

The Bullwhip: Time-to-Build and Sectoral Fluctuations*

Yan Leng	Ernest Liu	Yifei Ren	Aleh Tsyvinski
<i>UT Austin</i>	<i>Princeton and NBER</i>	<i>Princeton</i>	<i>Yale and NBER</i>

First draft: August 2024

This draft: July 2025

Abstract

We develop a theory of sectoral fluctuations driven by the propagation of demand shocks along supply chains with heterogeneous time-to-build production. We solve the model in closed form. Downstream producers respond directly to current demand. Upstream producers, due to time-to-build delays, respond to anticipated future demand (“news shocks”). Consequently, hump-shaped demand shocks to downstream goods propagate and amplify along the supply chain, generating pronounced volatility in upstream sectors and creating the bullwhip effect. Empirically and quantitatively, we show that the bullwhip is significant across downstream sectors that are important for final consumption.

*We thank Pol Antras, Ariel Burstein, Vasco Carvalho, Chad Jones, Aubhik Khan, Pete Klenow, Jean-Paul L’Huillier, Ezra Oberfield, Stephen Redding, Alireza Tahbaz-Salehi, Mathieu Taschereau-Dumouchel, Julia K. Thomas, Christopher Tonetti, our discussant Glenn Maggerman, and seminar participants at Cambridge, EAGLS, Princeton, Chicago Fed, Imperial College London, Penn State, Ohio State, UIUC, UCLA, UVA, and NBER ITI for very helpful comments and suggestions. We are particularly grateful to George-Marios Angeletos for a number of insightful suggestions. Ernest Liu acknowledges the support from the Sloan Foundation and the NSF.

1 Introduction

When a wrangler flicks a bullwhip, the motion propagates along its length, growing stronger and faster until the tip, or cracker, moves so fast that it breaks the sound barrier. Fluctuations in supply chains exhibit a similar phenomenon, with demand shocks originating downstream amplifying significantly as they propagate upstream.

We develop a theory of sectoral fluctuations driven by the dynamic propagation of demand shocks in a production network with heterogeneous time-to-build production delays, and we solve the model in closed form. Downstream producers respond immediately to current demand changes, whereas upstream producers, due to time-to-build, react based on expectations of future demand. As a result, downstream demand shocks with a hump-shaped time profile get amplified responses along the sectoral supply chain and result in large fluctuations in upstream sectors creating the bullwhip effect. We show empirically and quantitatively that the amplification of demand shocks through time-to-build along supply chains—the bullwhip—is a central feature of the economy’s response to demand shocks. Metaphorically, demand shocks originating in downstream sectors are like the wrangler’s initial flick: small initial movements get amplified dramatically as they travel upstream, culminating in pronounced volatility among upstream sectors.

Our model builds on two canonical models: the time-to-build environment of [Kydland and Prescott \(1982\)](#) as a theory of aggregate fluctuations, and the neoclassical sectoral model of [Long and Plosser \(1983a\)](#), in which all inputs require exactly one period to build, in the presence of either aggregate or sectoral productivity shocks. Methodologically, we develop and solve in closed form a production network model with sectoral demand shocks that explicitly incorporates heterogeneous time-to-build frictions. Substantively, our central focus is on how sectoral demand shocks propagate through and are amplified within supply chains, generating the bullwhip effect. This provides a new theory of sectoral fluctuations driven by demand shocks interacting with heterogeneous time-to-build dynamics.

Our theory of sectoral fluctuations features two core elements. The first is sectoral demand shocks, arising from consumers’ time-varying and stochastic preferences for different final goods. The second is heterogeneous time-to-build delays throughout the supply chain. We analyze the interplay between demand shocks, heterogeneous delays, and the cross-sectoral linkages in the economy. We fully characterize this model in a complete information setting, when the nature of shocks are known, and then extend the characterization to an incomplete information setting, in which agents cannot distinguish shocks with different time profiles.

The primary difficulty in characterizing the effects of sector-specific demand shocks within a time-to-build environment is as follows. Producers must choose inputs in anticipation of future demands across all time horizons. The expectation of future demands themselves arise from

both the exogenous impulse response of consumer demand shocks—at all time horizons—and endogenous direct and indirect responses to the consumer shocks by other producers—again, at all time horizons. The challenge is that each producer must account simultaneously for multiple future horizons and for the endogenous actions of producers situated at various positions in the supply chain, each potentially facing distinct time-to-build delays. The complexity of the producers’ endogenous response emerges from the interaction between the network structure and the heterogeneity in time-to-build delays across inputs, as well as the temporal structure of sectoral demand shocks. Theoretically, our problem is a path-dependent stochastic optimal control problem where the value functions depend on the input decisions at *all* previous heterogeneous time-to-build delays, which cannot be summarized by a low-dimensional state variable.

Our first contribution is an analytical characterization of the solution to this model for general sectoral demand shock processes, expressed explicitly in terms of the model’s primitives: the demand shocks, the production network structure, and the heterogeneity in time-to-build delays. The central object of the analytical characterization is the impulse response of sectoral revenue to sectoral demand shocks affecting various parts of the supply chain. Specifically, Theorem 1 provides a closed form solution for each sector’s revenue at a given time as an explicit function of expected consumer demands across all future time horizons and across all different goods that directly or indirectly rely on that sector as an input supplier.

Formally, sectoral revenue at a given time is a summation across future time horizons, with each term comprising two components explicitly derived from economic primitives. The first component is the conditional expectations of future consumer demands across goods, reflecting the temporal evolution of the demand shock process. The second component consists of horizon-specific weights applied to these conditional expectations. These weights reflect the direct and indirect input-output dependence across all possible network paths, along which the time-to-build delays cumulate to the specific time horizon of the expected future demand. These weights adjust the conditional expectations according to the time-to-build structure, resulting in different impacts of demand across sectors at various time horizons. Unlike the Leontief inverse—which aggregates direct and indirect linkages equally across all time horizons—the weights in our framework explicitly distinguish interactions at different time horizons, precisely due to heterogeneity in time-to-build delays between suppliers and input-users. The sequence of horizon-specific weights, and hence the impulse-response of sectoral revenue, is derived in closed form in terms of the economy’s primitives: the demand shocks, the sectoral linkages, and the time-to-build structure.

We then specialize the demand shock process into two distinct stochastic components. The first component captures transient changes in consumer preferences, modeled as monotone-decay AR(1) shocks, where a demand innovation today is expected to diminish monotonically

over time. The second component captures slower-developing changes, such as those arising from habit formation, consumer fads, or delayed adoption processes. These are modeled as hump-shaped AR(2) shocks, where a current increase in demand initially leads to increased demand at future horizons before eventually declining.

Theorem 2 explicitly characterizes the impacts of hump-shaped and monotone-decay demand shock components. For the monotone-decay component, we apply a similar approach as in the proof for Theorem 1, decomposing the impact into direct and indirect network effects at different time horizons, where each round of the network effect is discounted by both the discount rate and the rate at which the shock dissipates. The challenge in Theorem 2 lies in deriving the effect of the hump-shaped component because the impact of a hump-shaped shock does not decay monotonically. We show that the impulse response to a hump-shaped shock follows a non-monotone time path, represented explicitly as the product of an arithmetically increasing term and an exponentially decaying term.

In order to illustrate the intuition, consider an increase in consumer demand for automobiles. How producers respond to this shock depends crucially on their positions along the automobile supply chain. Downstream producers—those that are close in distance to the shock, such as automakers and manufacturers of tires or glass windows—respond by ramping up production to meet current or near-future expected demand. In contrast, producers located further upstream—such as manufacturers of semi-conductors or steel alloy, which are indirectly used in auto manufacturing as intermediate components—face significant time-to-build delays. As a result, these upstream producers adjust their production levels primarily based on anticipated future demand rather than current demand.

When a demand shock is hump-shaped—meaning higher automobile demand today signals even stronger demand in the future—upstream producers have significantly amplified responses, creating the bullwhip effect. This happens because distant suppliers respond more strongly to anticipated future peaks in demand, compared to downstream producers that mainly respond to current or short-term demand. Conversely, when demand shocks monotonically decline over time, downstream producers are always more responsive than upstream producers, preventing any amplified response along the supply chain.


The position of a producer in the supply chain, and the distance between them, play an important role in our analysis. A key feature of our model is the notion of delay—how long it takes for a given sector’s value-added to reach a specific user of its input along the supply chain. With time-to-build, the distance from an input supplier to a producer translates to the relative delays of producing goods between producers. We show that the ratio between a producer’s revenue response to a hump-shaped demand shock downstream and the response to a monotone-decay shock is a formal notion of supply chain distance in our model. This ratio captures the weighted

average of delays over all possible direct and indirect routes between the input supplier and input user, where the weights are given by the share of the supplier's output sent along each route. The remaining theoretical challenge is to characterize the weights, which are themselves endogenous. Our closed form solution of the impulse responses of producers actions to the hump-shaped shocks allows us to determine these endogenous weights and thus determine the supply chain distance measure in closed form over the primitives of the economy. Our definition of supply chain distance from one sector to another extends the influential upstream measures by [Antràs et al. \(2012\)](#) and [Alfaro et al. \(2019\)](#).

We further extend our analysis to an incomplete information environment, where both heterogeneous delays and learning are present. Specifically, we consider an environment in which agents observe only the realized consumer demand but cannot distinguish whether shocks exhibit a monotone-decay or a hump-shaped time profile. Theorem 3 extends our analytical characterization of the impulse responses from Theorem 2 to include this additional filtering problem faced by the suppliers. We find that supply chains connected to downstream sectors experiencing predominantly hump-shaped shocks exhibit pronounced amplification of demand fluctuations.

Empirically, we use our bilateral supply chain distance measure to document the bullwhip effect across U.S. sectors. As motivational evidence, we show in the cross section that volatility increases upstream within supply chains for goods important to final consumption. Figure 1 illustrates this for automobiles: each circle represents a sector supplying the auto industry (sized by input intensity), with supply chain distance on the Y-axis and value-added growth volatility on the X-axis. The bullwhip effect is evident—more upstream sectors exhibit greater volatility. In Section 4.2, we show that this pattern generalizes across all final goods supply chains.

Figure 1. Upstreamness and volatility for the supply chain of automobiles



Our theory’s predictions extend beyond the cross-sectional observation that upstream sectors often exhibit higher volatility, an empirical regularity that could be driven by factors such as supplying multiple downstream industries (and thus aggregating multiple shocks) or differences in inventory policies. Instead, our theory provides closed form predictions regarding impulse-response dynamics: specifically, under what conditions do hump-shaped demand shocks originating in a downstream sector can generate amplified volatility in upstream sectors.

We empirically validate the predictions by examining impulse-response functions for upstream value-added following downstream shocks. In order to do this, we first isolate downstream innovations in sectoral value-added—which respond to demand shocks—and classify downstream sectors according to the relative importance of AR(2) versus AR(1) shocks using partial autocorrelation function (PACF). We then show how sectoral responses propagate through supply chains using the estimated innovations.

The observed heterogeneity in impulse-response functions provides a sharp test of our theory, aligning with four key theoretical predictions under incomplete information. First, downstream sectors dominated by AR(2) shocks exhibit a distinctly hump-shaped response in their own value-added following demand shocks. Second, the value-added response along the supply chain is also hump-shaped and initially amplified as shocks propagate upstream, generating pronounced volatility. Third, this amplification diminishes in sectors located further upstream. Fourth and most importantly, the downstream hump-shaped responses and upstream amplification effects are markedly stronger in sectors where AR(2) shocks dominate, while such patterns are weak or absent in sectors characterized primarily by AR(1) shocks.

In the final part of the paper we conduct a model-based exercise. We infer the empirical process of demand shocks using the model. We perform variance decompositions for supply chain associated with each downstream good. We find that on average, hump-shaped AR(2) shocks account for 47.6 percent of supply chain volatility in the US economy. The fraction ranges from 9.2 percent on the lower end (for the supply chain associated with “Coffee and Tea”) to 92.2 percent on the higher end (for the supply chain associated with “Navigational, Measuring, Electromedical, and Control Instruments”). In a counterfactual world where all demand shocks are hump-shaped—while holding overall demand volatility constant—supply chain volatility significantly increases, by up to 145.7 percent for sectors supplying transportation goods. Conversely, when all demand shocks decay monotonically, supply chain volatility declines substantially, by more than 69 percent across all sectors. Furthermore, employing the standard static production network model, which neglects time-to-build dynamics, significantly understates supply chain volatility attributable to downstream demand shocks, by over 50 percent in certain consumption-good sectors.

