

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

Matthew L. Elliott and Matthew O. Jackson*

November 2025

Abstract

We introduce a parsimonious multi-sector model of international production and use it to study the impact of a disruption in the production of some goods propagates to other goods and consumers, and how that impact depends on the goods' positions in, and overall structure of, the production network. We show that the short-run impact of a disruption can be dramatically larger than the long-run impact. 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, lowering the chances for disruption but increasing the impact conditional upon disruption. We use the model to characterize the power that a country has over others via diversions of its production as well as quotas on imports and exports.

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 David Baqaee, 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. Alastair Langtry, Shane Mahen, Sarah Taylor and Yifan Zhang provided excellent research assistance. Disclosure: Jackson is a scientific advisor of (with stock options in) Altana AI, which is involved in collecting supply chain data and advising companies and governments on supply chain issues.

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, having global economic implications.³ How can we predict the economic impact of the disruption in the supply of one or more goods in a production network?

Despite repeated disruptions of supply chains due to disease, natural disasters, cyber failures, transportation failures, political crises, economic sanctions, changes in tariffs and regulations, and other events, there are no models that allow economists and others to analyze the short-run, “out-of-equilibrium” impact of different shocks on the production of goods and the shortages of final good production that ensue.⁴ The existing models that are used to estimate the impact of shocks are all equilibrium models—whether general equilibrium, imperfect competition, or some partial equilibrium model. Although these can be used to estimate the long-run impact of some shocks, they neglect important constraints on production that are needed to assess the short-run impact. We show that these constraints are important as they can significantly amplify shocks in the short-run and the magnitude of this amplification varies substantially the network position of the shocked industry.

It is important to emphasize that existing models cannot be adjusted to estimate short-run disruptions simply by reducing the elasticity of substitution between inputs and/or making labor immobile. Although this reduces the ability of the economy to adjust in some ways, it is conceptually and qualitatively distinct from an out of equilibrium analysis in which resources (especially labor) are wasted and production is temporarily halted. Constraints capture the reality that some shocks in the short run are such that alternative suppliers cannot be quickly

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 .

³There are a variety of studies documenting that disruptions propagate through supply chains Barrot and Sauvagnat (2016), Boehm, Flaaen, and Pandalai-Nayar (2019a) and Carvalho et al. (2021).

⁴Individual disruption incidents have impacts in the tens to hundreds of billions of dollars, and widespread events such as COVID have impacts in the trillions of dollars (Loh and Van Thai, 2015; Durugbo and Al-Balushi, 2023; Aldrighetti et al., 2021; Hägele, Grosse, and Ivanov, 2023; Kamiya et al., 2018; Queiroz et al., 2022). Trade Disruption Insurance (Business Interruption Insurance) is a rapidly growing industry that offers companies insurance against specific types of supply chain disruption they face, and supply chain management is an increasingly important part of most businesses (Durugbo and Al-Balushi, 2023). Despite the huge economic importance of this, the economic models that are available take a long-run approach that limits their usefulness for industry. Instead, industry relies on off-the-shelf tools from the operations management literature that take a localized approach and neglect systemic effects.

found and that supply relationships governed by explicit or implicit contracts prevent prices from adjusting and efficiently rationing goods. For illustration, consider the chip example discussed above, and suppose there is a product that must have a particular kind of fairly inexpensive chip to function (e.g., a car) and so the demand for that chip is very inelastic. If the car manufacturer sources that chip (directly or indirectly via its direct suppliers) from one particular plant or country and some event, for instance political, closes access to that plant then the production of the car is completely disrupted in the short run. This inefficient shutdown of production is a fundamentally out-of-equilibrium phenomenon. In an equilibrium model, there would instead be more expensive production of that chip, and that would drive up the cost of the chip, but as a percentage of the price of the car, the impact would be relatively small even with completely inelastic demand for the input. However, stockpiled inventories are quickly exhausted (even if optimally redirected) and it can take months if not years for new sources of production to arise. Indeed, during the pandemic, this is exactly what we saw.⁵ Thus, an equilibrium model that omits the misallocation of resources that occurs in the aftermath of disruptions, does not account for lags in building new production, and abstracts from limitations in reallocating existing production, can fail to account for much if not most of the short-run impact.⁶

In this paper, we introduce an international production and trade model that incorporates such short-run constraints on production, and can also evaluate full equilibrium adjustments. Thus, we can use it to characterize and contrast the extremes of short-run as well as long-run impacts of shocks to production. The extremes are both benchmarks, but both are of importance, as well as their contrast. We use the model to characterize the impact of production disruptions and how this depends on the 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 producers' productivities. We also show how the potential for short-term disruption increases linearly in the complexity of supply chains (for small disruption probabilities) and how this compares with the long run. In addition, we examine how the potential for disruption changes as transportation costs drop and supply chains evolve to be more specialized and involve more trade across countries. We also use the model to characterize the power that one country has over some other(s) by disrupting trade flows in and/or out of its borders. This has become increasingly important as international trade bargaining stakes have increased.

The model is important to have in hand given the increasingly complex supply chains that have emerged over time, and the fact that the production of some key inputs is highly consolidated,⁷ while others are not. Thus, the model also needs to be careful in tracing out supply

⁵See, for example, Reuters, (06-222021).

⁶Extensive anecdotal evidence of production being disrupted or halted because of difficulty sourcing inputs is presented in Section 2 of Elliott and Golub (2022).

⁷For 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/>).

chains and the full network of production in order to quantify the potential for disruptions to propagate.

We show that in the long run, when the change in productivity of an input can be completely compensated for by equilibrium adjustments of quantities and sources of all inputs, a version of Hulten’s (1978) 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.⁸ However, in the short run, the impact is much larger. A 5 percent shortage of integrated circuits propagates to prevent the production of 5 percent of all 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 proportional to the value of all downstream final products, while the long-run impact is proportional to the change in the cost of producing the input itself.

Central to our theory is that shocks propagate through the supply network and are amplified in the short-run. Much evidence is consistent with this. [Carvalho et al. \(2021\)](#) analyze the Great East Japan Earthquake and the resulting propagation of supply shocks. The localized impact of the disaster accounts for less than a 0.1 percentage point reduction in overall GDP growth, and this includes the propagation of disruptions within the affected areas. So, the direct impact of the shock is likely to be considerably below this. At the same time, there was a 0.4 percentage point reduction in the growth of GDP year-on-year and their counterfactual analysis estimates the actual impact at 0.47 percentage points. Further, tracing through affected customers and suppliers in the year following the shock, they find that the impact diminishes, but slowly. A further point of interest is that while they find inputs to be weak substitutes on the time frame of a year, [Barrot and Sauvagnat \(2016\)](#) find them to be complements on a quarterly basis. As we show, this is a prediction of our theory. Intuitively, as time progresses (which can be years, as in the [Carvalho et al. \(2021\)](#) case) constraints are relaxed. This allows alternative suppliers to increase their production where multisourcing is already present, and for new suppliers to be found. This makes inputs more substitutable over time.

The lack of a short-run model able to capture out-of-equilibrium propagation through a network, and inadequacy of a long-run equilibrium approach to substitute for such calculations has been lamented by Larry Summers:

“I always like to think of these crises as analogous to a power failure or analogous to what would happen if all the telephones were shut off for some time. Consider such an event. The

⁸This 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.

network would collapse. The connections would go away. And output would, of course, drop very rapidly. There would be a set of economists who would sit around explaining that electricity was only 4% of the economy, and so if you lost 80% of electricity, you couldn't possibly have lost more than 3% of the economy. Perhaps in Minnesota or Chicago there would be people writing such a paper, but most others would recognize this as a case where the evidence of the eyes trumped the logic of straightforward microeconomic theory. ... And we would understand that somehow, even if we didn't exactly understand it in the model, that when there wasn't any electricity, there wasn't really going to be much economy.” (Summers, 2013).⁹

We develop an algorithm to characterize how the impact of a short-run disruption depends on the structure of the supply chains involved, and show that this algorithm finds the solution to a well-defined constrained optimization problem. 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 constrained, but in which there is equilibrium adjustment so that goods flow to the highest-value 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, the prices of chips can adjust and redirect the chips to the uses where they are most valued, and this can improve production. We show that this mitigation can lead to an impact that falls anywhere between the effect of the short versus long run. We show that this medium-run impact depends on the diversity in the values of the downstream production using an input.

⁹See Carvalho and Tahbaz-Salehi (2019) for some further discussion around this quote and the impact of equilibrium disruptions when Hulten's theorem doesn't hold. Of particular note is Baqaee and Farhi (2019). See also Wei (2003) for a discussion of the timing of adjustments.

We use our characterizations to 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 lie downstream of the shocked inputs.

We then use this to examine some comparative statics in trade costs. When all trade costs are reduced sufficiently the local prices of intermediate goods are equalized across locations and are sourced from a lowest cost technology and so 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 or other directed trade disruptions and anticipating their effects.¹⁰ As such, the model serves as a foundation to understand the power that one country holds over another in terms of how much disruption the first can cause in the second by disrupting imports and/or exports—e.g., via sanctions, quotas, or tariffs. We measure this by examining the damage a disrupting country can cause to a target country per unit of damage to the disrupting country’s domestic industries. In measuring the impact of intentional disruptions we need to track the distribution of disrupted (idle) labor across countries. We develop an algorithm for evaluating the total disrupted labor, up and downstream from a given disrupted technology and show how the potentially combinatorially intensive exercise of identifying the disruptions that maximize one country’s power over another, simplifies. Intuitive aspects emerge from our characterization of power: one country wields substantial power over another when it can disrupt a relatively low value-added industry domestically and this has far-reaching up and/or downstream consequences in the target country, that do not cause significant further disruptions at home. An example of such an industry in practice is the industry for rare-earth metals and magnets in China, which are (directly and indirectly) important inputs into high-value production in the US among other countries.¹¹ Our analysis also shows how it is possible for two countries each to have substantial power over each other (via different supply chains), or for power to be very asymmetric, depending on the structure of the supply network.

It is worth emphasizing that short-run disruptions in supply chains add up. Although widespread short-run shocks that cause multiple percentage point losses in GDP such as COVID-19 are rare, shocks to industries that cause temporary shortages and price spikes to

¹⁰In this context, 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.

¹¹See, for example, Reuters (04-04-2025).

consumers, such as the car industry example discussed above, disruptions in shipping lanes, local wars, strikes, natural disasters, climate shocks, trade wars, etc., are common and can lead to death-by-a-thousand-cuts. Thus, it is important to understand and better manage such short-run disruptions. It could be that in an ideal world without any such disruptions, GDP would be percentage points higher.

