Inventories, Demand Shocks Propagation and Amplification in Supply Chains*

Alessandro Ferrari†

University of Zurich

First Version: February 2019

This version: October 2022

Abstract

I study the role of industries position in supply chains in shaping the transmission of final demand shocks. First, I use a shift-share design based on destination-specific final demand shocks and destination shares to show that shocks amplify upstream. Quantitatively, I find upstream industries respond to final demand shocks up to three times as much as final goods producers. To organize the reduced form results, I develop a tractable production network model with inventories and study how the properties of the network and the cyclicality of inventories interact to determine whether final demand shocks amplify or dissipate upstream. I test the mechanism both by directly estimating the model and in reduced form and I find evidence of the role of inventories in explaining heterogeneous output elasticities.

Keywords: production networks, supply chains inventories, shock amplification.

JEL Codes: C67, E23, E32, F14, F44, L14, L16

*This paper supersedes a previously circulated version under the title "Global Value Chains and the Business Cycle". I am thankful to my advisors Ramon Marimon and Philipp Kircher for their invaluable guidance and support. I thank Daron Acemoglu, Pol Antràs, Giacomo Calzolari, Vasco Carvalho, Russell Cooper, Juan Dolado, David Dorn, Matt Elliott, Matteo Escudé, Dalila Figueiredo, David Hémous, Nir Jaimovich, Andrei Levchenko, Michele Mancini, Isabelle Méjean, Ralph Ossa, Nitya Pandalai-Nayar, Mathieu Parenti, Armin Schmutzler, Alireza Tahbaz-Salehi for their feedback. This paper also benefited from comments in the seminars at EUI, Collegio Carlo Alberto, Paris School of Economics, CERGE-EI Prague, CSEF Naples, University of Zurich, Queen Mary University London, HEC Montreal, ECB Research Department, EIEF, World Trade Institute and presentations at ASSET, RIEF AMSE, WIEM, the EEA 2019 meeting, ETSG 2019, Econometric Society European Winter Meeting, GVC Conference of Bank of Italy, Large Firms Conference at Bank of France and the GEN Workshop. Lorenzo Arcà and Lorenzo Pesaresi provided excellent research assistance. All the remaining errors are mine.

†alessandro.ferrari@econ.uzh.ch

1 Introduction

The way goods and services are produced and distributed has changed significantly over recent decades. Today, production occurs along complex supply chains with goods crossing borders multiple times before reaching final consumers. A natural question, in light of these trends, is whether we live in a less volatile world than we used to, as shocks can be absorbed at multiple nodes of the chain and diversified away. On the other hand, they could travel longer distances and snowball across firms and countries. This question has recently gained even more salience as governments try to understand what are the causes and consequences of the recent supply chain crisis and whether a policy response is necessary. Part of the policy discussion has considered reshoring as an option to shorten supply chains and reduce the propagation of shocks.[1]

In this paper I address this question by studying how shocks travel in production networks. I investigate two fundamental forces. First, a standard effect in networks such that perturbations are diffused and absorbed at multiple stages. This implies that, for a given shock, longer and more complex chains are less volatile. Second, a force somewhat overlooked by economists called the *bullwhip effect*, according to which, in the presence of inventories, shocks can magnify as they travel across firms. The role of inventories as shocks amplifier or absorber is of particular interest, as it is one of the possible strategies to avoid prolonged disruptions in supply chains. If firms respond to uncertain demand and supply by building up larger stocks, they might contribute to higher volatility and shock propagation.

I start by providing 5 empirical observations that motivate this paper: in the last decades i) production chains have significantly increased in length, measured by the number of production steps goods undergo before reaching consumers; ii) and the spatial concentration of demand, measured by the Herfindal-Hirschman Index of destination sales shares has significantly declined. These two empirical facts suggest that the rise of complex supply chains may better insulate from final demand shocks as they are absorbed in multiple steps and diversified away. However, I also show that iii) inventories are procylically adjusted, and, as a consequence, iv) output is more volatile than sales. These two observations, in the context of a production network, imply upstream amplification, which may be strengthened by longer and more complex chains. Finally, I confirm a recent finding by Carreras-Valle (2021): v) inventory-to-sales ratios have been increasing since 2005 for US manufacturing firms. This suggests that facts iii), iv) and, therefore, upstream amplification, may have become even more salient.

Motivated by these empirical observations, I ask whether we observe a differential output response to final demand shocks depending on an industry's position in the supply chain. I do so by means of a shift-share design based on destination-specific foreign demand shifters and a measure of exposure accounting for direct and indirect linkages. The shift-share structure

[1] In June 2021 the Biden-Harris Administration instituted the *Supply Chain Disruption Task Force* "to provide a whole-of-government response to address near-term supply chain challenges to the economic recovery", see White House (2021). For the European context see Raza et al. (2021), a study commissioned by the International Trade Committee of the European Parliament considering reshoring options and the European Parliament resolution calling for "smart reshoring [to] relocate industrial production in sectors of strategic importance for the Union" and the creation of a program that "helps make our supply chains more resilient and less dependent by reshoring, diversifying and strengthening them", see European Parliament (2020).

allows me to estimate the causal effect of changes in final demand on output. With this design, I estimate a model in which I allow the elasticity of output to this exogenous change in final demand to vary by upstreamness, which measures the network distance between an industry and final consumers. I find that industries at different points of the value chains have significantly different responses. Quantitatively, an industry very close to consumption increases the growth rate of output by about .5pp for every 1pp increase in the growth rate of final demand. At the same time, industries very far from consumption (6 or more steps of production away) respond 1.2pp for every 1pp increase in the growth rate of final demand. I confirm the same result by instrumenting consumers' final demand changes with government expenditure, by using the China syndrome shock of Autor et al. (2013) and through the fiscal spending shocks for the US of Acemoglu et al. (2012). In all cases the most upstream industries respond between 2 and 3 times as much as the ones closest to consumers. Finally, I also find that more upstream industries have a significantly larger positive response of inventories when demand grows, with an up to 6-fold increase going from one to six production steps away from consumption.

These empirical findings are hard to rationalize in the context of current models of shock propagation in production networks. For this reason I lay out an extension of the workhorse production network model in which I introduce the role of inventories. I build this framework for two reasons. First, formalizing the problem allows me to shed light on the key features of the network and of the inventory problem that lead to upstream amplification vs dissipation of final demand shocks. Second, the model provides an estimating equation, which allows me to test the key mechanism directly and to compute counterfactuals.

Theoretically, I show that, in the special case of a simple line network without labor, procyclical inventory adjustment is a sufficient condition for upstream amplification. This is not the case anymore when I allow for a general network structure and the use of labor in production. Intuitively, at every step of the network, part of the shock absorption is done by labor, which is by nature irreproducible. As such the network naturally dampens final demand shocks as they travel upstream. It is then possible to characterize how longer chains, or different position in them, alter the output response of firms depending on how strong the inventory amplification versus the network dissipation forces are.

I conclude by estimating the model-implied relationship directly and performing counterfactuals. First, the estimating the main equation in the model, linking output changes to demand changes through the position in the network and inventories, generates the predicted signs from theory. I confirm these results also by estimating a model-free reduced form regression of the inventory channel. Second, I separately identify the role of changes in the network and inventories on observed industry responses to final demand shocks. Quantitatively, by comparing the same economy with the 2000 vs 2014 network, I find that this change brings about two opposing forces: i) as the length of supply chains increases, amplification forces intensify; ii) the increase dispersion of destination shares reduces the effective volatility of final demand by about 10%. Quantitatively the latter force dominates and output growth volatility declines. This masks the opposing effect of lower demand volatility and higher average output elasticities. In an alternative experiment, where I also increase inventories-to-sales ratio by 25%, similar to the recent trend discussed above, I find that the benefits of lower demand volatility are significantly undone by the higher responsiveness of output through inventory amplification.

More broadly, these results highlight a further trade-off element to the discussion on how to make supply chains more resilient. If firms increase their inventory buffer to prevent prolonged disruption, this might come at the cost of permanently higher volatility for the economy. This paper shows that longer supply chains might reinforce this effect, making the economy more uncertain and volatile.

Related Literature This paper relates, first, to the growing literature on shocks in production networks. From the theoretical standpoint, this line of research, which stems from Carvalho (2010), Acemoglu et al. (2012), and more recently Baqaee and Farhi (2019), studies the role of network structure in the propagation of idiosyncratic industry-level shocks. This paper builds a similar model, explicitly allowing for forces generating potential amplification in the network. This extension allows me to characterize theoretically the amplification patterns I find in the data and reconcile them with more aggregate empirical results on the relative volatility of final and intermediate goods sales. From an empirical standpoint I build on Acemoglu et al. (2016); Barrot and Sauvagnat (2016); Boehm et al. (2019); Carvalho et al. (2020); Dhyne et al. (2021); Korovkin and Makarin (2021), who study how shocks propagate in a production network. Relative to these contributions, by exploiting destination and time-specific shocks to consumption, I build industry-level exogenous variation in the spirit of Shea (1993). This approach enables me to evaluate the heterogeneous response to final demand shocks across industries at different points of the supply chain holding fixed the size of the shock itself. Further, I include the role of inventories as an additional channel contributing to the patterns of shock propagation.

Secondly, it relates to the literature on the role of inventories as a source of amplification and the *bullwhip effect*. This literature, stemming from Forrester (1961) and more recently Kahn (1987); Blinder and Maccini (1991); Metters (1997); Chen et al. (1999); Ramey and West (1999); Lee et al. (2004); Alessandria et al. (2010) suggested that when inventories are procyclically adjusted, they can amplify shocks upstream. I embed this mechanism in the network context to study the horse race between the role of inventories and the network dissipation effect. Lastly, from an empirical standpoint, the effect of inventories as an amplification device has been studied at several levels of aggregation by Alessandria et al. (2010), Altomonte et al. (2012) and Zavacka (2012). These papers all consider exogenous variation given by the 2008 crisis to study the responsiveness of different sectors or firms to the shock, depending on whether they produce intermediate or final goods. Using an indicator for exposure to the shock creates an identification problem as it is not possible to separate whether different sectors responded differently to the same shock or were hit by a different shock altogether. Relative to these works, the approach based on the shift-share design allows me to isolate the heterogeneity in the response to the same change in final demand depending on an industry's position in the supply chain.

Roadmap The rest of the paper is structured as follows: Section 2 provides the key empirical observations that motivate this paper. Section 3 provides the details on the data and the empirical strategy. Section 4 presents the reduced form results on upstream shocks propagation. Section 5 describes the model and provides the key comparative statics results and counterfactual exercises. Finally, Section 6 concludes.

2 Motivating Evidence

In this section I provide key empirical observations on production chains and inventories which are critical in motivating the empirical analysis and disciplining the paper's theoretical framework.

Fact 1: Production chains have increased in length

In the past few decades modes of production have changed markedly. As highlighted by the World Bank Development Report (2020), a growing share of production now occurs in many stages and crosses borders multiple times before reaching consumers. Figure A.1 shows, on average, how many production steps a good undergoes before it is finally consumed. The sector-sales weighted average of upstreamness steadily increased in the period covered by the World Input Output Database (WIOD 2016 release, see Timmer et al., 2015), from 2.6 in 2000 to 3.3 in 2014. [2] This change is driven in equal measure by an increase in the weight of already long chains (between component) and the increase in length of chains with large weights (within component). As mentioned in the discussion of the existing literature, a salient feature of current models of production networks is that shocks tend to dissipate as they travel away from their source. Taken together with the increasing length of production chains, this feature would imply that the network is becoming more resilient to demand shocks.

Fact 2: The spatial distribution of sales has become less concentrated

A second element related to the increasing role of international linkages in production is that of diversification. Using the WIOD data it is possible to construct exposure shares of each industry which account for intermediate linkages. I discuss this in larger details in Section 3.2, but in summary this measure represents how much of a given industry's output is eventually consumed in a given destination whether it is sold directly or indirectly. Figure A.2 shows the trend in the Herfindal-Hirschman Index of these sales shares. The key observation is that the HHI has been significantly declining over the period 2000-2014, whether I use a simple or a sales-weighted average. Quantitatively the unweighted HHI went from .51 in 2000 to .44 in 2014, while the sales weighted HHI went from .7 to .6 in the same period. This observation would suggest that as industries are now exposed to a wider array of destinations they should be less exposed to idiosyncratic shocks. In turn, this should reduce output volatility.

Facts 1 and 2 are likely driven by the rise of of global value chains and cross-border production

[2] This observation holds true when using a measure of supply chains length which sums both the average distance of an industry from final consumers and the distance of the industry from pure valued as shown in Figure B.1.

which implies both longer chains and a more diversified portfolio of customers across space. Furthermore, these trends, taken at face value, would suggest that we live in a less volatile world than we used to since network dissipation has more potential to occur and industries are generally less exposed to idiosyncratic demand shocks.

Fact 3: Inventories are adjusted procyclically

Inventories can, in principle, both amplify or absorb shocks propagating in production chains. Their effect depends fundamentally on whether they are adjusted procyclically or countercyclically. As the literature has noted, inventory investment is procyclical while the inventory-to-sales ratio is countercyclical. In terms of magnitude, the inventory-to-sales ratio based on the monthly data from the US Census Manufacturing & Trade Inventories & Sales is approximately 130% of monthly sales. To study its cyclicality, I estimate non-parametrically the mapping between inventories and sales (after applying the Hodrick-Prescott filter). I obtain an average derivative of .1 for annual data and 1.35 for monthly data, as reported in Table A.1. These estimates suggest that end-of-period inventories are an increasing function of sales. Figure A.3 provides the distribution of the sector-specific estimated derivative of inventories with respect to sales for both samples. The graph shows that the whole distribution lies above zero, which suggests that all the available sectors feature procyclical inventories over the sample period. The observed procylicality is robust to the filtering and the inclusion of sector and time-fixed effects as shown in Table A.1.[3]

Fact 4: Output is more volatile than sales

The presence of procylically adjusted inventories can make output more volatile than sales. In Figure A.4 I plot the distribution of the ratio between the standard deviation of output (σ_y) and sales (σ_q). Both series are HP-filtered and output is computed by summing sales and inventory changes. The graphs show the distribution of σ_y/σ_q across sectors. I plot these distributions for monthly data from the US Census Manufacturing & Trade Inventories & Sales data from January 1992 to August 2021 covering NAICS 3-digit industries. I also provide the same statistic after aggregating the data to quarters and years. Lastly the same distribution is reported for the yearly data in the NBER CES Manufacturing Industries Database from 1958 to 2018 for NAICS 6-digit industries. In all cases the distributions highlight how for most industries, particularly at the quarterly and yearly frequency, the ratio lies above 1, suggesting that output is more volatile than sales.

Facts 3 and 4, taken together, would suggest that when a chain experiences a demand shock from final consumers, the response of industries increases as the shock travels upstream. To test this simple intuition, in Figure A.5 I plot the correlation between the log of volatility of sales growth rates and the log of upstreamness, which measures the distance of an industry from

[3] The Census data provides breakdowns of inventories by finished products and materials. These broken down series has a significantly larger amount of missing data relative to the total inventories series. For this reason here I use the total inventory as the measure of inventories for a given industry. In Appendix B.1 I provide the same distribution of estimates separating inventories of final products and materials. The main conclusion that inventories are procyclically adjusted remains.

final consumers. This simple correlation is positive and statistically significant, suggesting that a more upstream position in the production chains is associated with higher volatility of sales.

Fact 5: Inventories-to-sales ratios are increasing

I conclude by reporting an empirical observation first uncovered by Carreras-Valle (2021): inventories-to-sales ratios have been increasing since 2005 for US manufacturing industries. I replicate this finding in Figure A.6 for both the annual data in the NBER CES database and the monthly data from the Census. The author suggests that this is largely driven by the substitution of domestic with foreign inputs, which are cheaper but entail higher delivery lags, to which firms optimally respond by holding more inventories.[4] This observation would suggest that the same forces strengthening Facts 1 and 2 can endogenously lead to larger inventories and thereby reinforce Facts 3 and 4.

Combining these empirical observations paints an inconclusive picture as to whether the rise of value chains should imply more or less volatile production. The goal of the rest of the paper is to focus on demand shocks and ask whether we observe their effects amplify or attenuate as they travel upstream in a production chain. Testing this formally requires a measure of where a given industry is positioned relative to final demand and a set of exogenous demand shocks. In the next section I lay out my approach to address both these measurement problems and the data used in the estimation.

3 Data and Methodology

3.1 Data

Input-Output Data

The main source of data in this paper is the World Input Output Database (WIOD) 2016 release, see Timmer et al. (2015). It contains the Input-Output structure of sector-to-sector flows for 44 countries from 2000 to 2014 at the yearly level. The data is available at the 2-digit ISIC revision 4 level. The total number of sectors in WIOD is 56, which amounts to 6,071,296 industry-to-industry flows and 108,416 industry-to-country flows for every year in the sample. The full coverage of the data in terms of countries and industries is shown in Table B.1 and B.2 in the Appendix. The structure of the WIOD data is represented in Figure 1.

The World Input-Output Table represents a world economy with J countries and S industries per country. The $(S \times J)$ by $(S \times J)$ matrix whose entries are denoted by Z represents flows of output used by other industries as intermediate inputs. Specifically, Z_{ij}^{rs} denotes the value of output of industry r in country i used as intermediate input by industry s in country j. In addition to the square matrix of input use, the table provides the flows of output used for final consumption. These are denoted by F_{ij}^{r}, representing the value of output of industry r in country

[4]A similar argument is also underlying the results in Alessandria et al. (2010, 2013) suggesting that firms more involved international activity tend to hold more inventories.

Figure 1: World Input Output Table

			Input use & value added							**Final use**			**Total use**
			Country 1			$\cdots$	Country J			Country 1	$\cdots$	Country J	
			Industry 1	$\cdots$	Industry S	$\cdots$	Industry 1	$\cdots$	Industry S				
		Industry 1	Z_{11}^{11}	$\cdots$	Z_{11}^{1S}	$\cdots$	Z_{1J}^{11}	$\cdots$	Z_{1J}^{1S}	F_{11}^{1}	$\cdots$	F_{1J}^{1}	Y_1^1
Intermediate	Country 1	$\cdots$	$\cdots$	Z_{11}^{rs}	$\cdots$	$\cdots$	$\cdots$	Z_{1J}^{rs}	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$
		Industry S	Z_{11}^{S1}	$\cdots$	Z_{11}^{SS}	$\cdots$	Z_{1J}^{S1}	$\cdots$	Z_{1J}^{SS}	F_{11}^{S}	$\cdots$	F_{1J}^{S}	Y_1^S
inputs	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	Z_{ij}^{rs}	$\cdots$	$\cdots$	$\cdots$	$\cdots$	F_{ij}^{r}	$\cdots$	Y_i^r
		Industry 1	Z_{J1}^{11}	$\cdots$	Z_{J1}^{1S}	$\cdots$	Z_{JJ}^{11}	$\cdots$	Z_{JJ}^{1S}	F_{J1}^{1}	$\cdots$	F_{JJ}^{1}	Y_J^1
supplied	Country J	$\cdots$	$\cdots$	Z_{J1}^{rs}	$\cdots$	$\cdots$	$\cdots$	Z_{JJ}^{rs}	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$
		Industry S	Z_{J1}^{S1}	$\cdots$	Z_{J1}^{SS}	$\cdots$	Z_{JJ}^{S1}	$\cdots$	Z_{JJ}^{SS}	F_{J1}^{S}	$\cdots$	F_{JJ}^{S}	Y_J^S
Value added			VA_1^1	$\cdots$	VA_1^S	VA_j^s	VA_J^1	$\cdots$	VA_J^S				
Gross output			Y_1^1	$\cdots$	Y_1^S	Y_j^s	Y_J^1	$\cdots$	Y_J^S	$\cdots$			

i consumed by households, government and non-profit organizations in country j. Following the literature, I denote $F_i^r = \sum_j F_{ij}^r$, namely the value of output of sector r in country i consumed in any country in the world.

For a subset of results I use the 2002 I-O Tables from the BEA, following Antràs et al. (2012). These are a one year snapshot of the US production network and cover 426 industries.

Inventory Data

The Input-Output data is complemented with information about sectoral inventories from the NBER-CES Manufacturing Industry Database. This dataset contains information about sales and end-of-the-period inventories for 473 6-digit 1997 NAICS US manufacturing industries from 1958 to 2011. I concord the inventory data to the level of aggregation of WIOD.

The second source of inventory data is the US Census Manufacturing & Trade Inventories & Sales. This dataset covers NAICS 3-digit industries monthly since January 1992. The data includes information for finished products, materials and work-in-progress inventories which I sum into a single inventory measure. Throughout I use the seasonally adjusted version of the data.

3.2 Measurement and Methodology

This section describes the empirical methodology used. I start by reviewing the existing measure of upstreamness as distance from final consumption proposed by Antràs et al. (2012). Next, I discuss the identification strategy based on the shift-share design. I show how to compute the sales share in the industry portfolio accounting for indirect linkages. This allows me to evaluate the exposure of industry sales to specific partner country demand fluctuations even when goods reach their final destination by passing through third countries. Then, I discuss the fixed-effect model used to extract and aggregate country and time-specific demand shocks from the final consumption data.

3.2.1 Measuring the Position in Production Chains

The measure of the upstreamness of each sector counts how many stages of production exist between the industry and final consumers, as proposed by Antràs et al. (2012). The measure is bounded below by 1, which indicates the entire sector output is used directly for final consumption. The index is constructed by assigning value 1 to the share of sales directly sold to final consumers, value 2 to the share sold to consumers after it was used as an intermediate good by another industry, and so on. Formally:

$$U_i^r = 1 \times \frac{F_i^r}{Y_i^r} + 2 \times \frac{\sum_{s=1}^{S} \sum_{j=1}^{J} a_{ij}^{rs} F_j^s}{Y_i^r} + 3 \times \frac{\sum_{s=1}^{S} \sum_{j=1}^{J} \sum_{t=1}^{T} \sum_{k=1}^{K} a_{ij}^{rs} a_{jk}^{st} F_k^t}{Y_i^r} + ... \quad (1)$$

where F_i^r is the value of output of sector r in country i consumed anywhere in the world and Y_i^r is the total value of output of sector r in country i. a_{ij}^{rs} is dollar amount of output of sector r from country i needed to produce one dollar of output of sector s in country j, defined as $a_{ij}^{rs} = Z_{ij}^{rs}/Y_j^s$. This formulation of the measure is effectively a weighted average of distance, where the weights are the distance specific share of sales and final consumption.[5]

Provided that $\sum_i \sum_r a_{ij}^{rs} < 1$, which is a natural assumption given the definition of a_{ij}^{rs} as input requirement, this measure can be computed by rewriting it in matrix form: [6]

$$U = \hat{Y}^{-1}[I - \mathcal{A}]^{-2}F, \quad (2)$$

where U is a $(J \times S)$-by-1 vector whose entries are the upstreamness measures of every industry in every country. $\hat{Y}$ denotes the $(J \times S)$-by-$(J \times S)$ diagonal matrix whose diagonal entries are the output values of all industries. The term $[I - \mathcal{A}]^{-2}$ is the power of the Leontief inverse, in which $\mathcal{A}$ is the $(J \times S)$-by-$(J \times S)$ matrix whose entries are all a_{ij}^{rs} and finally the vector F is an $(J \times S)$-by-1 whose entries are the values of the part of industry output that is directly consumed. Equation 1 shows, the value of upstreamness of a specific industry r in country i can only be 1 if all its output is sold to final consumers directly. Formally, this occurs if and only if $Z_{ij}^{rs} = 0, \forall s, j$, which immediately implies that $a_{ij}^{rs} = 0, \forall s, j$.

Table B.3 provides a list of the most and least upstream industries in the WIOD sample. Predictably, services are very close to consumption while raw materials tend to be distant.

3.2.2 Identification Strategy

The goal is to evaluate the responsiveness of output to changes in demand for industries at different positions in the supply chain. To estimate this effect, the ideal setting would be one in which I observe exogenous changes in final demand for each producing industry in the sample. As

[5] This discussion implicitly assumes that a sector's input mix is independent of the output use or destination. De Gortari (2019), using customs data from Mexico, shows that this can lead to mismeasurement. These concerns cannot be fully addressed in this paper due to the limitation imposed by the WIOD aggregation level.

[6] For this not to be true some industry would need to have negative value added since $\sum_i \sum_r a_{ij}^{rs} > 1 \Leftrightarrow \sum_i \sum_r Z_{ij}^{rs}/Y_j^s > 1$, meaning that the sum of all inputs used by industry s in country j is larger than the value of its total output. To compute the measure of upstreamness I apply the inventory correction suggested by Antràs et al. (2012), the discussion of the method is left to Appendix B.3.

this is not possible, I approximate this setting using a shift-share instrument approach to gauge the causal effect of interest. In the present application, using a shift-share design boils down to generating plausibly exogenous changes in final demand for producing industries as averages of destination-specific aggregate changes weighted by the appropriate measure of exposure.

The methodology, as described in Borusyak et al. (2022), Goldsmith-Pinkham et al. (2020) and Adão et al. (2019), requires exogeneity of either the shares or of the shocks. In the current case it is implausible to assume that the destination shares are exogenous as firms choose the destinations they serve. Identification can be obtained by plausibly as good as randomly assigned shifters (destination-specific shocks).

Define the shift-share changes in final demand for industry r in country i at time t as

$$\hat{\eta}^r_{it} = \sum_j \xi^r_{ij} \hat{\eta}_{jt}. \tag{3}$$

Where ξ^r_{ij} represents the fraction of the value of output of industry r in country i consumed directly or indirectly in destination j in the first sample period and $\hat{\eta}_{jt}$ the change in final consumption of country j at time t across all products from all origin countries.

As shown in Borusyak et al. (2022), the shift-share instrument estimator is consistent provided that the destination-specific shocks are conditionally as good as randomly assigned and uncorrelated. I discuss the plausibility of these identifying assumptions after describing how I compute demand shock shifters in Section 3.2.4.

Next I describe how I compute the destination shares and the destination shocks.

3.2.3 Sales Shares

The standard measure of sales composition uses trade data to compute the relative shares in a firm's sales represented by different partner countries (see Kramarz et al., 2020). However, such a measure may overlook indirect dependencies through third countries. As an example of this problem, take wood manufacturing in Canada. The output of this industry can be used both by final consumers and by firms as intermediate input. Assume that half of the country's production is sold directly to Canadian consumers and the other half to the furniture manufacturing industry in the US. The standard trade data-based sales share would state that the sales composition of the industry is split equally between Canada and the US. But this is not necessarily true because the US industry may sell its output back to Canadian consumers. Take an extreme example in which all US furniture industry output is exported back to Canada: the only relevant demand for the Canadian wood manufacturing industry, then, comes from Canadian consumers.

This example illustrates that, particularly for countries highly interconnected through trade, measuring portfolio composition only via direct flows may ignore a relevant share of final demand exposure.[7] The input-output structure of the data allows for a full accounting of these indirect

[7] Using granular production network data for Belgium, Dhyne et al. (2021) show that firms respond to changes in foreign demand even when they are only indirectly exposed to them.

links when analyzing sales portfolio composition. Formally, define the share of sales of industry r in country i that is consumed by country j as

$$\xi_{ij}^r = \frac{F_{ij}^r + \sum_s \sum_k a_{ik}^{rs} F_{kj}^s + \sum_s \sum_k \sum_t \sum_m a_{ik}^{rs} a_{km}^{st} F_{mj}^t + ...}{Y_i^r}. \tag{4}$$

The first term in the numerator represents sales from sector r in country i directly consumed by j; the second term accounts for the fraction of sales of sector r in i sold to any producer in the world that is then sold to country j for consumption. The same logic applies to higher order terms. By definition $\sum_j \xi_{ij}^r = 1$. Note, importantly, that this aggregation of destination specific demand changes is the one implied by the Input-Output framework as formalized in the Remark 1.