Our paper relates to the growing literature on production networks ([Carvalho, 2010](#); [Gabaix,](#)

2011; Acemoglu, Carvalho, Ozdaglar and Tahbaz-Salehi, 2012; Jones, 2011, 2013; Oberfield, 2018; Liu, 2019; Baqaee and Farhi, 2019, 2020; Bigio and La’O, 2020; Elliott et al., 2022; Taschereau-Dumouchel, 2020; Pellet and Tahbaz-Salehi, 2023; Kopytov et al., 2024; Liu and Ma, 2024; Liu and Tsyvinski, 2024; Nikolakoudis, 2024). In contrast to the existing literature, which includes static models (e.g., Acemoglu et al., 2012), steady-states of dynamic models with one-period time-to-build delays (e.g., Long and Plosser, 1983b), and transition dynamics with time-invariant consumer demand (e.g., Vom Lehn and Winberry, 2022; Liu and Tsyvinski, 2024), we provide a full characterization of shock propagation in a dynamic production network with time-varying consumer demand and arbitrary and heterogeneous time-to-build.

Our focus of time-to-build’s implications on industry dynamics is related to an influential literature that studies heterogeneity of fixed costs that firms face in their investment decisions (Khan and Thomas, 2008, 2013). Similar to that literature, the presence of time-to-build in a production network as the presence of the fixed costs leads to sophisticated industry dynamics and complex responses of the economy to shocks. Time-to-build delays represent periods during which inputs have been produced but not yet sold to consumers, or completed but not immediately contributing value to production. Empirically, these goods are captured as inventories (Antràs and Tubdenov, 2025). Unlike models with inventory management margins to hedge against shocks (Khan and Thomas, 2007), inventories in our model have specific shelf-lives. Our central theoretical results depend only on the first moments (expected values) of future demand and are thus invariant to higher-order moments, such as variance. Consequently, our theoretical characterizations remain valid even when analyzing unanticipated shocks where the variance of shocks approaches zero, thus eliminating any hedging role of inventories.

Two recent working papers, Bizzarri et al. (2025) and Schaal and Taschereau-Dumouchel (2025), offer complementary studies of how heterogeneous time-to-build affects the propagation and persistence of shocks in input-output networks. While their analysis focuses on productivity shocks, similar in spirit to Liu and Tsyvinski (2024), our study instead centers on demand shocks. Both papers highlight asynchronous and oscillatory dynamics, offering insights into the cyclical mechanisms and time-series properties of sectoral adjustment.

The bullwhip effect—the amplified upstream response to demand shocks—has been extensively studied in industrial engineering and operations management (Forrester, 1958; Lee et al., 1997, 2004), with a focus on individual firms’ production decisions and a body of empirical work consisting of case studies based on firms’ different positions within a supply chain. In contrast, the topic has received limited attention in economics. A recent exception is Ferrari (2025), who documents upstream amplification of foreign demand shocks and builds an inventory-based model in partial equilibrium, with fixed prices and perfectly elastic labor supply.¹ Our contribution is to

¹Unlike in standard inventory models (e.g., Khan and Thomas, 2007), inventories in Ferrari (2025) do not serve to

extend the standard time-to-build friction (à la [Kydland and Prescott, 1982](#)) to allow for pairwise heterogeneity in a general equilibrium production network setting, and our central results are on the dynamic impact of time-varying sectoral demand shocks. We show that the bullwhip effect may arise endogenously as upstream sectors respond to expected future demand due to the accumulation of time-to-build along the supply chain. We derive closed-form expressions for the network’s dynamic impulse responses to both current and anticipated demand shocks, and use them to generate and test empirical predictions that vary both across sectors and over time.

Our paper also relates closely to the literature studying how expectations about future economic fundamentals can generate economic fluctuations independently of current realizations of shocks. [Beaudry and Portier \(2006\)](#) show that news about future productivity leads to significant macroeconomic responses well before productivity itself changes. [Lorenzoni \(2009\)](#) develops a model where noisy signals about future productivity induce agents to adjust current consumption and employment, creating fluctuations without contemporaneous fundamental shocks. Our multi-sector general equilibrium framework relates to this literature by introducing heterogeneous time-to-build delays, and shows that upstream producers respond predominantly to anticipated future demand rather than current conditions. These delays generate a bullwhip effect, wherein hump-shaped demand shocks amplify volatility upstream. This mechanism aligns closely with a key insight highlighted in [Lorenzoni \(2011\)](#): real frictions, such as adjustment costs or production delays, are essential for anticipated shocks to drive present-day economic fluctuations. Empirically, our analysis provides detailed sector-level evidence consistent with the aggregate findings of [Blanchard et al. \(2013\)](#), who document the quantitative importance of expectation-driven shocks at the aggregate level. Specifically, we show how anticipation-driven sectoral demand shocks propagate through production networks and amplify due to heterogeneous delays.

2 Model

In this section we describe the environment and characterize the equilibrium. The model builds on [Long and Plosser \(1983a\)](#) and is a discrete-time dynamic production network with heterogeneous time-to-build, i.e., intermediate inputs are required to be supplied in advance. The main departures from [Long and Plosser \(1983a\)](#) are the introduction and analysis of sectoral demand shocks and of the heterogeneous time-to-build across inputs.

meet current demand—assumed to be always satisfied via just-in-time production—but instead impose a quadratic penalty when they deviate from next period’s demand. This reduced-form feature creates contemporaneous upstream amplification, but the analysis is silent on dynamic propagation of demand shocks.

2.1 Economic Environment

Preferences There is a representative consumer with preferences

$$V(\mathcal{S}_t) = \mathbb{E} \left[\sum_{s=t}^{\infty} \beta^{s-t} \left(\sum_{i=1}^N \theta_{is} \ln c_{is} - v(\bar{\ell}_s) \right) | \mathcal{S}_t \right] \quad (1)$$

where $\mathcal{S}_t$ is the set of state variables at the beginning of time t , c_{it} is consumption of good i , $\{\theta_{it}\}_{\geq 0}$ are the demand shocks, and $\bar{\ell}_t$ is total labor supplied.

Technology There are N sectors, each with constant-returns-to-scale production that uses labor and intermediate goods from other sectors. There is heterogeneous time-to-build across inputs: each input j must be supplied d_{ij} periods before production takes place in sector i , where $0 < d_{ij} < \infty$. The time-to-build is supplier-user specific; the mathematical formulation can be equivalently interpreted either as the shipping delay associated with delivering input j to producer i , or the duration required for input j to complete the production process for good i .

The production function of sector i is

$$y_{it} = z_{it} \ell_{it}^{\alpha_i} \prod_{j=1}^N m_{ij,t-d_{ij}}^{\omega_{ij}}, \quad \alpha_i + \sum_j \omega_{ij} = 1, \quad (2)$$

where z_{it} is sectoral productivity, ℓ_{it} is the labor input, and $m_{ij,t-d_{ij}}$ is the intermediate good j used in production of good i . The elasticities of intermediate inputs $\{\omega_{ij}\}$ define the production network $\Omega \equiv [\omega_{ij}]'$, where each column of Ω is an input user and each row is an input supplier.²

Market Clearing The goods and labor market clearing conditions are

$$y_{jt} = c_{jt} + \sum_{i=1}^N m_{ijt}, \quad \bar{\ell}_t = \sum_i \ell_{it}. \quad (3)$$

Sectoral Demand and Productivity Shocks We assume the economy is affected by sectoral demand shocks, as θ_{it} may fluctuate from period to period. For expositional simplicity, we assume that consumer demand is mean-reverting in the long run, i.e., $\lim_{s \rightarrow \infty} \mathbb{E}_t[\theta_{t+s}] = \bar{\theta}$ where $\bar{\theta}$ represents steady-state levels of consumer demand. The extension to time-varying long-run demand is conceptually straightforward but complicates exposition.

The economy is also affected by sectoral productivity shocks z_{it} .

Demand shocks are the key objects of this study. As we show below, most of our results are invariant to the process of productivity shocks.

²We assume that $I - \Omega$ is invertible, which is ensured when the production of any good eventually involves labor along the supply chain (i.e., for any i , $\sum_i \alpha_i [\Omega^k]_{ji} > 0$ some k).

2.2 Equilibrium

In order to analyze the competitive equilibrium, it is more transparent to study the social planner's problem. Because there are no distortions or externalities in the model—time-to-build is a technological feature of the economic environment and is also part of the social planner's problem—the competitive equilibrium is efficient: allocations are identical to the planner's solution, and equilibrium prices correspond to the Lagrange multipliers of the planner's problem up to a choice of numeraire (and this choice is inconsequential for our analysis). The concordance between the planner's Lagrange multipliers and equilibrium prices enables us to interpret the planner's solution as equilibrium objects involving prices, such as GDP, value-added, sales, and the Domar weight (i.e., sectoral size measured as the sales-to-GDP ratio). For expositional clarity, we analyze the planner's solution in the main text, and we formally demonstrate the mapping to a decentralized equilibrium in Appendix B.1.

The planner solves

$$V_t \left(\{m_{ij,t-q}\}_{q=1,\dots,d_{ij}} \right) = \max_{\{\ell_{it}, c_{it}, m_{ijt}, \bar{\ell}_t\}} \sum_i \theta_{it} \ln c_{it} - v(\bar{\ell}_t) + \beta \mathbb{E}_t \left[V_{t+1} \left(\{m_{ij,t+1-q}\}_{q=1,\dots,d_{ij}} \right) \right] \\ + w_t \left[\bar{\ell}_t - \sum_j \ell_{jt} \right] + \sum_j p_{jt} \left[z_{jt} \ell_{jt}^{\alpha_j} \prod_k m_{jk,t-d_{jk}}^{\omega_{jk}} - c_{jt} - \sum_i m_{ijt} \right],$$

where w_t is the Lagrange multiplier for the labor market clearing constraint, and p_{jt} is the Lagrange multiplier for the market clearing constraint of output from sector j .

We now discuss the difficulty of analyzing this problem. The key complication lies in the fact that this is an optimal control problem with endogenous state variables $\{m_{ij,t-q}\}_{q=1,\dots,d_{ij}}$ that are matrices of all of the past choices of the suppliers inputs at all horizons of the past time delays. That is, the optimal control problem is path-dependent through the dependence on the choices of the previous inputs at all previous delay horizons that cannot be summarized in a low-dimensional state variable.³ The key difficulty is that the delays at all horizons also affect all of the future value functions and represent the complicated endogenous responses of input suppliers to all future stochastic demand shocks at all time horizons. In Theorem 1, we show that the solution to this problem can be represented in a strikingly simple form.

The first-order conditions with respect to c_{jt} , ℓ_{jt} , and $\bar{\ell}_t$ are:

$$\{c_{jt}\} \quad \theta_{jt} = p_{jt} c_{jt} \tag{4}$$

$$\{\ell_{jt}\} \quad w_t \ell_{jt} = \alpha_j p_{jt} y_{jt} \tag{5}$$

$$\{\bar{\ell}_t\} \quad v'(\bar{\ell}_t) = w_t. \tag{6}$$

³The analysis of these problems, especially in stochastic environments, is in its nascency. See [Bayraktar and Keller \(2018\)](#) for a recent mathematical treatment of path-dependent Hamilton-Jacobi equations and [Boerma et al. \(2024\)](#) for the use of stochastic calculus of variations to provide a non-recursive approach.

The first-order condition with respect to $m_{jk,t-d_{jk}}$ and the envelope condition jointly imply

$$\{m_{jkt}\} \quad p_{kt}m_{jkt} = \omega_{jk}\beta^{d_{jk}}\mathbb{E}_t[p_{j,t+d_{jk}}y_{j,t+d_{jk}}] \quad (7)$$

As we show in Appendix B.1, the Lagrange multipliers correspond to prices at the time when input decisions are made in a decentralized competitive equilibrium when the numeraire is chosen such that the marginal utility of income in each period is normalized to one. We henceforth refer to p_{jt} as the price of good j and w_t as the wage rate.