Related Literature 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.¹² 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](#); [Baqae, 2018](#); [Oberfield, 2018](#); [Acemoglu and Tahbaz-Salehi, 2020](#); [Acemoglu and Azar, 2020](#); [Baqae and Farhi, 2021](#); [Kopytov et al., 2021](#); [König et al., 2022](#); [Chen et al., 2023](#); [Grossman, Helpman, and Lhuillier, forthcoming](#); [Pin, 2025](#)). 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 by introducing a first model that incorporates constraints into networked production.¹³

Perhaps closest to us is research that considers the impact of shocks in the presence of frictions, including limited substitutability of inputs, fixed factors of production, building delays and search and matching frictions (e.g., [Baqae and Farhi \(2019\)](#), [Bachmann et al. \(2022, 2024\)](#), [Bui et al. \(2022\)](#), [Pellet and Tahbaz-Salehi \(2023\)](#), [Bai et al. \(2024\)](#), [Baslandze and Fuchs \(2025\)](#), [Brancaccio, Kalouptsi, and Papageorgiou \(2025\)](#), [Arkolakis, Huneeus, and Miyauchi \(2025\)](#) and [Schaal and Taschereau-Dumouchel \(2025\)](#)). In terms of the typology of time horizons we develop in Section 2.1, that work relates to our long-run or medium-run analysis. Our main focus is on the short-run which fundamentally differs by, instead of being in (constrained) equilibrium, being out-of-equilibrium (through the inability of prices to adjust to clear markets).

There is also considerable work studying networks and fragility.¹⁴ 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

¹²See [Carvalho and Tahbaz-Salehi \(2019\)](#) and [Baqae and Rubbo \(2022\)](#) for recent surveys.

¹³There is a literature on “disequilibrium” models, beginning with [Galbraith \(1947\)](#); [Barro and Grossman \(1971\)](#); [Benassy \(1977\)](#). That literature is more interested in modeling macroeconomic dynamics under price controls or other frictions on convergence to equilibrium prices, rather than on supply chain disruptions, and hence those models are not closely related to ours.

¹⁴See [Elliott and Golub \(2022\)](#) for a recent survey, while [Baldwin and Freeman \(2022\)](#) reviews the literature on risks in global supply chains.

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.¹⁵ 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. They find that the empirical impacts are large in magnitude, consistent with the contrast that we draw between the short- and long-run.

Finally, there is related research on networks and international trade.¹⁶ 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 \(2019\)](#); [Grossman, Helpman, and Redding \(2021\)](#)). 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

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.¹⁷

¹⁵The 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.

¹⁶For recent surveys see [Bernard \(2018\)](#) and [Antràs and Chor \(2022\)](#).

¹⁷Simply 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 depend on costs of inputs and so identically-produced goods end up with the same prices in equilibrium.

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.¹⁸

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.

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.

Throughout we associate each technology with a unique producer, but nothing would change if we instead associated the production of a given good g in a given country n with a unique producer endowed with all technologies $\{\tau \in T_n : O(\tau) = g\}$.

An example of technologies appears in [Figure 1](#).

[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.

¹⁸This simply makes prices and examples easier to compute in our model, but would not affect our results (Hulten’s Theorem would still hold in the long run, and our constrained approach does not adjust labor inputs in the short run).




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




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

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 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.

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,¹⁹ 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.

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

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

defined on $\mathbb{R}_+^F$ that is increasing ($c \geq c', c \neq c'$ implies $U(c) > U(c')$) and strictly quasi-concave and homogeneous of degree 1, and continuously differentiable on the domain.

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

¹⁹These capture tariffs in terms of the calculation of equilibrium, but ignore the revenue that they generate which could impact welfare in ways ignored here.

Appendix B, our formulation admits a representative consumer with preferences represented by $U(\cdot)$.

Time horizons

We consider the response of the economy to a shock over four different time horizons, which we refer to as the "Short Run", "Medium Run", "Long Run", and "Very Long Run". In the very long run, the economy fully adjusts and technologies can be changed. In the long run, technologies are held fixed (so it is as if each economy is only endowed with the technologies that were active when the shock hit), but everything else adjusts.

If, for example, each country is endowed with a set of technologies for producing a given good that corresponded to a constant elasticity of substitution production function, Then the difference between the long run and very long run would be captured by having the elasticity of substitution be zero in our long run (so the production function becomes Leontief). Adjusting these elasticities to capture the difference between time horizons is an approach taken in some of the literature (e.g. [Bachmann et al. \(2024\)](#)). There is evidence consistent with the elasticity of substitution being approximately zero for quite considerable lengths of time.²⁰

However, this differs completely from our analysis of the short run: the adjustments are not equilibrium ones in which prices adjust, but shortages on existing contracts. In the short run, prices are unable to adjust and shortages are rationed proportionally to the initial flows. For example, an automobile manufacturer does not receive the steering wheels it has contracted for, and cannot simply buy alternatives. Finding a new supplier that could deliver the custom-designed steering wheels could take months if not years. The manufacturer can only produce as many cars as it has steering wheels. This short run perspective is relevant for most manufacturing, where plants are designed and optimized relative to particular recipes for producing parts and assembling them. It can take months to make small changes, and years for larger changes. Even in the service industry specific recipes are followed. Fast food franchises have very specific and detailed tasks that are allocated to different workers, with extremely explicit workflows and dedicated inputs. Shortages of inputs can lead some goods to be completely deleted from menus, or become sold out, for months at a time.

The medium run is a sort of hybrid: labor allocations to technologies are held fixed (so it takes time for workers, if made idle, to find alternative employment, and it takes firms time to hire and train new workers), but prices and flows of goods can adjust. It is equivalent to a constrained equilibrium in which labor allocations and technologies are fixed.

Table 2.1 summarizes these differences.

	Short Run	Medium Run	Long Run	Very Long Run
Suppliers/Prices	Fixed	Flexible	Flexible	Flexible
Labor	Fixed	Fixed	Flexible	Flexible
Technology	Fixed	Fixed	Fixed	Flexible

²⁰For example, [Barrot and Sauvagnat \(2016\)](#), [Atalay \(2017\)](#) and [Boehm, Flaaen, and Pandalai-Nayar \(2019b\)](#).

While the outcomes obtained in all the other time horizons are either an equilibrium of the economy (very long run) or an appropriately constrained equilibrium for the economy (long run and medium run), in the short run the economy is out of equilibrium, and thus it differs fundamentally from the other time horizons.

For simplicity, and in order to fully consider all the ways in which the economy adjusts, we consider shocks that are permanent (or at least are not anticipated to revert).²¹

Equilibrium

We assume the economy is in equilibrium when a shock hits. In this subsection we formally define an equilibrium (which corresponds to the very long run outcome) and also introduce constrained equilibria to capture the economy in the medium and long run. We define short run disruptions in Section 3.2.

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 a list of prices of each good at each point of sale, productions from each technology, the inputs (including labor) that it uses, and final consumptions of the consumers/laborers, such that

- laborers supply their endowment of labor inelastically and choose final goods to consume to maximize their preferences,
- the output produced by each technology is feasible given the inputs it uses and it maximizes profits over all potential choices of outputs calculated as the revenue from the output minus the costs of corresponding inputs (including transportation costs), 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.

²¹See Liu, Smirnyagin, and Tsyvinski (2024) for adjustments in response to temporary shocks.

In the long run, a constrained equilibrium is equivalent to an equilibrium defined above but with countries' technology endowments restricted to only include those technologies in use at the time of the shock. In the medium run, labor allocations across technologies are also fixed.²²

An equilibrium (or constrained 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 between technologies / producers, labor not pictured. Green technologies produce final goods. Nodes of the same color are at the same “level” (directed path distance from final goods) in the supply chain. Some of the technologies might or might not produce the same 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.

²²More formally, this is equivalent to looking for an equilibrium of the following adjusted economy: (i) Set the number of countries equal to the number of technologies currently in use; (ii) Assign each technology currently in use to a different country; (iii) Endow each country with the labor used by their technology prior to the shock; and (iv) set the iceberg cost to infinity for all flows of labor.

3.1 Hulten’s Theorem: The Long-Run Impact of a Change in the Supply Chain

Let World GDP (GDP for short, henceforth) denote the total expenditures on final goods:²³

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

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 B). 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 equilibrium of an economy, and a technology τ used in positive amounts in equilibrium to produce good $k = O(\tau)$ such that there is no other technology that can produce the good at the same (transportation-adjusted) price. 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,*

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

The condition that there is no alternative technology that can produce the good at the same (transportation-adjusted) price serves the same role as the condition that there are unique producers for each good (or that all producers of the same good are shocked and use the same equilibrium mix of inputs) that is usually implicitly or explicitly assumed in statements of Hulten’s Theorem. If we relax this condition and allow for duplicate production of the same good at the same effective price and shock only one of the technologies/producers, then this is an upper bound on the impact, and so the equality becomes less than or equal to in the above equation.²⁴ This condition rules out that there are exact duplicates for the shocked technologies and so complete substitution away from the shocks is costly.²⁵

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.

²³This is nominal GDP , but prices can be normalized so that the total price of the consumed bundle is 1, and then nominal and real GDP are equivalent. Rescaling all prices in our model (including labor) by some common factor has no effect on the equilibrium.

²⁴As an example, given costs of transportation, it can be optimal to produce the exactly same good using exactly the same technology separately in different regions to be closer to inputs. For instance, there are lumber mills using the same technology in many locations around the world to be close to the forests, and much of that lumber is later shipped as it is cheaper to ship the lumber than the trees. As such, having multiple suppliers of the same input can be a feature of some equilibria. Nonetheless, those could end up with different transportation-adjusted costs and hence not be perfect substitutes for each other.

²⁵Bahal 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.

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.

We note that there is no adjustment in the statement of Hulten's Theorem for transportation costs. This is because y_τ already includes the extra amount produced to account for shipping costs. That is, for any intermediate good, $y_\tau = \sum_{\tau''} x_{\tau\tau''}$ and each $x_{\tau\tau''}$ is $\theta_{\tau\tau''}$ times the amount that technology τ'' finally receives as an input. Technology τ 's price also includes the transportation costs that accumulated along its supply chain.

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.




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.

3.2 Bottlenecks and Disruptions: Short-Run Impacts of Shocks

Hulten's Theorem applies in the long run as an approximation for sufficiently small shocks, and also in the very long run (for small shocks²⁶) under additional conditions that limit how

²⁶And also for larger shocks if the set of technologies available to countries to produce a given output corresponds to a Cobb-Douglas production function, in which case the Theorem is exact for non-marginal shocks.

substitution works across technologies. Once one examines large changes, or the very long run when certain nonlinearities as technology substitution occurs, then adjustments are needed. That is well-studied (e.g., see Baqaee and Farhi (2019, 2020)), and thus we next examine the impact of the technologies when some aspects of production are fixed in the short and medium run.

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 now examine the short-run 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 3.4, we return to the medium run in which there are shortages but price adjustments and 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 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 impact of a short run shock is found by the following (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, $y_\tau^1 = y_\tau$ and $x_{\tau\tau'}^1 = x_{\tau\tau'}$ for all $\tau \notin T^{Shocked}$ and τ').
- Iteratively in k (starting from $k = 2$):





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.

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


- Iterate until $y_{\tau}^{k-1} = y_{\tau}^k$ (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). Note that this does not allow for the redistribution of labor and so some labor becomes idle. This reduces the labor income of the corresponding representative consumer.


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).




(a) 10 percent shock to τ_2




(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.

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 a 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 a minimum disruption problem. This is formalized in Proposition 2.

Although the short-run outcome is out of equilibrium, a variation on the First Welfare Theorem nevertheless applies to it: the flows that the shock propagation algorithm converge to solve a suitably constrained value maximization problem. We call this problem the minimum disruption problem.

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. Consider outputs and flows $(\hat{y}_\tau)_\tau$ 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

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

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

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

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

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.