Remark 1 (Model-Consistent Aggregation)
Suppose the economy is populated by agents with Cobb-Douglas preferences over varieties such that the expenditure share on variety i is β_i and by firms with constant returns to scale Cobb-Douglas production functions such that Input-Output linkages are summarized by an input requirement matrix $\mathcal{A}$. Then the growth rate of output of industry k is given by

$$\frac{\Delta Y_{it}^r}{Y_{it}^r} = \sum_j \xi_{ij}^r \eta_{jt}, \tag{5}$$

where $\xi_{ij}^r = \sum_k \sum_s \ell_{ik}^{rs} \beta_{kj}^s D_j / Y_i^r$ with ℓ_{ik}^{rs} element of the Leontief inverse $L = [I - \mathcal{A}]^{-1}$ and $\eta_{jt} = \frac{D_{jt} - D_{jt-1}}{D_{jt-1}}$ is the growth rate of total consumption expenditure in country j, D_{jt}.

Proof. See Appendix A.3. ■

I come back to this aggregation in Section 5 to estimate the model-implied relationships using the shift-share design.

3.2.4 Estimating Demand Shocks

To evaluate the total demand innovation that affects a specific industry one needs to estimate country-specific demand shocks. I do this via a fixed effects decomposition of final consumption.[8]

To build up intuition, define the value of output of industry r in country i that is consumed in country j at time t as F_{ijt}^r and denote f_{ijt}^r its natural logarithm. The simplest possible fixed effects model used to estimate demand innovations then takes the following form

$$\Delta f_{ijt}^r = \eta_{jt} + \nu_{ijt}^r. \tag{6}$$

Where ν_{ijt}^r is a normal distributed error term. The country and time-specific demand innovations would then be the series of $\hat{\eta}_{jt}$. This set of fixed effects extracts the change in consumption of destination market j at time t that is common to all sellers.

[8]A similar approach is used by Kramarz et al. (2020) and Alfaro et al. (2021).

A potential threat to identification through equation 6 would be if industry r is a sizeable fraction of j's consumption. In this case there would be reverse causality between industry and destination and one would not be able to claim exogeneity of η_{jt} to industry r in country i. Noting that in the WIOD sample the median domestic sales share is 67%, I partially account for this threat by estimating a different model for every industry r in producing country i, specifically

$$\Delta f^s_{kjt} = \eta_{jt}(i,r) + \nu^s_{kjt} \quad k \neq i, s \neq r. \tag{7}$$

For each industry r in country i, we need a shock that removes the possible reverse causality discussed above. I estimate country's j fixed effect using all industries of all countries except those of country i. This is tantamount to identifying the variation of interest through the trade of all other countries to the specific destination. To exemplify the idea behind the identification strategy, suppose that I observe US, Indian and Chinese producers of cars, textile and furniture. When estimating the change in the final demand faced by the US car manufacturing industry, I exclude the US as a production country. In principle this leaves me with identifying demand changes coming from sales of Indian and Chinese car, textiles and furniture producers to America, Indian and Chinese consumers. However if Chinese and American cars are very substitutable, then the observed sales of Chinese car might be related to supply shocks to American car manufacturers. To avoid this type of reverse causality I also restrict the analysis to sectors $s \neq r$, which in this example would be restricting to textile and furniture producers. In summary, when studying the observed change in final demand for US car manufacturer, I exploit variation coming from sales of Indian and Chinese textile and furniture manufacturers to US, Indian and Chinese consumers. This logic extends to the 56 sectors and 44 countries so that destination-time specific changes in final demand are estimated using (44-1)*(56-1) observations every year.[9] I provide robustness checks on this specification in Section B.9 in the Supplementary Material.

The estimated shifters can be aggregated as described above to create industry r in country i effective demand shocks at time t

$$\hat{\eta}^r_{it} = \sum_j \xi^r_{ij} \hat{\eta}_{jt}(i,r). \tag{8}$$

,Where the effective sales shares are evaluated at time $t = 0$ to eliminate the dependence of portfolio shares themselves on simultaneous demand innovations. This procedure implies that sales from i do not affect $\hat{\eta}_{jt}(i,r)$ and, therefore, $\hat{\eta}^r_{it}$.[10]

The identification of demand shocks relies on the rationale that the fixed effect model in equation 7 captures the variation common to all industries selling to a specific partner country in a given year. When producing industries are small relative to the destination, the estimated

[9]Further excluding all domestic flows does not change the results qualitatively or quantitatively. Formally it would imply additionally imposing $k \neq j$ in equation 7.

[10]In an alternative aggregation strategy in which I use time-varying sales shares, I follow Borusyak et al. (2022) and test pre-trends, namely that the sales shares are conditionally uncorrelated to the destination shifters. The results are reported in Online Appendix in Table B.6. I find no evidence of pre-trends.

demand shocks are exogenous to the producing industry, thereby providing the grounds for causal identification of their effects on the growth of sales.

Finally, note that, by construction, the aggregated demand shocks $\hat{\eta}^r_{it}$ already control for diversification potential. Suppose, for instance, that an industry delivered half its output to each of two countries, which had fully negatively correlated shocks. The industry would always have a realized $\hat{\eta}^r_{it} = 0$. Further, as noted in Figure B.6 in Section B.6 of the Online Appendix, the correlation between upstreamness and concentration of sales shares is negative, suggesting that the destinations of more upstream industries are more diversified.

4 Results

This section provides the results from the empirical analysis on how demand shocks propagate along the supply chain to industry output.

4.1 Demand Shock Amplification and Supply Chain Positioning

The goal of this section is to estimate if and how the output response to changes in final demand is heterogeneous depending on an industry's position in the supply chain. To this end I use the shift-share demand shocks aggregate according to eq. 8.In all the analyses in the remainder of this section I drop values of industry output growth rates larger than 69%, which is the 98^{th} percentile of the industry growth distribution. The results are consistent with different cuts of the data and without dropping any entry.

To first test that the estimated industry shocks are valid I run two preliminary regressions. First I check that output responds to these shocks with the expected sign and that they have explanatory power. To this end I regress the growth rate of industry output on the shocks and country-industry fixed effects. The estimated industry shocks generate a positive industry output growth response and they explain 43% of the variance, as shown in columns 1-2 of Table A.2. The estimation suggests that a 1 percentage point increase in the growth rate of final demand produces a .6pp increase in the growth rate of industry output. Secondly, as a further test for the shift-share demand shocks I check their effect on the industry deflator. Theory would suggest that positive demand shocks would increase the deflator.[11] Columns 3-4 in Table A.2 show that increases in the measured shock generate increases in the sectoral price index, suggesting that they are likely picking up shifts in final demand.

Having established that the estimated final demand shocks induce an increase in the growth rate of output, I study their heterogeneous effects depending on the industry's position in the supply chain. To do so, I estimate an econometric model in which the exogenous demand shocks can be considered a treatment and study the heterogeneity of the treatment effect along the upstreamness distribution.

[11] I compute the sectoral deflator by combining the baseline Input-Output data with the same dataset at previous year prices. I then compute the ratio of output in the two to obtain the change in the deflator.

Before discussing the results, it is important to note that given the shift-share structure which account for direct and indirect linkages, the demand shocks naturally account for diversification forces. As a consequence, as shown in Figure B.6 in the Online Appendix B.6.4, the standard deviation of measured demand shocks falls with upstreamness, suggesting that more upstream industries tend to be exposed to a lower demand volatility. The advantage of the empirical approach in this paper is that it allows me to isolate the differential output response fixing the size of demand shocks. To do so, I split the upstreamness distribution in bins through dummies taking values equal to 1 if $U^r_{it-1} \in [1,2]$ and $[2,3]$, and so on.[12] I use the lagged version of the upstreamness measure as the contemporaneous one might itself be affected by the shock and represent a bad control (see Angrist and Pischke, 2008). Formally, I estimate

$$\Delta \ln Y^r_{it} = \sum_j \beta_j \mathbb{1}\{U^r_{it-1} \in [j, j+1]\}\hat{\eta}^r_{it} + \nu^r_{it}, \qquad j = \{1, \ldots, 6\}. \tag{9}$$

The resulting coefficients are plotted in Figure 2, while the regression output is displayed in the first column of Table A.3 in the Appendix.

Figure 2: Effect of Demand Shocks on Output Growth by Upstreamness Level

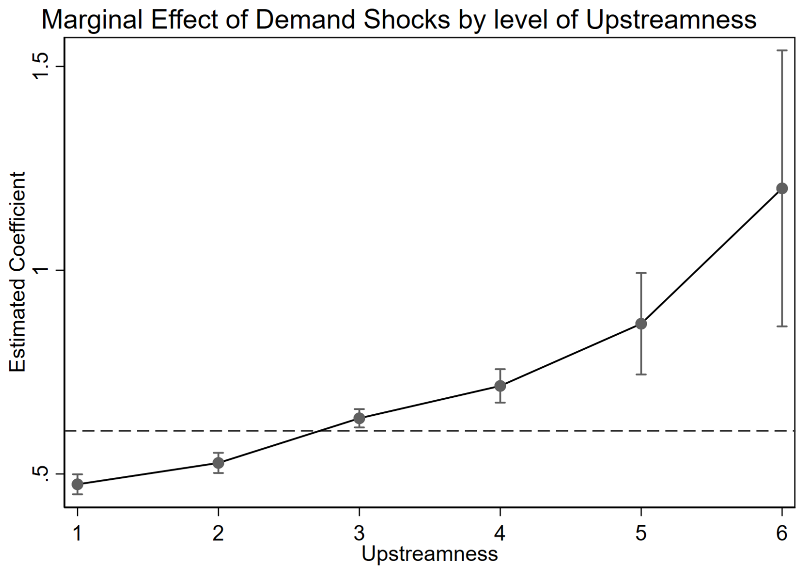

Note: The figure shows the marginal effect of demand shocks on industry output changes by industry upstreamness level. The dashed horizontal line represent the average coefficient as estimated in Table A.2. The vertical bands illustrate the 95% confidence intervals around the estimates. The regression includes country-industry fixed effects and the standard errors are cluster bootstrapped at the country-industry level. Note that due to relatively few observations above 7, all values above 7 have been included in the $U \in [6, 7]$ category. The full regression results are reported in the first column of Table A.3.

The results suggest that the same shock to the growth rate of final demand produces largely heterogeneous responses in the growth rate of industry output. In particular, industries between one and two steps removed from consumers respond approximately 60% less than industries six or more steps away. These results, which are robust across different fixed effects specifications,

[12]Since only 0.1% of the observations are above 7, I include them in the last bin, $\mathbb{1}\{U^r_{it-1} \in [6, \infty)\}$.

highlight how amplification along the production chain can generate sizable heterogeneity in output responses.This estimation also suggests that each additional unit of distance from consumption raises the responsiveness of industry output to demand shocks by approximately .09, which represents 14% of the average response.

It is important to note that these results are based on within producing-industry variation since I always include fixed effect at the country×industry level. This is noteworthy because it rules out alternative explanations such as that industries located more upstream tend to produce more durable goods and therefore be exposed to different intertemporal elasticity of substitution in household purchases. Secondly, given the construction of $\hat{\eta}^r_{it}$, these shocks include indirect exposure to other sectors' changes in demand. For example, the microprocessor industry is exposed to changes in the demand for both computers and financial services, insofar as computers are a key input in the financial services production function, where clearly computer and financial services have very different levels of durability.

A second important observation is that, as discussed above and shown in Figure B.6 in the Online Appendix, the volatility of the measured demand shocks $\hat{\eta}^r_{it}$ negatively correlates with upstreamness. This is intuitive as more upstream industries have a less concentrated sales distribution as shown in Figure B.5. Absent heterogeneous responses of output to changes in demand we should observe a declining volatility of sectoral output as move along the upstreamness distribution. As shown in the right panel of Figure B.6 this is counterfactual. The volatility of output growth positively correlates with upstreamness even if the effective volatility of demand declines.

In summary, more upstream industries face smaller fluctuations in their effective demand and yet have more volatility in their output growth. These differences are quantitatively important as the elasticity of output growth to changes in demand more than doubles along the upstreamness distribution. To assess the robustness of these results I run an extensive set of additional checks which I discuss in detail and report in the Online Appendix. These include using the re-centered instrument proposed by Borusyak and Hull (2020) to solve potential omitted variable bias; using an ordinal, rather than a cardinal, split of the upstreamness distribution; a more general model to estimate the demand shifters, allowing for supply side effects; using downstreamness as a potential source of heterogeneous effects; using deflated data to avoid confounding prices variation; using time-varying rather than time-invariant aggregation in the shift-share design; controlling for past output following Acemoglu et al. (2012).

4.2 Instrumenting Demand Shocks with Government Consumption

In the empirical specification in Section 4.1 I have used the variation arising from destination-time specific changes in foreign aggregate demand. To further alleviate concerns of endogeneity of this measure, in this section I use foreign government consumption as an instrument.

More specifically, the WIOD data contains information about the value of purchases of the government of country j of goods of industry r from country i in period t. Denote this G^r_{ijt}. I apply the same steps as in Section 3.2 replacing the consumers' purchases with the

the government ones. This procedure allows me to create a measure $\hat{\eta}_{it}^{rG} = \sum_j \xi_{ij}^r \hat{\eta}_{jt}^G$. With the understanding that, as in the main results, the estimated destination-time shifter $\hat{\eta}_{jt}^G$ is calculated by excluding all purchases of goods from country i or industry r.

I use this instrument in a control function approach. Formally, I estimate $\eta_{it}^r = \beta \hat{\eta}_{it}^{rG} + \epsilon_{it}^r$ and predict the residual $\hat{\epsilon}_{it}^r$. I then estimate equation 9 including the interactions with the estimated residuals.

$$\Delta \ln(Y_{it}^r) = \sum_j \beta_j \mathbb{1}\{U_{it-1}^r \in [j, j+1]\} \hat{\eta}_{it}^r + \gamma_j \mathbb{1}\{U_{it-1}^r \in [j, j+1]\} \hat{\epsilon}_{it}^r + \nu_{it}^r, \qquad j = \{1, \dots, 6\}.$$

As a further check, I allow for heterogeneity in the relationship between the consumer and government shocks. In particular, I estimate separately for each upstreamness bin $\hat{\varepsilon}_{it}^{rn}$ as the residual of the regression $\eta_{it}^r = \beta^n \hat{\eta}_{it}^{rG} + \varepsilon_{it}^{rn}$ if $U_{it-1}^r \in [n, n+1]$. I then estimate again equation 9, controlling for all $\hat{\varepsilon}_{it}^{rn}$. Figure 3 shows the estimated coefficient of interest in the two models.

Figure 3: Effect of Demand Shocks on Output Growth by Upstreamness Level - Government Consumption

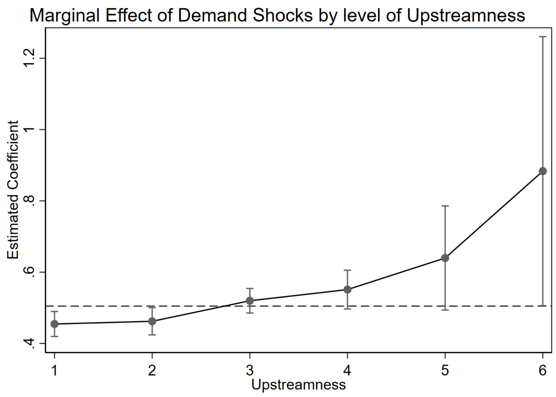

(a) Single First Stage

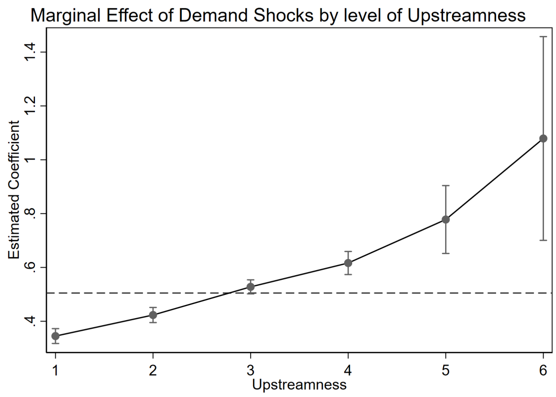

(b) Heterogeneous First Stage

Note: The figure shows the marginal effect of demand shocks on industry output changes by industry upstreamness level in the control function models. Panel (a) shows the single first stage result while Panel (b) shows the results when the first stage is heterogeneous by upstreamness bin. The vertical bands illustrate the 95% confidence intervals around the estimates. The regression includes country-industry fixed effects and the standard errors are cluster bootstrapped at the country-industry level. The dotted horizontal line represents the average coefficient. Note that due to relatively few observations above 6, all values above it have been included in the $U \in [6, 7]$ category. The regression results are reported in Table A.4.

Both approaches confirm the main results qualitatively and quantitatively. Allowing for heterogeneous dependence of demand shocks to government purchase shocks delivers estimates which are almost identical to main result in their shape.

4.3 China-Shock and Upstream Amplification

The two approaches used so far have the same structure of final demand shocks, generated through the I-O exposure and destination-specific aggregate changes. One might worry that the finding described so far could be due to using potentially related statistics of the network (upstreamness and exposure). To circumvent this possibility I use the China shock in Acemoglu et al. (2016) as a proxy for changes in US demand. The idea is that some US sectors suffered a drop in their domestic demand as Chinese import competition intensified. Note that this is an imperfect proxy for final demand as it confounds both demand from consumers and other firms.

Following Autor et al. (2013) I instrument the change in export from China to the US with the change from China to 8 other large economies. As I am using only US data I can employ the NBER CES Manufacturing dataset, in which I observe directly sales, inventories and valued added, which allows me to build output as sales plus the change in inventories. The merged sample consists of 312 NAICS industry for 21 years from 1991 to 2011. I estimate the regression in eq. 9 with two key differences. First, I estimate it with the growth rate of both value added (as in Acemoglu et al. (2016)) and output. Second, the US data has less than 1% of industries with upstreamness above 4, hence I aggregate all industries with upstreamness above 3 in the last group so that β_3 is for all industries r such that $U^r \in (3, \infty)$. Lastly, I apply the network transformation to the China shock to account for indirect exposure. As in the previous section, I use the control function approach to instrument the endogenous shock. The main results are plotted in Figure 4 and reported in Table A.5.

Figure 4: Effect of China Shock on Output Growth by Upstreamness Level

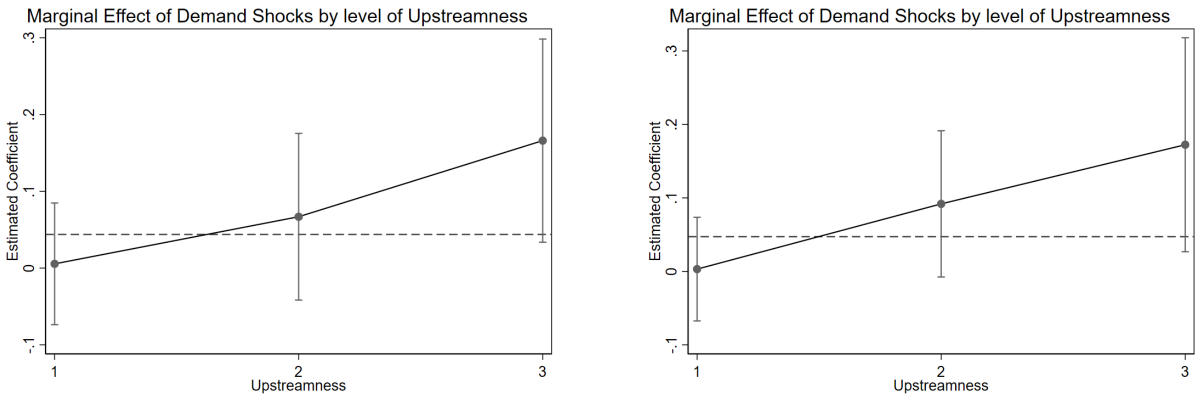

(a) Output Growth

(b) Value Added Growth

Note: The figure shows the marginal effect of the China shock on industry output changes by industry upstreamness level by the control function models. I apply the network transformation of the original shock to account for indirect exposure. Following Autor et al. (2013) I instrument the change in US imports from China with the change in other advanced economies. The vertical bands illustrate the 95% confidence intervals around the estimates. The regression includes country-industry fixed effects and the standard errors are cluster bootstrapped at the country-industry level. The dotted horizontal line represents the average coefficient. Note that due to relatively few observations above 4, all values above it have been included in the $U \in (3, 4)$ category. These estimates report the output and value added growth changes in response to a 1 standard deviation change in the networked China shock. The regression results are reported in Table A.5.

As in the previous analysis I find a positive gradient in the output growth response to the shock, meaning that output elasticities are increasing in upstreamness. These estimates are much noisier than the ones on the WIOD data but the slope is still statistically significant. Quantitatively, for a given 1 standard deviation increase in the shock, increasing upstreamness by 1 implies a 4.7% increase in the growth rate of output and 4.4% increase in the growth rate of value added. These magnitudes are comparable to the ones obtained by Acemoglu et al. (2016). They report the direct exposure to 1 std. dev. increase of the shock to generate a 3.4% drop in value added. I find that for the 3 upstreamness bins the same figures are 0.5%, 6.7% and 16.6%. These findings suggest that, underlying the large indirect effect, is a strong heterogeneity in the output elasticities, depending on industries' positions in the supply chain.[13]

4.4 Inventories Amplification

These reduced form results paint a consistent picture in which firms located further away from consumer respond significantly more to the same changes in demand. The operations literature on the *bullwhip effect* has suggested that procyclical inventory adjustment can lead to upstream amplification of shocks. To test this mechanism more directly I estimate the same regressions using the change of inventories as the dependent variable.

In particular, WIOD provides information on the change in inventories of a producing industry computed as the row residual in the I-O matrix. Intuitively, given the accounting identity that output equals sales plus change in inventory stock, the tables provide the net change in inventories as residual between output and sales to other industries or final good consumers. I standardize the change in inventories by dividing it by total output so that scale effects are at least partially accounted for. I then estimate equation 9 with $\Delta I^r_{it}/Y^r_{it}$ on the left hand side. I drop observations for which $\Delta I^r_{it}/Y^r_{it}$ is below -1 and above 3. The results are plotted in Figure 5 and reported in Table A.6.

Figure 5: Effect of Demand Shocks on Inventory Changes by Upstreamness Level

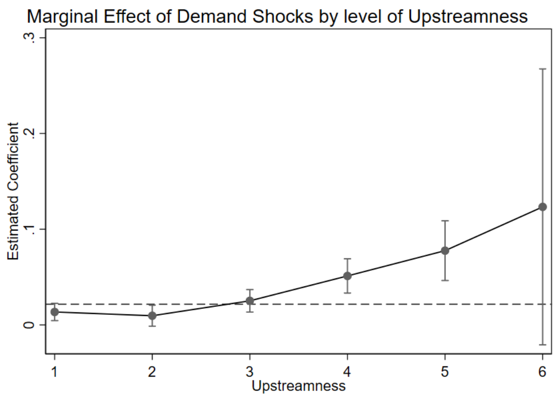

(a) Final Demand Shocks

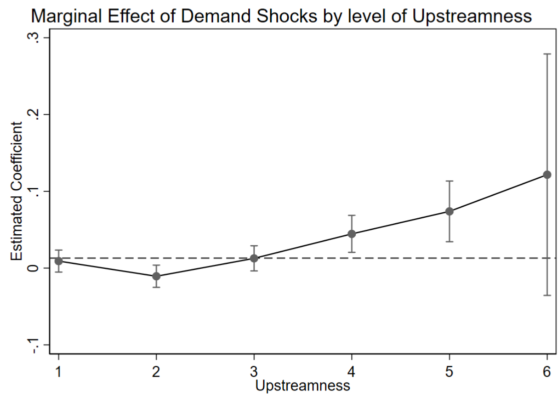

(b) Government Consumption Instrument

Note: The figure shows the marginal effect of demand shocks on industry inventory changes by industry upstreamness level. The left panel shows the estimation using the demand shocks described in section 3 while the right panel uses government consumption as an instrument. The dashed horizontal line represent the average coefficient. The vertical bands illustrate the 95% confidence intervals around the estimates. The regression includes country-industry fixed effects and the standard errors are cluster bootstrapped at the country-industry level. Note that due to relatively few observations above 7, all values above 7 have been included in the $U \in [6, 7]$ category. The full regression results are reported in the first column of Table A.6.

The left panel shows the baseline estimation with demand shocks, while the right panel

[13]In Figure B.3 in Online Appendix B.4 I report the results of the same procedure using Acemoglu et al. (2016) federal spending shocks as the metric of changes in demand. This exercise confirms a positive gradient over upstreamness in the response of output to changes in demand.

plots the results for the instrumented shocks using government consumption. The estimation suggests that the response of inventories, consistently with the one of output, increases along the upstreamness distribution. Quantitatively a 1pp increase in the growth of demand generates .02pp increase in the change in inventories over output for industries closest to final consumers. The same figure for industries at 6 or more steps of production away is .12pp. increase in the rescaled change in inventories. The baseline estimation results are confirmed when using government consumption as an instrument for final demand.

These results suggest that inventories can act as a force of upstream amplification in network economies. Following the insights of the operations literature on the *bullwhip effect*, in the next section I lay out a simple extension to the workhorse network model to include inventories. The goal is to characterize under which conditions on inventories and network structure we would observe upstream amplification or dissipation.

5 Conceptual Framework

I start by building an extended example of demand shock propagation in vertically integrated economies with inventories solely based on accounting identities. Secondly, I let firms optimally choose inventories and map their policy into sufficient conditions for amplification. I then combine the framework with the standard model of production network (see Acemoglu et al., 2012; Carvalho and Tahbaz-Salehi, 2019) to evaluate the conditions under which amplification or dissipation is observed, depending on the features of the network and the inventory response.

5.1 Vertically Integrated Economy

Consider an economy with one final good whose demand is stochastic, and $N-1$ stages sequentially used to produce the final good. Throughout I use industry, sector, and firm interchangeably. The structure of this production network is a line, where stage N provides inputs to stage $N-1$ and so on until stage 0, where goods are consumed.

The demand for each stage n in period t is D_t^n with $n \in \mathcal{N}$. Stage 0 demand, the final consumption stage, is stochastic and follows an AR(1) with persistence $\rho \in (-1, 1)$ and a positive drift $\bar{D}$. The error terms are distributed according to some finite variance distribution F on a bounded support. $\bar{D}$ is assumed to be large enough relative to the variance of the error so that demand is never negative.[14] Formally, final demand in period t is

$$D_t^0 = (1-\rho)\bar{D} + \rho D_{t-1}^0 + \epsilon_t,\ \epsilon_t \sim F(0, \sigma^2).$$

The production function is linear: for any stage n, if production is Y_t^n, it also represents the demand for stage $n+1$, D_t^{n+1}. This implies $Y_t^n = D_t^{n+1}$.

[14] Including the positive drift does not change the inventory problem since, for storage, the relevant statistic is the first differenced demand.

Stage 0 production is the sum of the final good demand and the change in inventories. Inventories at time t for stage n are denoted by I_t^n.

Firms at stage n form expectations on future demand $\mathbb{E}_t D_{t+1}^n$ and produce to end the period with target inventories $I_t^n = I(\mathbb{E}_t D_{t+1}^n)$. Where $I(\cdot)$ is some non-negative differentiable function that maps expectations on future demand into end-of-period inventories.