Equation (4) shows that total expenditure on consumption of good j equals the concurrent demand shock θ_{jt} . Equation (5) shows that each producer's labor expenditure is α_j fraction of the revenue, where α_j is the value-added intensity in the sector. Equation (7) states that producer j 's expenditure on input k at time t , $p_{kt}m_{jkt}$, equals a fraction ω_{jk} of the producer's expected revenue at future time $t + d_{jk}$, discounted at rate β per period. The expectation in (7) encapsulates the measurability constraint in the planner's solution.

Let $\gamma_{jt} \equiv p_{jt}y_{jt}$ denote the sector j 's revenue at time t . This is an important object for our analysis, as we can express the equilibrium allocation in shares (i.e., the fraction of workers in each sector $\ell_{it}/\bar{\ell}_t$; the share of each sector j 's output used for consumption, c_{jt}/y_{jt} ; and the share used as production inputs, m_{ijt}/y_{jt}) all as functions of the demand shocks and the current and expected future revenues:

$$\frac{\ell_{jt}}{\bar{\ell}_t} = \frac{\alpha_j \gamma_{jt}}{\sum_i \alpha_i \gamma_{it}}, \quad \frac{c_{jt}}{y_{jt}} = \frac{\theta_{jt}}{\gamma_{jt}}, \quad \frac{m_{ijt}}{y_{jt}} = \omega_{ij} \beta^{d_{ij}} \frac{\mathbb{E}_t[\gamma_{i,t+d_{ij}}]}{\gamma_{jt}}. \quad (8)$$

The GDP, i.e., the total value-added in the economy, is $\sum_i \alpha_i \gamma_{it}$. The Domar weight, defined as the sectoral sales-to-GDP ratio, is

$$(\text{Domar weight}) \quad \zeta_{it} \equiv \frac{p_{it}y_{it}}{\sum_i w_t \ell_{it}} = \frac{\gamma_{it}}{\sum_i \alpha_i \gamma_{it}}.$$

We now solve for sectoral sales $\{\gamma_{it}\}$ as a function of expected future demand. First multiply the goods market clearing condition (3) by p_{jt} on both sides and then substitute out c_{jt} and m_{ijt} using (4) and (7), we obtain

$$\begin{aligned} \gamma_{jt} &= \theta_{jt} + \sum_{i=1}^N \beta^{d_{ij}} \omega_{ij} \mathbb{E}_t[\gamma_{i,t+d_{ij}}] \\ &= \theta_{jt} + \sum_{d=1}^{\infty} \sum_{i=1}^N \beta^d \omega_{ij} \mathbf{1}_{d_{ij}=d} \mathbb{E}_t[\gamma_{i,t+d}]. \end{aligned}$$

That is, the revenue of sector j at time t consists of the concurrent time- t consumer demand for good j and the demand from other producers. Because of the heterogeneous time-to-build input delays by d_{ij} periods, the expenditure on input j at time t by producer i is equal to the input cost share ω_{ij} times the expected revenue at time $t + d_{ij}$, discounted by $\beta^{d_{ij}}$.

We then re-write the sectoral revenue equation in vector form. Throughout the paper, bold-face letters represent matrices (uppercase) and column vectors (lowercase). Let the $N \times N$ matrix $\Omega_d \equiv [\omega_{ij} \mathbf{1}_{d_{ij}=d}]'$ denote the input cost shares for inputs with delay d , and let $\theta_t \equiv [\theta_{it}]$ be the column vector of demand shocks. The collection of matrices $\{\Omega_d\}_{d=1}^\infty$ defines the time-to-build structure over the production network. Note that each column of Ω_d represents an input user and each row represents an input supplier.

Using the matrix notation, the vector of revenue can be written as

$$\gamma_t = \theta_t + \sum_{d=1}^{\infty} \beta^d \Omega_d \mathbb{E}_t [\gamma_{t+d}]. \quad (9)$$

The theorem below solves for the sectoral revenue at time t as a function of the expected future demand across sectors.

Theorem 1. *Let Φ_s denote the set of positive sequences $\phi \equiv (\phi_1, \dots, \phi_n)$ that sum to s (i.e., $\sum_{k=1}^n \phi_k = s$ with $\phi_k > 0$ for all k). Sectoral revenue follows*

$$\gamma_t = \theta_t + \sum_{s=1}^{\infty} G_s \mathbb{E}_t [\theta_{t+s}], \quad \text{where } G_s \equiv \beta^s \sum_{\phi \in \Phi_s} \prod_{\phi_j \in \phi} \Omega_{\phi_j}. \quad (10)$$

Proof. See appendix A.1. □

The infinite series (10) always converge given that demand is mean-reverting in the long run.

The main difficulty that the theorem overcomes is to express the complicated interactions of the suppliers' endogenous responses at all future time horizons to demand shocks at all future time horizons. The theorem shows that these interactions can be represented in a strikingly simpler form as a sum of products of only two components over the primitives of the economy, $\mathbb{E}_t [\theta_{t+s}]$ and G_s . We now discuss the object G_s in detail.

The ij -th entry of the matrix G_s captures how expected consumer demand for good j at time $t + s$ affects the time- t revenue of sector i through *all possible paths of direct and indirect linkages* that has a cumulative delay of s periods. To interpret, consider writing out G_s explicitly for $s = 1, 2, 3$. There is one positive sequence that sums to $s = 1$: $\Phi_1 = \{(1)\}$. There are two positive sequences that sum to $s = 2$: $\Phi_2 = \{(1, 1), (2)\}$. There are four positive sequences that sum to $s = 3$: $\Phi_3 = \{(1, 1, 1), (1, 2), (2, 1), (3)\}$. Hence, the formula in (10) implies that $G_1 = \beta \Omega_1$, $G_2 = \beta^2 (\Omega_2 + \Omega_1^2)$, and $G_3 = \beta^3 (\Omega_3 + \Omega_2 \Omega_1 + \Omega_1 \Omega_2 + \Omega_1^3)$. The power of the matrices captures the number of walks (i.e., the number of edges) associated with a network path: Ω_1 , Ω_2 , and Ω_3 capture one-walks (i.e., direct connection between producer pairs) with one, two, and three-period delays, respectively. Ω_1^2 captures two-walks (i.e., paths with two edges that connect producer pairs) where each walk has one-period delay, $\Omega_2 \Omega_1$ is a two-walk where the first walk has two-period delay and the second walk has one-period delay, and so on.

The matrix G_s summarizes how the vector of expected demand at $t + s$ affects the vector of revenue at time t , through *all possible paths of direct and indirect linkages* that have cumulative delays of s periods. While the concurrent consumer demand θ_t contributes to a sector's revenue directly, the one-period ahead demand θ_{t+1} contributes to a sector's revenue only indirectly through other producers to which the focal sector supplies with one period delay (captured by the matrix Ω_1 in G_1). The contribution of demand two-period ahead (θ_{t+2}) to a focal sector i 's revenue (captured by G_2) is through both (a) producers to which the focal sector supplies with a delay of two periods (captured by the matrix Ω_2), and (b) any producer k that uses an input j with one period delay, and where sector j uses input i also with one period delay (captured by the matrix Ω_1^2). The contribution of θ_{t+3} is captured by enumerating all possible paths that have three periods of cumulative delays.

Theorem 1 states that the time- t revenue of a sector is the summation across contributions from expected demand of different consumer goods at different time in the future, linked by *all possible paths of direct and indirect linkages* in the production network with different delays.

In contrast to the existing literature on production networks, which includes static models (e.g., Acemoglu et al., 2012), steady-states of dynamic models with one-period time-to-build delays (e.g., Long and Plosser, 1983b), and transition dynamics with time-invariant consumer demand (e.g., Liu and Tsyvinski, 2024), Theorem 1 provides a full characterization that accounts for time-varying consumer demand in a dynamic production network model with arbitrary and heterogeneous time-to-build. We now specialize Theorem 1 to time-invariant demand and discuss how this special case generalizes known results in the literature to heterogeneous time-to-build.

Implication: Time-Invariant Demand with Heterogeneous Time-to-Build Specializing Theorem 1 to time-invariant demand $\theta_t = \bar{\theta}$, we can simplify the summation over G_s and obtain the following result.

Corollary 1. *When demand is time-invariant, sectoral revenues are*

$$\bar{\gamma} = \left(\mathbf{I} - \sum_{d=1}^{\infty} \beta^d \Omega_d \right)^{-1} \bar{\theta}.$$

Proof. See appendix A.2. □

We can contrast this result with existing literature. When all input delays are exactly one period ($\Omega = \Omega_1$), Corollary 1 recovers as a special case the steady-state result of Long and Plosser (1983a): $\bar{\gamma} = (\mathbf{I} - \beta\Omega)^{-1} \bar{\theta}$. Similarly, the static model of Acemoglu et al. (2012), which yields sectoral revenue $\bar{\gamma} = (\mathbf{I} - \Omega)^{-1} \bar{\theta}$, emerges as a special case of our model when we allow inputs to have zero delays.⁴

⁴For expositional clarity, our baseline model features a time-to-build of at least one period. We could extend the

Corollary 1 extends these results on time-invariant sectoral size to a production network with an arbitrary time-to-build structure. The corollary is a special case of Theorem 1, which fully characterizes the sectoral revenue away from steady-state, given any sequences of expected future consumer demand across different goods.

Implication: When Sectoral Demand Follows $ARMA(p, q)$ Theorem 1 shows that the equilibrium revenue in sector i at time t reflects the expected future demand $\mathbb{E}_t[\theta_{k,t+s}]$ at each time horizon in every other sector k , through each possible walks of direct and indirect input-output linkages that delivers output of i to k , accounting for the time-to-build associated with each walk.

We now provide a characterization of equilibrium revenue when sectoral demand follows a stationary and invertible $ARMA(p, q)$. The process determines the expected future demand $\mathbb{E}_t[\theta_{k,t+s}]$ under rational expectations. The analysis reveals an interesting parallel between (1) how a current demand shock affects the expected future demand through the autoregression of the demand shock process, and (2) how the time-to-build and the network structure determines the importance of demand at different horizons.

Specifically, we consider an $AR(\infty)$ representation of the ARMA shock process:⁵

$$\theta_{it} = \sum_{s=1}^{\infty} \delta_s \theta_{it-s} + e_{it}. \quad (11)$$

where e_t has mean zero and independent of past shocks, and $\{\delta_s\}$ are autoregressive coefficients.

In this setting, the response of expected future demand to current demand innovation is

$$\frac{\partial \mathbb{E}_t[\theta_{i,t+s}]}{\partial e_{it}} = \sum_{\phi \in \Phi_s} \prod_{\phi_j \in \phi} \delta_{\phi_j}, \quad (12)$$

where, as in Theorem 1, Φ_s is the set of positive sequences $\phi \equiv (\phi_1, \dots, \phi_n)$ that sum to s . That is, e_{it} 's impact on $\mathbb{E}_t[\theta_{i,t+1}]$ is δ_1 , its impact on $\mathbb{E}_t[\theta_{i,t+2}]$ is $\delta_1^2 + \delta_2$, and on $\mathbb{E}_t[\theta_{i,t+3}]$ is $(\delta_1^3 + \delta_2\delta_1 + \delta_1\delta_2 + \delta_3)$, and so on. In other words, the impact of a demand shock e_{it} on the expected future demand at time $t + s$ has a recursive structure: it is the sum of its impact on demand $\mathbb{E}_t[\theta_{i,t+k}]$ for all $0 \leq k \leq s$, and then the impact of $\mathbb{E}_t[\theta_{i,t+k}]$ on $\mathbb{E}_t[\theta_{i,t+s}]$ through the

baseline model to incorporate static inputs without time-to-build (i.e., $d_{ij} = 0$), thereby nesting the static production network model as a special case. In that case, the formula in Theorem 1 needs to be modified as

$$\gamma_t = \theta_t + (\mathbf{I} - \mathbf{\Omega}_0)^{-1} \sum_{s=1}^{\infty} \mathbf{G}_s \mathbb{E}_t[\theta_{t+s}], \quad \text{where } \mathbf{G}_s \equiv \beta^s \sum_{\phi \in \Phi_s} \prod_{\phi_j \in \phi} [\mathbf{\Omega}_{\phi_j} (\mathbf{I} - \mathbf{\Omega}_0)^{-1}],$$

where $\mathbf{\Omega}_0 \equiv [\omega_{ij} \mathbf{1}_{d_{ij}=0}]'$ captures the component of the input-output network with zero delay (proof is available upon request). Corollary 1 extends naturally to $\bar{\gamma} = (\mathbf{I} - \sum_{d=0}^{\infty} \beta^d \mathbf{\Omega}_d)^{-1} \bar{\theta}$.