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.*

In the medium run prices, and thus the rationing, adjust. We explore the extent to which this mitigates short run disruptions in Section 3.4. An implication of Proposition 2 is that

allowing for flexible prices will (generically) reduce overall disruption because it relaxes the proportional rationing constraint.

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 3.4 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 because this solution maximizes output good by good over flows consistent with the constraints.*

It is worth reflecting on our model of the short-run and contrasting it with the existing literature. A typical way to capture the difference between the short-run and long-run is via the elasticity of substitution (e.g., [Bachmann et al. \(2022\)](#)). By making this more inelastic in the short run, less adjustment in the input mix occurs as one changes productivity parameters and compares across equilibria. However, this is a different approach that answers a fundamentally different question from our analysis that constrains production. Our assumption that firms are unable to change technologies and hence are stuck with a short-run production function that is

Leontief captures this lack of substitution, and one could use this to compare outcomes across different (long-run) equilibria as productivity parameters change. However, that is not what we are doing and is not what is driving the difference between the short- and long-run in our model. To emphasize this, consider an example where a pandemic shuts down production of some input for a quarter of the year. So, $1/4$ of the annual output of this input is lost. In our model that cascades through the supply chain and up to $1/4$ of the output of downstream goods using this as an input (directly or indirectly) are lost. If one uses does not treat this as a constraint, but recalculates equilibrium with a new productivity parameter of $3/4$ of what it used to be, the solution would be to scale up production of this good to compensate for its lower productivity. That would come at some cost (e.g., labor cost that is pulled from elsewhere in the economy), and would lead to lower production, but a (possibly much) smaller impact than the $1/4$ lost production that our model generates. But in a pandemic that scaling up would be impossible: the production is constrained by the labor restrictions of the pandemic and there is a fundamental cap put on the production function that operates differently from a parameter change and a new equilibrium. This applies to many short-term shocks such as natural disasters, sanctions, political crises, etc. It also captures situations where equilibrium adjustments and reallocation of labor, etc., across different parts of production take some time (which is true of many settings), so that the economy is at least temporarily out of equilibrium due to an unanticipated shock. This is not something that can be captured just by changing elasticities in a production function, but needs to be modeled as short-run constraints.

To further contrast the approaches, consider a temporary shock. A standard approach here would be to adjust to a new equilibrium in the short run for the duration of the shock, and then to adjust back to the original equilibrium in the long run, perhaps slowly if there are adjustment costs (see, for example, [Liu, Smirnyagin, and Tsyvinski \(2024\)](#)). However, our short-run analysis would let the economy be out of equilibrium, and adjustments back to the long-run equilibrium would occur as short-run constraints are relaxed and the economy can adjust on additional dimensions. Including the constraints leads to markedly different short-run outcomes, and identifies different points in the supply network as crucial.

Given the monotonicity of the algorithm, from it we can also deduce an upper bound on the impact of a shock that is 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

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

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 $F(T^{Shocked})$.




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.

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 $\mathcal{G}(T^{Shocked})$,²⁷ then the bound from Proposition*

²⁷Note that this permits the set of shocked technologies $T^{Shocked}$ to intersect both $R(T^{Shocked})$ and $F(T^{Shocked})$.

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 $T^{Shocked}$ 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 and 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. Similarly, one can easily construct examples where the bound holds that have extraneous cycles that are irrelevant to the cut set. Ruling out cycles is helpful in our partially constructive proof.

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 comparison is to a version of the long run in which we endow each firm with a unique (Leontief) technology and do not allow for substitutability in the input mix. Substitutability in the input mix would dampen the long-run impact of the shock further, making the contrast with the short-run even starker. Our approach highlights the difference between our short-run out-of-equilibrium analysis and models that compare the long-run and short-run to each other by reducing the elasticities of substitution between inputs.


The contrast between the long- and short-run 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.²⁸ To see this, consider an economy with two variations of the sets of technologies from 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 technology τ_1 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.


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.

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

²⁸Again, 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.




(a) Impact is 1/100th of GDP




(b) Impact is 1/100th of GDP

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 $\frac{p_{R1} y_{R1}}{GDP} = \frac{1}{10}$ 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.



(a) Impact is 1/20th of GDP



(b) Impact is 1/10th of GDP

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.

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.²⁹ 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.

Finally, consider a producer i downstream of a shocked firm and suppose, for the sake of illustration, that only its supply of one input from one direct supplier is affected. In the short run, i ’s unaffected suppliers of that same input cannot increase their production or re-route additional goods towards it. Thus, even if i is multi-sourcing the disrupted input, the flows of these alternative supplies remain fixed. For i ’s suppliers of different inputs, the amount used by i diminishes. Thus, in the short run, inputs are predicted to appear, in data, as complements. However, in the long run, prices adjust, flows reroute, productions change, and affected firms can adjust the technology they are using. All these adjustments reduce the observed complementarity between inputs. This is consistent with previous findings of short-run complementarity and longer-run substitutability between inputs. For example, on a quarter-long horizon [Barrot and Sauvagnat \(2016\)](#) present evidence for inputs being complements, while on a year-long horizon [Carvalho et al. \(2021\)](#) find inputs are weak substitutes.

²⁹If a good is unambiguously more downstream of another in that it uses all of the production of some good upstream (directly or indirectly), then the expenditures on the downstream good are strictly greater. If one can order goods in levels, so that goods at some level use all the production of goods further upstream, then this also holds across the levels of production. Thus, there is a well-defined way in which one can think of downstream expenditures being greater than upstream. Of course, this is having the number of goods in different layers be comparable. Having one raw material and many intermediate goods, could lead to any one intermediate good having lower expenditure than the raw material.

3.4 Flexible Prices: Shock Impacts in the Medium Run

While in the short run existing contracts prevent prices from adjusting (as embedded within our proportional rationing assumption), once prices can adjust existing production can reroute to higher value uses. This decreases the impact of the shock, yielding an outcome between the short-run and long-run impacts of the shock.³⁰ We now consider this medium-run situation, in which prices can adjust and shortages can be rationed efficiently, but in which production technologies are fixed, as is labor, which prevents any firms from increasing their output beyond pre-shock levels.

The medium run outcome can end up anywhere between the short- and long-run outcomes, depending on circumstances. Here we discuss some examples illustrating the forces, as a full characterization is simply an equilibrium statement that does not provide much further insight. We discuss these ideas more extensively, aided by numerical examples, in Appendix A.1.

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 create priorities in rationing and corresponding differences in prices. These ensure that electricity is allocated more efficiently towards higher value uses when there are future shortages. Again, it is an empirical question regarding the extent to which efficient contracts emerge, and which are more challenging to implement in other industries with supplier-specific inputs.

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 constraint (and effectively, reoptimizing). However, there are situations in which there is no difference between the medium and short run. For example, if the conditions of Proposition 5 hold, such that there is no technological diversity and there are industry specific shocks and each shocked technology is used—directly or indirectly—in the production of only one final good, then there is no scope for reducing the impact of the shock in the medium run. There is no scope for correlating the downstream instances of the disruption nor any scope for assigning the disruptions disproportionately to supply chains of lower-value final goods. On the other hand, it is possible to construct examples in which price flexibility makes a big difference. Specifically, the ratio of lost output in the short run to lost output in the medium run with price flexibility is unbounded (as we show in Supplementary Appendix II).

³⁰E.g., see [Iyoha et al. \(2024\)](#).

4 Disruption Centrality and Power

We now return to the short run, and provide further results that leverage the characterization to understand the impact of specific disruptions and how they sit in the supply chain. This allows us to not only define the “centrality” of a given technology based on short-run disruptions,³¹ but also to examine how much power a given country might have over others.

The two approaches in this section take different perspectives. One is of inadvertent disruption, where disruptions travel as in the short run proportionally along all relevant flows. The other is of deliberate and targeted disruptions for strategic purposes.

4.1 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 in the short run where there is proportional rationing on affected flows. 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' , and 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.

To keep the definition relatively intuitive, we examine acyclic supply chains. The definition easily extends, similarly to how we dealt with cycles above, but is no longer defined by a finite algorithm. 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 upstream in a clearly defined sense. This rules out (directed) cycles, but still allows upstream goods to be used in the production of multiple downstream goods and downstream goods to source any of their inputs from multiple upstream sources, and so for rich structures (and for the existence of undirected cycles in the network). 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.³²

When technologies are well-ordered in some equilibrium, then we say that τ is downstream from τ' (and τ' is upstream from τ) relative to $\succ$ if there is a sequence of technologies $\tau^1 \succ \tau^2 \succ \dots \succ \tau^L$ with $L \geq 2$, $\tau^1 = \tau'$ and $\tau^L = \tau$.

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

³¹The idea of showing that a measure of network centrality determines the impact of shocking different industries dates to [Leontief \(1936\)](#); [Acemoglu et al. \(2012\)](#). Recent work advancing this agenda includes [Bahal and Lenzo \(2023\)](#).

³²For 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.

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 the disruption centrality of τ is

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





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 $S(\tau', \tau'')$.

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 $\tau = \tau_2$. As an example, if the equilibrium prices of the final goods are all 1 (which depend on the labor inputs of the different technologies), then the disruption centrality of technology τ_2 for these equilibrium flows is

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

If the output of τ is reduced by some amount γ , then the total short-run fraction of GDP that is lost is $\gamma DC(\tau)$.

If all technologies are single-sourced (i.e., $S(\cdot, \cdot)$ is always 0 or 1) then

$$DC(\tau) = \sum_{f \in F(\tau)} p_f y_f,$$

which corresponds to the upper bound we identified before. More generally, however, $DC(\tau)$ 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.

4.2 Power

The perspective of “Disruption Centrality” is agnostic on why a particular technology might be disrupted and considers a total factor productivity shock (as is standard). There are however situations in which disruption, or the threat of disruption, is strategic. This could come either in the form of sanctions, or other deliberate disruptions as part of bargaining or conflict, including various quotas or other trade limitations. To capture how the structure of the supply network feeds into such situations we consider the power that one country wields over another via purposeful disruptions of its own technologies.

Although we use the terminology of “countries”, it should be clear that this can be any distinction of regions under consideration (e.g., parts of countries or collections of countries, etc.).

Total factor productivity shocks are not the best way to capture deliberate disruptions by an aggressor country. The aggressor can limit its own losses by selectively choosing which inputs it sources for some goods (e.g., retaining domestic inputs and eschewing inputs from a target country) or similarly from selectively choosing to whom it sells various goods (continuing to sell domestically, but reducing exports to a target country). We therefore study the impact of disruptions in which countries selectively withhold trade along some paths that involve their imports and exports.

To calculate the power that one country has over another (in the short run), it is no longer sufficient to simply track final goods disruptions, as we now care about the distribution of the impact of the disruption across countries. A country’s GDP (and welfare) can be measured via its labor income, and following a disruption some labor becomes idle throughout the supply chain. Indeed if two percent of a country’s labor becomes idle, then two percent of its GDP is lost.³³ Additional losses in terms of utility could also come via the final bundle that is available to purchase.³⁴

To measure the power that one country has over another we analyze what happens when a country disrupts the production of some of its technologies. We allow the country instigating

³³We work with homogeneous labor for simplicity, but more generally one can simply wage-weight labor to adjust for differences in productivity by labor in different tasks.

