000**
A2-B1	2.000	**1.000**	1.000	1.000	1.000	1.000	...	**1.000**
A2-B2	2.000	**1.000**	1.000	1.000	1.000	1.000	...	**1.000**
B1-C1	3.333	3.333	**2.917**	2.917	2.917	**2.542**	...	**1.667**
B1-C2	4.667	4.667	**4.083**	4.083	4.083	**3.558**	...	**2.333**
B2-C2	2.000	2.000	**1.000**	1.000	1.000	**1.000**	...	**1.000**
C1-E1	3.000	3.000	3.000	**2.625**	2.625	2.625	...	**1.500**
C1-E2	2.000	2.000	2.000	**1.750**	1.750	1.750	...	**1.000**
C2-D1	5.000	5.000	5.000	**3.813**	3.813	3.813	...	**2.500**
C2-E2	2.000	2.000	2.000	**1.525**	1.525	1.525	...	**1.000**
C2-E3	3.000	3.000	3.000	**2.288**	2.288	2.288	...	**1.500**
D1-B1	4.000	4.000	4.000	4.000	**3.050**	3.050	...	**2.000**
D1-B2	1.000	1.000	1.000	1.000	**0.763**	0.763	...	**0.500**

Table 1: Flow adjustments for the first five iterations of the while loop in the Shock Propagation Algorithm, which converges to the flows shown in Figure 3.

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 and contrasts starkly with Hulten's Theorem. Consider a shock to some subset of the technologies Ψ. Let $F(\Psi)$ denote the set of final goods that either (i) are in Ψ directly or (ii) use an input good from a technology (directly or indirectly) that is shocked. Thus $F(\Psi)$ gives the set of all final goods that are impacted by the shock.

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

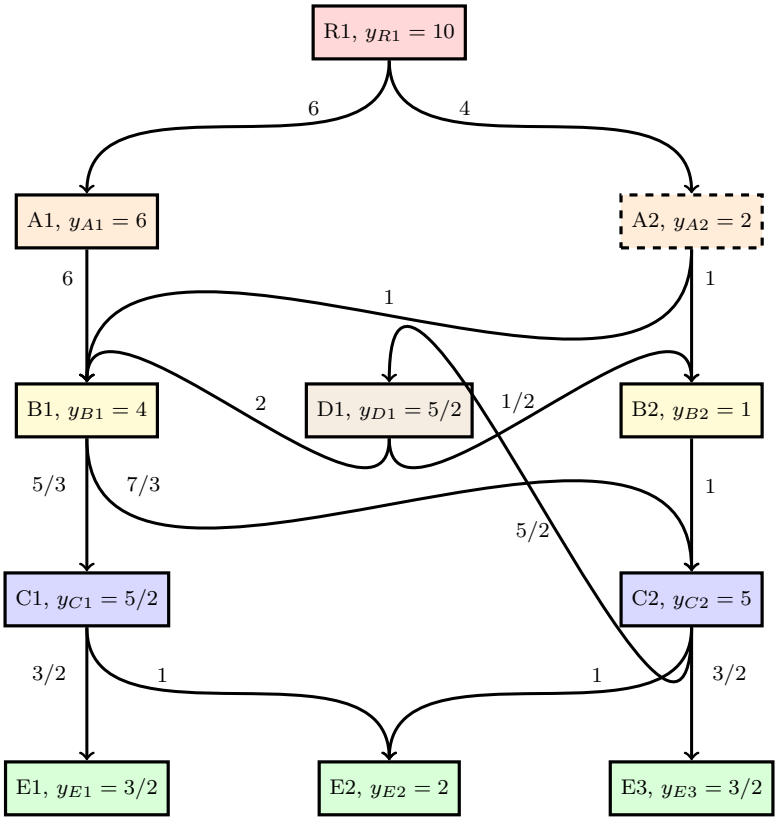
$$(1-\lambda)\left(\frac{\sum_{k' \in F(\Psi)} \sum_{\tau: O(\tau)=k'} p_{k'} y_\tau}{GDP}\right).$$

Proposition 3 provides an upper bound on the impact of a shock that is tied to the value of final goods related to the shocked goods, rather than the cost of the shocked goods, which is what matters in the long-run by Hulten's Theorem. As the disruption is just an upper bound, it is helpful to better understand when this bound is tight. We turn now to identifying some sufficient conditions under which the upper bound will be 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*, $\Omega(\Psi)$, is the sub-network induced on Ω by all nodes that are on a walk terminating at a final good node $F(\Psi)$.