⁵The representation exists for all stationary and invertible $ARMA(p, q)$ processes (see, e.g., Box et al., 2015).

autoregressive demand process.

Note in particular the parallel between equation (12) and the object G_s in Theorem 1. Both constructs enumerate all time gaps that sum to s . In G_s , the gap is due to the time-to-build delays; in equation (12), the gap is due to the autoregressive structure of demand.

Proposition 1. *When demand follows (11), the revenue response to a demand shock is*

$$\frac{\partial \gamma_t}{\partial e_t} = \mathbf{I} + \sum_{s=1}^{\infty} \beta^s \left(\sum_{\phi \in \Phi_s} \prod_{\phi_j \in \phi} \delta_{\phi_j} \right) \left(\sum_{\phi \in \Phi_s} \prod_{\phi_j \in \phi} \Omega_{\phi_j} \right),$$

where Φ_s is the set of positive, finite sequences $\phi \equiv (\phi_1, \dots, \phi_n)$ that sum to s .

Quantities and Aggregate Implications Theorem 1 analytically solves for sectoral revenue as a function of model primitives: the network structure, time-to-build, and the expected future demand. The focus of our paper is on sectoral dynamics over supply chains, and subsequent sections of this paper build on Theorem 1. For completeness, Proposition 2 characterizes equilibrium quantities and the aggregate implications. This characterization is the discrete-time analogue of Liu and Tsyvinski (2024).

Proposition 2. *Given the endogenous state variables $\{m_{ij,t-q}\}_{q=1,\dots,d_{ij}}$, exogenous productivities $\{z_{it}\}$, and the path of expected demand $\{\mathbb{E}_t[\theta_{t+s}]\}_{s=0}^{\infty}$, sectoral output $\{y_{it}\}$, and consumption $\{c_{it}\}$ at each time is*

$$y_{it} = z_{it} \left(\frac{\alpha_i \gamma_{it}}{\sum_k \alpha_k \gamma_{kt}} \bar{\ell}_t \right)^{\alpha_i} \prod_j m_{ij,t-d_{ij}}^{\omega_{ij}}, \quad c_{it} = \frac{\theta_{it}}{\gamma_{it}} y_{it}, \quad (13)$$

where γ_{it} is characterized by Theorem 1, and the total labor supplied $\bar{\ell}_t$ at time t satisfy:

$$v'(\bar{\ell}_t) \bar{\ell}_t = \sum_i \alpha_i \gamma_{it}. \quad (14)$$

The endogenous state variables evolve according to

$$m_{ijt} = \frac{\omega_{ij} \beta^{d_{ij}} \mathbb{E}_t[\gamma_{i,t+d_{ij}}]}{\gamma_{jt}} y_{it}. \quad (15)$$

The dynamics over the aggregate consumption index $c_t \equiv \prod_i c_{it}^{\theta_{it}}$ follows directly from (13).

Proof. See Appendix A.3. □

Inventories Our model features time-to-build delays in intermediate inputs, which implies periods during which inputs have been produced but not yet sold to consumers, or completed but not immediately contributing value to production. Empirically, inputs which have been produced but are not yet contributing value to production or sold to consumers are recorded as inventories. Unlike models with inventory management margins to hedge against shocks (Khan and

Thomas, 2007), inventories in our model have specific shelf-lives determined by the delay parameter d_{ij} . Nevertheless, our central theoretical results—Theorem 1 above and Theorem 2 in the next section—depend only on the first moments of future demand. Hence, our theoretical characterizations remain valid even in scenarios where the variance of shocks approaches zero, effectively eliminating any hedging role of inventories.

3 The Supply Chain Response to Demand Shocks

In this section, we analyze the equilibrium response to demand shocks along the supply chain. We begin by applying our characterization to a stochastic process for sectoral demand that has two components: a monotone-decay AR(1) component and a hump-shaped AR(2) component. This stochastic process, empirically motivated as shown later, defines the conditional expectation over future demand based on current information $\mathbb{E}_t[\theta_{t+s}]$. Specifically, we use Theorem 1 to analyze how different sectoral demand shocks and their implied paths of expected future demand interact with the heterogeneous time-to-build and shape the equilibrium response along the supply chain. Section 3.1 describes the demand shock process and derives the impulse responses. Section 3.2 discusses equilibrium implications. In Section 3.3 we characterize the model under incomplete information, i.e., when agents cannot distinguish between shocks with different time-profiles.

3.1 Two Components of Demand Shocks

We assume sectoral demand is mean-reverting and has two distinct components. The first component captures transient consumer preference changes, modeled as monotone-decay AR(1) shocks, where a demand innovation today is expected to monotonically diminish over time. The second component captures slower-developing changes, such as those arising from habit formation, consumer fads, or delayed adoption processes. These are modeled as hump-shaped AR(2) shocks, where a current increase in demand initially leads to further demand increases in future horizons before eventually declining. Specifically, sectoral demand θ_{it} follows:

$$\theta_{it} - \bar{\theta}_i = \rho (\theta_{it-1} - \bar{\theta}_i) + x_{it} + \epsilon_{it}, \quad (16)$$

$$x_{it} = \rho x_{i,t-1} + u_{it}. \quad (17)$$

The parameter $\rho < 1$ captures the rate at which sectoral demand mean-reverts to the steady-state or long-run level $\bar{\theta}_i$. There are two sources of demand shocks: ϵ_{it} and u_{it} , both are mean zero and mean independent over-time. These shocks may be correlated across sectors, and their variances may be time-varying (to ensure that sectoral demand stays non-negative). In our baseline specification, the two types of demand shocks share the same persistence parameter ρ . Appendix B.2

extends our results in this section to an environment with heterogeneous ρ 's for the two types of shocks.


The monotone-decay component of demand shocks is captured by ϵ_{it} , and the hump-shaped component is captured by u_{it} . In the absence of u_{it} , the monotone-decay shock ϵ_{it} is AR(1). In the absence of ϵ_{it} , the hump-shaped shock u_{it} is AR(2).

Consider the response of expected demand $\mathbb{E}_t [\theta_{i,t+s}]$ to an impulse (ϵ_{it} or u_{it}). We have:

$$\text{For all } s \geq 0 : \quad \underbrace{\frac{\partial \mathbb{E}_t [\theta_{i,t+s}]}{\partial \epsilon_{it}}}_{\text{response to a monotone-decay shock}} = \rho^s, \quad \underbrace{\frac{\partial \mathbb{E}_t [\theta_{i,t+s}]}{\partial u_{it}}}_{\text{response to a hump-shaped shock}} = (s+1) \rho^s. \quad (18)$$

The expected demand response to an ϵ_{it} shock declines monotonically and exponentially over time at rate ρ . In contrast, for any $\rho > 0.5$, the expected demand response to a u_{it} shock is hump-shaped: $(s+1) \rho^s u_{it}$ increases initially over time s , peaking at $\bar{s} \equiv \frac{2\rho-1}{1-\rho}$, and subsequently declines for $s > \bar{s}$. Figure 2 illustrates the impulse-response of demand following a hump-shaped shock at different persistence levels ρ . As shown, the impulse-response is hump-shaped, and greater persistence (larger ρ) results in longer-lasting effects with peaks occurring later in time.

Figure 2. The impulse response $\partial \mathbb{E}_t [\theta_{i,t+s}] / \partial u_{it}$ of future demand following a hump-shaped shock, for different levels of ρ



Impulse Responses under Complete Information We begin by assuming agents can separately observe the hump-shaped and monotone-decay components of demand. The conditional expectation is $\mathbb{E}_t [\cdot] \equiv \mathbb{E} [\cdot | \{\theta_{t-s}, \mathbf{x}_{t-s}\}_{s \geq 0}]$. Given the stochastic process of demand shocks, the information set consisting of the history of state variables $\{\theta_{t-s}, \mathbf{x}_{t-s}\}_{s \geq 0}$ is identical to the information set consisting of the history of shocks $\{\epsilon_{t-s}, \mathbf{u}_{t-s}\}_{s \geq 0}$, and we use the two information sets interchangeably.

We let $\tilde{\cdot}$ denote deviations from steady-state levels, e.g., $\tilde{\theta}_t \equiv \theta_t - \bar{\theta}$.⁶ We can then derive the characterization of conditional expectations as

$$\mathbb{E}_t \left[\tilde{\theta}_{t+s} \right] = \rho^s \left(\tilde{\theta}_t - \mathbf{x}_t \right) + (s+1) \rho^s \mathbf{x}_t. \quad (19)$$

On the right-hand side, the first term captures the monotone-decay component of demand deviation from the steady-state levels; this component decays exponentially at rate ρ . The second term captures the hump-shaped component: for $\rho > \frac{1}{2}$, $(s+1) \rho^s$ first increases and then decreases in s , peaking at time $s \equiv \lceil \frac{2\rho-1}{1-\rho} \rceil$ periods into the future.

Substituting the conditional expectation into Theorem 1, we obtain how the two types of demand shocks explicitly affect sectoral revenue. We first state and prove the result under general network and time-to-build structure and then provide additional interpretation through a series of examples.

Theorem 2. *Under complete information when agents separately observe the hump-shaped and monotone-decay components of demand shocks up to the current period, with demand characterized by equations (16) and (17), sectoral revenue at time t follows*

$$\tilde{\gamma}_t = \mathbf{G}_\infty^\epsilon \left(\tilde{\theta}_t - \mathbf{x}_t \right) + \mathbf{G}_\infty^x \mathbf{x}_t, \quad (20)$$

where $\mathbf{G}_\infty^\epsilon$ is the impact of demand's monotone-decay component on revenue, and $\mathbf{G}_\infty^x$ is the impact of demand's hump-shaped component, and

$$\mathbf{G}_\infty^\epsilon \equiv \left(\mathbf{I} - \sum_{d=1}^{\infty} (\rho\beta)^d \Omega_d \right)^{-1}, \quad \mathbf{G}_\infty^x \equiv \mathbf{G}_\infty^\epsilon \left(\sum_{d=1}^{\infty} d (\rho\beta)^d \Omega_d \right) \mathbf{G}_\infty^\epsilon + \mathbf{G}_\infty^\epsilon.$$

The impulse-response functions are:

$$\frac{\partial \tilde{\gamma}_{t+s}}{\partial \epsilon_t} = \rho^s \mathbf{G}_\infty^\epsilon, \quad \frac{\partial \tilde{\gamma}_{t+s}}{\partial \mathbf{u}_t} = s \rho^s \mathbf{G}_\infty^\epsilon + \rho^s \mathbf{G}_\infty^x.$$

Proof. See Appendix A.4. □

3.2 Implications of Theorem 2

We now discuss several implications of Theorem 2.

3.2.1 Upstream's Amplified Response to Expectations of Heightened Future Demand

We define the bullwhip effect as the phenomenon where a small downstream demand shock lead to amplified response in upstream sectors. Theorem 2 implies that the bullwhip effect does not occur when demand shocks decay monotonically, but it may arise when demand shocks are

⁶Since $\mathbf{x} = 0$ in steady state, we have $\mathbf{x}_t = \tilde{\mathbf{x}}_t$.

hump-shaped. Intuitively, the difference stems from producers' positions along the supply chain: those closer to the demand shock immediately increase production in response to current or near-future expected demand. On the other hand, because of time-to-build delays, upstream producers respond primarily to anticipated future demand. Therefore, when demand shocks monotonically diminish over time, upstream producers' responses are muted relative to downstream producers. In contrast, hump-shaped shocks generate anticipated future demand increases, eliciting an amplified response in upstream sectors and thereby creating a pronounced bullwhip effect.