5.1.1 An Accounting Framework

Given this setup it is possible to derive how output behaves at every step of production n by solving the economy upward from final demand. At stage n output is given by

$$Y_t^n = D_t^n + I(\mathbb{E}_t D_{t+1}^n) - I(\mathbb{E}_{t-1} D_t^n), \tag{10}$$

Where D_t^n is the demand for sector n's products. By market clearing this is also total production of sector $n-1$. In the context of this model, asking whether exogenous changes in final demand amplify upstream is effectively comparing $\frac{\partial Y_t^n}{\partial D_t^0}$ and $\frac{\partial Y_t^{n+1}}{\partial D_t^0}$. In particular, amplification occurs if $\frac{\partial Y_t^n}{\partial D_t^0} < \frac{\partial Y_t^{n+1}}{\partial D_t^0}$. Proposition 1 formalizes the sufficient condition for amplification in this economy.[15]

Proposition 1 (Amplification in Vertically Integrated Economies)
A vertically integrated economy with inventories features upstream amplification of positively autocorrelated final demand shocks if and only if the inventory function satisfies

$$0 < I' < \frac{1}{1-\rho}.$$

Proof. See Appendix A.3. ■

The first inequality requires that the inventory function is increasing. This ensures that, as demand rises, so do inventories. If inventories increase when demand rises then output increases more than one-to-one with demand. This, in turn implies that the demand change faced by the upstream firm is amplified relative to the one faced by the downstream firm. In other words, the demand shock amplifies upstream. The second inequality requires that the function is not "too increasing" relative to the persistence of the process. The second inequality arises because a positive change of demand today implies that the conditional expectation of demand tomorrow is lower than demand today, due to mean reversion. This condition ensures that the first effect dominates the second one. Intuitively, as shocks become arbitrarily close to permanent, the second condition is trivially satisfied, and it is enough for inventories to be increasing in expected demand.

An alternative way of summarizing the intuition is the following: in vertically integrated economies without labor and inventories, changes in final demand are transmitted one-to-one upstream, as no substitution is allowed across varieties. When such an economy features inventories, this result need not hold. If inventories are used to smooth production, meaning

[15] I generalize this result to the case in which firms have heterogeneous inventories in Proposition B.1 in section B.5.1 of the Online Appendix.

that $I(\cdot)$ is a decreasing function, shocks can be transmitted less than one-to-one as inventories partially absorb them. On the other hand when inventories are adjusted procyclically the economy features upstream amplification.

As discussed in Section 2 the estimated average derivative I' is approximately .1. Given an empirical estimate of the autocorrelation of HP-filtered sales at around .7, the data suggests that the condition in Proposition 1 condition is empirically verified.[16] In the next section I specify the problem of a firm to obtain an inventory policy which yields closed form solutions for output in the vertical network economy.

5.1.2 Endogenous Inventories

Suppose firms at a generic stage n have the following objective function:

$$\begin{aligned}&\mathbb{E}_t \sum_t \beta^t \left[D_t^n - c^n Y_t^n - \frac{\delta}{2}(I_t^n - \alpha D_{t+1}^n)^2\right] \quad st \\ &I_t^n = I_{t-1}^n + Y_t^n - D_t^n,\end{aligned} \tag{11}$$

where c^n is the marginal cost of production which, in equilibrium, is given by the price of goods at stage $n-1$ and $\delta, \alpha > 0$ govern the costs of holding inventories or facing stockouts and backlogs.[17] The optimal inventory policy is given by an affine function of the expected demand:

$$I_t^n = \max\{\mathcal{I}^n + \alpha \mathbb{E}_t D_{t+1}^n, 0\}, \tag{12}$$

with $\mathcal{I}^n := (\beta - 1)c^n/\delta$. [18] This optimal rule predicts that inventories are procyclically adjusted as is corroborated by the inventory data (see Figure A.3). This formulation of the problem, where the presence of inventories is motivated directly by the structure of the firm's payoff function, is a reduced form stand-in for stock-out avoidance motives. In Appendix B.5.2 I provide a simple dynamic model in which firms face stochastic production breakdowns and stochastic demand to show that the optimal dynamic policy implies procyclical inventory changes. Secondly, note that, in this setup, procylicality follows from the optimal target-rule adopted by firms. An alternative motive for holding inventories could be production smoothing, whereby a firm holds a stock of goods to avoid swings in the value of production between periods. In appendix B.5.3 I introduce a production smoothing motive and show that the firm optimally chooses procyclical inventories if the smoothing motive is not too strong. Furthermore, if the production smoothing motive were to dominate inventories would have to be countercyclical which is counterfactual based on the evidence in Section 2. The present formulation of the problem has two great advantages. First, the linear affine mapping between inventory holdings and future sales (as given by a

[16]Table A.1 provides the estimates of $I'(\cdot)$. Using the empirical version of the condition in Proposition 1, I find that the $I'(\cdot) \approx .1$, which satisfies the condition $\forall \rho > 0$.

[17]This model of inventory choice is a simplified version of the linear-quadratic inventory model proposed by Ramey and West (1999) following Holt et al. (1960) as a second order approximation of the full inventory problem. I discuss the more general version in Appendix B.5.2.

[18]Note that $\mathcal{I}^n < 0$ since, in the presence of time discounting or depreciation of inventories, the firm would ideally like to borrow output from the future and realize the sales today.

constant target rule) matches the data extremely well, as shown in Figure B.8 in the Online Appendix. Secondly, in this economy output has the following closed-form solution:

Lemma 1 (Industry Output in Vertical Economies)
In a vertical economy where the optimal inventory rule is given by 12, industry output for a generic sector at distance n from final consumption is

$$Y_t^n = D_t^0 + \alpha\rho \sum_{i=0}^{n} (1 + \alpha(\rho - 1))^i \Delta_t^0. \tag{13}$$

Where $\Delta_t^0 = D_t^0 - D_{t-1}^0$.

Proof. See Appendix A.3. ■

Using the insights of Proposition 1, note that $I' = \alpha > 0$ trivially satisfies the first inequality, while the second one is satisfied if $\alpha < 1/(1 - \rho)$.[19] I henceforth assume that these conditions are verified so that $1 + \alpha(\rho - 1) \in (0, 1)$. The reason for this assumption is twofold: first, it naturally follows from the empirical range of estimates of α and $\rho > 0$; second, assuming it is bounded above by 1 ensures that, if chains become infinitely long, the economy still features finite GDP.[20]

Using Lemma 1, I can characterize the responsiveness of industry n output to a change in final demand in the following proposition

Proposition 2 (Amplification in Vertically Integrated Economies)
The output response of a firm at stage n to a change in final demand is given by

$$\frac{\partial Y_t^n}{\partial D_t^0} = 1 + \alpha\rho \sum_{i=0}^{n} (1 + \alpha(\rho - 1))^i. \tag{14}$$

Furthermore, the economy features upstream amplification if

$$\frac{\partial^2 Y_t^n}{\partial D_t^0 \partial n} = \alpha\rho(1 + \alpha(\rho - 1))^n > 0. \tag{15}$$

Which is verified given the assumption $0 < \alpha < 1/(1 - \rho)$.

Proof. See Appendix A.3. ■

This result states that any shock to final demand traveling upstream gets magnified if $1 + \alpha(\rho - 1) > 0$, as assumed above. The operations literature labels this result the *bullwhip effect* or *Forrester effect* (see Forrester, 1961).

[19] Using the NBER CES Manufacturing Industries Database for the years 2000-2011, I find that these conditions are typically satisfied in the data as the yearly values of α range between 0 and 50% of next year sales, with an average of 12%.

[20] This assumption is not needed in the context of a vertically integrated production economy because, if the number of sectors is finite, so is the length of chains as shown in Appendix B.5.4. Assuming $1 + \alpha(\rho - 1) \in [0, 1]$ is useful in the context of a general network in which the presence of cycles may generate infinitely long chains.

Note that, from Lemma 1, it is apparent that if shocks are i.i.d. no amplification occurs as the second term is zero. Importantly, note that in this setting, due to production taking place on a line with only one endpoint, the structure of the network plays no role in determining the degree of amplification.

In the next section, I extend the model by including labor and allowing for a more general production structure, such that the network itself shapes the degree of propagation of demand shocks.

5.2 Network Structure and Amplification

In this section I extend the model to study how the structure of the production network interplays with the inventory amplification mechanism.

Suppose the economy is populated by domestic and foreign consumers. Domestic consumers have preferences over a homogeneous consumption good c_0, a differentiated bundle C and inelastically supply labor $\bar{l}$. The utility is given by

$$U = c_0 + \ln C. \tag{16}$$

They maximize utility subject to the budget constraint $\bar{l} = PC + c_0 - T$, where T is lump sum taxes that government uses to subsidize firms. The wage is the numeraire and the homogeneous good is produced linearly from labor, so that $w = p_0 = 1$. The household maximization yields a constant expenditure on the differentiated bundle equal to 1. Foreign consumers have a stochastic demand X which follows an $AR(1)$ process with some mean $\bar{X}$. Total demand faced by a firm is then given by

$$D_t = (1-\rho)(1+\bar{X}) + \rho D_{t-1} + \epsilon_t,\ \epsilon_t \sim F(0, \sigma^2).$$

I assume that the composition of the domestic and foreign consumption baskets is identical and generated through a Cobb-Douglas aggregator over varieties

$$C = \prod_{s \in S} C_s^{\beta_s},$$

where S is a finite number of available products, β_s the consumption weight of good s and $\sum_s \beta_s = 1$. This formulation implies that the expenditure on good s is $E_{s,t} = \beta_s D_t$ for $E_{s,t}$ solving the consumer expenditure minimization problem.

The network is characterized by an input requirement matrix $\mathcal{A}$, in which cycles and self-loops are possible.[21] I denote elements of $\mathcal{A}$ as $a_{rs} = [\mathcal{A}]_{rs}$. The network has a terminal node given by final consumption.

Firms produce using labor and a bundle of other sectors' output. This is generated through

[21] An example of a cycle is: if tires are used to produce trucks and trucks are used to produce tires. Formally, $\exists r : [\mathcal{A}^n]_{rr} > 0, n > 1$. An example of a self-loop is: if trucks are used in the production of trucks. Technically, such is the case if some diagonal elements of the input requirement matrix are positive, i.e. $\exists r : [\mathcal{A}]_{rr} > 0$.

Cobb-Douglas production functions

$$Q_{s,t} = Z_s l_{s,t}^{1-\gamma_s} M_{s,t}^{\gamma_s},$$

where l_s is the labor used by industry s, M_s is the input bundle and γ_s is the input share for sector s. Finally, Z_s is an industry specific normalization constant. The input bundle is aggregated as

$$M_{s,t} = \left(\sum_{r \in R} {a_{rs}}^{1/\nu} Q_{rs,t}^{\frac{\nu-1}{\nu}} \right)^{\frac{\gamma_s \nu}{\nu-1}},$$

where Q_s is the output of sector s, Q_{rs} is the output of industry r used in sector s production. ν is the elasticity of substitution and a_{rs} is an input requirement, in equilibrium this will also coincide with the expenditure amount Q_{rs} needed for every dollar of Q_s. R is the set of industries potentially supplying inputs to sector s. The aggregator function is assumed to have constant returns to scale

I maintain that competitive firms in each sector solve the problem in 11, where the marginal cost of production is given by the expenditure minimizing bundle of labor and inputs, and therefore follow the optimal policy in 12.[22]

Definition 1 (Equilibrium)
An equilibrium in this economy is given by a set of inventory and output policies for firms, the consumption policy of households and market clearing conditions for the homogeneous good, all differentiated varieties and labor.

I assume that the government raises lump sum taxes and subsidizes firms' production to cover the cost of inventories. These assumptions are extremely convenient to simplify the firms' problem and allow an analytical characterization of the equilibrium. In particular, they imply that firms price at the inventory-less marginal cost. As a direct consequence of this structure and the normalization $w = p_0 = 1$, the solution to the all firms pricing problem is $p_s = 1$, $\forall s$.[23] Effectively this is equivalent to households using labor in the production of the homogeneous good to pay for any deviation from marginal cost pricing of the differentiated good firms. It follows that there is no difference between value and quantity of output in this economy. I come back to these assumptions in the discussion part of this section. Output of final goods producers, denoted by the superscript 0, is

$$Y_{s,t}^0 = \beta_s [D_t + \alpha \rho \Delta_t].$$

[22]In particular the marginal cost of a firm in sector s is given by $c_s = Z_s^{-1}(1-\gamma_s)^{\gamma_s - 1} w^{1-\gamma_s} \gamma_s^{-\gamma_s} \left(\sum_r a_{rs} p_r^{1-\nu}\right)^{\frac{\gamma_s}{1-\nu}}$. The normalizing constant is then $Z_s := (1-\gamma_s)^{\gamma_s - 1} \gamma_s^{-\gamma_s \frac{\nu}{\nu-1}}$. See Carvalho and Tahbaz-Salehi (2019) for a similar treatment.

[23]To build such equilibrium one can conjecture the vector of prices $p_s = 1 \, \forall s$ and note that it is an equilibrium of the inventory-less economy since there are no productivity shocks. From there one can design government subsidies such that this equilibrium is enforced in the economy with inventories. This boils down to covering the cost of inventories for each firm so that a free entry condition holds at these conjectured prices.

This also represents the input expenditure of sector s to its generic supplier r, once it is rescaled by the input requirement $\gamma_s a_{rs}$. Hence output of producers one step of production removed from consumption obtains by summing over all final good producers s. Market clearing and the inventory policy imply $Y^1_{r,t} = \sum_s \gamma_s a_{rs} Y^0_{s,t} + \Delta I^1_{r,t}$ and therefore

$$Y^1_{r,t} = \sum_s \gamma_s a_{rs} \left[D_t + \alpha\rho \sum_{i=0}^{1} (1 + \alpha(\rho - 1))^i \Delta_t \right].$$

Denote $\gamma_s a_{rs} = \tilde{\mathcal{A}}_{rs}$, so that $\sum_s \gamma_s a_{rs} = \sum_s \tilde{\mathcal{A}}_{rs}$ is the weighted outdegree of a node r, namely the sum of the shares of expenditure of all industries s coming from input r. Iterating forward to generic stage n, and defining for industry k

$$\chi^n_k := \underbrace{\sum_v \tilde{\mathcal{A}}_{kv} \sum_q \tilde{\mathcal{A}}_{vq} \ldots \sum_r \tilde{\mathcal{A}}_{or} \sum_s \tilde{\mathcal{A}}_{rs}}_{n \text{ sums}} \beta_s = \tilde{\mathcal{A}}^n B_k,$$

we can write the value of production at stage n as

$$Y^n_{k,t} = \chi^n_k \left[D_t + \alpha\rho \sum_{i=0}^{n} (1 + \alpha(\rho - 1))^i \Delta_t \right]. \tag{17}$$

In equation 17 the structure of the network is summarized by χ^n_k, while the rest of the equation represents the inventory effect both directly and indirectly through network connections. In this setup, the effect of a change in contemporaneous demand on the value of production is

$$\frac{\partial Y^n_{k,t}}{\partial D_t} = \chi^n_k \left[1 + \alpha\rho \sum_{i=0}^{n} (1 + \alpha(\rho - 1))^i \right]. \tag{18}$$

Where the first term summarizes the network effect and the second term represents the inventory amplification. Observe that equation 18 is a generalization of equation 14, as it accounts for the network structure.

Finally, as firms operate at multiple stages of production, total output is $Y_{k,t} = \sum_{n=0}^{\infty} Y^n_{k,t}$. I can now characterize the value of sectoral output as a function of the inventory channel and the features of the network.

Lemma 2 (Sectoral Output)
The value of sectoral output for a generic industry k is given by

$$Y_{k,t} = \sum_{n=0}^{\infty} \chi^n_k \left[D_t + \alpha\rho \sum_{i=0}^{n} (1 + \alpha(\rho - 1))^i \Delta_t \right]. \tag{19}$$

This can be written in matrix form as

$$Y_{k,t} = \tilde{L}B_k D_t + \alpha\rho \left[\sum_{n=0}^{\infty} \tilde{\mathcal{A}}^n \sum_{i=0}^{n} (1+\alpha(\rho-1))^i\right]_k B\Delta_t, \tag{20}$$

where B is the $S \times 1$ vector of demand shares and $\tilde{L}_k$ is the k^{th} row of the Leontief inverse, defined as

$$\tilde{L} = [I + \tilde{\mathcal{A}} + \tilde{\mathcal{A}}^2 + ...] = [I - \tilde{\mathcal{A}}]^{-1}.$$

Where $\tilde{\mathcal{A}} \coloneqq \mathcal{A}\hat{\Gamma}$ and $\hat{\Gamma} = \mathrm{diag}\{\gamma_1, ..., \gamma_R\}$.
Sectoral output exists non-negative for any α, ρ such that $\alpha(\rho-1) \in [-1, 0]$.

Proof. See Appendix A.3. ■

A number of features of Lemma 2 are worth discussing. The first observation is that the model collapses to the standard characterization of output in production networks when there is no inventory adjustment, as the second term in equation 20 vanishes to recover $Y_{k,t} = \tilde{L}_k BD_t$. This occurs whenever there are no inventories ($\alpha = 0$) or when current shocks do not change expectations on future demand ($\rho = 0$). A second implication is that output might diverge as $n \to \infty$ if $\alpha(\rho-1) > 0$. Lastly, by the assumptions made on $\tilde{\mathcal{A}}$,[24] and the maintained assumption that $1+\alpha(\rho-1) \in [0, 1]$, additional distance from consumption implies ever decreasing additional output, so output converges.[25]

With Lemma 2 we are in the position to characterize the object of interest, namely how the change in output as a response to a change in final demand moves with a firm's position. Formally, that would require the characterization $\frac{\partial Y_{kt}}{\partial D_t \partial n_k}$, where n_k is a measure of distance from consumption for industry k, which would extend the result in Proposition 2 to a general network setting. Unfortunately this comparative statics is ill-defined in a general network as there is no such thing as n_k. For example, firms can be simultaneously at distance 1 and 5 from final consumers. To overcome this issue I proceed in two steps: first, I show that the natural candidate to calculate a firm's distance from final consumption is upstreamness; second, I engineer two simple comparative statics on primitives, designed to induce a marginal change in upstreamness. Remark 2 formalizes the first step.

Remark 2 (Upstreamness)
In a general production network characterized by the Leontief inverse discussed above, with $\alpha(\rho-1) \in [-1, 0]$, distance from consumption for some industry k is

$$U_k = \sum_{n=0}^{\infty} (n+1) \frac{Y_k^n}{Y_k}, \quad U_k \in [1, \infty). \tag{21}$$

[24] In particular the fact that $\sum_k \tilde{\mathcal{A}}^{kv} < 1$, i.e. the assumption that the firm labor share is positive.

[25] In Appendix B.5.4 I show that restricting the network to a Directed Acyclic Graph allows existence and non-negativity even if $\tilde{\mathcal{A}}^n \sum_{i=0}^{n} (1+\alpha(\rho-1))^i$ has a spectral radius outside the unit circle which is the sufficient condition used in Lemma 2.

Proof. See Appendix A.3. ■

The goal is characterize as closely as possible $\frac{\partial Y_{kt}}{\partial D_t \partial U_k}$ as this is what I measured in Section 4. However, U_k is measurement device rather than a primitive, so such comparative statics is poorly defined. In Proposition 3 I provide two well-defined comparative statics that generate a marginal increase in a firm's upstreamness. Both in the model and in the data, the position of an industry in the supply chain is determined by the composition of demand, governed by the vector of expenditure shares B, and by the Input-Output matrix defined by $\mathcal{A}$. Therefore the first comparative statics, denoted Δ_β, changes consumers' expenditure shares at the margin so that some firm k goes from U_k to $U_k + \epsilon$. The second one, denoted $\Delta_{\tilde{L}}$, changes the network structure by altering elements of the Input-Output matrix $\mathcal{A}$. An example of the latter thought experiment, which I discuss more in detail later, is taking all possible network paths linking a firm to consumers and just adding a step of production.

Proposition 3 (Comparative Statics)
This proposition formalizes the comparative statics on the responsiveness of output to final demand shocks. For ease of notation, denote $\omega = 1 + \alpha(\rho - 1)$.

a) *The effect of change in aggregate demand on sectoral production are given by*

$$\frac{\partial Y_{k,t}}{\partial D_t} = \tilde{L}B_k + \alpha\rho \sum_{n=0}^{\infty} \tilde{\mathcal{A}}_k^n \sum_{i=0}^{n} \omega^i B. \tag{22}$$

b) *Furthermore, a change in the composition of demand, defined as a marginal increase in the s^{th} element of the vector B (β_s), paired with a marginal decrease of the r^{th} element (β_r), changes output response to aggregate demand as follows:*

$$\Delta_\beta \frac{\partial Y_{k,t}}{\partial D_t} := \frac{\partial}{\partial \beta_s}\frac{\partial Y_{k,t}}{\partial D_t} - \frac{\partial}{\partial \beta_r}\frac{\partial Y_{k,t}}{\partial D_t} = \sum_{n=0}^{\infty}\left[\tilde{\mathcal{A}}_{ks}^n - \tilde{\mathcal{A}}_{kr}^n\right]\left[1 + \alpha\rho\sum_{i=0}^{n}\omega^i\right], \tag{23}$$

where $\tilde{\mathcal{A}}_{ks}^n, \tilde{\mathcal{A}}_{kr}^n$ are the elements of $\tilde{\mathcal{A}}$ in positions (k, s) and (k, r) respectively.

c) *Finally, a change of the structure of the network path from industry k to final consumption, denoted by a new I-O matrix $\tilde{\mathcal{A}}'$, implies a change in the responsiveness of production to aggregate demand given by*

$$\Delta_{\tilde{L}} \frac{\partial Y_{k,t}}{\partial D_t} = \sum_{n=0}^{\infty}\left[\tilde{\mathcal{A}'}_k^n - \tilde{\mathcal{A}}_k^n\right]\left[1 + \alpha\rho\sum_{i=0}^{n}\omega^i\right] B. \tag{24}$$

Proof. See Appendix A.3. ■

The first result in Proposition 3 shows that the effect of a change in final demand on sector output can be decomposed in two distinct terms in eq. 22. The first one, the standard term in production network economies, states that the change in output is a function of the structure of the network and, in particular, of the centrality of the sector. The second term states that an additional response is driven by the behavior of inventories. The more important inventories are in the economy and the more autocorrelated demand shocks are, the larger the additional effect of changes in demand on output.[26]

The second half of Proposition 3 characterizes how output responds differentially when we engineer changes in the position of firms. This is done through changes in the composition of demand in point 3b and through changes in the I-O matrix in point 3c. At this level of generality the model can feature both amplification or dissipation upstream of shocks. Which one prevails depends on the comparison between the network positions, as summarized by $\mathcal{A}^n_{ks} - \tilde{\mathcal{A}}^n_{kr}$ and the intensity of the inventory effect in $\alpha\rho\sum_{i=0}^{n}\omega^i$. For the purpose of sharpening the intuition and to come as close as possible to the ideal characterization of $\frac{\partial Y_{kt}}{\partial D_t \partial U_k}$, consider the following special case of the comparative statics in Proposition 3.

Example Suppose sector k has a vector of connections $\mathcal{A}_k$. Suppose further that all chains from k to consumers are increased by one link so that the upstreamness of sector k moves from U_k to $U'_k = U_k + 1$. As a practical example, suppose that the connection from tires to consumption used to be *tires* $\rightarrow$ *cars* $\rightarrow$ *consumption* and is now *tires* $\rightarrow$ *wheels* $\rightarrow$ *cars* $\rightarrow$ *consumption*. Applying Proposition 3c

$$\sum_{n=0}^{\infty}\left[\tilde{A'}^n_k - \tilde{\mathcal{A}}^n_k\right]\left[1+\alpha\rho\sum_{i=0}^{n}\omega^i\right]B = \sum_{n=0}^{\infty}\tilde{\mathcal{A}}^n_k\left[\left(1+\alpha\rho\sum_{i=0}^{n}\omega^i\right)(\tilde{\mathcal{A}}\mathbf{1}_k - 1) + \alpha\rho\omega^{n+1}\tilde{\mathcal{A}}\mathbf{1}_k\right]B.$$

This result leads to the following condition, determining whether, after having increased the distance from consumption of sector k, demand shocks will be further amplified or further dissipated:

$$\operatorname{sgn}\Delta_{\tilde{L}}\frac{\partial Y_{k,t}}{\partial D_t} = \operatorname{sgn}\left[\left(1+\alpha\rho\sum_{i=0}^{n}\omega^i\right)(\tilde{\mathcal{A}}\mathbf{1}_k - 1) + \alpha\rho\omega^{n+1}\tilde{\mathcal{A}}\mathbf{1}_k\right]. \tag{25}$$

Equation (25) shows the effect of moving marginally more upstream on the responsiveness of output to demand shocks. The first term on the right hand side states that moving more upstream implies exposure to potential dissipation by the network. To see this, note that $A\mathbf{1}_k = \sum_s a^{ks}$ is the outdegree of sector k which governs the additional connections for the industry after we increased the length of the chain. The second term, instead, represents the additional inventory amplification as can be seen by the $n+1$ exponent. Depending on which of the two forces prevails, the change in demand will be amplified or dissipated as it travels upstream in the network. This result states that if the inventory amplification effect dominates

[26] This result is close in spirit to Propositions 2 and 3 in Carvalho et al. (2020). In their setting, as inventories are absent, the term in the second square bracket in 3b and 3c is equal to 1. Recall that $\sum_{n=0}^{\infty} A^n_{ks} = \ell_{ks}$, where ℓ_{ks} is an element of the Leontief Inverse L. Furthermore, their assumptions on "pure" upstreamness and downstreamness imply that, if i is further removed from the source of the shock s than k, then $\ell_{ks} > \ell_{is}$. Hence, in their setting the network can only dissipate shocks upstream. I discuss the comparison between Carvalho et al. (2020) and my empirical results in the Online Appendix B.9.

the network dissipation effect, the change in the structure of supply chains implies that shocks will snowball upstream. This effect is driven by the increase in the sector's distance from final consumers.[27]

The model can therefore provide a potential rationale for the empirical results in Section 4. In the remainder of the section I proceed by testing the proposed mechanism directly and providing quantitative counterfactuals.

Multiple Destinations Before going back to the data and studying counterfactuals I lay out a simple extension to the multiple destination case. The key insight is that all of the above goes through, provided that firms optimize separately by destination. To see this note that it is possible to think of the model laid out so far as being the description of the problem for a single destination out of many. Indexing that destination as j, it is possible to write output of industry k as

$$Y_{k,t} = \sum_{j}^{J} \tilde{L} B_{kj} D_{jt} + \alpha\rho \left[\sum_{n=0}^{\infty} \tilde{\mathcal{A}}^n \sum_{i=0}^{n} \omega^i \right]_k B_{kj} \Delta_{jt}$$

The properties derived in Proposition 3 now hold for changes in the demand of a single destination. Trivially, the total variability will depend on the covariance between the destination-specific shocks. In particular note that if two destinations are hit by different shock realization it is now possible to reallocate labor across chains. As a consequence the exact pattern of propagation will depend on which destination is hit by which shock realization. This extension is exactly the one discussed in Remark 1, so the appropriate aggregation of destination specific shocks is given by the shift-share structure proposed in Section 3.

Discussion The model laid out in this section can rationalize the evidence on the cross-sectional distribution of output elasticities discussed in Section 4. However it does so under a set of strong assumptions which are worth discussing.

Two technical assumptions are required to deliver closed form results: i) linear-quadratic inventory problem and ii) pricing behaviour. The first assumption allows me to embed the inventory problem into the complex structure of the network economy. Absent this assumption it is hard to find a recursion such that the problem can be solved without assuming specific network structures (for example a Directed Acyclic Graph, see Appendix B.5.4). The linear-quadratic inventory can be seen as an approximation of a dynamic model in which firms with some probability are unable to produce in a given period, see Appendix B.5.2. Moreover, the specific formulation in this paper implies procyclical inventories due to the optimal target rule based on future demand. This property necessarily mutes any production smoothing motive, as I discuss and consider in Appendix B.5.3. If this force were to dominate we would observe inventories

[27] The exact same result obtains if we use a special case of the comparative statics described in Proposition 3b. In particular if take the case of tires and assume that they are consumed in two possible ways: i) *tires* $\rightarrow$ *cars* $\rightarrow$ *consumption*; ii) *tires* $\rightarrow$ *consumption*. Applying the comparative static $\Delta\beta^{cars} = -\Delta\beta^{tires} = \epsilon > 0$. This implies the upstreamness of tires increases at the margin and the result in 25 obtains.

moving countercyclically which, as discussed in Section 2, is counterfactual. Importantly, as shown in Figure B.8, the linearity assumption of inventories in sales fits the data extremely well.