When a *disrupted industries sub-network*, $\Omega(\Psi)$, is acyclic, there exist a set of nodes/technologies

Figure 4: Equilibrium Flow Network after a shock



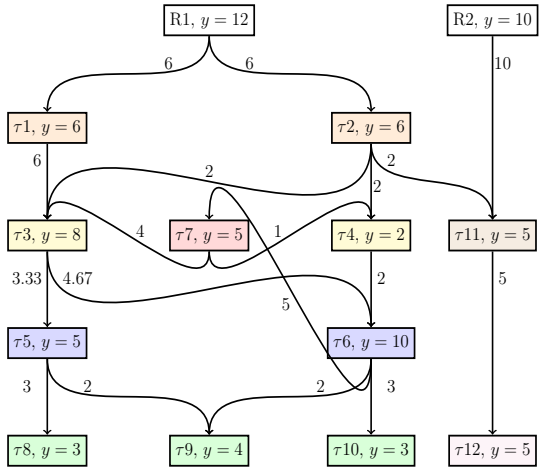
The equilibrium flow network from Figure 3 after a shock to $A2$ reduces its output and the impact of this propagates. Note that overall output of the final good E is cut in half, even through the shock only reduced output of good A by 1/8th.

that have no in-links. We denote these technologies $R(\Psi)$. 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.

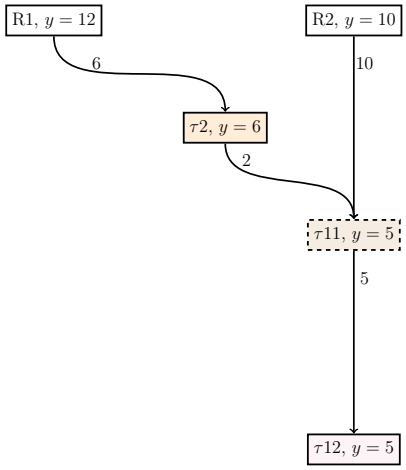
PROPOSITION **4.** *If the disrupted industries sub-network $\Omega(\Psi)$ is acyclic and the set of edges adjacent to the set of shocked technologies forms an $(R(\Psi), F(\Psi))$-cut set of the disrupted industries sub-network $\Omega(\Psi)$,*[14] *then the bound from Proposition 3 binds.*

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 Ψ constitutes a cut on the network that restricts attention to related goods (e.g.,

[14]Note that this permits the set of shocked technologies to intersect both $R(\Psi)$ and $F(\Psi)$.



(a) Equilibrium flows Ω



(b) Disrupted industries sub-network $\hat{\Omega}(\tau 11)$

Figure 5: Disrupted industries sub-networks: For the equilibrium flows in panel (a), panel (b) shows the disrupted industries subnetwork for $\Psi = \{\tau 11\}$. The disrupted industries sub-network for $\Psi = \{\tau 2\}$ is the network shown in panel (a), while the disrupted industries sub-network for $\Psi = \{\tau 1\}$ is the same as the equilibrium flows network shown in Figure 3.

disconnects raw materials from their related final goods). The example shown in Figure 5b illustrates a shock which, by Proposition 3, must obtain the bound.

Figure 4 shows that when the technology network is not acyclic, the bound can be obtained even when the edges adjacent to the shocked technologies do not constitute a cut set. For this example $A2$'s output is reduced to 1/2 its original value. As there is only one final good E, and $A2$ is used (indirectly) in the production of good E, Corollary 3 implies that an upper bound for the short run impact of the shock is 1/2 the value of GDP. Given that producer $A2$ only accounts for 2/5 of the production of good A, and only half of its output is lost, it might seem as though this upper bound is unlikely to bind. However, as we see in Figure 4, the upper bound does bind and the short-run impact of the shock is to reduce GDP by 50 percent. In contrast, the long-run impact of the shock, by Hulten's Theorem, is to reduce GDP by 20 percent.[15]

PROPOSITION **5.** *If all technologies for producing any given good use the same set of inputs and for any shocked technology $\tau \in \Psi$ all producers of good $O(\tau)$ are also in Ψ, then the bound from Proposition 3 binds.*