Formally, the revenue response to a concurrent monotone-decay shock is

$$\partial \tilde{\gamma}_t / \partial \epsilon_t = \mathbf{I} + \rho\beta\Omega_1 + (\rho\beta)^2 (\Omega_1^2 + \Omega_2) + (\rho\beta)^3 (\Omega_1^3 + \Omega_1\Omega_2 + \Omega_2\Omega_1 + \Omega_3) + \dots \quad (21)$$

and the response to a concurrent hump-shaped shock is

$$\partial \tilde{\gamma}_t / \partial u_t = 1 \times \mathbf{I} + 2 \times \rho\beta\Omega_1 + 3 \times (\rho\beta)^2 (\Omega_1^2 + \Omega_2) + 4 \times (\rho\beta)^3 (\Omega_1^3 + \Omega_1\Omega_2 + \Omega_2\Omega_1 + \Omega_3) + \dots \quad (22)$$

Each successive term in these summations captures the revenue response due to demand in a subsequent period. For a monotone-decay shock, the weights associated with each term form an exponentially decaying sequence: $1, \rho\beta, (\rho\beta)^2, \dots$. The further distant a sector is from the shock, the lower is the impact of the shock on sectoral revenue. By contrast, the weights for a hump-shaped shock have a time profile $1, 2\rho\beta, 3(\rho\beta)^2, \dots$, resulting in the potentially amplified response in upstream sectors.

In the special case where each input has one period delay only, $\mathbf{G}_\infty^\epsilon = (\mathbf{I} - \rho\beta\Omega_1)^{-1}$ and $\mathbf{G}_\infty^x = (\mathbf{I} - \rho\beta\Omega_1)^{-2}$, and equation (22) becomes

$$\partial \tilde{\gamma}_t / \partial u_t = (\mathbf{I} - \rho\beta\Omega_1)^{-2} = \mathbf{I} + 2\rho\beta\Omega_1 + 3(\rho\beta\Omega_1)^2 + 4(\rho\beta\Omega_1)^3 + \dots$$

where the Leontief-inverse-squared is the matrix analogous to the scalar summation formula $\sum_{t=0}^{\infty} (t+1) \rho^t = (1-\rho)^{-2}$.

The key takeaway from the characterizations is that how producers respond to a downstream demand shock depends on their positions along the supply chain. Downstream producers respond directly to current demand, while upstream producers, due to time-to-build delays, respond to anticipated future demand. This insight explains why hump-shaped shocks can propagate upstream and become significantly amplified, generating pronounced volatility in upstream sectors.

3.2.2 Example: A Vertical Network

In order to gain further intuition, consider a vertical supply chain with $N = 4$ sectors. Sector 1 produces with labor only; each successive sector $i \geq 2$ produces using input $i-1$ with intermediate elasticity $\omega_{i,i-1} = 1$. Each good i takes d_i periods to build before being used by sector $(i+1)$ as an input. All goods may be consumed. We say that sector j is upstream to sector i if

$j < i$, as sector j 's output eventually becomes sector i 's input.

Given the vertical network structure, we can write the sectoral revenue response to a monotone-decay demand shock as

$$\frac{\partial \tilde{\gamma}_t}{\partial \epsilon_t} \times \epsilon_t = \begin{bmatrix} 1 & (\rho\beta)^{d_1} & (\rho\beta)^{d_1+d_2} & (\rho\beta)^{d_1+d_2+d_3} \\ 0 & 1 & (\rho\beta)^{d_2} & (\rho\beta)^{d_2+d_3} \\ 0 & 0 & 1 & (\rho\beta)^{d_3} \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \epsilon_1 \\ \epsilon_2 \\ \epsilon_3 \\ \epsilon_4 \end{bmatrix}$$

and to a hump-shaped demand shock as

$$\frac{\partial \tilde{\gamma}_t}{\partial \mathbf{u}_t} \times \mathbf{u}_t = \begin{bmatrix} 1 & (1+d_1)(\rho\beta)^{d_1} & (1+d_1+d_2)(\rho\beta)^{d_1+d_2} & (1+d_1+d_2+d_3)(\rho\beta)^{d_1+d_2+d_3} \\ 0 & 1 & (1+d_2)(\rho\beta)^{d_2} & (1+d_2+d_3)(\rho\beta)^{d_2+d_3} \\ 0 & 0 & 1 & (1+d_3)(\rho\beta)^{d_3} \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \\ u_3 \\ u_4 \end{bmatrix}$$

The sensitivity of sector j 's revenue with respect to a monotone-decay demand shock in sector $i \geq j$ is $(\rho\beta)^d$, where d is the total periods of delay for any value-added in sector j to be used as inputs in sector i . By contrast, when the demand shock is hump-shaped, the sensitivity is $(d+1)(\rho\beta)^d$ and is scaled by the additional $(d+1)$ factor.

Consider, for example, a monotone-decay shock ϵ_4 to the most downstream sector four and examine the last column of the matrix $\partial \tilde{\gamma}_t / \partial \epsilon_t$ for the sectoral revenue response. Sector four's revenue responds one-for-one to the shock. The response of sector three is dampened by $(\rho\beta)^{d_3} < 1$. The response of sector two is further dampened to $(\rho\beta)^{d_2+d_3}$. The further a producer is from the monotone-decay shock—further upstream—the less is revenue response to the shock.

Now instead consider a hump-shaped shock u_4 to the most downstream sector four, and we examine the last column of $\partial \tilde{\gamma}_t / \partial \mathbf{u}_t$ for sectoral revenue response. Sector four's revenue responds one-for-one to the shock. The response of sector three is $(1+d_3)(\rho\beta)^{d_3}$ and may be amplified. The response of sector two, which is further upstream to sector three, is $(1+d_2+d_3)(\rho\beta)^{d_2+d_3}$ and may be further amplified relative to sector three.

Note also that, for a sector sufficiently upstream, the response will become dampened and eventually converges to zero (as $\lim_{s \rightarrow \infty} (s+1)(\rho\beta)^s = 0$), as the distance between the sector and the final demand shock could become large enough so that, due to the long time-to-build, the sector does not even respond to hump-shaped demand shocks.

3.2.3 A Bilateral Measure of Supply Chain Delays

We now introduce a measure of the pairwise input delay between an input supplier and an input user. Consider sector i 's revenue response to sector j 's hump-shaped demand shock relative to

that of a monotone-decay shock and define

$$\xi_{ij} \equiv \begin{cases} \left[\frac{\partial \tilde{\gamma}_t}{\partial \mathbf{u}_t} \right]_{ij} / \left[\frac{\partial \tilde{\gamma}_t}{\partial \epsilon_t} \right]_{ij} - 1 & \text{if } \left[\frac{\partial \tilde{\gamma}_t}{\partial \epsilon_t} \right]_{ij} \neq 0, \\ 0 & \text{otherwise.} \end{cases} \quad (23)$$

ξ_{ij} is the element-wise ratio between the matrices capturing the response to hump-shaped and monotone-decay shocks; $\xi_{ij} = 0$ if $\left[\frac{\partial \tilde{\gamma}_t}{\partial \epsilon_t} \right]_{ij} = 0$. In the vertical supply chain example, the matrix $\Xi \equiv [\xi_{ij}]$ is

$$\Xi = \begin{bmatrix} 0 & d_1 & d_1 + d_2 & d_1 + d_2 + d_3 \\ 0 & 0 & d_2 & d_2 + d_3 \\ 0 & 0 & 0 & d_3 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

The ξ_{ij} equals the delay between the production of input i and its eventual use in sector j , and is set to zero if input i is never used (even indirectly) by sector j .

We now show that the notion of ξ_{ij} as a bilateral measure of input delays extends naturally to a general network structure. When multiple paths with varying delays connect an input supplier i to an input user j , the measure ξ_{ij} defined in (23) calculates the weighted average delay across all possible paths from i to j . The weights reflect the value of inputs flowing along each route and are parametrized by ρ . A smaller value of ρ increases the importance of paths with shorter delays, whereas a larger value of ρ assigns greater weight to paths with longer delays.

Specifically, consider sector i 's revenue response to a monotone-decay shock in sector j . Whenever $\left[\partial \tilde{\gamma}_t / \partial \epsilon_t \right]_{ij}$ is non-zero, divide the ij -th entry on the right-hand side of equation (21) by the left-hand side to obtain

$$1 = \frac{[\mathbf{I}]_{ij}}{\left[\partial \tilde{\gamma}_t / \partial \epsilon_t \right]_{ij}} + \frac{[\rho \beta \mathbf{\Omega}_1]_{ij}}{\left[\partial \tilde{\gamma}_t / \partial \epsilon_t \right]_{ij}} + \frac{[(\rho \beta)^2 (\mathbf{\Omega}_1^2 + \mathbf{\Omega}_2)]_{ij}}{\left[\partial \tilde{\gamma}_t / \partial \epsilon_t \right]_{ij}} + \frac{[(\rho \beta)^3 (\mathbf{\Omega}_1^3 + \mathbf{\Omega}_1 \mathbf{\Omega}_2 + \mathbf{\Omega}_2 \mathbf{\Omega}_1 + \mathbf{\Omega}_3)]_{ij}}{\left[\partial \tilde{\gamma}_t / \partial \epsilon_t \right]_{ij}} + \dots \quad (24)$$

Sector i responds to sector j 's monotone-decay demand shock by raising production and sending its output along direct or indirect network routes to be used as an input in sector j . Each term on the right-hand side of (24) is the fraction of the output response that eventually becomes sector j 's input with a specific time delay. The first term is the direct response to i 's own demand shock. The second term is the revenue response to inputs sold with one-period delay to producers that experience the demand shock. The third term is sector i 's share of revenue response to demand shocks two periods later, and so on. In summary, each term in the infinite sequence reflects the fraction of sector i 's sales response due to sales at further distant positions in the supply chain, and the sequence sums to one.

We can write ξ_{ij} as


$$\xi_{ij} = 0 \times \frac{[\mathbf{I}]_{ij}}{[\partial \tilde{\gamma}_t / \partial \epsilon_t]_{ij}} + 1 \times \frac{[\rho \beta \mathbf{\Omega}_1]_{ij}}{[\partial \tilde{\gamma}_t / \partial \epsilon_t]_{ij}} + 2 \times \frac{[(\rho \beta)^2 (\mathbf{\Omega}_1^2 + \mathbf{\Omega}_2)]_{ij}}{[\partial \tilde{\gamma}_t / \partial \epsilon_t]_{ij}} + 3 \times \frac{[(\rho \beta)^3 (\mathbf{\Omega}_1^3 + \mathbf{\Omega}_1 \mathbf{\Omega}_2 + \mathbf{\Omega}_2 \mathbf{\Omega}_1 + \mathbf{\Omega}_3)]_{ij}}{[\partial \tilde{\gamma}_t / \partial \epsilon_t]_{ij}} + \dots$$

ξ_{ij} is therefore the average delay between when good i is produced and when it becomes an input in sector j , weighted by the additional output of i sent along each route in response to a monotone-decay demand shock to j .

Note that ξ_{ij} depends on ρ , the autocorrelation coefficient of demand shocks. As $\rho \rightarrow 1$, an AR(1) shock ϵ_{it} leads to a permanent shift in the level of demand.⁷ Hence, $\lim_{\rho \rightarrow 1} \partial \tilde{\gamma}_t / \partial \epsilon_t = \partial \bar{\gamma} / \partial \bar{\theta}$ provides the steady-state changes in revenue following steady-state demand changes, and $\lim_{\rho \rightarrow 1} \xi_{ij}$ captures the average delay from an input supplier i to an input user j along each network route using the steady-state response of goods flow along each route to demand changes in j as the weights.

Consider an example network below. The network has two supply chains labeled a and b , each with three sectors. Consider $3a$ as the focal sector. The delay from $2a$ to $3a$ is d_{2a} . The delay from $1a$ to $3a$ is $d_{1a} + d_{2a}$, as input $1a$ has to first go through producer $2a$. The delay from any sector along the b chain is unconnected to $3a$ and thus has a delay measure of zero. The matrix Ξ therefore reveals the bilateral supply chain distance measured by production delays.