³⁴If there is just one final good, then relative labor disruptions and corresponding reductions in relative purchasing power map exactly into GDP and utility losses. That is, with just one final good, the loss in labor wages (both up and down stream) equals the loss in consumption.

the disruption to reduce the output of any of its technologies, or combination of its technologies, and to direct these disruptions selectively. Thus a country can choose which customers receive how much less output from one of its disrupted technologies, and, when it sources the same input from multiple suppliers for a disrupted technology, how much less demand it sources from each supplier.

We start off by considering the short run in which the other countries don't have time to react strategically and the indirect disruptions propagate up and downstream proportionally across flows. In Section 4.5 we turn to time horizons over which the other countries have discretion over how to allocate their disrupted production and how to source the inputs they still need, and consider how this mitigates the power countries have over one another.

4.2.1 A Heuristic Definition of Power

We define the power of an aggressor country looking to disrupt production in a target country, as the biggest GDP losses that can be induced in the target country at the expense of the least GDP losses to itself. This is the disruption with the highest “bang-for-the-buck,” the disruption which maximizes the percentage GDP loss in the target country per percent of own GDP loss.

We denote the power that Country i has over Country j by Power_{ij} . For now, we define it heuristically (as we have not fully specified how to calculate total lost GDP as a function of disruptions yet). We offer a full definition further below, but this is useful to help fix ideas.

$$\begin{aligned} \text{Power}_{ij} &= \max_{\text{disruption of flows in/out of } i\text{'s technologies}} \frac{\% \text{ lost GDP in Country } j}{\% \text{ lost GDP in Country } i} \\ &= \max_{\text{disruption of flows in/out of } i\text{'s technologies}} \left[\frac{\text{lost GDP in Country } j}{\text{lost GDP in Country } i} \right] \bigg/ \left[\frac{\text{GDP in Country } j}{\text{GDP in Country } i} \right] \end{aligned}$$

Percentage losses capture the power that larger countries with more diversified economies have over small ones and reflects the losses per capita in the aggressor and target economies. An alternative measure of power, that is appropriate in some circumstances, is the ratio of lost GDP. However, as can be seen in the equation above, the same disruptions maximize the ratio of absolute losses in GDP also maximize the ratio of percentage losses in GDP.

If $\text{Power}_{ij} < 1$ then Country i has to be willing to forego a larger (percentage loss) in GDP itself than it imposes on Country j . More generally, when Power_{ij} is sufficiently low, no disruption by i may be credible and hence i may not wield any viable power over j .

4.2.2 An Illustrative Example

Consider the initial equilibrium shown in Panel (a) of Figure 14. There are two countries, Country i and Country j . Country i is on the right in Figure 14 and operates technologies $T_i = \{\tau 2, \tau 5\}$, while Country j operates technologies $T_j = \{\tau 1, \tau 3, \tau 4\}$. We normalize the price of the final good to be 1, and so the equilibrium wage is $1/3$ (a total of 120 units of labor consume 40 units of output). GDP in Country i is $70/3$ while it is $50/3$ in Country j , which

summed give the total consumption.

If Country i wants to reduce the GDP of Country j , it can only do so by limiting the amount of the output of technology τ_2 that it exports to Country j . Reducing the flow τ_2, τ_3 from 10 to 9, as shown in Panel (b) of Figure 14, reduces the amount that τ_3 can produce, which in turn reduces the amount that τ_4 can produce and the amount demanded from τ_1 . This is shown in Panel (c) of Figure 14. Thus, 1 unit of labor becomes idle in Country j whereas 5 units of labor become idle in Country i . Thus GDP in Country i falls by 1/70th while GDP in Country j falls by 1/10th.




Figure 12: A targeted withholding of the output of technology τ_2 by Country i (on the right) leads to one unit of lost labor to Country i and imposes a loss of 5 units of labor on Country j (on the left). Country i has a power of 5 over Country j if we measure it in absolute terms, and $7 = (5/50)/(1/70)$ if we measure it in percentage of GDP terms.

This example illustrates several points.

First, it is important to consider upstream disruptions, here to technology τ_1 , in order to account for how the idle labor is distributed across countries. This contrasts with our earlier analysis in which we were measuring overall (world) GDP, which can be captured via final goods reduction. Here, tracking changes in labor at each technology captures how the overall loss of 2/40 units of final goods are distributed. Country i loses 1/120 of the total labor in the economy while Country j loses 5/120 of the total labor. This leads to the ratio of 5, and correspondingly Country i loses 1/3 of a unit of consumption, while Country j loses 5/3 units of consumption in total.

Second, Country i exerts the biggest loss possible on Country j per unit of own loss by directing its disruption of τ_2 towards technology τ_3 and keeping the amount shipped to τ_5 intact. There is no other disruption available to Country i that would do as well as the one we have considered in terms of power. It could be scaled up, but would have the same ratio of losses, until it hits 10 units being withheld, and then further withholding would only hurt Country i . As such the relative power index, $\text{Power}_{ij} = 7$, and Country i has a relatively high amount of leverage over Country j .

Third, as we mentioned above, an alternative measure of power that it can be appropriate to use in some circumstances is the ratio of absolute GDP losses, and in this case the power of Country i over Country j would be 5.

Fourth, in this example, there is only one flow between the two countries and so if we want to examine the power of Country j over Country i , it could only disrupt things via the same flow. In this particular example, the power of j over i is just the reciprocal $1/7$ (or $1/5$ in absolute terms). In richer examples, with multiple flows across borders (see below), the best disruption for one country will generally not be the same as the best disruption for the other.

Fifth, we have discussed this as if $\tau 4$ and $\tau 5$ produce the same final good. If they are instead different goods, then the consumers will also be forced to purchase a worse combination of goods after the disruption (as production by $\tau 4$ is reduced but production by $\tau 5$ is not).³⁵ Finally, Country i is relatively powerful in this example. The only way Country j can disrupt production in Country i is to reduce its demand for the output of technology $\tau 2$. However, the impact of doing so is the same as shown in Panel (c) of Figure 14 and so it only has power $1/5$ over Country i . So, to hurt Country i , Country j has to hurt itself (considerably) more. The substantial power that Country i has over Country j is in part derived from its ability to insulate its own final good production from the shock it induces. Its disruption does not induce a further disrupted flow of goods back across borders affecting its own technologies.

4.3 A Formal Definition of Power

We now provide a formal definition of the power Country i has to disrupt Country j . To do this we must first specify how the full economy responds to a set of flow reductions by Country i .

This involves three complications compared to the downstream propagation algorithm that we defined previously:

- To account for all labor reductions, we also need to track how every directly or indirectly affected technology’s disruption propagates upstream.
- A disruption can involve multiple levels of discretion by Country i that have to be consistent with each other.
- Even without cycles in the network, there can be (finite) cyclical feedback as shocks propagate both upstream and downstream.³⁶

We attack these complications by defining a notion of a set of “consistent disruptions,” and showing it is without loss of generality to restrict attention to these. While this simplifies matters, the set of available sets of consistent disruptions can be large and complex. We show that under some basic conditions on the supply network, when looking for the optimal

³⁵Proportional rationing of the over-demanded final good will result in Country i ’s consumers consuming $18(69/114)$ and $20(69/114)$ units of the two final goods, while Country j ’s consumers will consume the remaining amounts. Thus the consumers in both countries will consume different amounts of the same bundle of final goods. As by assumptions preferences are homothetic, this is also what would also happen with flexible final good prices.

³⁶For instance, if a good 0 used in the production of good 1 is held back, and reduces production of good 2, but good 2 also uses the good 0 as an input, then reduced demand from good 2 further reduces the production of good 0.

disruption that determines the power one country has over another, attention can be restricted without loss of generality to sets of minimal consistent disruptions induced by a disruption to a *single* technology.

We consider supply networks with no (directed) cycles. Thus, as above there exists a well defined order $\succ$ and sequence of all technologies $\tau^1, \dots, \tau^L$ such that $\tau^\ell \succ \tau^{\ell'}$ implies τ^ℓ is not downstream from $\tau^{\ell'}$. We label the technologies from upstream to downstream so that $\tau^\ell \succ \tau^{\ell'}$ when $\ell' > \ell$.

Consider some equilibrium and let $x_{\tau\tau'}^*$ denote the equilibrium flow of output from technology τ to technology τ' and y_τ^*, L_τ^* be the equilibrium output and labor usage associated with τ .

A *consistent disruption* of an equilibrium $(x_{\tau\tau'}^*, y_\tau^*, L_\tau^*)_{\tau, \tau'}$ by Country i via technology $\tau^i \in T_i$ that is reduced to $\lambda < 1$ of its initial level³⁷ is a specification of inputs, outputs, and labor for each technology, $x_{\tau\tau'}, y_\tau, L_\tau$ for every τ, τ' , as follows. Let $\tau^i = \tau^\ell$ in our ordered sequence for some $\ell \in \{1, \dots, L\}$.

- First, specify new downstream flows $x_{\tau\tau'}^d$, such that $x_{\tau\tau'}^d \leq x_{\tau\tau'}^*$ for all τ' , and $\sum_{\tau'} x_{\tau\tau'}^d = \lambda y_{\tau^\ell}^*$. This specifies how i allocates the disrupted production of technology τ^ℓ .
- We then follow this iteratively for downstream technologies ℓ' (such that $\ell \succ \ell'$) starting with technology $\ell + 1$. It may be that ℓ' has had multiple cuts on its inputs from technologies further upstream when it is considered, and so the amount that it needs to ration is the max shortfall across its inputs. That is, the current (temporary) $y_{\tau^{\ell'}}^d = \min_{j \in I(\tau^{\ell'})} \frac{\sum_{\tau': O(\tau')=j} x_{\tau'\tau^{\ell'}}^d}{\sum_{\tau': O(\tau')=j} x_{\tau'\tau^{\ell'}}^*}$
 - and if $\tau^{\ell'} \in T_i$, then select any set $x_{\tau^{\ell'}\tau'}^d \leq x_{\tau^{\ell'}\tau'}^*$ for all τ' , and such that $\sum_{\tau'} x_{\tau^{\ell'}\tau'}^d = y_{\tau^{\ell'}}^d$. So i controls the disrupted flows through its own technologies.
 - and if $\tau^{\ell'} \in T_j$ for $j \neq i$, then set $x_{\tau^{\ell'}\tau'}^d \leq x_{\tau^{\ell'}\tau'}^*$ for all τ' , such that $\frac{x_{\tau^{\ell'}\tau'}^d}{x_{\tau^{\ell'}\tau'}^*} = \frac{y_{\tau^{\ell'}}^d}{y_{\tau^{\ell'}}^*}$. Thus there is proportional rationing of other countries' outputs.
- We do this until we hit $\ell' = L$, which is necessarily a final good. Note that this is well-ordered due to the lack of cycles. Set $y_{\tau^L}^u = y_{\tau^L}^d$, to start the upstream process.
- Next, we follow the consequences back upstream starting at technology τ^L and iterating until we get to technology τ^1 .³⁸ The iterative step is as follows.
- We ration $\tau^{\ell'}$'s inputs to levels $x_{\tau^{\ell'}\tau'}^u$ to match the production $y_{\tau^{\ell'}}^u$ from the previous steps, as follows:
 - if $\tau^{\ell'} \in T_i$, then select upstream choices $x_{\tau^{\ell'}\tau'}^u$ such that $x_{\tau^{\ell'}\tau'}^u \leq x_{\tau^{\ell'}\tau'}^d$ for all τ' and $\frac{\sum_{\tau': O(\tau')=j} x_{\tau^{\ell'}\tau'}^u}{\sum_{\tau': O(\tau')=j} x_{\tau^{\ell'}\tau'}^d} = \frac{y_{\tau^{\ell'}}^u}{y_{\tau^{\ell'}}^d}$ for all $j \in I(\tau^{\ell'})$

³⁷Note that any disruption of τ^i along any flows up or downstream reduces its output, and so this nests both cases.