[15]It is easy to write down with examples where this difference is starker. Indeed, we could adjust this example to shrink the long run impact below any strictly positive amount, while leaving the short run impact fixed. Suppose we adjust the equilibrium output of $A1$ to $10 - \varepsilon$ and the equilibrium output of $A2$ to ε for $\varepsilon \in (0, 4]$, while propagating these flows through the network to reach an alternative equilibrium (but with the same prices). Specifically, this increases the flows from $A1$ to $B1$, reduces the flows from $A2$ to $B1$ and $A2$ to $B2$, increases the production of $B1$ and the flow from $B1$ to $C2$, and reduces the equilibrium production of $B2$ and the flow from $B2$ to $C2$, while holding all other outputs and flows fixed. Then for all $\varepsilon < 4$, a shock that reduces the output of $A2$ to 1/2 its initial level also reduces GDP to 1/2 its initial level in the short run, but in the long run reduces GDP to $\varepsilon/20$ of its initial value, which can be made arbitrarily small.

Proposition 5 does not follow from Proposition 4 as it allows for cyclic production and also it allows for cases in which the technologies that produce some good do not form the basis for 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 here. Each final good in $F(\Psi)$ has a path to it from some shocked good. Tracing back from the affected input good into the final good, all of the technologies that it sources must involve paths that go back to some shocked good (given that technologies producing any given good use the same inputs and then all producers of any shocked good are shocked). Thus, all of those paths are reduced by $1-\lambda$, and so that input good has a shortage of $1-\lambda$.

Note under Proposition 5 is that if all production of a given good is located in the same country, then country-technology specific shocks become industry shocks, and the upper bound is obtained. 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, 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. 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 a tiny value if the chip is relatively cheap.

We emphasize that if the bound is obtained, the short-run impact of a disruption to some technology τ with output good k (in a world which technologies are all single-sourced) 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)[16] impact is, by Hulten's Theorem,

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

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 $\tilde{p}_\tau y_\tau$.

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 will be no impact on final good production. This again contrasts with the long-run impact of

[16] This emphasizes another contrast, in that the long-run impact is a marginal calculation, while our short run impact holds for larger changes.

such changes as captured by Hulten's Theorem, where the (marginal) impact of positive and negative shocks is symmetric.

3.3 Supply Chain Complexity and Increased Fragility

We find that shocks have more impact as production becomes more complex, all else held equal.

To make this precise, we consider economies such two technologies producing the same good use the same set of inputs (although in possibly different ratios) and shocks that affect goods rather than specific production locations (so all producers of good are affected). Thus Proposition 5 applies.

Let S denote the average number of different types of goods used directly or indirectly to produce a final good—averaged across final goods. Thus, S is a measure of the complexity of supply chains.

PROPOSITION **6.** *[Aggregate Supply Chain Disruption as a Function of Complexity] Consider a generic specification of an economy and one in which final goods prices are independent of the complexity of their supply chains. Consider independent chances of a fractional disruption of each good's production of size $-D$ with independent probability π and let ΔGDP denote the initial equilibrium GDP minus the GDP after any shocks. Then for small π small (so that the probability of two shocks on the same chain are vanishingly small relative to the probability of a single shock):*

$$\mathbb{E}\left[\Delta GDP\right] \approx -\pi DS.$$

The proof of Proposition 6 is straightforward and so we simply outline it here: Let $m = SF/(M+F)$ be the average number of supply chains that a given input lies on. Given that $\pi(M+F)$ is expected number of disruptions, and then each hits m chains, and that is a fraction of m/F of the total set of goods. Thus, $\pi(M+F)m/F = \pi S$ is the expected fraction of final goods shocked (here appealing to the fact that when π is small the limit probability is that a single technology is shocked at a time, and so we do not worry about multiple shocks to the same supply chain). A disrupted chain reduces consumption of the final good by D, and then disrupting each final good by D with a probability πS gives the result.

4 Globalization

Next, we study the impact of reduced transportation costs. The following proposition characterizes how specialization increases with reduced costs, which we then trace through to resulting changes in the equilibrium production network affects overall fragility.

PROPOSITION **7** (Specialization)**.** *Suppose that T_n is finite for each n and generic.*[17] *If iceberg costs are sufficiently low there is full specialization: There exists a threshold on iceberg cost*

[17] Generic indicates that the vectors differ for each technology.

$\bar{\theta} > 1$ *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.