Figure 3. Two production chains



It is important to note that our measure ξ_{ij} differs from the standard Leontief inverse. The latter captures a notion of network dependence—the network-adjusted value of each input used to produce a unit value of each good, whether in steady-state ($\partial \bar{\gamma} / \partial \bar{\theta} = (\mathbf{I} - \sum_{d=1}^{\infty} \beta^d \mathbf{\Omega}_d)^{-1}$) or in response to monotone-decay demand shocks ($\partial \tilde{\gamma}_t / \partial \epsilon_t = (\mathbf{I} - \sum_{d=1}^{\infty} (\rho \beta)^d \mathbf{\Omega}_d)^{-1}$). Instead, our bilateral measure ξ_{ij} captures a notion of network distance due to delays. Our measure ξ_{ij} can therefore be used to identify the ordering of inputs along each good-specific supply chains, as evident from the example in Figure 3 of two production chains. We implement this measure empirically in Section 4.

Related to ξ_{ij} is the influential upstreamness measure ψ_i of Antràs et al. (2012), which is

⁷A hump-shaped AR(2) shock u_t becomes a permanent change in the growth of demand as $\rho \rightarrow 1$.

defined as the average distance in terms of production stages from a sector i to the final consumer demand in a static model. Antràs et al. (2012)'s measure ψ_i is isomorphic to the average delay in our model from input i to all consumption goods j , weighing each j by the fraction of sector i 's value-added that eventually gets consumed through j in steady-state, in a special case without discounting and where all input delays are one period only (specifically, $\psi_i = \lim_{\beta, \rho \rightarrow 1} \sum_j \xi_{ij} \frac{[(I-\Omega)^{-1}]_{ij} \bar{\theta}_j}{\sum_k [(I-\Omega)^{-1}]_{ik} \bar{\theta}_k}$). Our delay measure ξ_{ij} also relates to the empirical measure of pair-wise upstreamness proposed by Alfaro et al. (2019). Our model microfound and generalizes these empirical upstreamness measures to accommodate different lengths of production delays.

3.2.4 Hump-shaped Demand Shocks and Amplified Volatility in Upstream Sectors

In sections 3.2.1 and 3.2.2 we have shown that upstream sectors' revenues are more responsive to hump-shaped demand shocks. Under these shocks, an increase in demand at time t implies a higher expected future demand, to which upstream is more responsive due to time-to-build.

In this section, we formally derive the volatility of sectoral revenue in terms of period-to-period changes and show that the same intuition implies that under hump-shaped shocks, upstream sectors tend to have higher volatility. First, we show that our derivation of Theorem 2 extends naturally for deriving the variance of one-period-ahead sectoral revenue conditional on the current information.

Proposition 3. *Under complete information, the sectoral revenue volatility is:*

$$\text{Var}_t(\tilde{\gamma}_{t+1}) = \mathbf{G}_\infty^\epsilon \Sigma_t^\epsilon \mathbf{G}_\infty^{\epsilon'} + \mathbf{G}_\infty^x \Sigma_t^u \mathbf{G}_\infty^{x'},$$

where Σ_t^ϵ and Σ_t^u are the variance-covariance matrices of ϵ_t and u_t , respectively, and $\mathbf{G}_\infty^\epsilon, \mathbf{G}_\infty^x$ are as defined in Theorem 2.

Proof. See Appendix A.5. □

As in Theorem 2, the matrix $\mathbf{G}_\infty^\epsilon$ characterizes the impact of demand's monotone-decay component on revenue, and $\mathbf{G}_\infty^x$ characterizes the impact of demand's hump-shaped component. Conditional on current period information, $\mathbf{G}_\infty^\epsilon \Sigma_t^\epsilon \mathbf{G}_\infty^{\epsilon'}$ and $\mathbf{G}_\infty^x \Sigma_t^u \mathbf{G}_\infty^{x'}$ characterize two demand components of the variance of the one-period-ahead revenue. In parallel to our analysis in sections 3.2.1 and 3.2.2, Proposition 3 shows that monotone-decay demand shocks from downstream sectors create dampened volatility in upstream revenue and hump-shaped demand shocks create amplified volatility.

Importantly, this result differs from the cross-sectional claim that upstream sectors may exhibit higher volatility—which could be driven by other factors such as supplying multiple downstream sectors and thereby accumulating many downstream shocks. Instead, Proposition 3 shows

that hump-shaped demand shocks in a single downstream sector can create amplified effects in upstream even when there is just a single input-output path connecting these sectors.

In order to demonstrate Proposition 3's implication as clearly as possible, consider the vertical supply chain example in Section 3.2.2. When only the most downstream sector four faces uncertain demand—with variance $(\sigma_4^\epsilon)^2$ and $(\sigma_4^u)^2$ for the monotone-decay and hump-shaped shocks, respectively—the conditional variance of one-period-ahead revenue $Var_t(\tilde{\gamma}_{i,t+1})$ in the most downstream sector ($i = 4$) and the most upstream sector ($i = 1$) are, respectively,

$$Var_t(\tilde{\gamma}_{4,t+1}) = (\sigma_4^\epsilon)^2 + (\sigma_4^u)^2$$

$$Var_t(\tilde{\gamma}_{1,t+1}) = (\rho\beta)^{2(d_1+d_2+d_3)} (\sigma_4^\epsilon)^2 + (1 + d_1 + d_2 + d_3)^2 (\rho\beta)^{2(d_1+d_2+d_3)} (\sigma_4^u)^2$$

The effect of monotone-decay demand shocks on upstream revenue is dampened by a factor of $(\rho\beta)^{2(d_1+d_2+d_3)}$, whereas the effect of hump-shaped demand shocks may be magnified.

3.3 Incomplete Information

Our analysis thus far has demonstrated that downstream sectors' revenue responds more to monotone-decay demand shocks, whereas upstream sectors' revenue may respond more to hump-shaped demand shocks. We now generalize our environment to the setting where agents only observe the realized demand $\tilde{\theta}_t$ but not the separate components. That is, when a demand shock arises at time t , it is not possible for agents to accurately assess its impact on future demand. In this setting, agents, including the social planner, need to solve a filtering problem and forecast future demand based on the history of demand shocks. We show that the impulse response of sectoral revenue to a demand shock is in general hump-shaped, that upstream sectors have more pronounced revenue response, and that this is especially true when demand shocks tend to be more hump-shaped.

Formally, agents observe only the history of demand $\left\{\tilde{\theta}_{t-s}\right\}_{s \geq 0}$ but not the two components separately. The conditional expectation operator is then defined as $\mathbb{E}_t[\cdot] \equiv \mathbb{E}\left[\cdot \mid \left\{\tilde{\theta}_{t-s}\right\}_{s \geq 0}\right]$. In this setting, agents, including the social planner, need to solve a filtering problem and forecast future demand based on the history of demand shocks.

The exogenous state variables at time t for consumer demand continue to be $(\tilde{\theta}_t, \mathbf{x}_t)$, as in the baseline model with complete information. Under incomplete information, only $\tilde{\theta}_t$ is observable; $\mathbf{x}_t$, the hump-shaped component of current demand, is a hidden state. There is an additional endogenous state variable that captures the agents' belief $\hat{\mathbf{x}}_t \equiv \mathbb{E}_t\left[\mathbf{x}_t \mid \left\{\tilde{\theta}_{t-s}\right\}_{s \geq 0}\right]$ of the hidden state based on the history of demand up to time t . We refer $\hat{\mathbf{x}}_t$ as the nowcast (see [Kalman, 1960](#), [Kalman and Bucy, 1961](#)).

The demand shock at time t is defined as the unexpected innovation in demand:

$$\mathbf{v}_t \equiv \tilde{\boldsymbol{\theta}}_t - \mathbb{E}_{t-1} [\tilde{\boldsymbol{\theta}}_t] = \tilde{\boldsymbol{\theta}}_t - \rho (\tilde{\boldsymbol{\theta}}_{t-1} + \hat{\mathbf{x}}_{t-1}).$$

Kalman Filter We posit that learning follows a Kalman filter, which provides the optimal nowcast under Gaussian shocks and the best linear nowcast (that minimizes the expected squared loss) under non-Gaussian shocks (Humpherys et al., 2012). Generally, demand shocks need to be non-stationary and non-Gaussian in order for demand to be non-negative; nevertheless, our analysis under incomplete information can be interpreted as analyzing the linearized model around a steady-state shocks approximated by Gaussian distributions. Hence, for expositional simplicity and also relevant for our empirical specification, we assume the covariance matrices Σ^u and Σ^ϵ are time-invariant, so that the Kalman filter is stationary.⁸

In this case, the demand innovation $\mathbf{v}_t$ updates the nowcast according to the law of motion:

$$\hat{\mathbf{x}}_t = \rho \hat{\mathbf{x}}_{t-1} + \mathbf{K} \mathbf{v}_t \quad (25)$$

The Kalman gain $\mathbf{K}$ is an $N \times N$ matrix as derived in Appendix A.6:

$$\mathbf{K} = (\rho^2 \mathbf{F} + \Sigma^u) (\Sigma^\epsilon + \rho^2 \mathbf{F} + \Sigma^u)^{-1}, \quad (26)$$

with Σ^u is the covariance matrix of hump-shaped shocks, Σ^ϵ is the covariance matrix of monotone-decay shocks, and $\mathbf{F} \equiv \text{Var}_t(\mathbf{x}_t - \hat{\mathbf{x}}_t)$, a measure of the accuracy of the nowcast, satisfies the following equation:

$$\mathbf{F} = \Sigma^\epsilon (\Sigma^\epsilon + \rho^2 \mathbf{F} + \Sigma^u)^{-1} (\rho^2 \mathbf{F} + \Sigma^u). \quad (27)$$

The following proposition derives the impulse-response of expected future demand $\mathbb{E}_t [\tilde{\boldsymbol{\theta}}_{t+s}]$ and revenue $\mathbb{E}_t [\tilde{\gamma}_{t+s}]$ to demand shocks.

Theorem 3. *Under incomplete information where learning follows (25), the response of the vector of expected future demand to a vector of demand shocks $\mathbf{v}_t \equiv \tilde{\boldsymbol{\theta}}_t - \mathbb{E}_{t-1} [\tilde{\boldsymbol{\theta}}_t]$ at time t is*

$$\frac{\partial \mathbb{E}_t [\tilde{\boldsymbol{\theta}}_{t+s}]}{\partial \mathbf{v}_t} = \rho^s \mathbf{I} + s \rho^s \mathbf{K}.$$

The response of the vector of expected future revenue to demand shocks is

$$\frac{\partial \mathbb{E}_t [\tilde{\gamma}_{t+s}]}{\partial \mathbf{v}_t} = \rho^s (\mathbf{G}_\infty^\epsilon (\mathbf{I} + (s-1) \mathbf{K}) + \mathbf{G}_\infty^x \mathbf{K}),$$

where $\mathbf{I}$ is the identity matrix, $\mathbf{G}_\infty^\epsilon$ and $\mathbf{G}_\infty^x$ are as defined in Theorem 2.

Proof. See Appendix A.7. □

⁸There is a large literature on nonlinear and nonstationary filtering (e.g., Hamilton, 1989; Ito and Xiong, 2000; Farmer, 2021); the extension to this case is straightforward but less expositionaly clear.

Theorem 3 characterizes the expected time path of each sector i 's consumer demand and sectoral revenue following an unexpected innovation in the demand for each good j . The impulse-response of expected demand is governed by the Kalman gain matrix $\mathbf{K}$ which characterizes learning. The impulse-response of expected revenue depends on both the network structure and the delay (characterized by $\mathbf{G}_\infty^\epsilon$ and $\mathbf{G}_\infty^x$), as well as learning (characterized by $\mathbf{K}$). When there are only monotone-decay shocks ϵ_t (i.e., the covariance matrix of $\mathbf{u}_t$ is zero), the Kalman gain is the zero matrix, and $\frac{\partial \mathbb{E}_t[\tilde{\gamma}_{t+s}]}{\partial v_t}$ coincides with the object $\frac{\partial \mathbb{E}_t[\tilde{\gamma}_{t+s}]}{\partial \epsilon_t}$ under complete information (c.f. Theorem 2). When there are only hump-shaped shocks, the Kalman gain is the identity matrix, and $\frac{\partial \mathbb{E}_t[\tilde{\gamma}_{t+s}]}{\partial v_t}$ coincides with the object $\frac{\partial \mathbb{E}_t[\tilde{\gamma}_{t+s}]}{\partial u_t}$ under complete information.