The second important assumption concerns the behaviour of prices. A key simplifying assumption in the current setup is that the behaviour of prices is muted by subsidies which enforce marginal cost pricing.[28] Without this assumption we do not have the analytical recursion that allows to characterize the equilibrium and the key comparative statics in the model. The practical implication is that the output responses I can characterize are necessarily an upper-bound as the movement of relative prices would dampen these effects. A second important implication is that by marginal-cost pricing and appropriate normalization the equilibrium features a vector of unitary prices. As a consequence there is no difference between output quantity and sales in the data other than the inventory change. Absent this assumption we would therefore need to separately keep track of the behaviour of quantities and sales, this would generate the further complication that the data only reports information on sales movements and it would therefore be significantly harder to discipline the model and the empirical observations.

In summary, this theoretical framework encompasses a richer pattern of propagation of final demand shocks in the network and highlights the key horse race between network features and inventories. The rest of the section recasts this framework to obtain a directly estimable relationship between observable quantities, which allows me to test the key mechanism directly.

5.3 Testing the Mechanism

The reduced form results in Section 4 suggest that firms further away from consumption have larger output responses to the same change in final demand. Furthermore, a similar behaviour is found for inventories, which motivated the model just described to provide a possible explanation for these cross-sectional patterns via the amplification generated by procyclical inventories. In this section I test this mechanism directly in two alternative ways. First, the framework in Section 5.2 provides a model-consistent estimating equation linking output growth to changes in demand through inventories and upstreamness. In particular, the object of interest $\Delta \log Y$ can be recovered by manipulating equation 22 as a function of observables. This is a direct test of the model, in that assumptions on the underlying parameters imply clear predictions on the signs of the coefficients to be estimated. Secondly, and more generally, I estimate a model-free specification combining shocks, inventories and network position to test the mechanism directly.

To recover a model-consistent estimating equation, I start from equation 22. The first observation is that it can be rewritten as

$$\frac{\partial Y_{kt}}{\partial D_t} = \tilde{L} B_k + \alpha \rho \psi \sum_{n=0}^{\infty} n \tilde{\mathcal{A}}^n B,$$

For some $\psi \in (0, 1)$. This follows from the maintained assumption that $\omega \in (0, 1)$, which implies that the term $\sum_{n=0}^{\infty} \tilde{\mathcal{A}}_k^n \sum_{i=0}^{n} \omega^i$ is bounded between 0 and $\sum_{n=0}^{\infty} n \tilde{\mathcal{A}}^n$. Next, recall that

[28] An alternative way of interpreting the assumption is that the shocks are small enough for firms not to be willing to change prices for any positive menu cost.

$U = \sum_{n=0}^{\infty}(n+1)\tilde{\mathcal{A}}^n B/\tilde{L}B$ and $\sum_{n=0}^{\infty}\tilde{\mathcal{A}}^n = \tilde{L}$. The local growth in output can be restated as

$$\frac{\Delta Y_{kt}}{Y_{kt}} = (1-\alpha\rho\psi)\frac{\tilde{L}B_k D_t}{Y_{kt}}\frac{\Delta_t}{D_t} + \alpha\rho\psi U_k \frac{\tilde{L}B_k D_t}{Y_{kt}}\frac{\Delta_t}{D_t}.$$

I can estimate this relationship directly through input-output and inventories data as

$$\Delta \ln Y_{it}^r = \delta_1 \hat{\eta}_{it}^r + \delta_2 \alpha_i^r U_i^r \hat{\eta}_t + \epsilon_{it}^r, \tag{26}$$

Upstreamness U is computed from the I-O data, α is given by the inventory data and $\hat{\eta}$ is the estimated demand shock discussed in Section 3.2. In this estimating equation $\delta_1 = 1 - \alpha\rho\psi$ and $\delta_2 = \rho\psi$.

To directly measure the inventory-to-sales ratio, I use the NBER CES Manufacturing.[29] This data only covers the US and a subset of the industries in the WIOD data. For this reason I maintain throughout the assumption that $\alpha_i^r = \alpha_{US}^r$, $\forall i$, namely that within an industry all countries have the same inventory-to-sales ratios. If one thinks that inventories increase in the level of frictions and that these decrease with the level of a country's development, then the US represents a lower bound in terms of inventory-to-sales ratios. With this assumption I can estimate this regression on a sample with all manufacturing industries in the WIOD data.

Before discussing the results, note that if the model was misspecified and inventories played no role, we should expect $\hat{\delta}_2 = 0$ and $\hat{\delta}_1 = 1$. If inventories smoothed fluctuations upstream, we should have $\hat{\delta}_2 < 0$. Finally, if the network dissipation role were to dominate we should also expect $\hat{\delta}_2 < 0$ as it would capture differential responses based on the position relative to consumers, as measured by U.

Table 1 reports the results of the estimation using different sets of fixed effects. I use time 0 versions of both inventories and I-O measures to avoid their contemporaneous response to the shocks. The results are consistent using either set of fixed effects. Using the most parsimonious specification in column 1 the model suggests that, ceteris paribus, increasing upstreamness by 1 to a sector with average inventories implies a .053pp increase in the output elasticity to demand shocks. When using time fixed effects this number drops to .027pp. These are respectively 9.1% and 4.6% of the average elasticity. Conversely, raising inventories-to-sales ratios by 1 standard deviation to industries at the average level of upstreamness yields a higher elasticity by .075pp or .038pp with time fixed effects, respectively 12.9% and 6.9% of the average. In all cases, consistently with the model, I estimate $\hat{\delta}_1 < 1$ and $\hat{\delta}_2 > 0$.

The estimates provided in Table 1 are subject to risk of model misspecification from the theoretical framework. I check that these results are robust to a more general specification of the empirical model by estimating a saturated version with the interactions between the demand shocks, upstreamness and α as proxied by the inventory-to-sales ratio. I discuss the details of this estimation in the Appendix and report the results in Table B.7. I find that the main results

[29] As discussed earlier in the paper, WIOD contains information on the changes in the inventory stocks which are computed as a residual in the I-O table. To recover the inventory-to-sales ratio I need the level of inventory stock as well as the correct allocation of the industries these inventories are used by. For these reasons I use the NBER CES data which provides reliable values for the inventory stock for US manufacturing industries.

Table 1: Model-Consistent Estimation of the Role of Inventories and Upstreamness

	(1)	(2)	(3)
	$\Delta \ln Y_{it}^r$	$\Delta \ln Y_{it}^r$	$\Delta \ln Y_{it}^r$
$\hat{\eta}_{it}^r$	0.527***	0.504***	0.256***
	(0.0271)	(0.0285)	(0.0249)
$\alpha_i^r \times U_i^r \times \hat{\eta}_{it}^r$	0.500***	0.513***	0.350***
	(0.0845)	(0.0870)	(0.0731)
Constant	0.0802***	0.0800***	0.0767***
	(0.00139)	(0.00155)	(0.00153)
Industry FE	No	Yes	Yes
Time FE	No	No	Yes
N	12098	12098	12098
R^2	0.352	0.425	0.513

Cluster bootstrapped standard errors in brackets.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: This table shows the results of the regressions in 26. Specifications in columns 2 and 3 include country-industry fixed effects. Column 3 also include time fixed effects. Standard errors are cluster bootstrapped at the country-industry pair.

are confirmed and in particular that the triple interaction between shocks, upstreamness and inventories is always positive and statistically significant.

5.4 Model Performance and Counterfactuals

I conclude by using a simulated version of the model to study the trends discussed in Section 2. As a first step, I use the actual WIOD input-requirement matrix $\tilde{\mathcal{A}}$ as the I-O matrix in the model. I do so using the data in 2000 and 2014. I simulate 24000 cross-sections of demand shocks for the J countries assuming that they follow an iid AR(1) process with volatility σ and persistence ρ. I use the volatility parameter σ to match the 2000 dispersion of $\hat{\eta}$ and use $\rho = .7$ as estimated in an AR(1) on the empirical demand shocks. For the inventory parameter there are two options: first, recall that the model is derived under the assumption of symmetric inventories, hence a first option is to use the US manufacturing average inventory-to-sales ratio as α. Alternatively, I can use the relative volatility of output growth and demand in 2000 to compute the level of inventories in 2000. Using the US data underestimates the relative volatility by about 40%, suggesting that the US inventories figure is too low for the rest of the world. Using an inventory-to-sales ratio of .3 I can match the relative volatility of 1.21. These moments are reported in Table A.7.

Following the discussion in Section 5 I report the results for two version of the model. The first one features a unique consumption destination, such that in every period there is only one change in final demand. The second version instead allows for multiple destinations. In the latter case I draw independently J changes in final demand and apply them to actual final demand data in the I-O table. For both cases I also report the results for a model without inventories, $\alpha = 0$, for comparison.

First, the model can replicate the reduced form evidence. In particular, estimating the main specification in equation 9 on the generated data, I obtain Figure 6. I find that, while

Figure 6: Model Data Regression

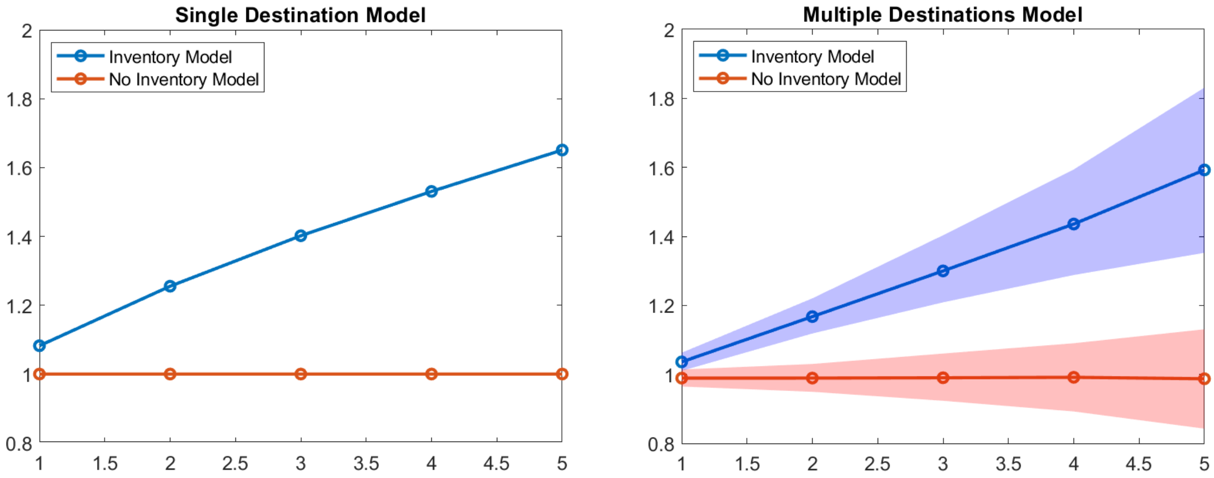

(a) Single Destination

(b) Multiple Destinations

Note: The figures show the model equivalent of Figure 2. Panel (a) shows the result of regression 9 for a model with a single consumer. Panel (b) shows the same estimation for economies with multiple destinations. In the latter case the propagation pattern is not deterministic as it matters which destination receives which shock. Therefore I build confidence intervals by simulating the economy 24000 times and reporting the 10th and 90th percentile of the estimated coefficient distribution as the bounds of the shaded area. In both plots the blue line represent the result in an economy with inventories, while the red line is for economies without inventories (i.e. with α set to 0 for all industries).

both models fail at matching the scale of the coefficient, the inventory model can replicate the slope. Namely the increasing response across different upstreamness bins. The model without inventories cannot generate the positive gradient found in Figure 2.

Note that, as discussed in the previous section, the single destination model does not have any uncertainty around the effect of a demand shock by upstreamness. On the other hand, allowing for iid shocks in a multiple destinations model implies that there is residual uncertainty depending on which country suffers which shock and, given the I-O matrix, which sector is then affected. Therefore, this setting has the additional important feature of allowing for diversification forces to operate. Quantitatively the model matches the slope found in the empirical analysis.

Counterfactuals

In the context of this model I study three alternative counterfactual scenarios.

The first counterfactual experiment replaces the network in 2000 with the one in 2014. As discussed in Section 2 these networks have two salient differences: i) the concentration of sales shares decreased; ii) the average distance from consumers increased. As a consequence one should expect that, fixing the variance of shocks η_j, lower sales share concentration implies more diversification and therefore lower variance in the demand shocks η_i^r that each industry faces. Secondly, the higher distance from consumption should reinforce the inventory amplification channel, generating an increase in the relative volatility of output to demand.

The second counterfactual, motivated by the recent trends in the inventory-to-sales ratio discussed in Section 2, consists of a 25% increase in inventories, holding fixed the network features. Intuitively this should increase the output response for a given change in demand.

Finally, I combine these experiments and allow for both the 25% increase in the inventory-to-sales ratio and a change in the global input-output network from the 2000 to the 2014 WIOD. This should bring about two opposing forces: i) the higher diversification should reduce changes in demand each sector is exposed to; ii) for a given level of changes in demand we should observed increased output responses as both inventories and upstreamness increased.

I report the results of these counterfactual exercises in Table 2. The firsts two columns in the baseline section report targeted moments for the multiple destination model. The statistics of interest are reported in the counterfactual columns.

The table reports three key moments of the model economy which I compute as follows. The standard deviation of demand $\sigma_\eta = \left(\sum_{i,r}(\eta_i^r - \bar{\eta})^2\right)^{\frac{1}{2}}$ where $\eta_i^r = \sum_j \xi_{ij}^r \eta_j$, where r is the industry, i the origin country, j a destination country and $\eta_j = \Delta \log D_j$. Output dispersion σ_y is computed the same way on the growth rate of output. Finally, $\frac{\Delta \log Y_i}{\Delta \log \eta_i}$ is the ratio between the output growth of industry i and the change in final demand industry i is exposed to. For each simulation I compute the dispersion measures and the median $\frac{\Delta \log Y}{\Delta \log \eta}$ and then average across simulations.[30]

Starting from the single destination economy, recall that here there is no scope for diversification forces as there is only one demand shock. The changes in the network imply an increase in the average length of chains which, in turn, generates an increase in the output response to changes in demand from 1.31 to 1.35 as shown in the first row of Table 2. Increasing the inventory-to-sales ratio predictably generates a significant increase of the change in output triggered by a change in demand from 1.31 to 1.37. When I combine these two changes the model predicts a reinforcing effects of the two forces as increasing chain length and inventories are complementary in generating upstream amplification. As a consequence the output response to changes in demand moves from 1.31 to 1.41. Note that, by construction, when changing the network in the single destination economy we only account for the role of increased chain lengths, not for the increased diversification.

The multiple destination model allows for diversification effects. The first observation is that when moving from the 2000 to the 2014 network the model predicts a decline in industries' effective demand shocks. As the destination exposure becomes less concentrated, for a given level of volatility of destination shocks, the cross-sectional dispersion of η_r^i declines from 0.115 to 0.104. At the same time changing the network implies increasing chain length so that the response of output to changes in demand increases from 1.22 to 1.24. In the second counterfactual,

[30]I take the median of $\frac{\Delta \log Y}{\Delta \log \eta}$ rather than the average because, as the denominator is at times very close to zero, the ratio can take extreme values and therefore significantly affect the average response.

Table 2: Counterfactual Moments

	Baseline			Counterfactual		
	σ_η	σ_y	$\frac{\Delta \log Y}{\Delta \log \eta}$	σ_η	σ_y	$\frac{\Delta \log Y}{\Delta \log \eta}$
$[A]$ - Single Destination Model						
$\tilde{A}_{2000} \to \tilde{A}_{2014}$	0	0	1.31	0	0	1.35
$\alpha \to 1.25\alpha$	0	0	1.31	0	0	1.37
$\tilde{A}_{2000} \to \tilde{A}_{2014}$, $\alpha \to 1.25\alpha$	0	0	1.31	0	0	1.41
$[B]$ - Multiple Destinations Model						
$\tilde{A}_{2000} \to \tilde{A}_{2014}$	0.115	0.136	1.22	0.104	.123	1.24
$\alpha \to 1.25\alpha$	0.115	0.136	1.22	0.115	0.14	1.27
$\tilde{A}_{2000} \to \tilde{A}_{2014}$, $\alpha \to 1.25\alpha$	0.115	0.136	1.22	0.104	0.126	1.28

Note: The Table presents the results of baseline and counterfactual estimation. The first 3 columns refer to the baseline model calibrated to 2000 while the last 3 show the counterfactual results. Panel [A] shows the results for the single destination setting while Panel [B] for the multiple destination model. In each model I perform 4 counterfactuals: i) keeping inventories constant I use the I-O matrix of 2014 instead of the one of 2000; ii) keeping the I-O matrix constant I increase inventories by 25%; iii) 25% increase of inventories and changing the I-O matrix from the one in 2000 to the one of 2014; changing the IO matrix in an economy without inventories. Each counterfactual is simulated 4800 times.

increasing inventories while fixing the network structure implies no change in demand exposure and therefore a significant increase of the output volatility as firms respond to demand shocks significantly more, from 1.22 to 1.27 for a 1pp increase in the growth rate of demand. The last counterfactual, allowing for both changes in the network and increasing inventories, suggests that the reduction in effective demand volatility is largely offset by the increase in inventories so that the volatility of output is mostly unchanged. There is, however, a significant increase in the output change triggered by a change in demand, from 1.22 to 1.28 largely driven by the increase in inventories.

These counterfactual experiments suggest that the reshaping of the network is generating opposing forces in terms of output volatility. First, the decreasing exposure to specific destination is reducing the effective volatility of final demand for each industry. At the same time the increase in chain length would imply a higher responsiveness of output to changes in final demand. The latter force seems to be quantitatively very small, in isolation. When combining these network changes with an increase of inventories from an inventory-to-sales ratio of 30% to 37.5% the benefits of the changes in the network are partially undone. In particular output dispersion drops by significantly less when inventories are allowed to increase.

6 Conclusions

Recent decades have been characterized by a significant change in the way goods are produced due to the rise of global value chains. In this paper I ask whether these trends trigger stronger or weaker propagation of final demand shocks. To answer this question, I start by asking whether we observe a higher output response to demand shocks by firms further away from consumption. Using a shift-share design based on global Input-Output data and I find that upstream firms respond up to three times more strongly than their downstream counterparts to the same final demand shock.

I build a theoretical framework embedding procyclical inventories in a network model to study the key features determining upstream amplification vs dissipation patterns. I then estimate the model and, in counterfactual exercises, I find that in absence of the inventory amplification channel we would observe significantly lower output responses to demand shocks. This last result becomes particularly salient in light of the recent trends of increasing inventories and lengthening production chains. The estimated model suggests that between 2001 and 2014 the output response to demand shocks has increased by about a third.

To conclude, this paper represents a first attempt at the study of the interactions between the rise of global supply chains and the role of inventories in propagating shocks. As such it ignores a number of important elements. Two examples of these are given by the role of re-pricing as an absorption mechanism and the dependence of inventory policies on the position in supply chains. This topic represents a promising avenue for both empirical and theoretical research given the recent supply chains disruptions in the Covid-19 crisis.

References

Acemoglu, Daron, Ufuk Akcigit, and William Kerr, "Networks and the Macroeconomy: An Empirical Exploration," *NBER Macroeconomics Annual*, 2016, *30* (1), 273–335.

_ , Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, "The Network Origins of Aggregate Fluctuations," *Econometrica*, 2012, *80* (5), 1977–2016.

Adão, Rodrigo, Michal Kolesár, and Eduardo Morales, "Shift-Share Designs: Theory and Inference," *The Quarterly Journal of Economics*, 2019, *134* (4), 1949–2010.

Alessandria, George, Joseph Kaboski, and Virgiliu Midrigan, "Trade wedges, inventories, and international business cycles," *Journal of Monetary Economics*, 2013, *60* (1), 1–20.

_ , Joseph P Kaboski, and Virgiliu Midrigan, "The Great Trade Collapse of 2008-09: An Inventory Adjustment?," *IMF Economic Review*, December 2010, *58* (2), 254–294.

Alfaro, Laura, Manuel García-Santana, and Enrique Moral-Benito, "On the direct and indirect real effects of credit supply shocks," *Journal of Financial Economics*, 2021, *139* (3), 895–921.

Altomonte, Carlo, Filippo Di Mauro, Gianmarco I. P. Ottaviano, Armando Rungi, and Vincent Vicard, "Global Value Chains During the Great Trade Collapse: A Bullwhip Effect?," CEP Discussion Papers dp1131, Centre for Economic Performance, LSE February 2012.

Angrist, Joshua D and Jörn-Steffen Pischke, *Mostly harmless econometrics: An empiricist's companion*, Princeton university press, 2008.

Antràs, Pol and Davin Chor, "On the Measurement of Upstreamness and Downstreamness in Global Value Chains," NBER Working Papers 24185, National Bureau of Economic Research, Inc January 2018.

_ , _ , Thibault Fally, and Russell Hillberry, "Measuring the Upstreamness of Production and Trade Flows," *American Economic Review*, May 2012, *102* (3), 412–416.

Autor, David, David Dorn, and Gordon H Hanson, "The China syndrome: Local labor market effects of import competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–68.

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

Barrot, Jean-Noël and Julien Sauvagnat, "Input specificity and the propagation of idiosyncratic shocks in production networks," *The Quarterly Journal of Economics*, 2016, *131* (3), 1543–1592.

Blinder, Alan S and Louis J Maccini, "The resurgence of inventory research: what have we learned?," *Journal of Economic Surveys*, 1991, *5* (4), 291–328.

Boehm, Christoph E, Aaron Flaaen, and Nitya Pandalai-Nayar, "Input linkages and the transmission of shocks: Firm-level evidence from the 2011 Tōhoku earthquake," *Review of Economics and Statistics*, 2019, *101* (1), 60–75.

Borusyak, Kirill and Peter Hull, "Non-random exposure to exogenous shocks: Theory and applications," Technical Report, National Bureau of Economic Research 2020.

_ , _ , and Xavier Jaravel, "Quasi-experimental shift-share research designs," *The Review of Economic Studies*, 2022, *89* (1), 181–213.

Carreras-Valle, Maria-Jose, "Increasing Inventories: the Role of Delivery Times," Technical Report, Working Paper 2021.

Carvalho, Vasco M, "Aggregate Fluctuations and the Network Structure oF Intersectoral Trade," *mimeo*, 2010.

_ and Alireza Tahbaz-Salehi, "Production networks: A primer," *Annual Review of Economics*, 2019, *11*, 635–663.

_ , Makoto Nirei, Yukiko Saito, and Alireza Tahbaz-Salehi, "Supply chain disruptions: Evidence from the great east Japan earthquake," *Quarterly Journal of Economics*, 2020.

Chen, Frank, Zvi Drezner, Jennifer Ryan, and David Simchi-levi, "Quantifying the Bullwhip Effect in a Simple Supply Chain: The Impact of Forecasting, Lead Times, and Information," *Management Science*, 07 1999, *46.*

Dhyne, Emmanuel, Ayumu Ken Kikkawa, Magne Mogstad, and Felix Tintelnot, "Trade and domestic production networks," *The Review of Economic Studies*, 2021, *88* (2), 643–668.

European Parliament, "A New Industrial Strategy for Europe," Technical Report November 2020.

Fally, Thibault, "Production Staging: Measurement and Facts," *mimeo*, 01 2012.

Forrester, Jay W, "Industrial dynamics. 1961," *MT Press*, 1961.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik instruments: What, when, why, and how," *American Economic Review*, 2020, *110* (8), 2586–2624.

Gortari, Alonso De, "Disentangling global value chains," Technical Report, National Bureau of Economic Research 2019.

Holt, Charles C, Franco Modigliani, John Muth, and Herbert Simon, *Planning Production, Inventories, and Work Force.* 1960.

Kahn, James A, "Inventories and the volatility of production," *The American Economic Review*, 1987, pp. 667–679.

Korovkin, Vasily and Alexey Makarin, "Production Networks and War," *Available at SSRN*, 2021.

Kramarz, Francis, Julien Martin, and Isabelle Mejean, "Volatility in the small and in the large: The lack of diversification in international trade," *Journal of International Economics*, 2020, *122*, 103276.

Lee, L. Hau, V Padmanabhan, and Seungjin Whang, "Information Distortion in a Supply Chain: The Bullwhip Effect," *Management Science*, 12 2004, *43*, 546–558.

Los, Bart, Reitze Gouma, Marcel Timmer, and Pieter IJtsma, "Note on the construction of wiots in previous years prices," Technical Report 2014.

Metters, Richard, "Quantifying the Bullwhip Effect in Supply Chains," *Journal of Operations Management*, 05 1997, *15*, 89–100.

Ramey, Valerie A and Kenneth D West, "Inventories," *Handbook of macroeconomics*, 1999, *1*, 863–923.

Raza, W, J Grumiller, H Grohs, J Essletzbichler, and N Pintar, "Post Covid-19 value chain: options for reshoring production back to Europe in a globalised economy," *Brussels: European Union–Directorate General for External Policies of the Union*, 2021.

Shea, John, "The input-output approach to instrument selection," *Journal of Business & Economic Statistics*, 1993, *11* (2), 145–155.

Timmer, Marcel P, Erik Dietzenbacher, Bart Los, Robert Stehrer, and Gaaitzen J De Vries, "An illustrated user guide to the world input–output database: the case of global automotive production," *Review of International Economics*, 2015, *23* (3), 575–605.

White House, "Biden-Harris Administration Announces Supply Chain Disruptions Task Force to Address Short-Term Supply Chain Discontinuities," Technical Report June 2021.

World Bank, *World Development Report 2020* number 32437. In 'World Bank Publications.', The World Bank, 12-2019 2020.

Zavacka, Veronika, "The bullwhip effect and the Great Trade Collapse," 2012.

Appendix

A.1 Motivating Evidence

Fact 1: Production chains have increased in length

Figure A.1: Upstreamness Dynamics

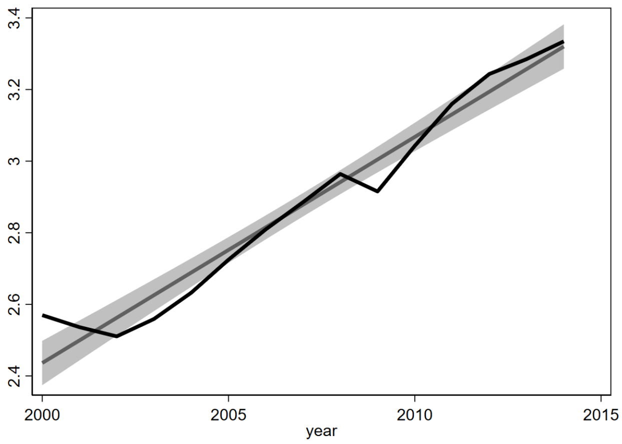

(a) Upstreamness Dynamics

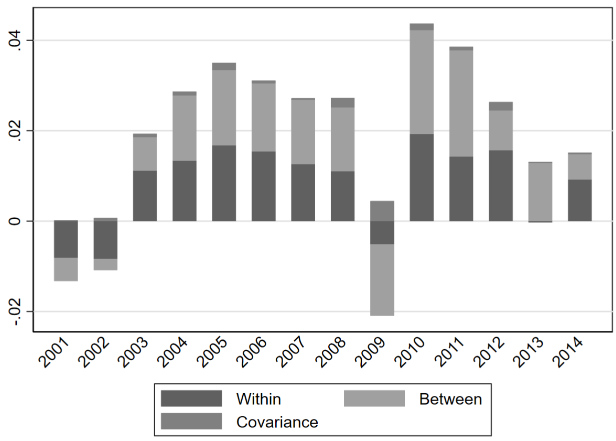

(b) Dynamics Decomposition

Note: The figure shows the dynamics of the weighted upstreamness measure computed as $U_t = \frac{\sum_i \sum_r y^r_{it} U^r_{it}}{\sum_i \sum_r y^r_{it}}$. The left panel shows the average over time and it includes the estimated linear trend and the 95% confidence interval around the estimate. The right panel shows the decomposition of these changes into the stacked contributions (in levels) of the different components of the changes in the weighted average upstreamness measure. The components are given by $\Delta U_t = \sum_i \sum_r \underbrace{\Delta U^r_{it} w^r_{it-1}}_{Within} + \underbrace{U^r_{it-1} \Delta w^r_{it}}_{Between} + \underbrace{\Delta U^r_{it} \Delta w^r_{it}}_{Covariance}$.