Note that it could be that more final goods are produced in the limit, as some technologies might be too expensive to be part of an equilibrium in the face of higher transportation costs. So, the proposition does not have implications about the number of goods produced in the world, but it does say that there is a unique technology used for each good in the limit, and is produced in just one country.

An implication of Propositions 3 and 7 is that moving to a frictionless economy results in increased fragility in a certain sense.

COROLLARY **1** (Fragility)**.** *Suppose that* T^n *is finite for each* n *and generic. There exists a threshold on iceberg 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 is binding.*

The fact that the bound from Proposition 3 is binding can be seen as follows. Every affected final good has a path from some shocked good to it. Since goods are all single-sourced, each good on the path between the final good and the shocked good are fully reduced by λ.

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.

5 Concluding Remarks

We comment on several further explorations for which the model can serve as a foundation.

5.1 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]

[18] Bachmann et al. (2022) 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, and find substantial differences.

5.2 Flexible Pricing and Proportional Rationing

There is a spectrum of alternatives to the extreme short and long runs that we have analyzed for here. For example, a natural third case to examine is one in which prices of goods in shortage adjust to direct them to the most valuable production chains. For example, when there is a shortage of chips, if their price rises then they might, for instance, continue to be used in the manufacture of expensive goods, but the production of cheap novelty toys might almost entirely stop, which would then change the ultimate impact on total utility. This is a sort of "medium-run" analysis where the production technologies and chains are not adjustable, but relative proportional mixes of inputs can be repriced and reallocated. It is an important next step that can build on this model.

5.3 Inventories and Endogenous Robustness

In our a benchmark of a short-run disruption, disrupted goods are fully missing from production. Firms can maintain 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.

References

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

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.

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, 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, 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.

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.

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* .

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.

Kalemli-Özcan, ?ebnem, Alvaro Silva, Muhammed A Yildirim, and Julian di Giovanni. 2022. "Global Supply Chain Pressures, International Trade, and Inflation." Tech. rep., Federal Reserve Bank of New York.

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.

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.

A Omitted proofs

LEMMA **2.** *In any equilibrium aggregate consumption is equivalent to the consumption choice of a single representative consumer with preferences represented by the utility function* $U(c_1, \ldots, c_F)$ *and with wealth* $\sum_n L_n p_{Ln}$.

Proof of Lemma 2:

The country n consumer solves

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

As the utility function is $U(\cdot)$ is 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 iceberg 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 will be a re-scaling of the same bundle of goods. Moreover, it follows that a representative consumer with utility function $U(c_1,\ldots,c_F)$ and with wealth $\sum_n L_n p_{Ln}$ would also choose a re-scaling of the same bundle of goods, and that re-scaling must equal aggregate consumption. ∎

A.1 Proof of Lemma 1

We prove existence of an equilibrium by mapping our economy into one without iceberg costs, and then applying standard results. In order to do this we use the iceberg costs to map each technology $\tau \in I\!R^{1+M+F}$ into source-specific technologies $t \in I\!R^{N+T}$. Specifically, for each technology τ that required k inputs, we create N^k source-specific technologies allowing for each possible combination of sourcing choices 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 iceberg costs. So, if technology τ requires k units of good g, and the iceberg costs associated with sourcing good g from country n are 2, then corresponding source-specific technologies t that source good g from country n, will require $2k$ units of good g.

Replacing iceberg costs and technologies with source-specific technologies, our environment is a special case of that considered in 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 2 our economy admits a representative consumer with utility $U(c_1,\ldots,c_F)$, all Pareto efficient allocations must maximize this utility function subject to feasibility constraints. We will 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 $c \in I\!R^F_+$ 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 conditions 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 will 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 $(1-\lambda)c'$ will be 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 is 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) will be unique in equilibrium. As utility is strictly 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.

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

Sketch of the 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 $(c_1, \ldots, c_F) \in \mathbb{R}^F_+$:

$$U(c_1, \ldots, c_F) = \sum_f c_f \frac{\partial U}{\partial c_f}. \tag{5}$$

By Lemma 2 total world consumption choices are as if there is a representative consumer with preferences represented by the homogeneous of degree 1 utility function $U(c_1, \ldots, c_F)$ 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