When demand shocks are uncorrelated across sectors, the Kalman gain matrix becomes diagonal, with the i -th entry κ_i increasing in the relative variance of hump-shaped AR(2) shocks compared to monotone-decay AR(1) shocks.⁹ As a result, sectors dominated by hump-shaped shocks exhibit larger values of κ_i . Empirically, since demand shocks are largely sector-specific and weakly correlated across sectors, the Kalman gain matrix $\mathbf{K}$ is diagonally dominant, as we confirm in the data (Section 4.4).

A key prediction of our model under complete information is that more upstream sectors have amplified response to hump-shaped demand shocks but not monotone-decay shocks (Theorem 2). Theorem 3 shows that under incomplete information and when both types of shocks are present, the impulse-response of sectoral revenue to demand innovations are in general hump-shaped, that as one moves from downstream to upstream the response first becomes more pronounced before it dampens, and that the initial amplification to upstream is especially true when demand shocks tend to be more hump-shaped.

To see this, consider again the vertical supply chain example as in Section 3.2.2, where sector i supplies to $i+1$ for $i = 1, 2, 3$. Suppose demand shocks are uncorrelated across sectors. Theorem 3 implies that, following a demand innovation in the most downstream sector 4, the expected time path of sectoral revenue is

$$\frac{\partial E_t[\tilde{\gamma}_{t+s}]}{\partial v_{4t}} = \rho^s \times \begin{bmatrix} (1 + s\kappa + d_1\kappa + d_2\kappa + d_3\kappa) (\rho\beta)^{d_1+d_2+d_3} \\ (1 + s\kappa + d_2\kappa + d_3\kappa) (\rho\beta)^{d_2+d_3} \\ (1 + s\kappa + d_3\kappa) (\rho\beta)^{d_3} \\ 1 + s\kappa \end{bmatrix},$$

where κ is the Kalman gain associated with the demand innovation for good 4, the most downstream sector.

The sectoral revenue responses to a demand innovation in the incomplete information environment exhibit three features similar as the responses to hump-shaped shocks under complete


⁹Specifically, $\kappa_i = \frac{-(1+(\sigma_i^u/\sigma_i^\epsilon)^2-\rho^2)+\sqrt{(1+(\sigma_i^u/\sigma_i^\epsilon)^2-\rho^2)^2+4\rho^2(\sigma_i^u/\sigma_i^\epsilon)^2}}{2\rho^2}$.

information. First, sectoral responses are hump-shaped over time: i.e., as s increases, the term $s\kappa$ is increasing, the term ρ^s is decreasing, and their product first increases over time (for $\rho > 1/2$) but eventually decreases as $s \rightarrow \infty$. Second, upstream sectors may exhibit amplified response relative to downstream, and third, the amplified response eventually dies out as one goes further upstream (for sufficiently high ρ , $(1 + s\kappa + x\kappa)(\rho\beta)^x$ is initially increasing but eventually decreasing in x).

The new implications under incomplete information are two-fold. First, the hump-shaped impulse-response becomes more pronounced along the supply chain when the Kalman gain κ is larger. This occurs precisely when demand shocks tends to more hump-shaped—that is, when the relative standard deviation of hump-shaped shocks is higher. Second, higher Kalman gains also lead to greater amplification of the sectoral response as shocks propagate upstream along the supply chain.

We illustrate these implications in Figure 4, which shows the expected response of sectoral revenue to a downstream shock at different positions along the supply chain. The left panel displays the downstream sector's response, the middle panel shows a relatively upstream sector's response, and the right panel illustrates the response of a sector located further upstream. The blue curves represent impulse responses under a high Kalman gain, corresponding to sectors where demand shocks tend to be more hump-shaped, while the red curves represent sectors with lower Kalman gains. Consistent with our theory, the sectoral response is initially amplified as shocks propagate upstream (from left to middle panels), but this amplification diminishes for sectors situated further upstream (right panel). Crucially, when demand shocks are more likely to be hump-shaped (higher Kalman gain), both the hump-shaped downstream responses and the initial upstream amplification effects become notably more pronounced. We test these predictions empirically in Section 4.3.

Figure 4. The expected revenue paths under high κ and low κ for a vertical supply chain



Finally, we formally derive the sectoral revenue volatility under incomplete information.

Proposition 4. *Under incomplete information, the volatility of demand is*

$$Var_t(\tilde{\theta}_{t+1}) = \rho^2 \mathbf{F} + \Sigma^\epsilon + \Sigma^u, \quad (28)$$

and the volatility of sectoral revenue is:

$$Var_t(\tilde{\gamma}_{t+1}) = (\mathbf{G}_\infty^\epsilon - \mathbf{G}_\infty^\epsilon \mathbf{K} + \mathbf{G}_\infty^x \mathbf{K}) Var_t(\tilde{\theta}_{t+1}) (\mathbf{G}_\infty^\epsilon - \mathbf{G}_\infty^\epsilon \mathbf{K} + \mathbf{G}_\infty^x \mathbf{K})', \quad (29)$$

where Σ^ϵ and Σ^u are the variance-covariance matrices of ϵ_t and u_t , respectively, and $\mathbf{G}_\infty^\epsilon, \mathbf{G}_\infty^x$ are as defined in Theorem 2. $\mathbf{K}$ is the matrix of Kalman gains as in (26), and $\mathbf{F} \equiv Var_t(\mathbf{x}_t - \hat{\mathbf{x}}_t)$ as in (27).

Proof. See Appendix A.8. □

Demand volatility under incomplete information consists of three parts. The term $\rho^2 \mathbf{F}$ captures the forecast error of the hidden state. Σ^ϵ and Σ^u captures the uncertainty due to the two types of demand shocks. Demand volatility then translates into sectoral volatility along the supply chain. When the Kalman gain is the identity matrix ($\mathbf{K} = \mathbf{I}$), sectoral volatility simplifies to $Var_t(\tilde{\gamma}_{t+1}) = \mathbf{G}_\infty^x \Sigma^u \mathbf{G}_\infty^{x'}$, which corresponds to the complete information case with only hump-shaped shocks. Similarly, when the Kalman gain is the zero matrix ($\mathbf{K} = \mathbf{0}$), sectoral volatility simplifies to $Var_t(\tilde{\gamma}_{t+1}) = \mathbf{G}_\infty^\epsilon \Sigma^\epsilon \mathbf{G}_\infty^{\epsilon'}$, which corresponds to the complete information case with only monotone-decay shocks.

4 Empirical Evidence and Quantification

We now turn to the empirical analysis with three goals. First, in Section 4.2 we provide evidence of the bullwhip effect in the cross section. That is, within a sectoral supply chain, upstream sectors tend to exhibit higher volatility than downstream sectors. Second, we extract innovations to downstream value-added and estimate the impulse-response functions of upstream sectors along the supply chain in Section 4.3. We demonstrate a number of features of the estimated impulse-response functions that are consistent with our model's predictions, thereby providing a sharp test of our theory. Third, in Section 4.4 we conduct a model-based exercise. We infer the empirical process of demand shocks using the model and demonstrate that hump-shaped demand shocks for downstream sectors account for a quantitatively significant share of sectoral fluctuations along the supply chains.

4.1 Data

We use data from three main sources. First is the Industrial Production (IP) data published by the Federal Reserve Board. The data set has panel information on the sectoral value-added for manu-

facturing industries over the period 1972-2023, with industries classified by the North American Industry Classification System (NAICS). To smooth out measurement errors and high-frequency shocks, we average the monthly data to construct quarterly sectoral value-added in our baseline analysis, while monthly specifications are presented as robustness checks in Appendix D.8.

Second, to construct our bilateral measure of supply chains, we need information on input-output linkages and input time-to-build delays. For input-output data, we utilize the 2007 BEA input-output (IO) use table,¹⁰ from which we construct input cost shares ω_{ij} and sectoral value-added intensities (value-added to revenue ratio).

We follow Foerster et al. (2011)’s algorithm to concord between the IP and IO datasets. This procedure results in 114 manufacturing industries appearing in both datasets, as well as 127 non-IP sectors present only in the IO table. Among the IP industries, the majority (65 out of 114) corresponds to four-digit NAICS codes, with the remaining mapping into either three- (4 industries) or five-digit (45 industries) NAICS codes. We format the IP data (“Value-Added Proportions”) to represent sectoral value-added shares—each sector’s value-added as a proportion of the total value-added across all 114 IP sectors. This format aligns with our analysis, which centers on sectoral and supply chain dynamics rather than aggregate effects.¹¹

Third, we need information on time-to-build delays. Our baseline strategy is to use the backlog ratio for each input, i.e., the ratio between the stock value of unfilled orders and the flow value of goods delivered, as a measure of supply chain delays (Liu and Tsyvinski, 2024). The strategy yields a supplier-specific measure of delay. We rely on the U.S. Census M3 survey of manufacturers’ shipments, inventories, and orders. The data set provides broad-based, monthly statistical data on economic conditions in the manufacturing sector. For each industry, we compute the average backlog ratio between the years 2015 and 2019.¹² We match industries in the M3 survey to the finest partition possible in the IO table. We impute the backlog ratio using the sample averages respectively for durable goods sectors and non-durable goods sectors that are not in the M3 survey. Because the backlog ratio captures only the average time between orders placed and orders received but not delays due to the production process of the input-buyer, we round up the backlog ratios when converting from months to quarters. Based on this procedure, twenty sectors have time-to-build delays of two quarters, two sectors have delays of three quarters, four sectors have delays of four, and one sector (“Ship and Boat Building”) has five-quarter delays.

As robustness checks, we implement two alternative strategies for measuring time-to-build

¹⁰The BEA IO tables are available every five years. We use the 2007 version as it is in the second half of our sample but before the Great Financial Crisis. Our results are not sensitive to this choice.

¹¹We impute missing data according to the Federal Reserve Board’s recommended methodology: if sector C is made up of sector A less sector B, then the value added of sector C is the value-added of A less than that of B.

¹²Most of the variation in the backlog ratio is cross-sectional. Regressing the backlog ratio between 1992 and 2019 on sector fixed effects yields an R^2 of 0.88, indicating limited variation over time.

delays. The first alternative strategy follows the recent work of [Antràs and Tubdenov \(2025\)](#), which infers an input-buyer-specific measure of time-to-build based on the ratio between the value of inventories and cost of goods sold (COGS) using the COMPUSTAT data and then aggregate to the sector level using a cost-weighted average across firms. This measure captures production delays by input-buyers but does not reflect shipping or production delays originating from input-suppliers. As such, it is complementary to our baseline measure. The second alternative strategy specifies the total time-to-build delay from an input producer j to input user i is the sum of the backlog ratio of good j and the inventory-to-COGS ratio of good i . Our main results remain robust across these alternative measures of time-to-build—network structure consistently plays a more significant role, as detailed in the Appendix [D.4](#).

4.2 Higher Volatility in Upstream Sectors Across Supply Chains

We first provide cross-sectional evidence for the bullwhip effect across supply chains of the U.S. production network: upstream sectors within a supply chain tend to have higher volatility.

We measure sectoral volatility as the standard deviation of the year-on-year (YoY) growth in sectoral value-added shares (calculated relative to total value-added across IP industries) over our sample. Using YoY growth in value-added offers several advantages: it removes seasonal fluctuations and normalizes for intrinsic size differences across sectors, allowing for meaningful volatility comparisons. Moreover, YoY growth serves as a convenient reduced-form proxy for the conditional variance, as last year’s value-added levels effectively summarize the relevant time-varying information set. Finally, although we have formulated our theory in terms of sectoral revenue, the YoY growth measure removes the distinction between revenue and value-added to the extent that sectoral value-added intensity stays constant.

it can be readily restated in terms of value-added (observed from the IP data), which equals revenue multiplied by the sectoral value-added intensity (obtained from the BEA IO table).

The main challenge in documenting the bullwhip effect is to identify the supply chains in the US production network. To do so, we rely on our bilateral measure of supply chain delays ξ_{ij} , following the procedure described in Section [3.2.3](#).

We start by identifying a set $\mathcal{D}$ of 25 downstream sectors that are important in the consumption bundle, following [Antràs et al. \(2012\)](#).¹³ These sectors are listed in Table [1](#); they supply most of their output directly to the final consumer instead of other producers. Broadly speaking, these sectors fall into the categories of transportation, food, and other household consumption goods.