Fact 2: Sales shares are becoming less concentrated

Figure A.2: Herfindahl Index of Sales Shares

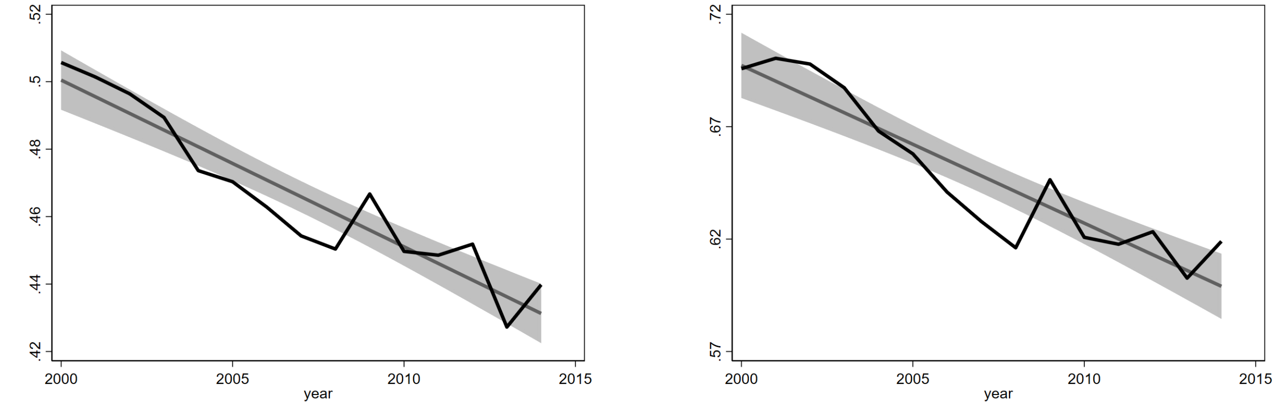

(a) Simple Average HHI

(b) Weighted Average HHI

Note: The figure shows the behavior of the Herfindahl Index of destination shares over time. Destination shares are described as in equation 4 in Section 3.2. The Herfindahl Index is computed at the industry level as $HHI_t^r = \sum_j {\xi_{ij}^r}^2$. The left panel shows the simple average across industry, i.e. $HHI_t = R^{-1} \sum_r HHI_t^r$. The right panel shows the weighted average using industry shares as weights: $HHI_t = \sum_r \frac{Y_t^r}{Y_t} HHI_t^r$. The plots include the estimated linear trend and the 95% confidence interval around the estimate.

Fact 3: Inventories are adjusted procyclically

Figure A.3: Distribution of estimated I'

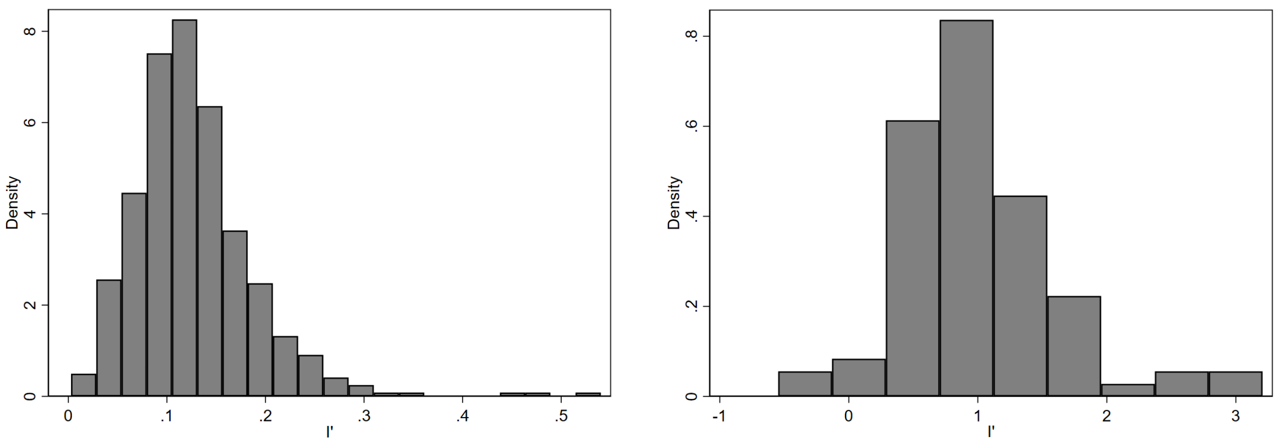

(a) NBER

(b) Census

Note: The graph shows the distribution of estimated $I'(\cdot)$, namely the derivative of the empirical inventory function with respect to sales. The sample is the full NBER CES sample of 473 manufacturing industries. The estimation is carried out sector by sector using time variation. The graph shows the sector-specific estimated coefficient. The left panel shows the same statistics based on the monthly data from the Manufacturing & Trade Inventories & Sales data of the US Census.

Table A.1: Estimation of $I'(\cdot)$

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
y_t^i	$\tilde{I}_t^i$	$\tilde{\alpha}_t^i$	$\tilde{I}_t^i$	$\tilde{\alpha}_t^i$	I_t^i	α_t^i	I_t^i	α_t^i
NBER CES Sample - Annual Data								
Mean of y_t^i	-11.36*	-0.00000270	-11.20***	-0.0000200	966.8***	0.127***	982.4***	0.127***
	(5.939)	(0.000170)	(4.154)	(0.000409)	(33.34)	(0.000891)	(16.43)	(0.000570)
$\frac{\partial y_t^i}{\partial Sales_t^i}$	0.0815***	-0.00000218***	0.0745***	-0.00000251***	0.106***	-0.00000129***	0.101***	-0.000000641***
	(0.0107)	(0.000000412)	(0.00994)	(0.000000205)	(.0013713)	(6.53e-08)	(0.00555)	(9.71e-08)
Industry FE	YES	YES	NO	NO	YES	YES	NO	NO
N	5658	5658	5659	5659	5664	5676	5676	5676
R^2	0.395	0.0239	0.307	0.0163	0.759	0.140	0.551	0.0617
US Census Sample - Monthly Data								
Mean of y_t^i	-11.53	0.00123	-3.475	0.00128	52760.8***	1.598***	52731.2***	1.620***
	(20.83)	(0.00135)	(34.10)	(0.00144)	(682.3)	(0.00515)	(764.1)	(0.00481)
$\frac{\partial y_t^i}{\partial Sales_t^i}$	0.337***	-0.00000139***	0.341***	-0.00000139***	1.353***	-0.00000218***	1.344***	-0.00000111***
	(0.0143)	(0.000000104)	(0.00841)	(0.000000118)	(0.00702)	(7.94e-08)	(0.00998)	(3.34e-08)
Industry FE	YES	YES	NO	NO	YES	YES	NO	NO
N	30479	30479	30479	30479	30479	30479	30479	30479
R^2	0.325	0.0120	0.313	0.00661	0.975	0.769	0.961	0.0678

Bootstrapped standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: This table displays the results of the non-parametric kernel estimations of $I'(\cdot)$ and $\alpha'(\cdot)$, the derivative of the inventory function and the inventory to sales ratio function with respect to current sales. The estimation is based on the data of the NBER CES Manufacturing Industries Dataset for 2000-11 for the top Panels and on the US Census Manufacturing & Trade Inventories & Sales data for the bottom Panels. Standard errors are bootstrapped. Variables with $\sim$ denote HP-filtered data. Columns (1), (2), (5) and (6) include industry fixed effects.

Fact 4: Output is more volatile than sales

Figure A.4: Relative Volatility of Output and Sales

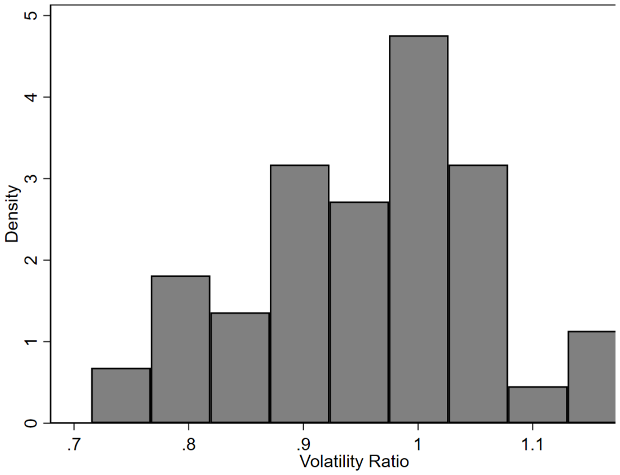

(a) Monthly

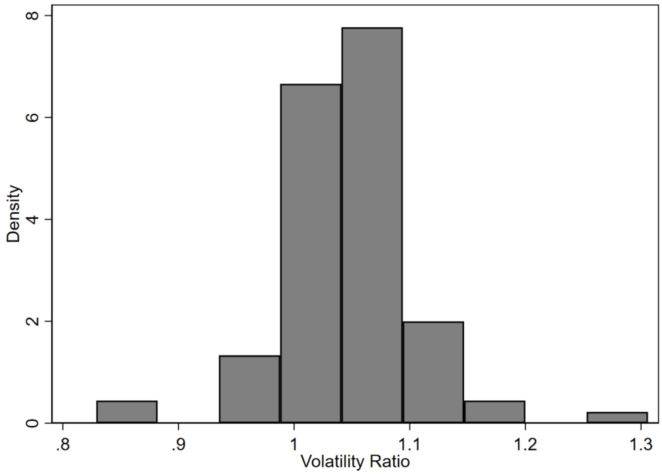

(b) Quarterly

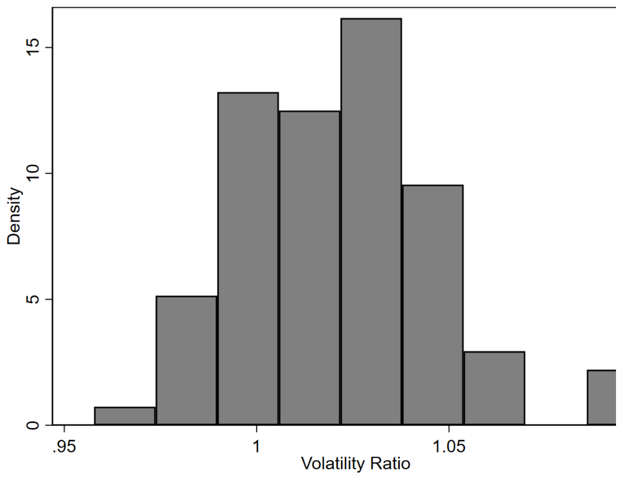

(c) Yearly

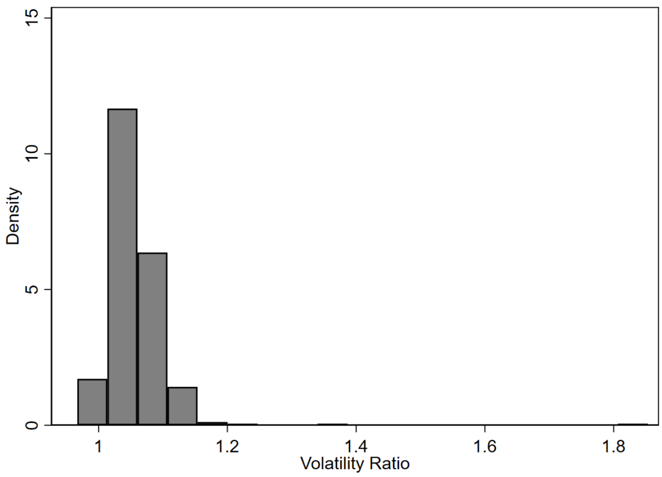

(d) Yearly NBER

Note: The graphs show the distribution of the ratio of volatility of HP-filtered output to HP-filtered sales across sectors. Panels (a), (b) and (c) represent data from the Manufacturing & Trade Inventories & Sales data of the US Census, while Panel (d) shows data from the NBER CES Manufacturing data. Both sources are described in the data Section. Panels (a) and (d) have no aggregation, while for Panels (b) and (c) I sum monthly output and sales to get quarterly and yearly output and sales.

Figure A.5: Correlation between Volatility and Upstreamness

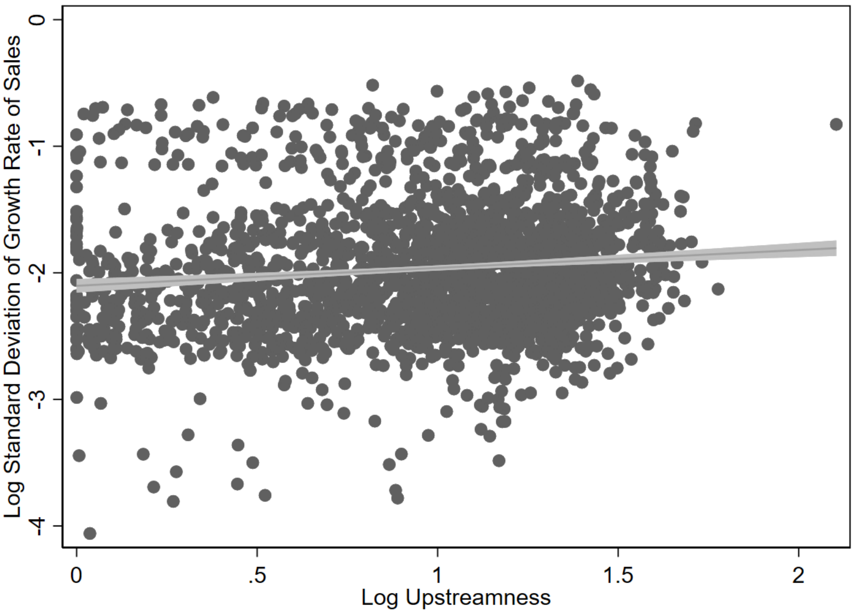

Note: The graph shows the simple correlation between the log of the standard deviation of the growth rate of output and the log of upstreamness. The black line represents the linear fit and the grey area is the 95% confidence interval.

Fact 5: : Inventories-to-sales ratios are increasing

Figure A.6: Trends in Inventory-to-Sales ratios

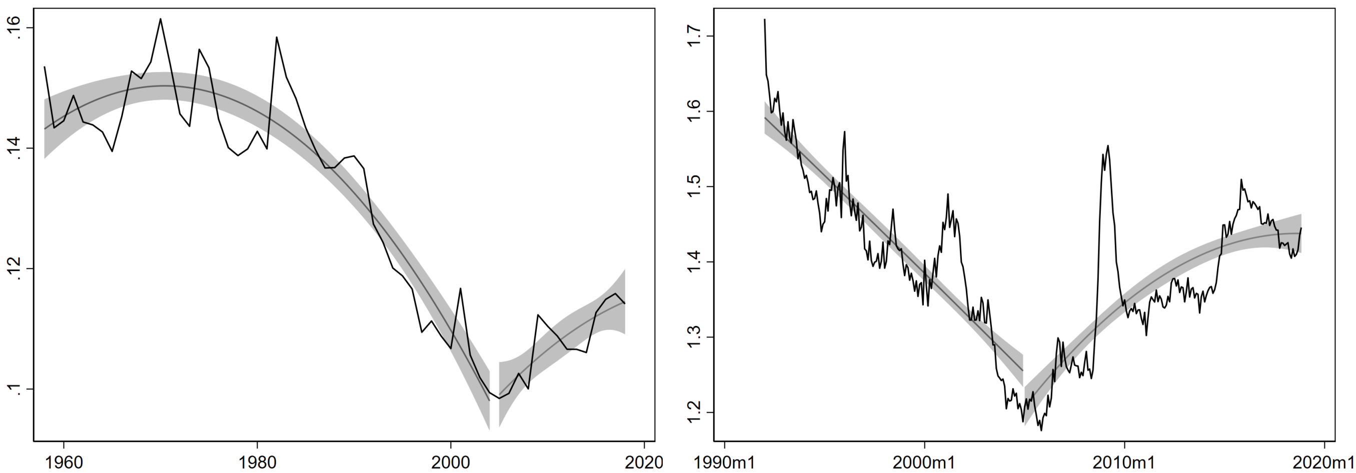

(a) Yearly NBER

(b) Monthly Census

Note: The graphs replicate the key finding in Carreras-Valle (2021). Panel (a) shows the inventory-to-sales ratio from 1958 to 2018 from the NBER CES Manufacturing Database. Panel (b) reports the same statistic from the Census data from Jan-1992 to Dec-2018. Both graphs include non-linear trends before and after 2005. I estimate separate trends as Carreras-Valle (2021) suggests that 2005 is when the trend reversal occurs.

A.2 Results

Table A.2: Industry Output Growth, Price Indices and Demand Shocks

	(1)	(2)	(3)	(4)
	$\Delta \ln Y_{it}^r$	$\Delta \ln Y_{it}^r$	$\Delta \ln P_{it}^r$	$\Delta \ln P_{it}^r$
$\hat{\eta}_{it}^r$	0.598***	0.332***	0.433***	0.231***
	(0.00673)	(0.00977)	(0.00535)	(0.00769)
Constant	0.0841***	0.0820***	0.0846***	0.0833***
	(0.000874)	(0.000861)	(0.00126)	(0.00125)
Industry FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	No	Yes
N	32371	32371	31911	31911
R^2	0.428	0.492	0.554	0.617

Cluster bootstrapped standard errors in brackets.

Standard errors are clustered at the producing industry × country level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table shows the regressions of the growth rate of industry output and the changes in the sectoral price index on the weighted demand shocks that the industry receives. Columns 1 and 2 regress output growth rates on demand shocks with industry and time fixed effects. Columns 3 and 4 show the same regression with the change in the deflator as outcome. This is computed by taking the ratio of the I-O tables at current and previous year prices to obtain the growth rate of the deflator from year to year.

Table A.3: Effect of Demand Shocks on Output Growth by Upstreamness Level

	(1) $\Delta \ln Y^r_{it}$	(2) $\Delta \ln Y^r_{it}$	(3) $\Delta \ln Y^r_{it}$
Upstreamness in [1,2]	0.473***	0.263***	0.262***
	(0.0116)	(0.0115)	(0.0122)
Upstreamness in [2,3]	0.544***	0.319***	0.319***
	(0.0125)	(0.0130)	(0.0127)
Upstreamness in [3,4]	0.663***	0.396***	0.395***
	(0.0108)	(0.0130)	(0.0136)
Upstreamness in [4,5]	0.763***	0.454***	0.456***
	(0.0201)	(0.0207)	(0.0217)
Upstreamness in [5,6]	0.911***	0.591***	0.606***
	(0.0630)	(0.0517)	(0.0496)
Upstreamness in [6,∞)	1.146***	0.749***	0.783***
	(0.210)	(0.188)	(0.213)
Constant	0.0852***	0.0829***	0.0830***
	(0.000838)	(0.000901)	(0.000887)
Time FE	No	Yes	Yes
Level FE	No	No	Yes
Country-Industry FE	Yes	Yes	Yes
N	32371	32371	32371
R^2	0.439	0.497	0.497

Cluster bootstrapped standard errors in brackets.

Standard errors are clustered at the producing industry × country level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The Table shows the results of the regression in equation 9. In particular I regress the growth rate of output on the estimated demand shocks interacted with dummies taking value 1 if upstreamness is in the $[1, 2]$ bin, $[2, 3]$ bin and so on. Observations with upstreamness above 7 are included in the $[6, \infty)$ bin. All regressions include producing industry-country fixed effects and columns 2 and 3 progressively add time fixed effects and upstreamness bin fixed effects. Standard errors are clustered at the producing industry-country level.

Table A.4: Effect of Demand Shocks on Output Growth by Upstreamness Level - Government Consumption Instrument

	(1)	(2)
	$\Delta \ln Y^r_{it}$	$\Delta \ln Y^r_{it}$
Upstreamness in [1,2]	0.455***	0.345***
	(0.0179)	(0.0141)
Upstreamness in [2,3]	0.462***	0.423***
	(0.0195)	(0.0143)
Upstreamness in [3,4]	0.520***	0.528***
	(0.0176)	(0.0133)
Upstreamness in [4,5]	0.551***	0.616***
	(0.0277)	(0.0219)
Upstreamness in [5,6]	0.640***	0.778***
	(0.0745)	(0.0643)
Upstreamness in [6,∞)	0.883***	1.079***
	(0.193)	(0.193)
First Stage Residual U in [1,2]	0.0461***	
	(0.0174)	
First Stage Residual U in [2,3]	0.157***	
	(0.0236)	
First Stage Residual U in [3,4]	0.230***	
	(0.0217)	
First Stage Residual U in [4,5]	0.326***	
	(0.0310)	
First Stage Residual U in [5,6]	0.463***	
	(0.0908)	
First Stage Residual U in [6,∞)	0.503**	
	(0.208)	
First Stage Residual		0.233***
		(0.0120)
Constant	0.0814***	0.0816***
	(0.0000958)	(0.0000900)
Industry FE	Yes	Yes
N	31653	31441
R^2	0.433	0.438

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: This table displays the results of the regression of growth rate of industry output on instrumented demand shocks interacted with dummies taking value 1 if the upstreamness level of the industry is in a given interval, e.g. [1,2]. Column 1 shows the results for the case with a single first stage. Column 2 displays the result when using multiple first stages, in other words allowing for a a different relationship between instrument and demand shocks by upstreamness bin. Both columns including country-industry pair fixed effects and standard errors are clustered at the country-industry pair level.

Table A.5: Effect of China Shock on Output Growth by Upstreamness Level

	(1) $\Delta \ln Y_t^r$	(2) $\Delta \ln VA_t^r$
Upstreamness in [1,2]	0.00549	0.00308
	(0.0404)	(0.0360)
Upstreamness in [2,3]	0.0669	0.0919*
	(0.0554)	(0.0508)
Upstreamness in (3, ∞)	0.166**	0.172**
	(0.0675)	(0.0743)
First Stage Residual U in [1,2]	0.0110	0.00915
	(0.0190)	(0.0162)
First Stage Residual U in [2,3]	-0.0243	-0.0662
	(0.0473)	(0.0439)
First Stage Residual U in (3, ∞)	-0.119**	-0.128**
	(0.0491)	(0.0617)
Constant	0.230***	0.215***
	(0.0152)	(0.0158)
Industry FE	Yes	Yes
Time FE	Yes	Yes
N	5871	6180
R^2	0.266	0.182

Cluster bootstrapped standard errors in brackets.

Standard errors are clustered at the producing industry level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The Table reports the results of estimating equation 9 using the China shock IV in Autor et al. (2013). I instrument the change in US imports from China with the change in other advanced economies. I apply the network transformation to the shocks to account for indirect linkages and standardize them so that the coefficient are changes to a 1 standard deviation of the shocks. I estimate the model with the control function approach, first estimating the endogenous variable on the instrument and fixed effects and controlling for the residual in the second stage. Both columns include producing industry and time fixed effects. Column 1 estimates the model using as outcome the growth rate of output while column 2 uses the growth rate of value added.

Table A.6: Effect of Demand Shocks on Inventory Changes by Upstreamness Level

	(1) $\frac{\Delta I_{it}^r}{Y_{it}^r}$	(2) $\frac{\Delta I_{it}^r}{Y_{it}^r}$	(3) $\frac{\Delta I_{it}^r}{Y_{it}^r}$	(4) $\frac{\Delta I_{it}^r}{Y_{it}^r}$
Uptreamness in [1,2]	0.0205***	0.0141**	0.0159***	0.00909
	(0.00547)	(0.00548)	(0.00520)	(0.00755)
Uptreamness in [2,3]	0.00703	0.000604	-0.000989	-0.0107
	(0.00580)	(0.00598)	(0.00618)	(0.00773)
Uptreamness in [3,4]	0.0275***	0.0187***	0.0171***	0.0127
	(0.00652)	(0.00623)	(0.00600)	(0.00841)
Uptreamness in [4,5]	0.0554***	0.0431***	0.0386***	0.0445***
	(0.00969)	(0.00976)	(0.00985)	(0.0120)
Uptreamness in [5,6]	0.0858***	0.0713***	0.0647***	0.0739***
	(0.0189)	(0.0183)	(0.0167)	(0.0192)
Uptreamness in [6, ∞)	0.117*	0.0923	0.0693	0.122
	(0.0667)	(0.0751)	(0.0589)	(0.0807)
First Stage Residual				0.0215***
				(0.00507)
Constant	0.0833***	0.0833***	0.0832***	0.0829***
	(0.00351)	(0.00343)	(0.00349)	(0.00361)
Time FE	No	Yes	Yes	No
Level FE	No	No	Yes	No
Country-Industry FE	Yes	Yes	Yes	Yes
N	32432	32432	32432	31306
R^2	0.917	0.917	0.918	0.918

Cluster bootstrapped standard errors in brackets.

Standard errors are clustered at the producing industry × country level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The Table shows the results of the regression in equation 9 with inventory changes as the dependent variable. In particular I regress the inventory changes over output on the estimated demand shocks interacted with dummies taking value 1 if upstreamness is in the $[1, 2]$ bin, $[2, 3]$ bin and so on. Observations with upstreamness above 7 are included in the $[6, \infty)$ bin. All regressions include producing industry-country fixed effects and columns 2 and 3 progressively add time fixed effects and upstreamness bin fixed effects. Column 4 shows the result instrumenting demand shocks with government consumption. Standard errors are clustered at the producing industry-country level.

A.3 Proofs

Proof of Remark 1. Starting by the I-O matrix we have that, in matrix form, output is given by $Y = [I - \mathcal{A}]^{-1}F$. Under the Cobb-Douglas assumption $F_{kj}^r = \beta_{kj}^r D_j$, $\forall r, j, k$. Therefore the change in output of sector k sold directly or indirectly to country j is given by

$$\Delta Y^r_{ijt} = \sum_k \sum_s \ell^{rs}_{ik}\beta^s_{kj}\Delta_{jt},$$

summing over destinations to recover the total output change for industry r in country i

$$\begin{aligned}\Delta Y^r_{it} = \sum_j \Delta Y^r_{ijt} &= \sum_j \sum_k \sum_s \ell^{rs}_{ik}\beta^s_{kj}\Delta_{jt} \\ &= \sum_j \sum_k \sum_s \ell^{rs}_{ik}\beta^s_{kj}D_{jt}\frac{\Delta_{jt}}{D_{jt}},\end{aligned}$$

dividing by total output to obtain the growth rate

$$\begin{aligned}\frac{\Delta Y^r_{it}}{Y^r_{it}} &= \sum_j \frac{\sum_k \sum_s \ell^{rs}_{ik}\beta^s_{kj}D_{jt}}{Y^r_{it}}\frac{\Delta_{jt}}{D_{jt}} \\ &= \sum_j \xi^r_{ij}\eta_{jt},\end{aligned}$$

where the last equality follows from the definition of ξ^r_{ij} and η_{jt}. ■

Proof of Proposition 1. The goal is to prove that if $0 < I'(x) < \frac{1}{1-\rho}$ then $\frac{\partial Y^n_t}{\partial D^0_t} > \frac{\partial Y^{n-1}_t}{\partial D^0_t}$, $\forall n, t$. The proof starts by characterising $\frac{\partial Y^n_t}{\partial D^0_t}$.
Evaluating 10 at stage 0

$$\frac{\partial Y^0_t}{\partial D^0_t} = 1 + \frac{\partial I(\mathbb{E}_t D^0_{t+1})}{\partial D^0_t} = 1 + \rho I'$$

Similarly at stage 1

$$\begin{aligned}\frac{\partial Y^1_t}{\partial D^0_t} &= \frac{\partial Y^0_t}{\partial D^0_t} + \frac{\partial I(\mathbb{E}_t D^1_{t+1})}{\partial D^0_t} \\ &= 1 + \rho I' + I'\left[\frac{\partial \mathbb{E}_t}{\partial D^0_t}\left[D^0_{t+1} + I(\mathbb{E}_{t+1}D^0_{t+2}) - I(\mathbb{E}_t D^0_{t+1})\right]\right] \\ &= 1 + \rho I' + \rho I'[1 + \rho I' - I']\end{aligned}$$

Similarly for stage 2

$$\begin{aligned}\frac{\partial Y^2_t}{\partial D^0_t} &= \frac{\partial Y^1_t}{\partial D^0_t} + \frac{\partial I(\mathbb{E}_t D^2_{t+1})}{\partial D^0_t} \\ &= 1 + \rho I' + \rho I'[1 + \rho I' - I'] + \rho I'[1 + \rho I' - I']^2\end{aligned}$$

From the recursion

$$\frac{\partial Y^n_t}{\partial D^0_t} = \frac{\partial Y^{n-1}_t}{\partial D^0_t} + \rho I'[1 + \rho I' - I']^n$$

Given $\rho > 0$, if $0 < I'(x) < \frac{1}{1-\rho}$ then the last term is positive and $\frac{\partial Y_t^n}{\partial D_t^0} > \frac{\partial Y_t^{n-1}}{\partial D_t^0}$. To show the opposite implication, note that if n is even, then $\frac{\partial Y_t^n}{\partial D_t^0} > \frac{\partial Y_t^{n-1}}{\partial D_t^0}$ implies $I' > 0$. If n is odd then $\frac{\partial Y_t^n}{\partial D_t^0} > \frac{\partial Y_t^{n-1}}{\partial D_t^0}$ implies either $I' > 0$ and $1 + \rho I' - I' > 0$ or $I' < 0$ and $1 + \rho I' - I' < 0$. The first case is true if $I' > 0$ and $I' < 1/(1-\rho)$. The second case would require $I' < 0$ and $I' > 1/(1-\rho) > 0$ which is a contradiction.