³⁸Note that upstream propagation has no further downstream consequences.

– and if $\tau^{\ell'} \in T_j$, then set $x_{\tau', \tau^{\ell'}}^u$ such that $x_{\tau', \tau^{\ell'}}^u \leq x_{\tau', \tau^{\ell'}}^d$ for all τ' and if $x_{\tau', \tau^{\ell'}}^d > 0$ then $\frac{x_{\tau', \tau^{\ell'}}^u}{x_{\tau', \tau^{\ell'}}^d}$ is equal for all such τ' and $\frac{\sum_{\tau': O(\tau')=j} x_{\tau', \tau^{\ell'}}^u}{\sum_{\tau': O(\tau')=j} x_{\tau', \tau^{\ell'}}^*} = \frac{y_{\tau^{\ell'}}^u}{y_{\tau^{\ell'}}^*}$ for all $j \in I(\tau^{\ell'})$.

- By construction, these upstream disruptions eliminate flows of inputs goods that are unused in production, but cutting an upstream flow has no further downstream consequences (given the absence of directed cycles).

We are left with a well defined (individual technology) disruption. We can now treat this $(x_{\tau\tau'}, y_{\tau}, L_{\tau})$ specification as if it were a new “equilibrium” and then apply a new disruption to some other $\tau^{i'} \in T_i$.

We let $\mathcal{D}_i$ denote the set of consistent disruptions available to Country i via sequences of such disruptions, i.e., those disruptions that can be obtained from any finite sequence of individual technology disruptions as described above.³⁹

For any $(x_{\tau\tau'}, y_{\tau}, L_{\tau})_{\tau, \tau'} \in \mathcal{D}_i$ denote the corresponding GDP of Country i (and similarly for j) by

$$GDP_i((x_{\tau\tau'}, y_{\tau}, L_{\tau})_{\tau, \tau'}) = w_i^* \sum_{\tau \in T_i} L_{\tau},$$

where w_i^* is the original wage level (as by assumption this is not adjusted in the short run). Letting GDP_i^* denote the initial equilibrium GDP of Country i ,

$$\begin{aligned} \text{Power}_{ij} &= \max_{(x_{\tau\tau'}, y_{\tau}, L_{\tau})_{\tau, \tau'} \in \mathcal{D}_i} \left[\frac{GDP_j^* - GDP_j((x_{\tau\tau'}, y_{\tau}, L_{\tau})_{\tau, \tau'})}{GDP_i^* - GDP_i((x_{\tau\tau'}, y_{\tau}, L_{\tau})_{\tau, \tau'})} \right] \bigg/ \left[\frac{GDP_j^*}{GDP_i^*} \right] \\ &= \max_{(x_{\tau\tau'}, y_{\tau}, L_{\tau})_{\tau, \tau'} \in \mathcal{D}_i} \left(\frac{\sum_{\tau \in T_j} L_{\tau}^* - L_{\tau}}{\sum_{\tau \in T_j} L_{\tau}^*} \right) \bigg/ \left(\frac{\sum_{\tau \in T_i} L_{\tau}^* - L_{\tau}}{\sum_{\tau \in T_i} L_{\tau}^*} \right) \end{aligned}$$

The space of consistent disruptions $\mathcal{D}_i$ can be large. We first consider the case in which disruptions are such that no technology is completely shut down. Once the new supply chains hit corners where some technology flows are completely disrupted, power calculations become more involved as we detail below. In most applications, countries are not completely shutting down cross border flows, and then things are much easier to analyze. In particular, if none of i 's technological flows are completely disrupted then, as we show, maximum power is achieved via a consistent individual technology disruption. We call these partial disruptions.

PROPOSITION 6. *If the initial equilibrium supply network has no (directed) cycles and only partial disruptions are considered, then there exists a technology $\tau^i \in T_i$ such that the power of Country i over Country j is maximized by a consistent individual technology disruption τ^i . Moreover, the disruption can be rescaled arbitrarily (subject to staying partial) while maintaining the same power.*

³⁹In applying sequences of disruptions, the marginal proportional rationing of Country j is done using the relative size of the original equilibrium flows. We show in the appendix that the order of disruptions is not consequential as long as no flows have been cut to zero.

Proposition 6 shows that when the supply network is acyclic there is a disruption available to i that hurts country j the most per amount of disruption in country i , and it involves country i initially disrupting a single technology in T_i . So, the problem of finding the set of disruptions available to i that maximizes i 's power over j , as defined above, simplifies to searching over disruptions to a single technology in T_i . This makes such disruptions relatively easy to find and implement.

The proof of Proposition 6 is in Appendix B. It works by showing that complicated disruptions by i can be broken down into separate independent disruptions, such that each independent disruption is induced by an initial disruption to a single technology in T_i . The power associated with the complicated disruption is then a weighted average of the power associated with these separate disruptions, and so one of the separate disruptions must be weakly more powerful than the combination.

We remark that an algorithm looking for the technologies to disrupt can actually further restrict attention to the technologies that are closest to the border in a well-defined sense. That is, one can concentrate on disruptions of $\tau \in T_i$ such that

- there is at least one downstream path from τ to a final good technology that involves no technologies in T_i anywhere on the path, or
- there is at least one upstream path from τ to an intermediate good technology that uses only labor in its production that involves no technologies in T_i anywhere on the path.

When supply networks become more involved, the importance of restricting attention to a smaller set of potential technologies to disrupt becomes more combinatorially important. We illustrate that with a more involved example in Appendix A.2.

That more involved example also shows how one country can have much more power over another than vice versa. In doing so it helps identify the characteristics of the supply network that grant Country i power over Country j . Country i will have a substantial amount of power over Country j when: (i) Country j 's output is relatively highly concentrated in one supply chain (measured by its proportion of wage income associated with that supply chain); (ii) Country i has the ability to disrupt this supply chain; and (iii) Country i has relatively diversified production (measured by its proportion of wage income associated with other supply chains).

An example of a powerful disruption is one to a key good that is relatively monopolized by the disrupting country and not costly in labor itself, but is present in high-labor-usage technologies (up and/or downstream) that lie significantly within the target country.

Note that it is possible for Countries i and j to both have substantial power over each other (potentially via completely separate parts of the supply network), and for the power relationship to be very asymmetric—so that one country has substantial power over the other but not vice versa.

4.4 The Disruption Possibility Frontier

If countries consider large disruptions that completely cut ties across borders, then that can alter the marginal impact of other technologies as it changes the ways in which subsequent flows are rationed. To capture such extreme interventions, one can no longer limit attention to just individual technology disruptions and combinations and orderings of disruptions become important. We can still define power, and it has a single well-defined value, but the algorithm for finding it when allowing for complete disruptions is more involved.

We ask, what is the maximum percentage reduction in GDP in Country j that Country i can impose for an $x\%$ loss in own GDP, varying from 0 to 100 percent own loss. This is defined by solving the following problem for all $\lambda \in [0, 1]$.

$$\min_{\{(x_{\tau\tau'}, y_{\tau}, L_{\tau})_{\tau, \tau'} \in \mathcal{D}_i\}} \frac{GDP_j((x_{\tau\tau'}, y_{\tau}, L_{\tau})_{\tau, \tau'})}{GDP_j^*}$$

subject to $GDP_i((x_{\tau\tau'}, y_{\tau}, L_{\tau})_{\tau, \tau'}) \geq (1 - \lambda)GDP_i^*$.

Solving this problem for different values of λ identifies the trade-off Country i faces between disrupting their own output and disrupting the output of Country j . As long as Country i prefers larger disruptions to Country j and smaller disruptions to itself, it will choose disruptions that implement some point on this disruption possibility frontier.

Figure 13 illustrates the disruption possibility frontier for the example from Figure 18. The disruption possibility frontier is piecewise linear. $Power_{ij}$ is given by the slope of the disruption possibility frontier at $\lambda = 0$; which corresponds to the optimal disruptions that fall short of completely disrupting any of i 's flows in the supply network.

While our power measure is a useful single statistic and an accurate measure of maximum power for disruptions that do not zero out any trade flows, for large disruptions that do zero out some trade flows, complex disruptions (that sequentially disrupt multiple technologies) need to be considered and the full disruption possibility frontier needs to be mapped.

In Figure 14, we illustrate how the marginal impact of a disruption of one flow can actually increase as another flow is zeroed out. This illustrates how the disruption possibility frontier need not be concave, and why when considering complete disruptions, the power analysis becomes combinatorially complex.

4.5 Strategic Power When All Countries Optimize Flows

The definitions above work with non-strategic decisions at all flows except those by the country that is doing the disruption. This does not provide scope for other countries responding to mitigate the damage or retaliate. While in the short run other countries may not have time to respond, by the medium run a response should be expected. Moreover, if the disruption has been anticipated by the target country then it might be in a position to immediately respond.

It is straightforward to develop a definition of power that includes a strategic response. It is important to define since, as we show via examples in the appendix, having a country react




Figure 13: Disruption Possibility Frontier for Country i disrupting Country j for the example from Figure 18.

strategically rather than proportionally can end up rerouting disruptions back to the country that is initiating the disruptions. This can lead to substantial changes to the power calculation. While this is important to account for, it is analytically straightforward and so we provide the definitions and examples in Appendix A.3.

5 Complexity, Fragility, and Globalization

5.1 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 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.




Figure 14: For the initial equilibrium shown in Panel (a), Panels (b) and (c) show the impact of the different marginal disruptions available to Country i (on the right in each panel). The power of both these disruptions is the same and constant up to the point where flows zero out. Panel (d) shows the impact of both these aforementioned disruptions being sequentially implemented (in either order). The power of this combined disruption is greater than the power of either marginal disruption.

PROPOSITION 7. [Supply Chain Disruption as a Function of Complexity] *Consider an 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):*

$$\text{Short-Run: } \mathbb{E} \left[\frac{\Delta \text{GDP}}{\text{GDP}} \right] \approx -(1 - \lambda)\pi S,$$

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

Some intuition behind Proposition 7 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 7 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 15. We see 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.




Figure 15: 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 $p_{vertical} = (\frac{1}{5}, \frac{1}{5}, \frac{2}{5}, \frac{3}{5}, \frac{4}{5}, 1)$, $p_{horizontal} = (\frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5}, 1)$, $p_{parallel} = (\frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$.

5.2 Globalization

We next discuss what happens when globalization, modeled as reduced transportation costs, leads to increased specialization and the conditions hypothesized in Proposition 5 apply. The idea is simple: As transportation costs become sufficiently low, technological diversity decreases and each good becomes single-sourced by the least costly technology, and technology shocks become industry shocks.