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

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 (5) it follows that

$$U(c_1, \ldots, c_F) = \lambda \sum_f c_f p_f. \tag{6}$$

Next, consider an allocation that is Pareto efficient subject to production and market clearing constraints (1)-(4), which is true of any equilibrium. Given the representative consumer, this is equivalent to maximizing the total utility subject to (1)-(4). In particular, the constraint of (4) is the one containing c_f and in the context of the representative consumer and now allowing τ_f to change from one becomes:

$$c_f - \sum_{n,\tau \in T^n : O(\tau) = f} \tau_f y^{\tau n} = 0 \tag{7}$$

The Envelope Theorem and first order conditions imply that for a given production technology τ with $O(\tau) = f$,

$$\frac{\partial U}{\partial \tau_f} = \mu_f y^{\tau n} \tag{8}$$

where μ_f is the Lagrange multiplier on the final good clearing constraint (7)

Then from (8),

$$\frac{\partial \log(U)}{\partial \log(\tau_f)} = \frac{\partial U / U}{\partial \tau / \tau} = \mu_f y^{\tau n} \frac{\tau}{U}$$

and so from (6) it follows that

$$\frac{\partial \log(U)}{\partial \log(\tau_f)} = \frac{\mu_f \tau_f y^{\tau n}}{\lambda \sum_{f'} p_{f'} c_{f'}} = \frac{\mu_f y^{\tau n}}{\lambda \sum_{f'} p_{f'} c_{f'}}. \tag{9}$$

Next, let us show that μ_f is equal to λp_f. From the first order conditions on utility maximization of a consumer

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

From the Pareto efficiency maximization problem,

$$\frac{\partial U}{\partial c_f} = \mu_f,$$

and so $\mu_f = \lambda p_f$. Therefore by (9)

$$\frac{\partial \log(U)}{\partial \log(\tau_f)} = \frac{\lambda p_f y^{\tau n}}{\lambda \sum_{f'} p_{f'} c_{f'}} = \frac{p_f y^{\tau n}}{\sum_{f'} p_{f'} c_{f'}}.$$

Note that as $GDP = \sum_f p_f c_f$, by (6) $U = \lambda GDP$, and so we also have

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

To extend this to intermediate goods, note that it must be that for each n, τ used in the production of some final good f and an intermediate good k used in τ

$$y^{\tau n} = \sum_{n'} \frac{x_{kn'}^{\tau n}}{-\tau_k \theta_{kn'}^n}.$$

Then considering $\tau' \in T^{n'}$s such that $O(\tau') = k$ we can write this as

$$y^{\tau n} = \sum_{n'} \frac{\tau'_k x_{kn'}^{\tau n}}{-\tau_k \theta_{kn'}^n}. \tag{10}$$

Then the equivalent of (8) is

$$\frac{\partial U}{\partial \tau'_k} = \phi_k^f \frac{x_{kn'}^{\tau n}}{-\tau_k \theta_{kn'}^n}, \tag{11}$$

where ϕ_k^f is the multiplier on the constraint (10).

Next, note that from the Pareto problem

$$0 = \frac{\partial U}{\partial y^{\tau n}} = 0 = -\mu_f + \sum_k \phi_k^f,$$

and since $\mu_f = \lambda p_f$ it follows that

$$\lambda p_f = \sum_k \phi_k^f.$$

Noting that $p_f = \sum_k -\tau_k \theta_{kn'}^n \widehat{p}_k^n$, [[fixme: here presuming each is singly sourced, but must hold for each - insert footnote]]

$$\lambda \sum_k -\tau_k \theta_{kn'}^n \widehat{p}_k^n = \sum_k \phi_k^f.$$

These are unique multipliers and a solution is $\phi_k^f = -\lambda\tau_k\theta_{kn'}^n\widehat{p}_k^n$, and then rewriting (11)

$$\frac{\partial U}{\partial \tau_k'} = -\lambda\tau_k\theta_{kn'}^n\widehat{p}_k^n\frac{x_{kn'}^{\tau n}}{-\tau_k\theta_{kn'}^n} = \lambda\widehat{p}_k^n x_{kn'}^{\tau n}$$

Then, following the same argument as for final goods given above, it follows that

$$\frac{\partial \log(U)}{\partial \log(\tau_k')} = \frac{\lambda\widehat{p}_k^n x_{kn'}^{\tau n}}{\lambda\sum_{f'} p_{f'}c_{f'}} = \frac{\widehat{p}_k^n x_{kn'}^{\tau n}}{\sum_{f'} p_{f'}c_{f'}}.$$

which is the desired conclusion. This is for an intermediate good used in the production of a final good, but one can then iterate on this argument for further goods, including labor. Summing this across all the n in which the production of τ' is used in, leads to the expression $\frac{\widehat{p}_k^n y^{\tau n}}{GDP}$. ∎

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 $S = \prod_{j=1,\ldots,|\Phi|}[0,\phi_j]$. 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 $\Gamma : S \to S$ 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. Then by the definition of Γ any fixed point of it is a solution. Given that Γ is continuous on a compact space, it is direct to check that any limit point is a fixed point.