■

Proof of Lemma 1. From equation 10 for stage 0 and the optimal rule in 12

$$Y_t^0 = D_t^0 + \alpha \mathbb{E}_t D_{t+1}^0 - \alpha \mathbb{E}_{t-1} D_t^0 = D_t^0 + \alpha\rho\Delta_t$$

Using the market clearing condition $D_t^1 = Y_t^0$, the definition of Y_t^1 as a function of demand at stage 0 and inventory adjustment

$$Y_t^1 = D_t^0 + \alpha\rho(2 - \alpha + \alpha\rho)\Delta_t = Y_t^0 + \alpha\rho(1 - \alpha(\rho - 1))\Delta_t.$$

Similarly, for stage 2,

$$\begin{aligned} Y_t^2 &= D_t^0 + \alpha\rho(3 + 3\alpha\rho - 3\alpha + \alpha^2 - 2\alpha^2\rho + \alpha^2\rho^2)\Delta_t \\ &= Y_t^1 + \alpha\rho(1 + \alpha(\rho - 1))^2\Delta_t. \end{aligned}$$

It follows from the recursion that

$$Y_t^n = Y_t^{n-1} + \alpha\rho(1 + \alpha(\rho - 1))^n\Delta_t,$$

or, as a function of final demand,

$$Y_t^n = D_t^0 + \alpha\rho\sum_{i=0}^{n}(1 + \alpha(\rho - 1))^i\Delta_t.$$

As stated in the Lemma.

■

Proof of Proposition 2. The proof of the first statement follows immediately by taking the partial derivative with respect to D_t of equation 13. The second part of the statement follows by taking the second derivative and noting that it is equal to $\alpha\rho(1 + \alpha(\rho - 1))^n$, which is always positive if $0 < \alpha < 1/(1-\rho)$. ■

Proof of Lemma 2. The first part of the Lemma follows immediately from the definition of output at a specific stage n and total sectoral output as the sum over stage-specific output. The proof of the second part requires the following steps: first, using the definition of χ_k^n and denoting $\omega = 1 + \alpha(\rho - 1)$, rewrite total output as

$$\begin{aligned}
Y_{k,t} &= \sum_{n=0}^{\infty} \chi_k^n \left[D_t + \alpha\rho \sum_{i=0}^{n} \omega^i \Delta_t \right] \\
&= \left[\tilde{\mathcal{A}}^0 + \tilde{\mathcal{A}}^1 + \dots \right]_k BD_t + \alpha\rho \left[\tilde{\mathcal{A}}^0\omega^0 + \tilde{\mathcal{A}}^1(\omega^0 + \omega^1) + \dots \right]_k B\Delta_t \\
&= \tilde{L}_k BD_t + \alpha\rho \left[\sum_{n=0}^{\infty} \tilde{\mathcal{A}}^n \sum_{i=0}^{n} \omega^i \right]_k B\Delta_t.
\end{aligned}$$

The equality between the second and the third row follows from the convergence of a Neumann series of matrices satisfying the Brauer-Solow condition. To show that $Y_{k,t}$ exists non-negative for $\omega - 1 = \alpha(\rho - 1) \in [-1, 0]$, note that if $\omega - 1 = -1$ then $\omega = 0$, the second term vanishes and existence and non-negativity follow from $\tilde{L}$ finite and non-negative. If $\omega - 1 = 0$, then $\omega = 1$ and

$$\begin{aligned}
Y_{k,t} &= \tilde{L}_k BD_t + \alpha\rho \left[\sum_{n=0}^{\infty} (n+1) \tilde{\mathcal{A}}^n \right]_k B\Delta_t \\
&= \tilde{L}_k BD_t + \alpha\rho \left[\tilde{\mathcal{A}}^0 + 2\tilde{\mathcal{A}}^1 + 3\tilde{\mathcal{A}}^2 + \dots \right]_k B\Delta_t \\
&= \tilde{L}_k BD_t + \alpha\rho \tilde{L}_k^2 B\Delta_t,
\end{aligned}$$

where the last equality follows from $\sum_{i=0}^{\infty} (i+1)\mathcal{A}^i = [I - \mathcal{A}]^{-2}$ if $\mathcal{A}$ satisfies the Brauer-Solow condition. Existence and non-negativity follow from existence and non-negativity of $[I - \tilde{\mathcal{A}}]^{-2}$. If $\omega - 1 \in (-1, 0)$, then $\omega \in (0, 1)$. As this term is powered up in the second summation and as it is strictly smaller than 1, it is bounded above by $n + 1$. This implies that the whole second term

$$\sum_{n=0}^{\infty} \tilde{\mathcal{A}}^n \sum_{i=0}^{n} \omega^i < \sum_{n=0}^{\infty} (n+1) \tilde{\mathcal{A}}^n = \tilde{L}^2 < \infty.$$

Alternatively, note that the second summation is strictly increasing in ω, as $\omega \leq 1$ the summation is bounded above by $n + 1$. Which completes the proof.

■

Proof of Remark 2. An economy with general input-output structure can be thought of as an infinite collection of vertical production chains with length $n = 0, 1, 2, ...$ Upstreamness is defined as

$$U_k = \sum_{n=0}^{\infty} (n+1) \frac{Y_k^n}{Y_k}.$$

To prove that this metric is well defined first, recall $Y_k = \sum_{n=0}^{\infty} Y_k^n$. Secondly, by Lemma 2 the following holds

$$Y_k = \tilde{L}_k BD_t + \alpha\rho \left[\sum_{n=0}^{\infty} \tilde{\mathcal{A}}^n \sum_{i=0}^{n} \omega^i\right]_k B\Delta_t,$$

and

$$Y_k^n = \tilde{\mathcal{A}}_k^n BD_t + \alpha\rho\tilde{\mathcal{A}}_k^n \sum_{i=0}^{n} \omega^i B\Delta_t.$$

Then

$$U_k = \left[\tilde{L}_k BD_t + \alpha\rho \left[\sum_{n=0}^{\infty} \tilde{\mathcal{A}}^n \sum_{i=0}^{n} \omega^i\right]_k B\Delta_t\right]^{-1} \left[\sum_{n=0}^{\infty}(n+1)\left[\tilde{\mathcal{A}}_k^n BD_t + \alpha\rho\tilde{\mathcal{A}}_k^n \sum_{i=0}^{n} \omega^i B\Delta_t\right]\right].$$

To show that U_k is finite, first note that $\sum_{n=0}^{\infty}(n+1)\tilde{\mathcal{A}}_k^n BD_t = [I-\tilde{\mathcal{A}}]_k^{-2} BD_t$ which is finite. Hence, I am left to show that the last term is finite. Following similar steps to the proof of Lemma 2, note that if $\omega = 0$ then the last term is 0. If $\omega = 1$ then $\sum_{n=0}^{\infty}(n+1)\alpha\rho\tilde{\mathcal{A}}_k^n \sum_{i=0}^{n} \omega^i B\Delta_t = \alpha\rho\sum_{n=0}^{\infty}(n+1)^2\tilde{\mathcal{A}}_k^n B\Delta_t$. Note that

$$\begin{aligned}
&\sum_{n=0}^{\infty}(n+1)^2\tilde{\mathcal{A}}_k^n = \sum_{n=0}^{\infty} n^2\tilde{\mathcal{A}}_k^n + 2\sum_{n=0}^{\infty} n\tilde{\mathcal{A}}_k^n + \sum_{n=0}^{\infty}\tilde{\mathcal{A}}_k^n \\
&= \sum_{n=0}^{\infty} n^2\tilde{\mathcal{A}}_k^n + 2\sum_{n=0}^{\infty}(n+1)\tilde{\mathcal{A}}_k^n - \sum_{n=0}^{\infty}\tilde{\mathcal{A}}_k^n \\
&= \sum_{n=0}^{\infty} n^2\tilde{\mathcal{A}}_k^n + 2[I-\tilde{\mathcal{A}}]_k^{-2} - [I-\tilde{\mathcal{A}}]_k^{-1}.
\end{aligned}$$

To show that the first term is bounded, totally differentiate

$$\begin{aligned}
&\frac{\partial}{\partial\tilde{\mathcal{A}}}\sum_{n=0}^{\infty}(n+1)\tilde{\mathcal{A}}_k^n = \frac{\partial}{\partial\tilde{\mathcal{A}}}[I-\tilde{\mathcal{A}}]_k^{-2} \\
&\sum_{n=0}^{\infty} n^2\tilde{\mathcal{A}}_k^{n-1} + \sum_{n=0}^{\infty} n\tilde{\mathcal{A}}_k^{n-1} = 2[I-\tilde{\mathcal{A}}]_k^{-3} \\
&\sum_{n=0}^{\infty} n^2\tilde{\mathcal{A}}_k^n = 2\tilde{\mathcal{A}}[I-\tilde{\mathcal{A}}]_k^{-3} - \tilde{\mathcal{A}}[I-\tilde{\mathcal{A}}]_k^{-2}.
\end{aligned}$$

As both terms on the right hand side are bounded, so is the term on the left hand side. This implies that $\sum_{n=0}^{\infty}(n+1)^2\tilde{\mathcal{A}}_k^n$ is bounded. As the term is bounded for $\omega = 1$ and it is strictly increasing in ω, U_k is well defined for any $\omega \in [0,1]$. Finally, note that $U_k = 1$ iff $Y_k = Y_k^0$.

■

Proof of Proposition 3. The result in part a follows from the partial derivative of output from Lemma 2. The statement in part b can be shown as follows

$$
\begin{aligned}
\Delta_\beta \frac{\partial Y_{k,t}}{\partial D_t} \equiv \frac{\partial}{\partial \beta_s}\frac{\partial Y_{k,t}}{\partial D_t} - \frac{\partial}{\partial \beta_r}\frac{\partial Y_{k,t}}{\partial D_t} = \\
&= \tilde{L}_{ks} + \alpha\rho \sum_{n=0}^{\infty} \tilde{\mathcal{A}}_{ks}^{n} \sum_{i=0}^{n} \omega^i - \tilde{L}_{kr} - \alpha\rho \sum_{n=0}^{\infty} \tilde{\mathcal{A}}_{kr}^{n} \sum_{i=0}^{n} \omega^i \\
&= \sum_{n=0}^{\infty} \left[\tilde{\mathcal{A}}_{ks}^{n} - \tilde{\mathcal{A}}_{kr}^{n}\right] \left[1 + \alpha\rho \sum_{i=0}^{n} \omega^i\right].
\end{aligned}
$$

Where the last equality follows from the definition of $\tilde{L}$.
Finally, the result in part c can be derived analogously

$$
\begin{aligned}
\Delta_{\tilde{L}} \frac{\partial Y_{k,t}}{\partial D_t} \equiv \frac{\partial Y_{k,t}|\tilde{\mathcal{A}}'}{\partial D_t} - \frac{\partial Y_{k,t}|\tilde{\mathcal{A}}}{\partial D_t} \\
&= \Delta \tilde{L}_k B + \alpha\rho \sum_{i=0}^{n} \omega^i \left[\tilde{A'}_k^n - \tilde{\mathcal{A}}_k^n\right] B \\
&= \sum_{n=0}^{\infty} \left[\tilde{A'}_k^n - \tilde{\mathcal{A}}_k^n\right] \left[1 + \alpha\rho \sum_{i=0}^{n} \omega^i\right] B.
\end{aligned}
$$

Where the last equality follows from the definition of $\tilde{L}$.

■

A.4 Quantitative Model

Table A.7: Targeted moments and model counterparts

	Data	Model
σ_η	0.11	0.115
σ_y	0.133	0.136

Note: The Table reports the targeted moments in the data and in the model. σ_η is the cross-sectional dispersion of η_r^i in 2000, while σ_y is the cross-sectional dispersion of output growth rates in 2000.

Supplementary Material

Not for Publication

B.1 Additional Results on Motivating Evidence

Total Length of Supply Chains

Figure B.1: Dynamics of Supply Chains Length

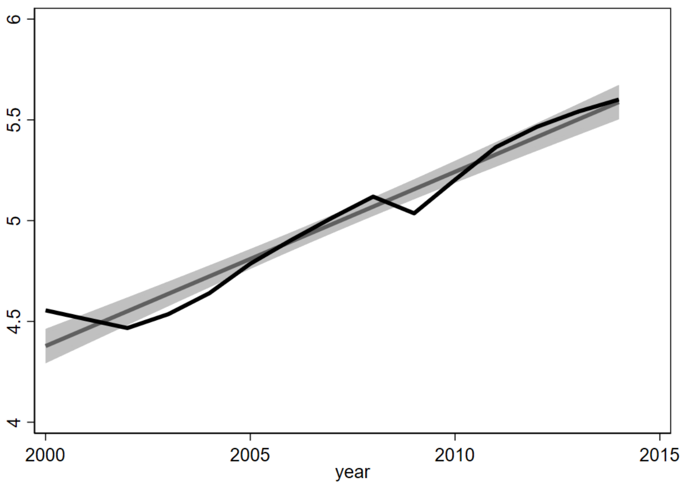

Note: Note: The figure shows the dynamics of the weighted length of chains measure computed as $L_t = \frac{\sum_i \sum_r y^r_{it} L^r_{it}}{\sum_i \sum_r y^r_{it}}$, here $L^r_{it} \coloneqq U^r_{it} + D^r_{it}$, namely the sum of upstreamness and downstreamness to count the total amount of steps embodied in a chain from pure value added to final consumption. The figure shows the average over time and it includes the estimated linear trend and the 95% confidence interval around the estimate.

Inventory Procyclicality by Inventory Type

Figure B.2: Distribution of estimated I' by inventory type

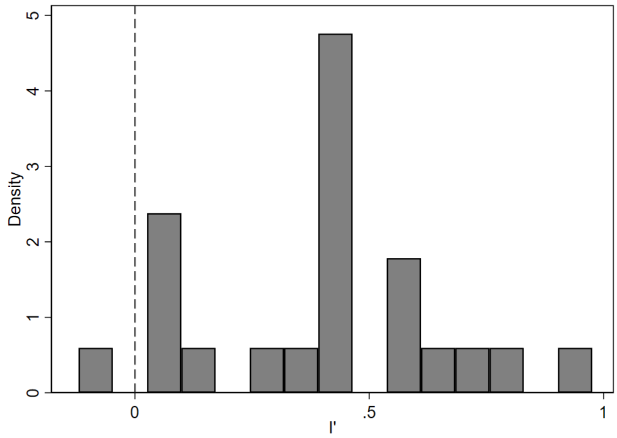

(a) Final Goods Inventories

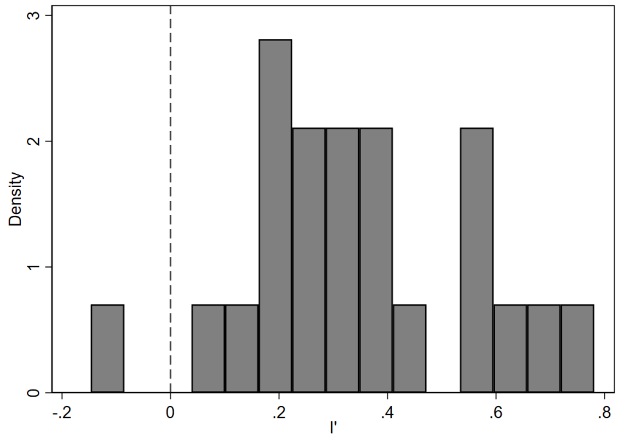

(b) Materials Inventories

Note: The graph shows the distribution of estimated $I'(\cdot)$, namely the derivative of the empirical inventory function with respect to sales. The left panel shows the same statistics based on the final goods monthly inventories data from the Manufacturing & Trade Inventories & Sales data of the US Census while the right panel shows the estimates using materials inventories.

B.2 WIOD Coverage

Table B.1: Countries

Country			
Australia	Denmark	Ireland	Poland
Austria	Spain	Italy	Portugal
Belgium	Estonia	Japan	Romania
Bulgaria	Finland	Republic of Korea	Russian Federation
Brazil	France	Lithuania	Slovakia
Canada	United Kingdom	Luxembourg	Slovenia
Switzerland	Greece	Latvia	Sweden
China	Croatia	Mexico	Turkey
Cyprus	Hungary	Malta	Taiwan
Czech Republic	Indonesia	Netherlands	United States
Germany	India	Norway	Rest of the World

Table B.2: Industries

Industry	Industry
Crop and animal production	Wholesale trade
Forestry and logging	Retail trade
Fishing and aquaculture	Land transport and transport via pipelines
Mining and quarrying	Water transport
Manufacture of food products	Air transport
Manufacture of textiles	Warehousing and support activities for transportation
Manufacture of wood and of products of wood and cork	Postal and courier activities
Manufacture of paper and paper products	Accommodation and food service activities
Printing and reproduction of recorded media	Publishing activities
Manufacture of coke and refined petroleum products	Motion picture
Manufacture of chemicals and chemical products	Telecommunications
Manufacture of basic pharmaceutical products and pharmaceutical preparations	Computer programming
Manufacture of rubber and plastic products	Financial service activities
Manufacture of other non-metallic mineral products	Insurance
Manufacture of basic metals	Activities auxiliary to financial services and insurance activities
Manufacture of fabricated metal products	Real estate activities
Manufacture of computer	Legal and accounting activities
Manufacture of electrical equipment	Architectural and engineering activities
Manufacture of machinery and equipment n.e.c.	Scientific research and development
Manufacture of motor vehicles	Advertising and market research
Manufacture of other transport equipment	Other professiona activitiesl
Manufacture of furniture	Administrative and support service activities
Repair and installation of machinery and equipment	Public administration and defence
Electricity	Education
Water collection	Human health and social work activities
Sewerage	Other service activities
Construction	Activities of households as employers
Wholesale and retail trade and repair of motor vehicles and motorcycles	Activities of extraterritorial organizations and bodies

B.3 Inventory Adjustment

Antràs et al. (2012) define the measure of upstreamness based on the Input-Output tables. This measure implicitly assumes the contemporaneity between production and use of output. This is often not the case in empirical applications since firms may buy inputs and store them to use them in subsequent periods. This implies that, before computing the upstreamness measure, one has to correct for this possible time mismatch.

The WIOD data provides two categories of use for these instances: net changes in capital and net changes in inventories. These categories are treated like final consumption, meaning that the data reports which country but not which industry within that country absorbs this share of output.

The WIOD data reports as Z_{ijt}^{rs} the set of inputs used in t by sector s in country j from sector r in country i, independently of whether they were bought at t or in previous periods. Furthermore, output in the WIOD data includes the part that is stored, namely

$$Y_{it}^{r} = \sum_{s} \sum_{j} Z_{ijt}^{rs} + \sum_{j} F_{ijt}^{r} + \sum_{j} \Delta N_{ijt}^{r}. \tag{27}$$

As discussed above the variables reporting net changes in inventories and capital are not broken down by industry, i.e. the data contains ΔN_{ijt}^{r}, not ΔN_{ijt}^{rs}.

This characteristic of the data poses a set of problems, particularly when computing bilateral upstreamness. First and foremost, including net changes in inventories into the the final consumption variables may result in negative final consumption whenever the net change is negative and large. This cannot happen since it would imply that there are negative elements of the F vector when computing

$$U = \hat{Y}^{-1}[I - \mathcal{A}]^{-2}F.$$

However, simply removing the net changes from the F vector implies that the tables are no longer balanced, which is also problematic. By the definition of output in equation 27, it may be the case that the sum of inputs is larger than output. When this is the case $\sum_i \sum_r a_{ij}^{rs} > 1$, which is a necessary condition for the convergence result, as discussed in the Methodology section.

To solve these problems I apply the the inventory adjustment suggested by Antràs et al. (2012). It boils down to reducing output by the change of inventories. This procedure, however, assumes inventory use. In particular, as stated above, the data reports ΔN_{ijt}^{r} but not ΔN_{ijt}^{rs}. For this reason, the latter is imputed via a proportionality assumption. Namely, if sector s in country j uses half of the output that industry r in country i sells to country j for input usages, then half of the net changes in inventories will be assumed to have been used by industry s. Formally:

$$\Delta N_{ijt}^{rs} = \frac{Z_{ijt}^{rs}}{\sum_s Z_{ijt}^{rs}} \Delta N_{ijt}^{r}.$$

Given the inputed vector of ΔN_{ijt}^{rs}, the output of industries is corrected as

$$\tilde{Y}_{ijt}^{rs} = Y_{ijt}^{rs} - \Delta N_{ijt}^{rs}.$$

Finally, whenever necessary, Value Added is also adjusted so that the the columns of the I-O tables still sum to the corrected gross output.

These corrections ensure that the necessary conditions for the matrix convergence are always satisfied. I apply these corrections to compute network measures while I use output as reported when used as an outcome.

B.4 Additional Results

Figure B.3: Effect of Federal Spending Shock on Output Growth by Upstreamness Level

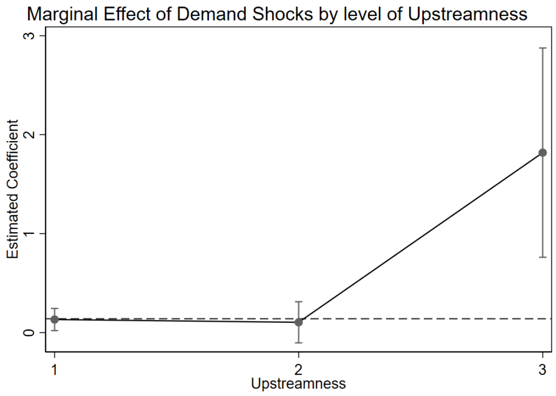

(a) Output Growth

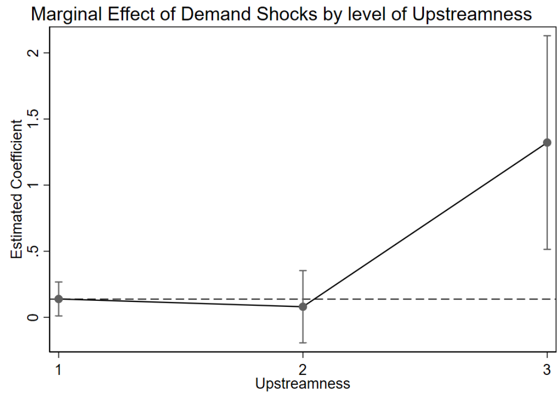

(b) Value Added Growth

Note: The figure shows the marginal effect of the federal spending shock on industry output changes by industry upstreamness level by the control function models. I use the changes in federal spending from Acemoglu et al. (2016) and apply the network transformation through the Leontief inverse, so that changes in federal spending are accounted for both directly and indirectly. The vertical bands illustrate the 95% confidence intervals around the estimates. The regression includes country-industry fixed effects and the standard errors are cluster bootstrapped at the country-industry level. The dotted horizontal line represents the average coefficient. Note that due to relatively few observations above 4, all values above it have been included in the $U \in (3, 4)$ category. The regression results are reported in Table A.5.

B.5 Model Extensions and Additional Theoretical Results

B.5.1 Heterogeneous Inventory Policies

I extend the model of section 5 to allow for heterogeneous inventory policies in this section. Denote $I_i \geq 0$ the inventory policy of sector $i \in \{0, \dots, N\}$.[31] The following generalization of Proposition 1 holds.

Proposition B.1 (Amplification with Heterogeneous Inventory Policies)
A vertically integrated economy with heterogeneous inventory policies features upstream amplification between sectors m and $n > m$ if $\exists\ k \in [m+1, n]$ such that $0 < I'_k$ and $\nexists j \in [m+1, n] : I'_j > \frac{1}{1-\rho}$.

Proof of Proposition B.1. I start by constructing the recursion that links the response of sector n to that of sector $m < n$. Starting with sector zero it is immediately evident that

$$\frac{\partial Y_t^0}{\partial D_t^0} = 1 + \rho I'_0.$$

[31]As shown in Figure A.3 this assumption is supported by the empirical evidence on the procyclicality of inventories.

Similarly

$$\frac{\partial Y_t^1}{\partial D_t^0} = \frac{\partial Y_t^0}{\partial D_t^0} + \rho I_1' \omega_0.$$

with $\omega_0 \equiv 1 + \rho I_0' - I_0'$. Following the recursion

$$\frac{\partial Y_t^n}{\partial D_t^0} = \frac{\partial Y_t^{n-1}}{\partial D_t^0} + \rho I_n' \prod_{i=0}^{n-1} \omega_i.$$

Substituting in

$$\frac{\partial Y_t^n}{\partial D_t^0} = 1 + \rho I_0' + \sum_{i=1}^{n} \rho I_i' \prod_{j=0}^{i-1} \omega_j.$$

Then

$$\begin{aligned}\frac{\partial Y_t^n}{\partial D_t^0} - \frac{\partial Y_t^m}{\partial D_t^0} &= 1 + \rho I_0' + \sum_{i=1}^{n} \rho I_i' \prod_{j=0}^{i-1} \omega_j - \left(1 + \rho I_0' + \sum_{i=1}^{m} \rho I_i' \prod_{j=0}^{i-1} \omega_j\right) \\ &= \sum_{i=m+1}^{n} \rho I_i' \prod_{j=0}^{i-1} \omega_j.\end{aligned}$$

Given the maintained assumptions that $\rho > 0$ and $I_i' \geq 0$, $\forall i$ and $\nexists j \in [m+1, n] : I_j' > \frac{1}{1-\rho}$ it follows that $\omega_j \geq 0$, $\forall j$. This immediately implies that $\prod_{j=0}^{i-1} \omega_j \geq 0$, $\forall i$. Further, $\exists\, k \in [m+1, n]$ such that $0 < I_k'$ implies $\omega_k > 0$, which in turn implies $\sum_{i=m+1}^{n} \rho I_i' \prod_{j=0}^{i-1} \omega_j > 0$. The statement follows. ■

The sufficient condition to observe amplification between two sectors is that at least one sector in between has to amplify shocks through inventories while no sector can dissipate them. This condition can be relaxed only by requiring that, while some sectors absorb shocks via countercyclical inventory adjustment, they do not so in such a way as to fully undo the upstream amplification of procyclical inventories.