In particular, suppose that the set of world technologies is such that if there were no transportation costs, then no two distinct technologies that produce the same good use exactly the same amount of labor (in total, directly plus indirectly in terms of their most efficient equilibrium input mix). This is effectively a generic property in terms of sets of technologies T 's, as it is a knife-edge case that two different technologies end up using exactly the same amount of labor.

In this case, once transportation costs drop sufficiently, any good that is produced in equilibrium ends up being sourced from the same technology.

The fact that this implies 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, 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 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.⁴⁰ Nonetheless, Proposition 8 is important to state. It has as a corollary that if a lowering of trade costs leads to the consolidation of supply chains and—via improved efficiency—larger outputs of final goods from some single technology, then 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 due to globalization the production network evolves and supply chains can become more complex and cross more borders. This can expose them more to shocks such as political disruptions and issues that disrupt shipping. Moreover, the overall distribution of shocks can co-evolve with the supply network. This opens the possibility that one could see both an increase in the probability of a disruption *and* the size of the disruption.

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.⁴¹

⁴⁰The main 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.

⁴¹See [Bachmann et al. \(2022\)](#) for an analysis of the impact of Russian energy sanctions for the German economy.

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. 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.⁴²

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.

⁴²For 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).

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.
- Aldrichetti, Riccardo, Daria Battini, Dmitry Ivanov, and Ilenia Zennaro. 2021. “Costs of resilience and disruptions in supply chain network design models: a review and future research directions.” *International Journal of Production Economics* 235:108103.
- 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.
- Arkolakis, Costas, Federico Huneus, and Yuhei Miyauchi. 2025. “Production Network Formation, Trade, and Welfare.” Tech. rep., CESifo Working Paper.
- Arrow, Kenneth J. and Gerard Debreu. 1954. “Existence of an equilibrium for a competitive economy.” *Econometrica* :265–290.
- Atalay, Enghin. 2017. “How important are sectoral shocks?” *American Economic Journal: Macroeconomics* 9 (4):254–280.
- 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.
- . 2024. “What if? The macroeconomic and distributional effects for Germany of a stop of energy imports from Russia.” *Economica* 91 (364):1157–1200.
- Bahal, Girish and Damian Lenzo. 2023. “Beyond Domar Weights: A New Measure of Systemic Importance in Production Networks.” .
- Bai, Xiwen, Jesús Fernández-Villaverde, Yiliang Li, and Francesco Zanetti. 2024. “The causal effects of global supply chain disruptions on macroeconomic outcomes: evidence and theory.” .

- 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.
- Baqae, David and Elisa Rubbo. 2022. “Micro propagation and macro aggregation.” Tech. rep., National Bureau of Economic Research. Working Paper No. 30538.
- Baqae, David R. 2018. “Cascading failures in production networks.” *Econometrica* 86 (5):1819–1838.
- Baqae, David R. and Emmanuel Farhi. 2019. “The macroeconomic impact of microeconomic shocks: beyond Hulten’s Theorem.” *Econometrica* 87 (4):1155–1203.
- . 2020. “Productivity and misallocation in general equilibrium.” *Quarterly Journal of Economics* 135 (1):105–163.
- . 2021. “Entry vs. rents: Aggregation with economies of scale.” Tech. rep., National Bureau of Economic Research. Working Paper No. 27140.
- Barro, Robert J and Herschel I Grossman. 1971. “A general disequilibrium model of income and employment.” *The American Economic Review* 61 (1):82–93.
- 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.
- Baslandze, Salomé and Simon Fuchs. 2025. “The Price of Delay: Supply Chain Disruptions and Pricing Dynamics.” .
- Benassy, Jean-Pascal. 1977. “A neokeynesian model of price and quantity determination in disequilibrium.” In *Equilibrium and Disequilibrium in Economic Theory: Proceedings of a Conference Organized by the Institute for Advanced Studies, Vienna, Austria July 3–5, 1974*. Springer, 511–544.
- 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. 2019a. “Input linkages and the transmission of shocks: firm-level evidence from the 2011 tōhoku earthquake.” *Review of Economics and Statistics* 101 (1):60–75.
- . 2019b. “Input linkages and the transmission of shocks: Firm-level evidence from the 2011 Tōhoku earthquake.” *Review of Economics and Statistics* 101 (1):60–75.
- Brancaccio, Giulia, Myrto Kalouptsi, and Theodore Papageorgiou. 2025. “Rigidities in Transportation and Supply Chain Disruptions.” In *AEA Papers and Proceedings*, vol. 115. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 543–550.
- 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.
- Durugbo, Christopher M and Zainab Al-Balushi. 2023. “Supply chain management in times of crisis: a systematic review.” *Management Review Quarterly* 73 (3):1179–1235.
- 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.
- Galbraith, John K. 1947. “The disequilibrium system.” *The American Economic Review* 37 (3):287–302.
- Giovanni, Julian di and Andrei A Levchenko. 2009. “Trade openness and volatility.” *The Review of Economics and Statistics* 91 (3):558–585.
- Grossman, Gene, Elhanan Helpman, and Stephen Redding. 2021. “When tariffs disrupt global supply chains.” *American Economic Review* .
- Grossman, Gene M, Elhanan Helpman, and Hugo Lhuillier. forthcoming. “Supply Chain Resilience: Should Policy Promote International Diversification or Reshoring?” *Journal of Political Economy* .
- Hägele, Sascha, Eric H Grosse, and Dmitry Ivanov. 2023. “Supply chain resilience: a tertiary study.” *International Journal of Integrated Supply Management* 16 (1):52–81.
- Hulten, Charles R. 1978. “Growth accounting with intermediate inputs.” *The Review of Economic Studies* 45 (3):511–518.
- Iyoha, Ebehi, Edmund Malesky, Jaya Wen, Sung-Ju Wu, and Bo Feng. 2024. *Exports in Disguise?: Trade Rerouting during the US-China Trade War*. Harvard Business School Boston, MA, USA.

- 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>.
- Kamiya, Shinichi, Jun-Koo Kang, Jungmin Kim, Andreas Milidonis, and René M Stulz. 2018. “What is the impact of successful cyberattacks on target firms?” Tech. rep., National Bureau of Economic Research.
- 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.
- Liu, Ernest, Vladimir Smirnyagin, and Aleh Tsyvinski. 2024. “Supply Chain Disruptions and Supplier Capital in US Firms.” .
- Loh, Hui Shan and Vinh Van Thai. 2015. “Cost consequences of a port-related supply chain disruption.” *The Asian Journal of Shipping and Logistics* 31 (3):319–340.
- 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* .

- Pin, Paolo. 2025. “Network Effects of Tariffs.” *arXiv preprint arXiv:2504.04816* .
- Queiroz, Maciel M, Dmitry Ivanov, Alexandre Dolgui, and Samuel Fosso Wamba. 2022. “Impacts of epidemic outbreaks on supply chains: mapping a research agenda amid the COVID-19 pandemic through a structured literature review.” *Annals of operations research* 319 (1):1159–1196.
- Schaal, Edouard and Mathieu Taschereau-Dumouchel. 2025. “Echoes and delays: time-to-build in production networks.” .
- 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 Analyses and Examples Omitted from Main Text

A.1 The Medium Run

While in the short run existing contracts prevent prices from adjusting, as embedded within our proportional rationing assumption, once prices can adjust existing production can reroute which can decrease the impact of the shock, and end up between the short run and long run impacts.⁴³ We now consider this medium run situation, in which prices can adjust and shortages can be rationed efficiently, but in which production technologies are fixed, as is labor, which prevents any firms from increasing their output beyond pre-shock levels.

As the medium run can vary and end up anywhere between the short and long run outcomes, depending on circumstances, we provide examples and illustrate the forces, as a full characterization is simply an equilibrium statement that does not provide much further insight.

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 $\tau 1$ (without flexible prices) shown in panel (a) of Figure 16.

Here a disruption to $\tau 1$ affects the production of the final goods $\tau 4$ through two potential channels. First $\tau 1$ is used directly as an input, and secondly it is used to produce $\tau 3$ which is an input for $\tau 4$. In this example $\tau 1$ and $\tau 2$ produce the same good, and so by redirecting some of the output of $\tau 1$ to $\tau 3$ there is less disruption of that good which is critical to the production of $\tau 4$, while there is a substitute for $\tau 1$ in the direct input for $\tau 4$.

A second example is illustrated in Figure 17. Suppose labor is priced at 1, one unit of labor is needed to make one unit of the intermediate good $R1$, 1 unit of $R1$ and no units of labor are

⁴³E.g., see [Iyoha et al. \(2024\)](#).




Figure 16: 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.

needed to make final good F_1 , while $1/9$ th of a unit of R_1 and $8/9$ ths of a unit of labor are needed to make a unit of final good F_2 . Thus both F_1 and F_2 are priced at 1. Suppose there are 10 units of labor, and at these prices, labor demands 9 units of good F_2 and 1 unit of F_1 . Thus the overall production by F_2 is much more valuable than the overall production by F_1 : F_2 's output contributes 9 to a GDP of 10, while F_1 's output contributes 1.

Consider now a 10% total factor productivity shock to R_1 . 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 F_2 will not be affected at all, while output of F_1 will decrease by 20%. However, because production of good F_2 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.

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 constraint (and effectively, reoptimizing). However, there are situations in which there is no difference between the medium and short run. For example, if the conditions of Proposition 5 hold, such that there is



Figure 17: 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)).

no technological diversity and there are industry specific shocks and each shocked technology is used—directly or indirectly—in the production of only one final good, then there is no scope for reducing the impact of the shock in the medium run. There is no scope for correlating the downstream instances of the disruption nor any scope for assigning the disruptions disproportionately to supply chains of lower-value final goods. On the other hand, it is possible to construct examples in which price flexibility makes a big difference. Specifically, the ratio of lost output in the short run to lost output in the medium run with price flexibility is unbounded (as we show in Supplementary Appendix II).

In this section we have considered the medium run relaxation of the short run impact of a shock that allows prices to adjust. There are further variations that live between the medium and long run. For example, one can relax the no increased output constraints for some technologies. In a somewhat different setting, this is the 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.

A.2 A More Involved Example of Power

Consider five countries labeled 1 – 5, three raw materials (intermediate goods that only use labor in their production) $R1$ to $R3$, six intermediate goods $I1$ to $I6$ and four final goods $F1$ to $F4$. The initial equilibrium flows are shown in Figure 18. Labor is immobile and we index each active technology by the good it produces and the country in which it is located. Countries 3-5 each operate a single technology, but both Country 1 and Country 2 operate multiple technologies. Units of output are defined such that one unit of labor is required to produce one unit of output (which can be done when each good is produced by one technology

(possibly in different countries, but with identical recipes).

Country 1 and Country 2 are both endowed with 100 units of labor, which are fully deployed across their active technologies (illustrated by the respective shaded regions in Figure 18). The GDP of Country 1 is therefore $100w_1$, where w_1 is wage in country 1, and correspondingly, the GDP of Country 2 is $100w_2$.




Figure 18: There are five countries labeled 1 to 5. Country 1 is in the blue rectangle and Country 2 is in the yellow rectangle. Each technology is indexed by the good it produces and then the country in which it is located. So technology $I1-2$ produces intermediate good $I1$ and is located in country 2. Countries 3-5 operate a single technology. However Countries 1 and 2 each operate multiple technologies, as illustrated by the shaded areas.