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 Ω. This proves that for any path $\mathcal{P}$ starting at an affected node i we have that $\omega_{jk}^* < \omega_{jk}\ \forall j,k \in \mathcal{P}$.

To show maximality with respect to network structure notice that as all flows ω^* are feasible

by construction of the algorithm, these flows represent the maximum that can be produced final good by final good after the disruption. □

A.4 Proof of Proposition 3

Proof. The upper bound follows from the algorithm directly, noting that if no good's production is ever below $1-\lambda$ of its initial level, then given the structure of the production functions, no output is ever below $1-\lambda$ of its initial level. Thus, given that all goods start with at least $1-\lambda$ of their initial levels, any limit point of the algorithm has production of all goods at least $1-\lambda$ of initial levels. □

A.5 Proof of Proposition 4

Proof. Suppose that the edges adjacent to Ψ constitute a $(R(\Psi), F(\Psi))$-cut. We can attribute the equilibrium flows (before the shock) to a set of 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 Ψ to be an $(R(\Psi), F(\Psi))$-cut every path in the network leading to every final good $k \in F(\Psi)$ must contain at least one node in Ψ; and (ii), that for any path that passes through a node in Ψ, 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 $F(\Psi)$ are reduced by a proportion $1-\lambda$ of their initial level, and the upper bound is obtained. □

A.6 Proof of Proposition 7

Proof. Suppose $\max_{\tau,\tau'} \theta_{\tau\tau'} = \max_{n,\tau'} \theta_{n\tau'} = 1$ (so there are no iceberg cost). By Lemma 1, there are generically unique equilibrium outputs and, by Lemma **??**, we can look for an equilibrium of the induced direct technology economy $\hat{\mathcal{E}}$. Moreover, as it is costless to reallocate labor across countries, the input requirements for each direct technology $\hat{\tau} \in D(T)$ can be represented by a single number $\xi(\hat{\tau}) > 0$ representing the number of units of labor (which are not country specific) required to produce one unit of $O(\hat{\tau}) \in F$. Thus any equilibrium outputs, by the proof of Lemma 1, must solve:

$$\max_{\{y_{\hat{\tau}} \geq 0\}_{\hat{\tau}}} \sum_{\hat{\tau}: O(\hat{\tau}) \in F} p_{\hat{\tau}} y_{\hat{\tau}} \qquad \text{subject to} \quad y_{\hat{\tau}} = \ell_{\hat{\tau}} \xi(\hat{\tau}),$$

where $\ell_{\hat{\tau}}$ is the amount of labor used by direct technology $\hat{\tau}$.

This implies that only the most efficient direct technologies for producing a given final good can be used—i.e., $q_{\hat{\tau}} > 0$ only if $\xi(\hat{\tau}) \leq \xi(\hat{\tau}')$ for all $\hat{\tau}'$ such that $O(\hat{\tau}') = O(\hat{\tau})$. Generically, there is a unique such (simple) technology τ^* for producing a given final good f. Moreover, consider

the corresponding supply chains in the economy $\mathcal{E}$. For each such supply chain producing final good f, the cost of sourcing each required input must be the same, but generically there will be a unique lowest cost producer of each input, so there must be a unique such supply chain. Moreover, if any other final good is produced using one of these intermediate goods, the same lowest cost supplier must be used. Thus, in generic economies there will be full specialization.

Now suppose the maximum iceberg cost is greater than 1. This changes the technologies available in the corresponding simple economy so we work with the corresponding economy but consider the used supply chains S^* from the frictionless economy. First note that for all iceberg cost any combination of supply chains must produce less than was produced in the frictionless economy using supply chains S^*. Generically, as shown above, the supply chains S^* are the unique ones that could be used in frictionless economy to maximize GDP. Thus, 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 $\max_{n,\tau'} \theta^{n\tau'} < \bar{\theta}$ then the supply chains S^* are still the unique ones that maximized GDP. Hence these are the unique supply chains that are used in equilibrium and there is full specialization. □