B.5.2 A dynamic model of optimal procyclical inventories

In this section I show that optimally procyclical inventories obtain as the policy for a firm subject to production breakdowns. Consider a price taking firm facing some stochastic demand $q(A)$ where A follows some cdf Φ. The firm produces at marginal cost c and with probability $\chi > 0$ is unable to produce in a given period. The problem of the firm is described by the value functions for the "good" state where it can produce and the "bad" state where production is halted. The firm can store inventories I between periods. Inventories follow the law of motion $I' = I + y - q(A)$, where y is output and $q(A)$ is, by market clearing, total sales. Suppose further that firms do not face consecutive periods of halted production.

$$
\begin{aligned}
V^G(I,A) &= \max_{I'} pq(A) - cy + \beta \mathbb{E}_{A'|A}\left[\chi V^B(I',A') + (1-\chi)V^G(I',A')\right], \\
V^B(I,A) &= p\min\{q(A), I\} + \beta \mathbb{E}_{A'|A} V^G(I',A').
\end{aligned}
$$

The first order condition for next period inventories is then given by

$$
\frac{1-\beta(1-\chi)}{\beta\chi}c = \frac{\partial \mathbb{E}_{A'|A} V^B(I',A')}{\partial I'}.
$$

Note that the LHS is a positive constant and represents the marginal cost of producing more today relative to tomorrow. This is given by time discounting of the marginal cost payments, which the firm would prefer to backload. Note trivially that if the probability of halted production goes to zero the firm has no reason to hold inventories. The marginal benefit of holding inventories is given by relaxing the sales constraint in the bad state. Denote $P(A, I')$ the probability that the realization of A' implies a level of demand larger than the firm's inventories, which implies that the firm stocks-out. This probability depends on the current state since demand realizations are not independent. Denote $P_{I'}(A, I') = \partial P(A, I')/\partial I'$. Then

$$
\begin{aligned}
&\frac{\partial \mathbb{E}_{A'|A} V^B(I',A')}{\partial I'} = \\
&pP(A,I') + \beta P_{I'}(A,I')\mathbb{E}_{A''|A}\left[V^G(0,A') - V^G(I'-q(A'),A')\right] + \beta(1-P(A,I'))c > 0.
\end{aligned}
$$

This states that extra inventories in the bad state imply marginal revenues equal to the price in the event of stockout. The last two terms state that it makes it less likely that the firm will have to start next period without inventories and that it will be able to save on marginal cost for production if it does not stockout.

Note that it is immediate that the value of both problems is increasing in the level of inventories the firm starts the period with. It is also straightforward to see that if $\partial \mathbb{E}_{A'|A}/\partial A > 0$, namely if shocks are positively autocorrelated, then the expected value in the bad state is non-decreasing in A. The optimality condition for inventories shows that the LHS is a constant while the RHS increases in inventory holdings and decreases in the level of demand. Evaluating the first order condition at different levels of A, it has to be that $I'^{\star}(I, A_1) > I'^{\star}(I, A_2)$, $\forall A_1 > A_2$. In other words a the firm will respond to a positive demand shock by increasing output more than 1-to-1 as it updates inventories procyclically. The reason is that a positive shock today increases the conditional expectation on demand tomorrow. As a consequence the likelihood of a stock-out for a given level of inventories increases, which implies that the RHS of the first order condition increases as the benefit of an additional unit of inventories rises.

B.5.3 Production smoothing motive

In the main body of the paper I assume the inventory problem is defined by a quadratic loss function $(I_t - \alpha D_{t+1})^2$. This assumption is a stand-in for the costs of holding inventories or stocking-out. However it imposes two possibly unrealistic restrictions: i) it implies a symmetry between the cost of holding excess inventories and the cost of stocking-out; ii) it excludes any production smoothing motive as it implies an optimal constant target rule on expected future sales. In this section I extend the problem to eliminate these restrictions following Ramey and West (1999) more closely. Formally, consider the problem of a firm maximizing

$$\mathbb{E}_t \sum_t \beta^t \left[D_t^n - Y_t^n \left(c^n + \frac{\nu}{2} Y_t^n \right) - \frac{\delta}{2} (I_t^n - \alpha D_{t+1}^n)^2 - \zeta I_t^n \right] \quad st$$
$$I_t^n = I_{t-1}^n + Y_t^n - D_t^n,$$

Where the term $Y_t^n \left(c^n + \frac{\nu}{2} Y_t^n \right)$ includes a convex cost of production, which in turns generate a motive to smooth production across periods. The term ζI_t^n implies a cost of holding inventories which breaks the symmetry between holding excessive or too little inventories. In what follows I drop the stage and time indices and denote future periods by $'$. The first order condition with respect to end-of-the-period inventories implies

$$I = (\nu(1+\beta)+\delta)^{-1} \left[c(\beta-1) - \zeta + \delta\alpha \mathbb{E}D' - \nu(D - I_{-1}) + \nu\beta \mathbb{E}(D' + I') \right],$$

Define $\mathcal{B} := (\nu(1+\beta)+\delta)^{-1}$, taking a derivative with respect to current demand implies

$$\frac{\partial I}{\partial D} = \mathcal{B} \left(\delta\alpha\rho - \nu(1-\beta\rho) + \nu\beta \frac{\partial \mathbb{E} I'}{\partial D} \right).$$

Define $\mathcal{C} := \delta\alpha\rho + \nu(1-\beta\rho)$ then

$$\frac{\partial \mathbb{E} I'}{\partial D} = \mathcal{B} \left(\rho \mathcal{C} + \nu \frac{\partial I}{\partial D} + \nu\beta \frac{\partial \mathbb{E} I''}{\partial D} \right).$$

Iterating forward and substituting the following obtains

$$\frac{\partial I}{\partial D} = \left(1 - \sum_{i=1}^{\infty} \left(\mathcal{B}^2 \nu^2 \beta \right)^i \right)^{-1} \mathcal{B}\mathcal{C} \sum_{j=0}^{\infty} \left(\mathcal{B}\nu\rho\beta \right)^j.$$

If both $\mathcal{B}^2\nu^2\beta$ and $\mathcal{B}\nu\rho\beta$ are in the unit circle then

$$\frac{\partial I}{\partial D} = \frac{\mathcal{B}\mathcal{C}}{1 - \mathcal{B}\nu\beta\rho} \frac{1 - 2\mathcal{B}^2\nu^2\beta}{1 - \mathcal{B}^2\nu^2\beta} \lesseqgtr 0.$$

This states intuitively that if the production smoothing motive is strong enough then inventories respond countercyclically to changes in demand. This is immediate upon noting that when $\xi = 0$ then $\mathcal{B} = \delta^{-1}$, $\mathcal{C} = \delta\alpha\rho$ and therefore $\frac{\partial I}{\partial D} = \alpha\rho > 0$, while if $\nu > \mathcal{B}(\beta/2)^{1/2}$ then $\frac{\partial I}{\partial D} < 0$. Had the latter effect dominated then the empirical estimates of the response of inventories to changes in sales would be negative which is counterfactual given the findings discussed in Section 2.

B.5.4 Directed Acyclic Graphs Economies

The model derived in the section 5.2 applies to economies with general networks defined by the input requirement matrix A, a vector of input shares Γ and a vector of demand weights B. As discussed in the main body this economy features finite output under some regularity condition on the intensity of the inventory channel. I now restrict the set of possible networks to Directed Acyclic Graphs (DAGs) by making specific assumptions on A and Γ. This subset of networks feature no cycle between nodes.

Definition 2 (Directed Acyclic Graph)
A Directed Acyclic Graph is a directed graph such that $[A^n]_{rr} = 0, \forall r, n$.

The next trivial lemma provides a bound for the maximal length of a path in such a graph.

Lemma 3 (Longest Path in Directed Acyclic Graph)
In an economy with a finite number of sectors R, whose production network is a Directed Acyclic Graph, there exists an $N \leq R$ such that ${}^n\tilde{A}^n = [0]_{R\times R}, \forall n \geq N \wedge {}^n\tilde{A}^n \neq [0]_{R\times R}, \forall n < N$. Such N is the longest path in the network and is finite.

Proof. A path in a graph is a product of the form $\tilde{a}^{rs} \dots \tilde{a}^{uv} > 0$. A cycle in such graph is a path of the form $\tilde{a}^{rs} \dots \tilde{a}^{ur}$ (starts and ends in r). The assumption that there are no cycles in this graph implies that all sequences of the form $\tilde{a}^{rs} \dots \tilde{a}^{ur} = 0$ for any length of such sequence. Suppose that there is a finite number of industries R such that the matrix A is $R \times R$. Take a path of length $R + 1$ of the form $\tilde{a}^{rs} \dots \tilde{a}^{uv} > 0$, it must be that there exists a subpath taking the form $\tilde{a}^{rs} \dots \tilde{a}^{ur}$, which contradicts the assumption of no cycles. Hence the longest path in such graph can be at most be of length R. ■

With this result it is straightforward to show that output is finite even if $\tilde{\mathcal{A}}^n \sum_{i=0}^{n} \omega^i$ has a spectral radius outside the unit circle. If the network is a DAG with R sectors then output is given by

$$Y_k = \tilde{L}_k B D_t + \alpha\rho \left[\sum_{n=0}^{R} \tilde{\mathcal{A}}^n \sum_{i=0}^{n} \omega^i\right]_k B\Delta_t,$$

Which is naturally bounded since the second term is a bounded Neumann series of matrices.

B.6 Descriptive Statistics

This section provides additional descriptive statistics on the World Input-Output Database (WIOD) data.

B.6.1 Upstreamness

Table B.3: Highest and Lowest Upstreamness Industries

Industry	Upstreamness
Activities of extraterritorial organizations and bodies	1
Human health and social work activities	1.14
Activities of households as employers	1.16
Education	1.22
Public administration and defence	1.22
Accommodation and food service activities	1.66
⋮	⋮
Construction	3.96
Manufacture of wood and of products of wood and cork, except furniture	4.22
Manufacture of fabricated metal products, except machinery and equipment	4.27
Manufacture of machinery and equipment n.e.c.	4.28
Manufacture of other non-metallic mineral products	4.39
Mining and quarrying	4.52
Manufacture of basic metals	5.13

B.6.2 Destination Shares

The distribution of sales portfolio shares is computed as described in the methodology section. Table B.4 reports the summary statistics of the portfolio shares for all industries and all periods. Importantly, the distribution is very skewed and dominated by the domestic share. On average 61% of sales is consumed locally. Importantly, the median export share is .16% and the 99^{th} percentile is 12%. These statistics suggest that there is limited scope for diversification across destinations.

Table B.4: Portfolio Shares Summary Statistics

	count	mean	sd	min	max	p25	p50	p90	p95	p99
portfolio share	1522475	.0227	.1029	3.33e-13	.9999	.0004	.0017	.026	.0659	.7143
domestic portfolio share	34632	.6146	.2744	.0001	.9999	.4176	.6674	.9442	.9793	.9974
export portfolio share	1487843	.0089	.0273	3.33e-13	.962	.0003	.0016	.0199	.0418	.1224

Note: The table displays the summary statistics of the sales portfolio shares. Shares equal to 0 and 1 have been excluded. The latter have been excluded because they arise whenever an industry has 0 output. No industry has an actual share of 1.

B.6.3 Degree Distributions

After calculating the input requirement matrix A, whose elements are $a_{ij}^{rs} = Z_{ij}^{rs}/Y_j^s$. One can compute the industry level in and outdegree

$$indegree_i^r = \sum_i \sum_r a_{ij}^{rs}, \tag{28}$$

$$outdegree_i^r = \sum_j \sum_s a_{ij}^{rs}. \tag{29}$$

The indegree measures the fraction of gross output attributed to inputs (note that $indegree_i^r = 1 - va_i^r$ where va_i^r is the value added share).
The weighted outdegree is defined as the sum over all using industries of the fraction of gross output of industry r in country i customers that can be attributed to industry r in country i. This measure ranges between 0, if the sector does not supply any inputs to other industries, and $S * J$, which is the total number of industries in the economy, if industry r in country i is the sole supplier of all industries. In the data the average weighted outdegree is .52.
The distributions of these two measures are in Figure B.4.

Figure B.4: Degree Distributions

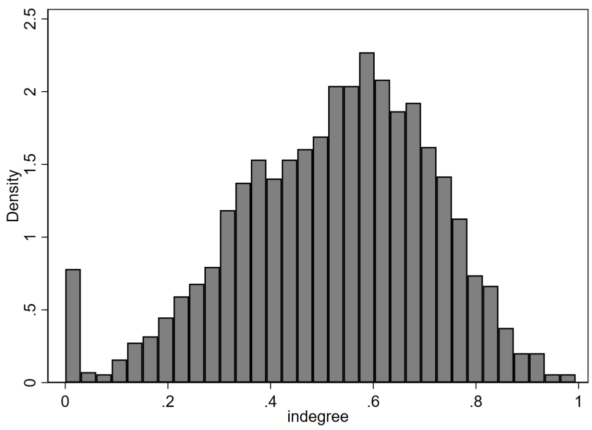

(a) Indegree

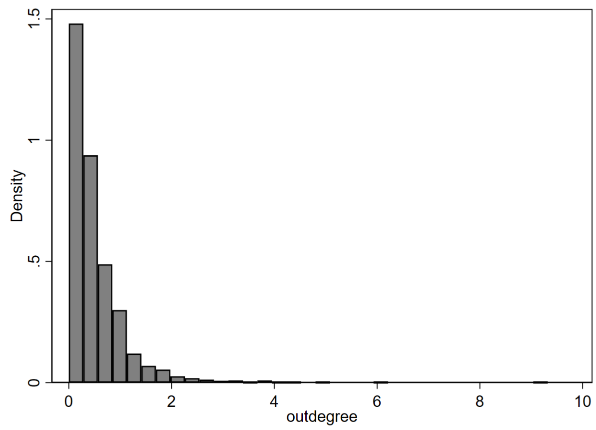

(b) Outdegree

Note: The figure depicts the distributions of the indegree and outdegree across all sectors and years in the WIOD data.

In the WIOD sample, industries' outdegree positively correlates with upstreamness, which suggests that industries higher in production chains serve a larger number (or a higher fraction) of downstream sectors. This relationship is shown in Figure B.5.

Figure B.5: Outdegree and Sales HHI

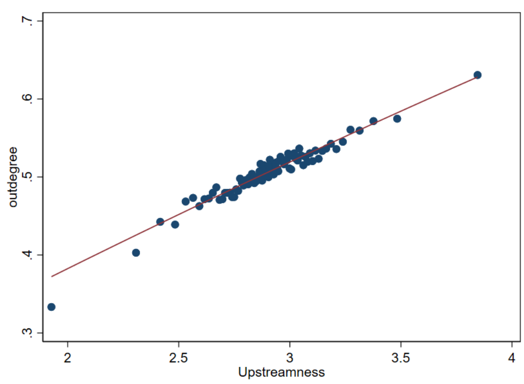

(a) Outdegree and Upstreamness

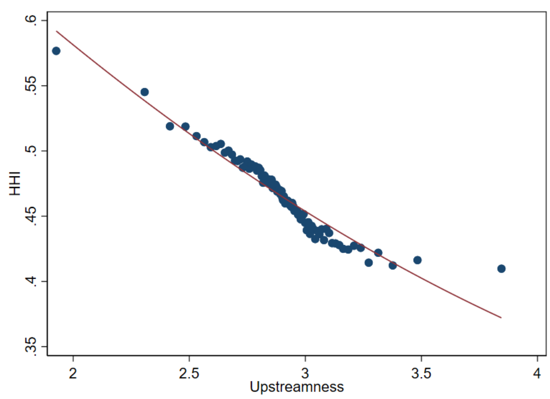

(b) HHI of sales and Upstreamness

Note: the figure plots the binscatter of industries' outdegree and upstreamness and of the Herfindahl-Hirschman Index of sales and upstreamness, controlling for country-industry fixed effects.

B.6.4 Demand Shocks and Output Growth Volatility by Upstreamness

Figure B.6: Demand Shocks and Output Growth Volatility by Upstreamness

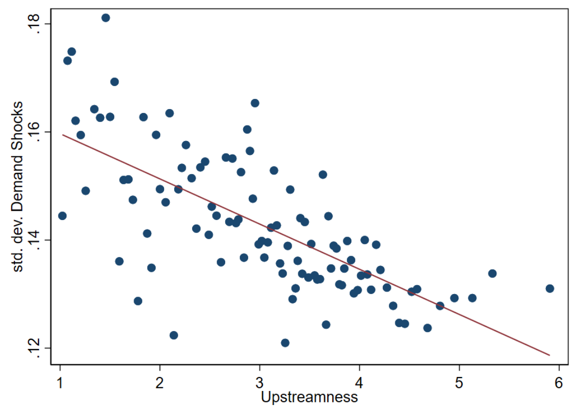

(a) Demand Shocks Volatility

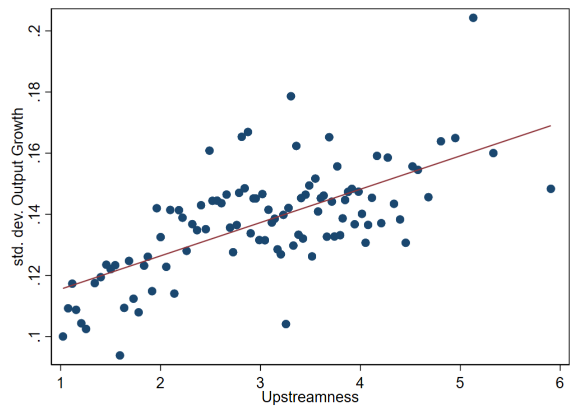

(b) Output Growth Volatility

Note: The graph shows the binscatter of the standard deviation of demand shocks and output growth within industry across time versus the industry average upstreamness across time.

B.6.5 Inventories

In the model presented in this paper the potential amplification is driven by procylical inventory adjustment. The WIOD data does not provide industry-specific inventory stock or change, eliminating the possibility of a direct test of the mechanism.

To provide partial evidence of the behavior of inventories I use NBER CES Manufacturing Industry data. This publicly available dataset covers 473 US manufacturing industries at the six-digit NAICS from 1958 to 2011. The data contains industry-specific information about sales and end-of-period inventories.

As mentioned in the main body of the paper, computing the parameter $\alpha \equiv I_t/\mathbb{E}_t D_{t+1}$ as $\alpha_t = I_t/D_{t+1}$ provides a set of numbers between 0 and 1, with an average of approximately 15%. Figure B.7 shows the distribution of α across all industries and years.

Figure B.7: Distribution of α

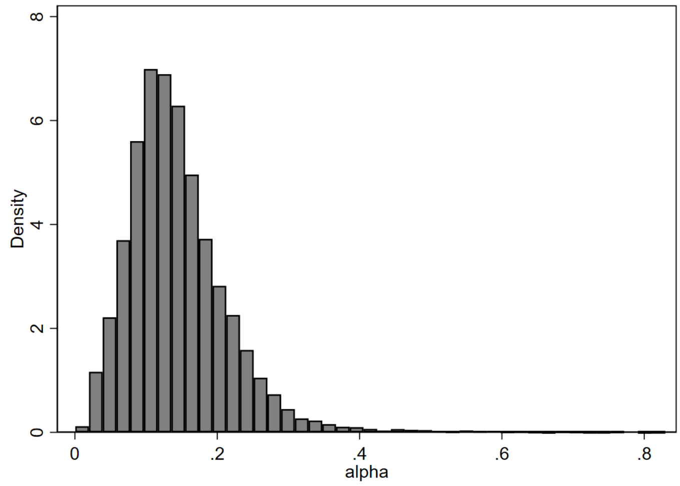

Note: The graph shows the distribution of $\alpha_t = I_t/Y_{t+1}$ across the 54 years and 473 industries in the NBER CES Manufacturing Industry data.

In the model the key assumption is that α is a constant across industries and time. This suggests that inventories are a linear function of sales. Figure B.8 shows the augmented component-plus-residual plot of the end-of-period stock of inventories as a function of current sales (the same picture arises for next-period sales). The underlying regression includes time and sector fixed effects. The graph is useful for detecting deviations from linearity in the relationship.

Figure B.8: Inventories and Sales

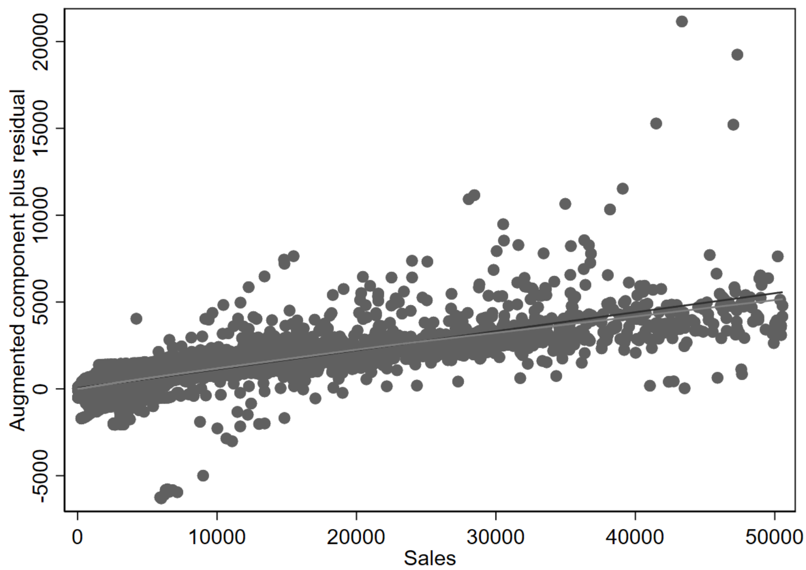

Note: The figure depicts the augmented component-plus-residual plot of the regression of inventories over sales, including time and industry fixed effects. The black line represents the linear fit of the model. The grey line is a local weighed smoothing fit. If the data presented significant deviations from linearity the two lines would be very different.

Figure B.8 suggests that linearity assumption is very close to the data. The function deviates from linearity only at high sales levels. This suggests that the inventory-to-sales ratio is mostly constant, other than in particularly high sales periods, when it starts to decline.

Table B.5 provides the correlation between sector position and inventory sales ratios. The two measures are positively correlated, which suggests that, in the data, more-upstream sectors tend to hold a larger fraction of future sales as inventories.

Table B.5: Inventories and Upstreamness

	(1)	(2)
	$\alpha^r_{i,t}$	$\alpha^r_{i,t}$
$U^r_{i,t}$	0.0121***	0.0241**
	(0.00355)	(0.0106)
Constant	0.0783***	0.0428
	(0.0108)	(0.0313)
Industry FE	No	Yes
N	210	210
R^2	0.0524	0.863

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: This table shows the results of the estimation of α against upstreamness. Column (1) reports the result of the OLS estimate while Column (2) includes industry fixed effects. change in sales.

B.7 Test of Uncorrelatedness of Instruments

As discussed in the main text, the identifying assumption for the validity of the shift share design is conditional independence of shocks and potential outcomes. Since this assumption cannot be tested, I provide evidence that the shares and the shocks are uncorrelated to alleviate endogeneity concerns. I test the conditional correlation by regressing the shares on future shocks and industry fixed effect. Formally

$$\xi^r_{ijt} = \beta\hat{\eta}_{jt+1} + \gamma^r_{it} + \epsilon^r_{ijt}.$$

This estimation results reported in Table B.6 suggest that the two are uncorrelated.

Table B.6: Test of Uncorrelatedness of Instruments

	ξ^r_{ijt}
$\hat{\eta}_{jt+1}(i)$	-0.0121
	(0.00762)
N	1517824
R^2	0.00284

Clustered standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B.8 Additional Empirical Results

To check that the empirical findings in Section 5.3 are robust to misspecification in the theoretical framework and the implied estimating equation. I estimate a saturated model with the interactions between the demand shocks, upstreamness and α as proxied by the inventory-to-sales ratio. I proceed in two alternative ways. In the first version I use a firm's inventory-to-sales ratio directly. In a second empirical model, I note that the model suggests that the inventory channel faced by industry r does not only depend on industry r's inventories, but also on the ones of all its downstream connected industries. To allow for this more flexible dependence I build an alternative measure of inventories along the chain. An empirical measure of this notion is given by $\underline{\tilde{\alpha}} := \tilde{L}\underline{\alpha}$, with $\tilde{L}$ being the Leontief inverse and $\underline{\alpha}$ being the vector of inventory-to-sales ratio. This allows me to use sectors whose inventories cannot be directly observed but that are connected to sectors whose inventories are. As a conservative approach, I assume that all industries whose inventories are not observed are zero. Finally, note that these two measures should not be directly compared since, by construction $\underline{\tilde{\alpha}} \geq \underline{\alpha}$, with equality in the limit case in which the sector does not belong to any production chain, while empirically $\underline{\tilde{\alpha}}$ takes values up to 10 times a for sectors with high $\tilde{\ell}$. The empirical model is given by

$$\Delta \ln Y_{it}^r = \beta_1 \hat{\eta}_{it}^r + \beta_2 U_i^r \times \hat{\eta}_{it}^r + \beta_3 \alpha_i^r \times \hat{\eta}_{it}^r + \beta_4 U_i^r \times \alpha_i^r \times \hat{\eta}_{it}^r + \epsilon_{it}^r, \tag{30}$$

Where the main coefficient of interest is β_4 and the theoretical model prediction is that it should be positive.[32] Table B.7 shows the results of the estimation with both inventory measures. The first two columns show the results for the direct measure of inventory-to-sales ratio. Columns 3 and 4 provide the estimates for the networked inventory measure $\underline{\tilde{\alpha}}$ while still keeping the same sample as the first two columns. Finally, the last two columns use the networked inventory measure on all industries. The key result on $\hat{\beta}_4$ is consistent with the model prediction of a positive interaction between inventories and the position in the production chain in amplifying shocks upstream. When using the direct measure of inventories I estimate $\hat{\beta}_2 = 0$, which suggests that all the positive effect from the position in the supply chain is driven by its interaction with inventories.. As a whole these estimates provide direct evidence of the inventory amplification channel, both based on the model estimating equation and on a reduced form specification. In the next section I use a simple calibration of the model to provide quantitative predictions on the volatility of the economy based on different counterfactuals for inventories and the network structure.

B.9 Robustness Checks

In this section I provide a set of robustness checks. First I apply the correction proposed by Borusyak and Hull (2020) to correct for potential omitted variable bias. Next I re-estimate the

[32] Note that the remaining interaction terms are subsumed in the fixed effects since they are industry-specific and time-invariant.

Table B.7: Reduced Form Estimation of the Role of Inventories and Upstreamness

	Inventories Manufacturing		Chain Inventories Manufacturing		Chain Inventories All Industries	
	(1)	(2)	(3)	(4)	(5)	(6)
	$\Delta \ln Y_{it}^r$	$\Delta \ln Y_{it}^r$	$\Delta \ln Y_{it}^r$	$\Delta \ln Y_{it}^r$	$\Delta \ln Y_{it}^r$	$\Delta \ln Y_{it}^r$
$\hat{\eta}_{it}^r$	0.589***	0.391***	0.431***	0.292***	0.367***	0.205***
	(0.149)	(0.130)	(0.0769)	(0.0659)	(0.0212)	(0.0188)
$U_i^r \times \hat{\eta}_{it}^r$	0.00506	-0.0156	0.0887***	0.0484**	0.0841***	0.0563***
	(0.0532)	(0.0451)	(0.0222)	(0.0188)	(0.00766)	(0.00675)
$\alpha_i^r \times \hat{\eta}_{it}^r$	-2.251*	-2.365**	-0.640	-1.213***	-0.271	-0.517***
	(1.318)	(1.178)	(0.417)	(0.378)	(0.198)	(0.172)
$U_i^r \times \alpha_i^r \times \hat{\eta}_{it}^r$	0.950**	0.886**	0.142	0.252***	0.105**	0.144***
	(0.475)	(0.407)	(0.102)	(0.0901)	(0.0527)	(0.0449)
Constant	0.0802***	0.0769***	0.0799***	0.0765***	0.0845***	0.0824***
	(0.000156)	(0.000226)	(0.000159)	(0.000234)	(0.0000698)	(0.0000917)
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	No	Yes	No	Yes	No	Yes
N	12098	12098	12098	12098	32371	32371
R^2	0.429	0.516	0.428	0.516	0.437	0.496

Clustered standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: This table shows the results of the regression in 30. Columns 1 and 2 show the results using the readily observed measure of inventories from the NBER-CES sample. Columns 3 and 4 use the network-based measure of inventories $\tilde{\alpha}$ but still restrict the sample to industries for which the direct measure is observed. Finally, columns 5 and 6 use the network-based measure of inventories $\tilde{\alpha}$ for the whole sample. All specifications include country-industry fixed effects while columns 2, 4 and 6 also include time fixed effects. Standard errors are cluster bootstrapped at the country-industry pair.

reduced form result after discretizing the network data to compare my estimates with existing ones in the recent literature. I also reproduce the main result in the ordinal, rather than cardinal, binning and under alternative fixed effects models to estimate the final demand shifters. I conclude by showing that the results are robust to estimation on deflated data, to account for potential price effects, using time varying aggregation shares and controlling for past output as suggested by Acemoglu et al. (2016).