Consider the power Country 1 has to disrupt production in Country 2. By Proposition 6 we just need to consider individual technology disruptions. And, as we remarked above, not even all of these need to be considered. In this case it is sufficient to consider individual disruptions to $\{\tau^{I2-1}, \tau^{I3-1}, \tau^{I6-1}\}$. As Country 1 operates 9 technologies and several of these technologies have multiple flows in and/or out of them there are many different combinations of disruptions that Country 1 could potentially take and so this simplifies matters considerably.

First consider a disruption to τ^{I2-1} . Reducing the flow $\tau^{I2-1}\tau^{F1-2}$ to 0 reduces Country 1's GDP by $10w_1$ and reduces Country 2's GDP by $45w_2$. As τ^{F1-2} is no longer able to source input $I2$ it shuts down, leaving 35 units of labor in Country 2 idle, and this reduces τ^{F1-2} 's demand for input $I1$ leaving a further 10 units of labor idle in Country 2. Thus this disruption grants Country 1 power 4.5 over Country 2.

An alternative disruption to τ^{I2-1} available to Country 1 is to eliminate the flow $\tau^{I2-1}\tau^{F1-5}$. This will reduce the demand for τ^{I1-2} 's output, causing Country 2 to lose $10w_2$ in GDP at a cost of $10w_1$ in GDP for Country 1. Thus this disruption is less effective, and as other disruptions to τ^{I2-1} are some convex combination of these two, the most power that disrupting τ^{I2-1} can give Country 1 over Country 2 is 4.5.

The second technology that Country 1 might want to disrupt is τ^{I3-1} . If Country 1 eliminated the flow $\tau^{I3-1}\tau^{I5-2}$, then propagating the shock both upstream and downstream, 15 units of labor become idle in Country 1 at a cost to GDP of $15w_1$, while GDP in Country 2 is reduced by $45w_2$. This grants Country 1 power 3 over Country 2 and is the most powerful disruption of technology τ^{I3-1} available to Country 1.

The final technology that Country 1 might want to disrupt is τ^{I6-1} . However, when the disruption to the flow $\tau^{I3-1}\tau^{I5-2}$ considered above is propagated it eliminates the flow $\tau^{I6-1}\tau^{F2-2}$ and disrupting either of these flows turn out to be equivalent. So again, this only gives Country 1 power 3 over Country 2.

Thus, the power that Country 1 has over Country 2 is $\text{Power}_{12} = 4.5$.

It is also instructive to consider the power that Country 2 has over the other countries. It has 0 power over Country 4 because it cannot impact its GDP at all. It has power 10 over Country 5, because it can completely eliminate its GDP at a cost of 10% of its own GDP by withholding the flow $\tau^{I1-2}\tau^{F1-5}$. The same disruption gives it power 2.5 over Country 3, because by eliminating its flow $\tau^{I1-2}\tau^{F1-5}$ demand for τ^{R1-3} 's output falls to 1 causing a 25% loss in GDP for Country 3 for the 10% loss in own GDP. Finally, Country 2 has power 1 over Country 1, which comes again from disrupting the flow $\tau^{I1-2}\tau^{F1-5}$. This is effectively no power, unless Country 2 is willing to forgo an equal percentage loss in GDP as it imposes on Country 2.

The example helps identify the characteristics of the supply network that grant Country i power over Country j . Country i will have a substantial amount of power over Country j when: (i) Country j 's output is relatively highly concentrated in one supply chain (measured by its proportion of wage income associated with that supply chain); (ii) Country i has the ability to disrupt this supply chain; and (iii) Country i has relatively diversified production (measured by its proportion of wage income associated with other supply chains).

An example of a powerful disruption is one in which a key good is relatively monopolized by one country, but the supply chain in which it is involved (up and/or downstream) lies nontrivially in another country.

Note that it is possible for Countries i and j to both have substantial power over each other (potentially via completely separate parts of the supply network), and for the power relationship to be very asymmetric—so that one country has substantial power over the other but not vice

versa.

A.3 Strategic Power

The definitions above work with non-strategic decisions at all flows except those by the country that is doing the disruption. This does not provide scope for other countries responding to mitigate the damage or retaliate. While in the short run other countries may not have time to respond, by the medium run a response should be expected. Moreover, if the disruption has been anticipated by the target country then it might be in a position to immediately respond. We now develop a definition of power that includes a strategic response.

To see some of the basic issues, consider the example in Figure 19. Here, by responding strategically with its intermediate good, Country j can drastically blunt Country i 's power.




Figure 19: Calculation of Country i 's power over Country j . Country j is on the left and Country i is on the right. If Country j is non-strategic, then Country j loses 7 units of labor to Country i 's 6 units of labor and Country i 's power is 7/6. In contrast, if Country j is strategic, then Country j loses 2 units of labor and Country i loses 11, and Country j 's power is only 2/11. (In this example, that is actually the inverse of Country i 's power over Country j which is 11/2).

Accounting for strategic reactions can actually simplify power calculations on the one hand, and complicate them on the other. The simplification is that optimal strategic disruptions throughout the supply network on the margin can be calculated via “threads”—that have one entrance and one exit from each technology on the thread and go between a raw material (i.e., an intermediate good with no incoming flows) and a final good. A complication is that the only relevant threads are those that are solutions to an equilibrium and nested max (or min) problem, where each country at each node is minimizing its domestic disruption anticipating the consequences given other countries' strategies at every node, and potentially maximizing the disruption of some other countries. A second complication is that we now need to specify what third-party countries do when they have choices of whom to harm, holding their own harm constant.

To keep things uncluttered, we focus on the case of two countries, but one can extend this to put third-party countries in one of the two blocks (with some complications if countries care about which country within the other block they do the most harm to).

We first note a complication that requires us to specify an equilibrium definition, where actions at each technology are specified, and each has to be a best response given the others. To see why this is true, consider the following example that shows that the optimal strategy at one technology depends on what is being done at others.

Suppose that the labor endowment of both countries is the same and that Country j is seeking to minimize the power that Country i has over it. Country j has the choice at a given technology to either disrupt 3 units of labor in each country, or else to disrupt 2 units of labor in the other country and 1 unit of its own labor. If all other paths involve only minor labor disruptions, then the second option results in i having a power of approximately $1/2$ and is preferable. However, if 5 units are already disrupted in j and none in i , then the first gives i a power of $8/3$ while the second gives i a power of $6/2$, and so now the first disruption minimizes i 's power.

To define an equilibrium, we can adjust the definition of a consistent disruption above to find the *strategic power* of Country i over Country j from disrupting some initial technology. The only difference is that we now allow Country j to have discretion over its flows as well. Note that if Country j acts to minimize i 's power, then this is a zero-sum game (Country i 's payoff is its power and Country j acts to minimize that power and so has a payoff equal to minus Country i 's power). Thus, it has a unique Nash equilibrium outcome and associated power.

That is, to define the strategic power of Country i over j from some initial technology, we adjust the previous definition of a consistent disruption to eliminate the requirement that if $\tau^{\ell'} \in T_j$ then there is proportional rationing. Everything else in the definition is as it was. Then we define a game for each $\tau^i \in T_i^{max}$ such that i and j choose flows that satisfy consistency as defined above (but not requiring j to equalize disruptions), and i maximizes its power while j minimizes it. As this is a zero sum game it has a unique equilibrium power for each $\tau^i \in T_i^{max}$, and i 's strategic power over j is the max of these powers over $\tau^i \in T_i^{max}$.

Computing the Nash equilibrium outcome given any initial $\tau^i \in T_i^{max}$ can be challenging, given that even in the absence of (directed) cycles a technology can be reached via multiple paths from another technology—as it could be reached from both upstream and downstream from paths originating at the other technology. If we focus on supply chains with no *undirected* cycles in the supply chain, then a technology can only be reached from one direction from some other technology, and so it has a well-defined direction of disruption to consider if it is reached. From any starting technology, we can then calculate the equilibrium power as follows. Consider all technologies that are on undirected paths from this technology, and thus lie on some combination of paths going up and downstream. Given the absence of undirected cycles, this means that each of these technologies is either reached via some path coming from upstream or coming from downstream, but not both. Not all such technologies will be reached, but we have to specify what they would do if reached in order to find the equilibrium disruption

consequence. We can work by backward induction. We specify what a country will do at maximum distance for every given set of choices of technologies at lower distance. We then move steps back iteratively, given what we have found at greater distances and for every possible set of choices at technologies of smaller distances. Equilibrium requires that decisions maximize the power of Country i when the technology is in Country i and minimize the power of i when the technology is in Country j . Note that in this case, one can find a pure strategy equilibrium by selecting a pure best response at each node in the backward induction.

This analysis of strategic power is summarized as follows.

- The strategic power of Country i over Country j is well-defined and the (unique) value of a zero sum game in which i picks some technology to disrupt and chooses consistent flows at each technology in that supply chain that it controls (as defined above), and j also chooses consistent flows at each technology.
- In the case of undirected cycles that power can be found in pure strategies and via backward induction.

Our notion of strategic power gives a maximum power for any given disruption, presuming that the two countries stay within the part of the supply chain associated with the initial disruption. There is a richer game in which j could respond with some other disruptions of its own, and so forth. To solve that game, one has to introduce some more widely encompassing preferences, and we leave that for further research. Our analysis provides an important building block for any further extension.

B 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 $\sum_n L_n p_{Ln}$.*

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. ■

B.1 Proof of Lemma 2

LEMMA 2. *There exists an equilibrium. Moreover, the same bundle of final goods is produced in all equilibria.*

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 $t \in \mathbb{R}^{N+T}$. For each technology τ that requires inputs $I(\tau)$, we create $\prod_{k \in I(\tau)} |\{\tau : O(\tau) = k\}|$ source-specific technologies allowing for each possible combination of sourcing choices from different technologies across the different inputs (where $|\cdot|$ denotes cardinality). 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. $\square$

B.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 $(c_1, \dots, c_F) \in \mathbb{R}_+^F$:

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

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 (12) holds. Allowing τ_f to change from one this implies that:⁴⁴

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

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.

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

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

⁴⁴One 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.

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

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

Iteratively substituting in these expressions for all intermediate good technologies except the shocked one, $\hat{\tau}$, the obtained expression converges 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. ■

B.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 $\omega_{jk}^* < \omega_{jk} \forall j, k \in \mathcal{P}$. $\square$

B.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. $\square$

B.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. $\square$

B.6 Proof of Proposition 6

Proof. Before presenting the proof, let us mention the logic. Given that no technological flow is sent to 0, and all of j 's proportional rations are done relative to the original equilibrium flows (regardless of the sequence of disruptions), it follows that the marginal impact of a disruption consistent with disruption of technology $\tau_i \in T_i$, does not depend on what other disruptions are conducted before or after it. Thus, the total value of a sequence of disruptions can be written as a weighted sum of the values of the powers of marginal disruptions. It follows that the total value can be no more than the maximum power across the marginal disruptions, which comes from a single $\tau_i \in T_i$.

More formally, consider a disruption $(x_{\tau\tau'}, y_\tau, L_\tau)_{\tau,\tau'} \in \mathcal{D}_i$ that is obtained via a sequence of K individual technology disruptions by Country i given an initial equilibrium $(x_{\tau\tau'}^*, y_\tau^*, L_\tau^*)_{\tau,\tau'}$. Let $(x_{\tau\tau'}^{\tilde{K}}, y_\tau^{\tilde{K}}, L_\tau^{\tilde{K}})_{\tau,\tau'}$ be the outcome reached after $\tilde{K} \leq K$ of the individual technology shocks. Thus $(x_{\tau\tau'}^{\tilde{K}}, y_\tau^{\tilde{K}}, L_\tau^{\tilde{K}})_{\tau,\tau'} = (x_{\tau\tau'}, y_\tau, L_\tau)_{\tau,\tau'}$. As all such outcomes are consistent (by construction)

$$y_\tau^{\tilde{K}} = \sum_{\tau'} x_{\tau\tau'}^{\tilde{K}} = \sum_{\tau': O(\tau')=g} \frac{x_{\tau'\tau}^{\tilde{K}}}{\tau_g},$$

for all $g \in I(\tau)$, all τ and all $\tilde{K} \leq K$.

Thus, if we let d_τ^k be the output reduction of technology τ , and let $d_{\tau\tau'}^k$ be the flow reduction from technology τ to τ' , caused by the k th individual technology disruption then

$$y_\tau^* - \sum_{k=1}^{\tilde{K}} d_\tau^k = \sum_{\tau'} (x_{\tau\tau'}^* - \sum_{k=1}^{\tilde{K}} d_{\tau\tau'}^k) = \sum_{\tau': O(\tau')=g} \frac{1}{\tau_g} \left(x_{\tau'\tau}^* - \sum_{k=1}^{\tilde{K}} d_{\tau'\tau}^k \right), \quad (8)$$

for all $g \in I(\tau)$, all τ and all $\tilde{K} \leq K$.

This implies that for $\tilde{K} = 1$ the disruption $(x_{\tau\tau'}^1, y_{\tau'}^1, L_{\tau'}^1)_{\tau, \tau'}$ is a consistent individual technology disruption (as it must be because in the first step we start from the equilibrium outcome). Now consider $\tilde{K} = k > 1$. As equation (8) holds for $\tilde{K} - 1$, this implies that

$$-d_{\tau}^k = -\sum_{\tau'} d_{\tau\tau'}^k = -\sum_{\tau': O(\tau')=g} \frac{d_{\tau'\tau}^k}{\tau_g},$$

for all $g \in I(\tau)$ and all τ , so using consistency of the initial equilibrium

$$y_{\tau}^* - d_{\tau}^k = \sum_{\tau'} x_{\tau\tau'}^* - d_{\tau\tau'}^k = \sum_{\tau': O(\tau')=g} \frac{1}{\tau_g} (x_{\tau\tau'}^* - d_{\tau\tau'}^k),$$

for all $g \in I(\tau)$ and all τ . Thus each disruption in the sequence is by itself a consistent individual technology disruption with respect to the initial equilibrium.

As by assumption $x_{\tau\tau'}^* > 0$ implies the final flow associated with the disruption $x_{\tau\tau'}^K > 0$ and, absent any zero-outed flows, rationing that occurs at all of the K steps in the sequence of disruptions is done relative to the initial equilibrium flows, the k th disruption is a valid and available individual technology disruption for all k , for the initial equilibrium.

Without loss of generality label these valid individual technology disruptions $\zeta_1, \dots, \zeta_K$. Each such consistent disruption in this set is associated with a change in employed labor in country i and in country j . Let $\Delta L_i(\zeta_1)$ denote the change in employed labor in country i due to disruption ζ_1 . The power of the disruption for country i is then a strictly increasing function of $\sum_{k=1}^K \Delta L_j(\zeta_k) / \sum_{k=1}^K \Delta L_i(\zeta_k)$. Moreover,

$$\frac{\sum_{k=1}^K \Delta L_j(\zeta_k)}{\sum_{k=1}^K \Delta L_i(\zeta_k)} \leq \max_k \frac{\Delta L_j(\zeta_k)}{\Delta L_i(\zeta_k)},$$

and so there exists a valid individual technology disruption to a technology in T_i that is weakly more powerful. $\square$

Supplementary Appendix for “Supply Chain Disruptions, the Structure of Production Networks, and the Impact of Globalization,” by Matthew L. Elliott and Matthew O. Jackson

I Equilibrium

It is helpful to define the transportation-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 $\hat{p}_{\tau f} = \min_{\tau': O(\tau')=f} p_{\tau'} = p_f$.

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 use 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 $(c_{1n}, \dots, c_{Fn})$ to maximize

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

- **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

$$\sum_{\tau': O(\tau')=k} \frac{x_{\tau'\tau}}{\theta_{\tau'\tau}} = -\tau_k y_\tau, \quad (9)$$

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:

$$L_n = \sum_{\tau} x_{n\tau}. \quad (10)$$

- 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:

$$\sum_{\tau'} x_{\tau\tau'} = y_\tau. \quad (11)$$

- 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:

$$\sum_n c_{fn} = \sum_{\tau \in T: O(\tau)=f} y_\tau. \quad (12)$$

Given constant returns to scale, in equilibrium active technologies 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.

I.1 Existence and Uniqueness of Consumption

To ensure equilibrium existence, it is necessary that the available technologies avoid money pumps. We assume that each technology uses a positive amount of labor, and given the limited supply of labor, that is sufficient for avoiding money pumps and is quite natural, although

weaker conditions could be used.⁴⁵

We also assume that T is such that every final good has a viable supply chain, meaning that there is some combination of technologies that produce a positive amount of it. This ensures that every final good has a finite price in equilibrium.

Lemma 2 in Appendix B shows that there exists an equilibrium and the same bundle of final goods is produced in all equilibria. We show that by expanding the space of technologies to let them take source-specific inputs, transportation costs can be incorporated into them. Then, with this transformation in hand, standard existence results apply.

We note that every final good is consumed and has a finite, positive price in equilibrium, and that the same is true of each country's labor. These follow from the fact that every final good has a viable supply chain, each country has a positive endowment of labor that is completely supplied in equilibrium, and preferences are increasing. This ensures that all final goods are produced and consumed, and so must have positive, finite prices (given increasing preferences), and labor cannot have either a 0 price (or some technology would use an infinite amount of it⁴⁶), or an infinite price (or some consumer's problem would result in infinite demand and markets would not clear).

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 transportation costs are sufficiently large that labor never moves, but in contrast intermediate goods have transportation 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.⁴⁷ If instead, labor transportation

⁴⁵We 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.

⁴⁶Note that this follows from the proof in that there are equivalent primitive technologies in terms of labor units, and 0 profits ensure that it is used in

⁴⁷Suppose 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 positive vector summing to 1 such that the wage in country i is

costs are sufficiently low that labor in one country can substitute for that in another, then that begins to pin down relative wages, until with sufficient mobility wages are equal across countries.

II Potential impact of price flexibility

In other cases, the ability for prices to adjust can make a big difference. One way to measure this difference is to consider the ratio of lost output when prices can adjust, to lost output when prices cannot adjust (and the proportional rationing constraint must be satisfied). Given an initial economy $\mathcal{E}$ and associated pre-shock equilibrium, and letting y_τ^{ip} be a solution to the minimum disruption problem (with inflexible prices), we define the *loss to price rigidity (LPR)* as

$$LPR(\mathcal{E}, T^{Shocked}, \lambda) = \frac{\sum_{\tau: O(\tau) \in F} (y_\tau^* - y_\tau^{ip})}{\sum_{\tau: O(\tau) \in F} (y_\tau^* - y_\tau^{fp})} \geq 1.$$

PROPOSITION 9. *The potential loss to price rigidity is unbounded: there exists a sequence of economies $\{\mathcal{E}_t\}_t$ such that $\lim_{t \rightarrow \infty} LPR(\mathcal{E}_t, T^{Shocked}, \lambda) = \infty$.*

We prove Proposition 9 by example. Consider an economy $\mathcal{E}_t$ with initial equilibrium flows as depicted in Figure 20 for $t = 3$. For this example it is helpful to define intermediate good levels. There are two levels of intermediate goods. Level 1 intermediate goods use only labor as an input. Level 2 intermediate good use labor and intermediate level 1 goods as inputs. Finally, there are final goods which use intermediate level 2 goods and labor as inputs.

Intermediate level 1 consists of a single type of good, good 0, produced by two countries 1 and 2. Both countries require 1 unit of labor to produce one unit of output. Intermediate level 2 consists of t goods. All t goods are produced in all t countries. Each good intermediate level 1 good i requires one unit of good 0 and one unit of labor. In equilibrium, all goods are produced in all countries and for the production of good i by country i , and for $i = 1, \dots, t$, input good 0 is sourced from country 1, while otherwise input good 0 is sourced from country 2. Finally, each country also produces the unique final good. Production of this final good in all countries requires one unit of all intermediate goods $1, \dots, t$ and one unit of labor to produce one unit of output.

We assume that no transportation costs for any goods except labor. Country 1, is endowed with $2t + 1$ units of labor, country 2 is endowed with $t^2 + 1$ units of labor and the other countries are endowed with $t + 1$ units of labor. In equilibrium: Country 1 produces t units of good 0, one unit of goods $1, \dots, t$ and one unit of the final good, Country 2 produces $t(t - 1)$ units of good 0, one unit of goods $1, \dots, t$ and one unit of the final good. The other countries produce one unit of goods $1, \dots, t$ and one unit of the final good. There is then an equilibrium in which all final goods are produced using only domestically produced inputs, and all labor works in its home country.

weakly greater than $1/k_i$.




Figure 20: Initial equilibrium of economy $\mathcal{E}_3$, with red lines indicating flows that will be disrupted with inflexible pricing following a shock to technology 01.




Figure 21: Rerouted flows with flexible pricing following a shock to technology 01.

We consider a shock to the output of good 0 production in country 1 that reduces its output to $(1 - \lambda)t$ units. With inflexible prices (such that there is proportional rationing), by Proposition 4, this reduces the aggregate output of final good production to $(1 - \lambda)$ of its initial level (i.e., it achieves the bound from Proposition 3). Intuitively, output of good i in country i is reduced to $(1 - \lambda)$ of its initial level, because it sources good 0 from country 1, and as all of good i used by country i in the production of its final good is sourced domestically, output of its final good production also reduced to $(1 - \lambda)$ of its initial level.

In contrast, with flexible prices, it is possible to maintain final good production in countries $2, \dots, t$. As shown in Figure 21 for $t = 3$, this can be done by rerouting the output of input good i in country 1 to country i , for $i = 2, \dots, t$, for use in country i 's final good production. Thus, letting τ_{01} be the technology used to produce good 0 in country 1, $LPR(\mathcal{E}_t, \tau_{01}, \lambda) = \frac{\lambda}{\lambda/t}$, and so

$$\lim_{t \rightarrow \infty} LPR(\mathcal{E}_t, \tau_{01}, \lambda) = \lim_{t \rightarrow \infty} t \left(\frac{\lambda}{\lambda} \right) = \infty.$$

However, if instead both level 0 intermediate good producers has been shocked, then by Proposition 5 we would have $LPR(\mathcal{E}_t, \{\tau_{01}, \tau_{02}\}, \lambda) = \lambda/\lambda = 1$, for all t .