Re-centered Instrument

In a recent paper Borusyak and Hull (2020) show that when using shift-share design there is a risk of omitted variable bias arising from potentially non-random shares. They also suggest to re-center the instrument to prevent such bias by using the average counterfactual shock.
I apply this methodology by permuting N=1000 times, within year, the distribution of destination shocks $\hat{\eta}_{jt}$. After the permutation I compute the average for each treated unit and demean the original demand shock to create $\tilde{\eta}_{it}^r = \hat{\eta}_{it}^r - \mu_{it}^r$. Where $\mu_{it}^r \equiv \frac{1}{N}\sum_n \sum_j \xi_{ij}^r \tilde{\eta}_{jt}$ and $\tilde{\eta}_{it}^r$

is the permuted shock. I then re-estimate the main specification in equation 9 with the re-centered shocks. The results, shown in Table B.8 are unchanged both qualitatively and quantitatively.

Table B.8: Re-centered Instrument Estimation

	(1) $\Delta \ln Y_{it}^{r}$	(2) $\Delta \ln Y_{it}^{r}$
Upstreamness in [1,2]	0.475***	0.473***
	(0.0125)	(0.0125)
Upstreamness in [2,3]	0.527***	0.525***
	(0.0127)	(0.0126)
Upstreamness in [3,4]	0.637***	0.635***
	(0.0115)	(0.0114)
Upstreamness in [4,5]	0.716***	0.714***
	(0.0209)	(0.0209)
Upstreamness in [5,6]	0.869***	0.868***
	(0.0635)	(0.0633)
Upstreamness in [6, ∞)	1.201***	1.204***
	(0.173)	(0.174)
Constant	0.0673***	0.0674***
	(0.0000819)	(0.0000824)
Country-Industry FE	Yes	Yes
N	31921	31921
R^2	0.364	0.364

Clustered standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The Table shows the results of the regression in equation 9 in column 1 and the re-centered instrument approach in Borusyak and Hull (2020). The latter is done by computing permutations of the shocks to demean the shift-share shock. Standard errors are clustered at the producing industry-country level.

Discrete Network

The empirical analysis in this paper is applied to a network in which the graph is weighted and directed. This is possible because I use aggregate data from Input-Output tables. More granular production network data often only allows to construct unweighted graphs. To compare my empirical results to previous work I therefore discretize my network so that I only observe connections to be ones or zeros and re-estimate my analysis.

Discretizing the network requires the choice of a cutoff $\tilde{a}$ such that, if the connection between industries i and j is given by the element a_{ij} of the input requirement matrix, the discrete connection is encoded as a 1 only if $a_{ij} \geq \tilde{a}$. Similarly, to fully compare my results to the ones of Carvalho et al. (2020), I also discretize the shock by transforming the destination specific changes as follows

$$\tilde{\eta}_{jt} = \begin{cases} -1 & \text{if } \eta_{jt} \leq \eta^* \\ 0 & \text{if } \eta_{jt} \in (\eta^*, \eta^{**}) \\ missing & \text{if } \eta_{jt} \geq \eta^{**} \end{cases} \tag{31}$$

with $\eta^* < 0 < \eta^{**}$. I do so because Carvalho et al. (2020) have a large negative shock, therefore this encoding allows me to consider as treated the destinations receiving a large negative shock. I use as control destinations with a shock around zero and I drop the ones with a large positive shock. The negative encoding is to preserve the sign of the shock, so that the result is directly comparable with the ones in Figure 2. I provide extensive robustness checks on both $\tilde{a}$ and the thresholds η^* and η^{**}. As shown in Figure B.9 I recover the result in Carvalho et al. (2020) such that these shocks have ever smaller effects when moving upstream in the production chain. Throughout I maintain $\eta^{**} = 0.05$.

A result consistent with the one in Section 4 would be a positive coefficient, increasing in upstreamness. Conversely, if dissipation forces were to dominate, we would expect a positive effect which is decreasing (towards zero) in upstreamness. As shown in Figure B.9 most of the discrete network estimates are decreasing in upstreamness. The takeaway from this robustness test is that discretizing (or observing only and unweighted version of) the network can significantly alter the conclusions on the how shocks propagate along production chains. I attribute the difference between the results in the weighted and unweighted graphs to the way measurement error moves with upstreamness.

Figure B.9: Effect of Demand Shocks on Output Growth by Upstreamness Level - Discrete Network

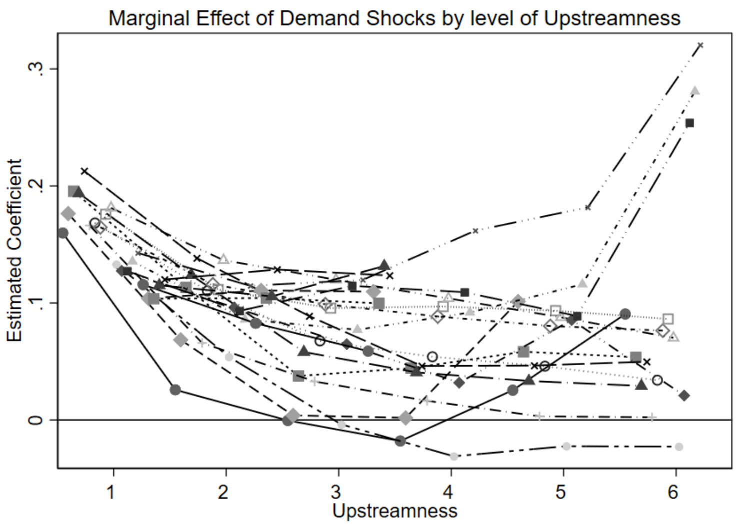

Note: The figure shows the marginal effect of demand shocks on industry output changes by industry upstreamness level in the control function models. Every line represents the result of a regression with thresholds $\tilde{\alpha} = \{.01, .05, .1, .2\}, \eta^* = \{-0.2, -.15, -0.1, -0.05, 0\}, \eta^{**} = .05$. Note that due to relatively few observations above 6, all values above it have been included in the $U \in [6, 7]$ category.

To further inspect the comparison, I make use the model to test the limitations of using discrete networks. I simulate the calibrated model and apply the same discretization procedure. I estimate eq. 9 on the true data, the discretized network with continuous shocks and the the discretized network with dichotomous shocks and report the resulting graphs in Figure B.10. The left panel represents the estimation on the economy without inventories and the right panel the one with inventories. The economy with $\alpha > 0$ has built-in upstream amplification, while the one with $\alpha = 0$ should have no gradient.
Estimating 9 on the economy with amplification successfully detects the increasing responsiveness of upstream industries as evidenced by the blue line on the right panel. The same estimation procedure on the discretized network does not estimate any positive gradient in the output response upstream. The conclusion from this exercise is that if the econometrician only observes a discrete version of both networks and shocks the specification can fail at detecting upstream amplification.
A similar result obtains by estimating the interaction version of this empirical model:

$$\Delta \log Y_{it}^r = \beta_0 + \beta_1 \hat{\eta}_{it}^r + \beta_2 U_{it-1}^r \hat{\eta}_{it}^r + \epsilon_{it}^r$$

Figure B.11 shows the distribution of 2500 estimates of $\hat{\beta}_2$ in generated data when $\alpha = \{.3, 0\}$, standing for yearly inventories and no inventories. Each figure shows the estimates for the three cases discussed above, namely when the econometrician observes a continuous network and a continuous shock, a discrete network and a continuous shock, and finally when only discrete network and discrete shocks are observable. The discretization is carried out as discussed above. The empirical model applied to discrete data does considerably worse at detecting even strong upstream amplification. When applying this specification to a discrete graph the estimated distribution moves leftward as α increases, suggesting that this form of amplification is even less likely to be detected when the inventory channel is strong.

Figure B.10: Model Based Estimates of Continuous and Discrete Graph

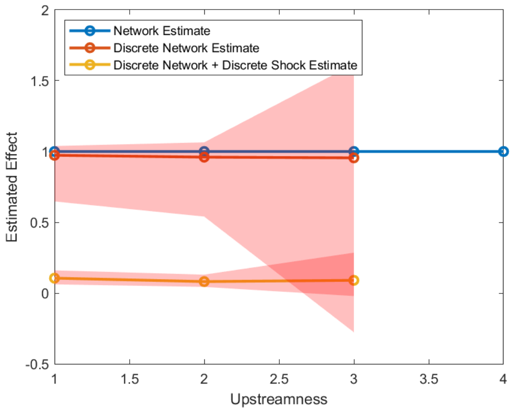

(a) $\alpha = 0$

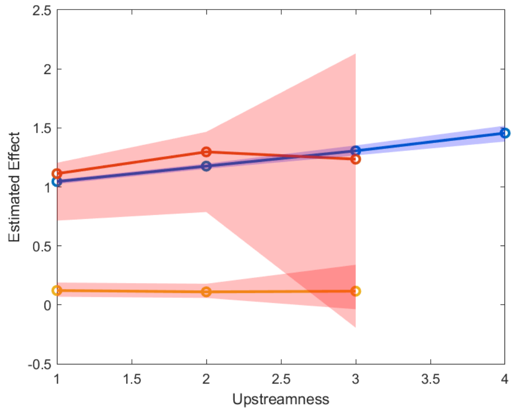

(b) $\alpha = .3$

Note: The figures show the estimates of the empirical model in equation 9. Panel (a) shows the estimates of an economy without inventories, while Panel (b) shows the estimates for an economy with inventory to sales ratio of 30%. The blue lines describe the estimates on the true network and shock data while the red line represent the estimates on a discrete network with the true shocks. Finally the yellow line is the estimate for a discrete network with dichotomous shocks. The bands in both cases represent the min-max range of estimates across the 2500 simulations.

Figure B.11: Model Based Estimates of Continuous and Discrete Graph

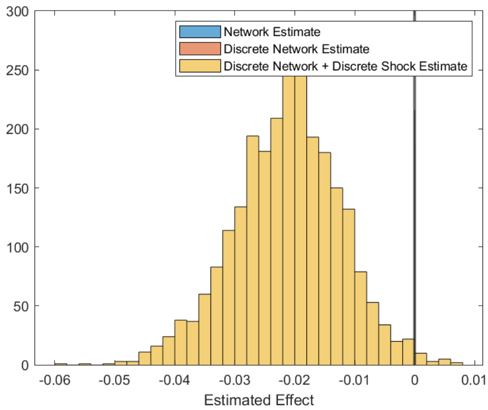

(a) $\alpha = 0$

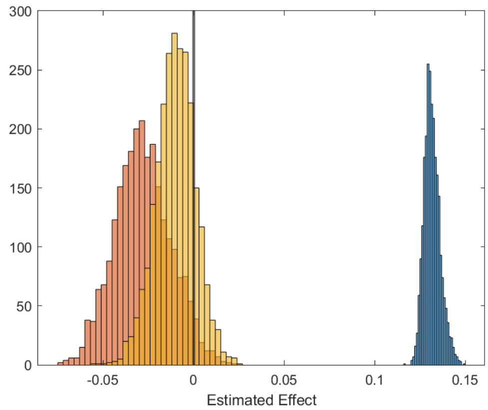

(b) $\alpha = .3$

Note: Note: The figures show the distribution of estimates of $\hat{\gamma}$ from the regression $\Delta \log Y_{it} = \beta\eta_{it} + \gamma U_i \eta_{it}$ on model generated data. Panel (a) shows the estimates of an economy without inventories, while Panel (b) shows the estimates for an economy with inventory to sales ratio of 30%. The blue lines describe the estimates on the true network and shock data while the red line represent the estimates on a discrete network with the true shocks. Finally the yellow line is the estimate for a discrete network with dichotomous shocks. Note that in the left panel the distribution of estimates on the true data and the discrete network are degenerate at zero.

Downstreamness

The conceptual framework built in Section 5 suggests that upstreamness is the key determinant of the inventory amplification across industries. A natural test is to check that alternative measures of positions do not have the same ability to explain the observed variation. Following Antràs and Chor (2018) I compute the measure of downstreamness which counts the the average number of production stages embodies in a sector's output. Formally, following Fally

(2012) this is defined recursively as $D_i^r = 1 \sum_j \sum_s a_{ji}^{sr} D_j^s$. Intuitively the sum of upstreamness and downstreamness measures the full length of a supply chain from pure value added to final consumption. As a consequence this is not necessarily negatively correlated with upstreamness since more complex goods might feature high upstreamness and high downstreamness. To estimate whether downstreamness can account for part of the cross-sectional variation in output responses I split the distribution following the same steps leading to equation 9 and estimate the regression with both interactions of shocks, upstreamness and downstreamness. The result is displayed in Figure B.12. I use industries with downstreamness between 1 and 2 as the reference category. The estimation suggests that the inclusion of downstreamness does not change the conclusion on the positive gradient of output responses with upstreamness and that along the downstreamness distribution there is no significant difference in the estimated responses to demand shocks. The same conclusion holds when using the WIOD inventory change as the outcome.

Figure B.12: Effect of Demand Shocks on Output Growth and Inventory Changes by Upstreamness and Downstreamness Levels

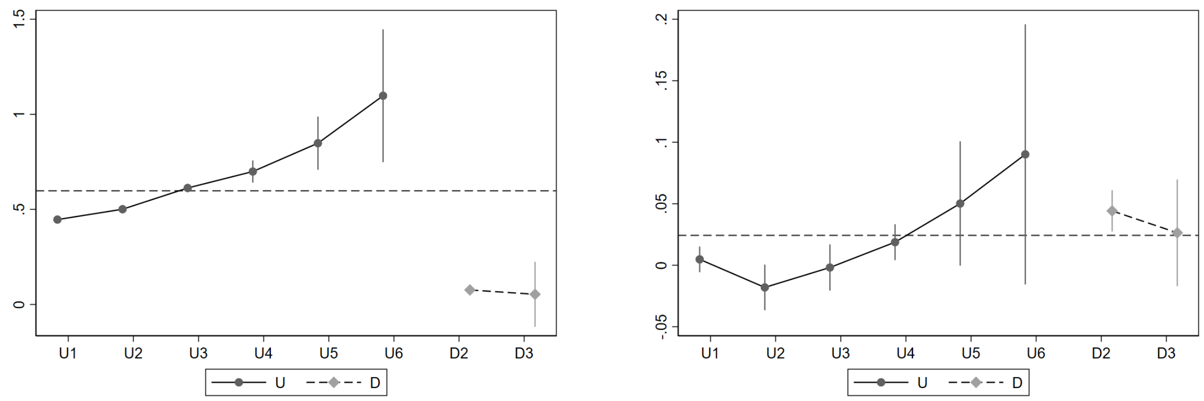

(a) Output Growth

(b) Changes in Inventories

Note: The figure shows the marginal effect of demand shocks on industry output growth and inventory changes by industry upstreamness and downstreamness levels. The left panel shows the estimation using the demand shocks described in section 3 on output growth while the right panel uses inventory changes as the outcome. The dashed horizontal line represent the average coefficient. The vertical bands illustrate the 95% confidence intervals around the estimates. The regression includes country-industry fixed effects and the standard errors are cluster bootstrapped at the country-industry level. Note that due to relatively few observations above 7, all values above 7 have been included in the $U \in [6, 7]$ category and similarly for the category $D \in [3, 4]$.

Ordinal Effects of Upstreamness

The results presented in section 4 are based on an split of the sample into industries whose upstreamness is between 1 and 2, 2 and 3, and so forth. To confirm that this sample split is not driving the results, I estimate a similar model to the main specification in section 4.1 using ordinal measures from the upstreamness distribution. Namely, I interact the industry-level shocks with dummies taking value 1 if an industry belongs to an upstreamness decile. Formally the estimated model is

$$\Delta \ln(Y_{it}^r) = \sum_j \beta_j \mathbb{1}\{U_{it-1}^r \in Q_j\}\hat{\eta}_{it}^r + \nu_{it}^r \quad , j = \{1...10\}, \tag{32}$$

where Q_j denotes the j^{th} deciles of the upstreamness distribution. The results are shown in Table B.9 in the Appendix. The estimation suggests that moving upstream in production chains increases the responsiveness of output to final demand shocks. The effect increases by 80% when moving from the first to the last decile. This corresponds to moving from 1.17 to 4.37 production stages away from final demand.

As in the main specification, the results suggest that the output response to demand shocks increases with distance from consumption. Ordinally the estimation states that moving from the first to the last decile of the distribution implies an increase in the output response from .47 to .81 of a percentage point. Note that all the results in this section are robust to the inclusion of industry, country, and upstreamness decile fixed effects.

Table B.9: Effect of Demand Shocks on Output Growth by Upstreamness Decile

	(1)	(2)
	$\Delta \ln Y_{it}^r$	$\Delta \ln Y_{it}^r$
decile 1	0.461***	0.209***
	(0.0303)	(0.0260)
decile 2	0.504***	0.233***
	(0.0178)	(0.0181)
decile 3	0.534***	0.238***
	(0.0202)	(0.0201)
decile 4	0.551***	0.263***
	(0.0203)	(0.0200)
decile 5	0.559***	0.267***
	(0.0209)	(0.0203)
decile 6	0.633***	0.314***
	(0.0193)	(0.0200)
decile 7	0.649***	0.301***
	(0.0212)	(0.0221)
decile 8	0.689***	0.330***
	(0.0213)	(0.0231)
decile 9	0.710***	0.323***
	(0.0239)	(0.0231)
decile 10	0.830***	0.416***
	(0.0301)	(0.0284)
Constant	0.0708***	0.0679***
	(0.000108)	(0.000124)
Industry FE	Yes	Yes
Time FE	No	Yes
N	31921	31921
R^2	0.391	0.442

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The Table shows the results of the ordinal version of the regression in equation 9. In particular I estimate e different coefficient for each decile of the upstreamness distribution. Both regressions include producing industry-country fixed effects and columns 2 adds time fixed effects. Standard errors are clustered at the producing industry-country level.

Alternative Shifter Estimation

In section 3.2.4 I used the fixed effect model to gauge the idiosyncratic demand shocks. Such a model may confound other sources of variation such as supply shocks, along with the object of interest. To investigate this possibility I use an alternative econometric model to extract the demand shocks. Following Kramarz et al. (2020) more closely, I include producer fixed effects: γ_{it}^r is the fixed effect for the producing industry r in country i at time t, namely

$$\Delta f_{kjt}^s = \eta_{jt}(i,r) + \gamma_{jt}^s + \nu_{kjt}^s \quad k \neq i, s \neq r. \tag{33}$$

Where the conditions $k \neq i, s \neq r$ ensure that domestically produced goods used for final consumption are not included in the estimation and neither are the goods within the same sector. The result for the main specification (equation 9) is presented in Table B.10. The findings confirm the main results in Section 4.1 both qualitatively and quantitatively.

Table B.10: Effect of Demand shocks by level of Upstreamness - Alternative Shifter Estimation

	(1) $\Delta \ln Y_{it}^r$	(2) $\Delta \ln Y_{it}^r$	(3) $\Delta \ln Y_{it}^r$
Upstreamness in [1,2]	0.476***	0.135***	0.192***
	(0.0126)	(0.0151)	(0.0141)
Upstreamness in [2,3]	0.528***	0.160***	0.221***
	(0.0127)	(0.0165)	(0.0151)
Upstreamness in [3,4]	0.638***	0.209***	0.282***
	(0.0116)	(0.0174)	(0.0157)
Upstreamness in [4,5]	0.715***	0.231***	0.329***
	(0.0212)	(0.0237)	(0.0221)
Upstreamness in [5,6]	0.865***	0.363***	0.480***
	(0.0636)	(0.0512)	(0.0537)
Upstreamness in [6,∞)	1.213***	0.568***	0.786***
	(0.169)	(0.167)	(0.178)
Constant	0.0673***	0.0661***	0.0664***
	(0.0000816)	(0.0000788)	(0.0000781)
Time FE	No	Yes	Yes
Level FE	No	No	Yes
Country-Industry FE	Yes	Yes	Yes
N	31921	31921	31921
R^2	0.364	0.434	0.481

Clustered standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The Table shows the results of the regression in equation 9 under the alternative shifter estimation. All regressions include producing industry-country fixed effects and columns 2 and 3 progressively add time fixed effects and upstreamness bin fixed effects. Standard errors are clustered at the producing industry-country level.

Deflated Data

As an additional robustness check I use the deflated version of the WIOD dataset (see Los et al., 2014) and replicate the entire industry level empirical analysis to test whether price movements could possibly be responsible for the results discussed above. The results of this check are displayed in Table B.11 in the Appendix. The findings in section 4.1 are confirmed both qualitatively and quantitatively.

Table B.11: Effect of Demand Shocks on Output Growth by Upstreamness Level - Deflated Data

	(1) $\Delta \ln Y_{it}^r$	(2) $\Delta \ln Y_{it}^r$	(3) $\Delta \ln Y_{it}^r$
Upstreamness in [1,2]	0.606***	0.269***	0.264***
	(0.0191)	(0.0205)	(0.0205)
Upstreamness in [2,3]	0.717***	0.339***	0.332***
	(0.0229)	(0.0254)	(0.0253)
Upstreamness in [3,4]	0.852***	0.412***	0.404***
	(0.0198)	(0.0234)	(0.0234)
Upstreamness in [4,5]	0.960***	0.490***	0.486***
	(0.0308)	(0.0314)	(0.0315)
Upstreamness in [5,6]	1.114***	0.633***	0.635***
	(0.0794)	(0.0741)	(0.0757)
Upstreamness in [6,∞)	1.113***	0.846***	0.847***
	(0.285)	(0.253)	(0.256)
Constant	0.0733***	0.0763***	0.0764***
	(0.0000872)	(0.000132)	(0.000131)
Time FE	No	Yes	Yes
Level FE	No	No	Yes
Country-Industry FE	Yes	Yes	Yes
N	29809	29809	29809
R^2	0.281	0.336	0.351

Clustered standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The Table shows the results of the regression in equation 9 on the deflated version of the WIOD Data. All regressions include producing industry-country fixed effects and columns 2 and 3 progressively add time fixed effects and upstreamness bin fixed effects. Standard errors are clustered at the producing industry-country level.

Time Varying Sales Share Aggregation

In the main specification I use a shift-share instrument in which the shares are fixed using the base year of the sample network data. Here I report the results of estimating equation 9 using the following shocks: $\hat{\eta}_{it}^r = \sum_j \xi_{ijt-1}^r \hat{\eta}_{jt}(i,r)$, where I aggregate using lagged sales share. The results are reported in Table B.12 and confirm the main results both qualitatively and quantitatively.

Table B.12: Effect of Demand Shocks on Output Growth by Upstreamness Level - Time Varying Sales Shares

	(1) $\Delta \ln Y_{it}^r$	(2) $\Delta \ln Y_{it}^r$	(3) $\Delta \ln Y_{it}^r$
Upstreamness in [1,2]	0.478***	0.134***	0.196***
	(0.0127)	(0.0156)	(0.0143)
Upstreamness in [2,3]	0.536***	0.160***	0.229***
	(0.0129)	(0.0171)	(0.0154)
Upstreamness in [3,5]	0.655***	0.210***	0.291***
	(0.0117)	(0.0186)	(0.0165)
Upstreamness in [4,6]	0.736***	0.234***	0.343***
	(0.0215)	(0.0250)	(0.0230)
Upstreamness in [5,6]	0.887***	0.360***	0.491***
	(0.0650)	(0.0543)	(0.0565)
Upstreamness in [6,∞)	1.204***	0.535***	0.802***
	(0.173)	(0.164)	(0.180)
Constant	0.0673***	0.0661***	0.0664***
	(0.0000836)	(0.0000837)	(0.0000823)
Time FE	No	Yes	Yes
Level FE	No	No	Yes
Country-Industry FE	Yes	Yes	Yes
N	31921	31921	31921
R^2	0.365	0.434	0.481

Clustered standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The Table shows the results of the regression in equation 9 using time-varying aggregation via sales share. All regressions include producing industry-country fixed effects and columns 2 and 3 progressively add time fixed effects and upstreamness bin fixed effects. Standard errors are clustered at the producing industry-country level.

Past Output

Finally, as discussed in previous work studying the effect of demand shocks and their propagation in the network (see Acemoglu et al., 2016), I include lags of the output growth rate. The results of the estimation are shown in Table B.13. This robustness check confirms the results of the main estimation both qualitatively and quantitatively.

Table B.13: Effect of Demand Shocks on Output Growth by Upstreamness Level - Output Growth Lags

	(1)	(2)	(3)	(4)	(5)
	$\Delta \ln Y_{it}^r$	$\Delta \ln Y_{it}^r$	$\Delta \ln Y_{it}^r$	$\Delta \ln Y_{it}^r$	$\Delta \ln Y_{it}^r$
Upstreamness in [1,2]	0.475***	0.488***	0.482***	0.456***	0.489***
	(0.0147)	(0.0156)	(0.0154)	(0.0175)	(0.0192)
Upstreamness in [2,3]	0.527***	0.544***	0.555***	0.550***	0.592***
	(0.0168)	(0.0174)	(0.0172)	(0.0212)	(0.0238)
Upstreamness in [3,4]	0.637***	0.653***	0.670***	0.671***	0.739***
	(0.0151)	(0.0154)	(0.0168)	(0.0213)	(0.0238)
Upstreamness in [4,5]	0.716***	0.748***	0.775***	0.813***	0.886***
	(0.0279)	(0.0282)	(0.0287)	(0.0345)	(0.0414)
Upstreamness in [5,6]	0.869***	0.934***	0.949***	0.975***	1.079***
	(0.114)	(0.123)	(0.121)	(0.130)	(0.115)
Upstreamness in [6,∞)	1.201***	1.335***	1.413***	1.392***	1.415***
	(0.128)	(0.144)	(0.145)	(0.148)	(0.148)
$\Delta \ln Y_{it-1}^r$		-0.0637***	-0.0856***	-0.0431**	-0.0616***
		(0.00986)	(0.0122)	(0.0173)	(0.0199)
$\Delta \ln Y_{it-2}^r$			-0.0248**	-0.0422***	-0.0422***
			(0.0118)	(0.0106)	(0.0111)
$\Delta \ln Y_{it-3}^r$				0.00932	-0.00143
				(0.0100)	(0.0103)
$\Delta \ln Y_{it-4}^r$					-0.0298***
					(0.00904)
Constant	0.0673***	0.0860***	0.0963***	0.0969***	0.106***
	(0.000128)	(0.000803)	(0.00181)	(0.00304)	(0.00343)
Country-Industry FE	Yes	Yes	Yes	Yes	Yes
N	31921	29077	26390	23843	21503
R^2	0.364	0.415	0.459	0.446	0.451

Clustered standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: This table contains the results of the regression of industry output growth rates on demand shocks interacted with dummies taking value 1 if the upstreamness level of the industry is in a given interval, e.g. [1,2]. The first column of the table includes the first lag of the dependent variable, the other columns progressively add lags up $t-4$.