¹³Specifically, we first select the 30 most downstream sectors according to the upstreamness measure by [Antràs et al. \(2012\)](#). We then drop “Other Miscellaneous Manufacturing”, “Construction Machinery”, “Agricultural Implements”, “Industrial Machinery” and “Support Activities for Mining”, as these five sectors do not account for significant consumer expenditure and are thus not important for the consumer demand.

Table 1. The list of twenty-five downstream sectors

Apparel	Coffee and Tea
Ship and Boat Building	Other Transportation Equipment
Newspaper Publishers	Heavy Duty Trucks
Tobacco	Office And Other Furniture
Soft Drinks and Ice	Aerospace Products and Parts
Bakeries and Tortilla	Breweries
Motor Vehicle Bodies and Trailers	Other Food Except Coffee and Tea
Carpet and Rug Mills	Pharmaceuticals and Medicines
Soap, Cleaning Compounds, and Toilet Preparation	Navigational/Measuring/Electromedical Instruments
Periodical, Book, and Other Publishers	Sugar and Confectionery Products
Automobiles and Light Duty Motor Vehicles	Animal Slaughtering and Processing
Major Electrical Household Appliances	Medical Equipment and Supplies
Fruit and Vegetable Preserving and Specialty Foods	

We then identify the supply chain for each of these downstream sectors. First, for each downstream sector $j \in \mathcal{D}$, we keep a set of sectors $\mathcal{U}_j$ that are considered as direct or indirect suppliers to j : a sector i is in the set $\mathcal{U}_j$ if producing 200 dollar of j 's output requires at least one dollar of input i , directly or indirectly. Formally, this means that the ij -th entry of the Leontief inverse $L_{ij} \equiv (I - \Omega)_{ij}^{-1}$ is greater than 0.5 percent. This threshold is chosen so that the each downstream good has on average 25 (median 26) number of sectors along its supply chain.

Next, for each downstream sector $j \in \mathcal{D}$, we order the set of suppliers to j using our bilateral supply chain delay measure ξ_{ij} , following the procedure described in Section 3.2.3. In the baseline, we set $\rho = \beta = 1$ when calculating ξ_{ij} , which depends on the product $\rho\beta$. Our results are robust to using alternative values of these parameters—as Appendix Figure A.9 shows, ξ_{ij} remains highly correlated across a wide range of $\rho\beta$ —including a realistic quarterly discount factor $\beta = 0.98$ and the rate of decay parameter $\rho = 0.818$ estimated from the data (see Section 4.4). The resulting measure captures the average number of quarters in the time-to-build delays for input i to reach producer j , accounting for all the direct and indirect network linkages from i to j . Within the supply chain of each good j , we interpret sector $i \in \mathcal{U}_j$ as more upstream to $k \in \mathcal{U}_j$ if there is a longer delay from input i to producer j than from input k .

Take “Automobiles and Light Duty Motor Vehicles” (“automobiles” for short) and “Motor Vehicle Bodies and Trailers” (“trailers” for short) as two examples of downstream sectors. Despite the fact that these two sectors share many common suppliers—such as “Machine Shops; Turned Products; and Screws, Nuts, and Bolts”, “Paperboard Containers”, “Semiconductors”, “Iron and Steel Products”, and “Electric Power Generation, Transmission and Distribution”—our measure can isolate the differences in their supply chains. Specifically, our measure identifies that the sectors most immediately upstream to automobiles are “Motor Vehicle Bodies and Trailers”, “Motor

Vehicle Parts” and “Tires”. By contrast, those that are most immediately upstream to trailers include “Major Electrical Household Appliances” and “Millwork”.

Figure 5 shows the relationship between upstreamness, as measured by supply chain delays, and volatility, along the sector-specific supply chain for 9 of the 25 downstream sectors (we produce the same figure for each of the remaining 16 downstream sectors in Appendix Figure A.2).¹⁴ Specifically, each panel represents the supply chain associated with a downstream sector, as indicated by the subtitle (e.g., the top-left panel is for the supply chain of automobiles). Within each panel, a circle represents an input supplier along the supply chain (e.g., the bottom circle in the top-left panel is “Motor Vehicle Parts”). The Y-axis captures the time-to-build delays from the supplier to the downstream producer, i.e., a notion of upstreamness. The X-axis is the volatility in the YoY value-added growth of the supplier. The size of each circle represents the share of the total downstream cost that is spent on each input, directly or indirectly.

The key takeaway from Figure 5 is the significant positive relationship between sectoral volatility and upstreamness, an empirical pattern consistently observed across supply chains for all 25 downstream sectors. In terms of magnitude, pooling across 25 supply chains we find that going from 25 to 75 percentile in terms of upstreamness is associated with 5.5 percentage points higher volatility in value-added growth. This effect is pronounced across transportation, food, households consumption goods, and all other categories of downstream goods.

One potential concern is that the observed positive relationship between sectoral volatility and upstreamness might be driven by the volatility of commodity prices, as four commodity sectors occupy upstream positions across many supply chains.¹⁵ However, we find that this not the case. Appendix Figure A.3 shows that the positive relationship between sectoral volatility and upstreamness remains qualitatively unchanged even after excluding these commodity sectors.


4.3 Supply Chain Response to Downstream Value-Added Shocks

Crucially, our theory’s predictions extend beyond the cross-sectional observation that upstream sectors often exhibit higher volatility, an empirical regularity that could be driven by other factors such as supplying multiple downstream industries and thereby aggregating multiple shocks. Instead, our theory provides precise predictions regarding impulse-response dynamics. Specifically, Theorem 3 implies that following an innovation to value-added in downstream sectors dominated by AR(2) shocks:

¹⁴We include only non-durable inputs—goods for which no more than 50% of its final use goes into fixed capital formation in the BEA IO table—when investigating this cross-sectional relationship, as the volatility of durable goods may be driven by important factors not considered in the model.

¹⁵These sectors are “Oil and Gas extraction”, “Coal mining”, “Copper, Nickel, Lead, and Zinc Mining”, and “Iron, Gold, Silver and Other Metal Ore Mining”.

Figure 5. Along supply chains, upstream sectors exhibit higher volatility



1. The downstream sectors' own value-added exhibits a hump-shaped response over time.
2. Value-added responses along the supply chain are also hump-shaped and initially amplified as shocks propagate upstream.
3. Moving along the supply chain, the initial amplification diminishes in sectors located further upstream.
4. The downstream hump-shaped responses and upstream amplification effects should be stronger in sectors where AR(2) shocks dominate, while such patterns could be absent in

sectors characterized primarily by AR(1) shocks.

We empirically test these predictions by first removing linear time trends from each sector’s value-added series, standardizing the detrended data, and using these standardized z-scores in our analysis to account for heterogeneous sectoral sizes.

Although in our theoretical framework, fluctuations in value-added (i.e., revenue multiplied by sectoral value-added intensity) originate exclusively from demand shocks—since productivity shocks have offsetting effects on prices and quantities—our empirically testable predictions directly concern fluctuations in sectoral value-added.

In the main text, we directly test the predictions based on extracting shocks to downstream value-added and estimating their impulse response functions along supply chains. Additionally, in Appendix D.9, we discuss a robustness check using a generalized model with constant-elasticity-of-substitution (CES) production functions to explicitly control for sector-specific cost shocks. Specifically, we solve the CES model with one-period delays, linearize it around a steady-state, and derive separate implications of cost shocks and demand shocks. This approach enables us to nonparametrically control for cost shocks when estimating impulse-response functions. Our empirical findings remain robust after incorporating these controls.

Downstream Value-Added Has Both AR(1) and AR(2) Shocks We begin by using a reduced-form test to show that the time-series of detrended downstream value-added is well captured by a stochastic process with AR(1) and AR(2) shocks. Specifically, we follow the econometrics literature and use partial autocorrelation function (PACF) to select appropriate lags in an $AR(p)$ model (Box et al., 2015). The PACF coefficient ϕ_k of a time series is defined as the autocorrelation at lag k after conditioning on all previous lags 1 through $k - 1$. According to the Box–Jenkins method (Box et al., 2015), insignificance of the PACF coefficient at lag $p + 1$ indicates that the time-series can be modeled as an $AR(p)$ process.

We estimate the PACF coefficient sector-by-sector for the 25 downstream sectors $i \in \mathcal{D}$. We find that in our sample, all twenty-five downstream sectors have statistically significant (5%) PACF coefficients at lag 1 for the time-series of their value-added. Twenty-one out of twenty-five downstream sectors have statistically significant PACF coefficients at lag 2. Only two sectors (“Tobacco”, “Other Food Except Coffee and Tea”) have statistically significant PACF coefficients at lag 3. These findings suggest that downstream value-added is well approximated by a stochastic process with AR(1) and AR(2) shocks.

Moreover, our model implies that sectors with larger absolute values of PACF coefficients at lag 2 ($|\phi_{i,2}|$) are precisely those in which hump-shaped AR(2) shocks play a more dominant role.

Specifically, $|\phi_{i,2}|$ is increasing in the relative variance of AR(2) shocks, $(\sigma_i^u/\sigma_i^\epsilon)^2$.¹⁶ Conversely, sectors with smaller $|\phi_{i,2}|$ predominantly experience monotone-decay AR(1) shocks.

To test our model's prediction on the differential supply chain response following shocks to these sectors, we define P_i as the dummy variable for whether $|\phi_{i,2}|$ is above median among the 25 downstream sectors. For expositional simplicity we refer to these sectors as having relatively large AR(2) components.

Extracting Innovations in Downstream Value-Added We use local projections to estimate the impulse-response along the supply chain follows to downstream value-added shocks. To estimate the value-added shock of the 25 downstream sectors, we project each sector's value-added on its own lags, as well as sector fixed effects μ_i and time fixed effects μ_t , and we extract the residuals. Specifically, the value-added shocks r_{it} in a downstream sector i is the innovation in the following:

$$VA_{it}^{down} = \sum_{s=1}^p \gamma_s VA_{i,t-s}^{down} + \mu_i + \mu_t + r_{it}, \quad i \in \mathcal{D}. \quad (30)$$

In the baseline specification, we include $p = 2$ lags, consistent with our model specification of AR(1) and AR(2) shocks. Our results are not sensitive to the choice of lags; Appendix D.6 replicates our results for $p = 1$ and $p = 3$ as robustness checks.

Impulse-Response of Downstream Value-Added to Its Own Innovations Equipped with downstream's value-added shocks, we first estimate the impact of a shock on the downstream sector's own value-added in subsequent periods:

$$VA_{i,t+h}^{down} - VA_{i,t-1}^{down} = \alpha_0^h r_{it} + \alpha_1^h r_{it} \times P_i + X_{it} + \zeta_{it}, \text{ with controls } X_{it} \in \left\{ \{VA_{i,t-s}^{down}\}_{s=1}^p, \mu_i, \mu_t \right\}. \quad (31)$$

In the set of regressions, value-added in each quarter $t + h$ is regressed on the estimated value-added shock in quarter t , as well as controls X_{it} . Recall P_i is the dummy for whether sector i has a dominant AR(2) component (i.e., the PACF coefficient $|\phi_{i,2}|$ is above the median among all downstream sectors). The estimate α_0^h represents the impulse response function at horizon h for downstream sector i 's value-added when the dummy P_i is turned off, and the estimate $\alpha_0^h + \alpha_1^h$ represents the impulse response function when the dummy is turned on.

The impulse response functions displayed in all figures show 95% confidence interval. Because the second stage regression (31) uses the residuals from the first stage (30) as an independent variable, we adopt a bootstrap procedure that resamples the original data with replacement at the buyer-supplier-sector-pair level to obtain the confidence interval in the second stage.


The top-left panel (a) in Figure 6 reports the coefficients α_0^h (solid line) and $\alpha_0^h + \alpha_1^h$ (dashed line), displaying the estimates from equation (31) of the impact of a demand shock from down-

¹⁶Formally, $\phi_{i,1} = \rho + \frac{\frac{\rho}{(1-\rho^2)^2}(\sigma_i^u/\sigma_i^\epsilon)^2}{\frac{1+\rho^2}{(1-\rho^2)^3}(\sigma_i^u/\sigma_i^\epsilon)^2 + \frac{1}{1-\rho^2}}$ and $\phi_{i,2} = \left(\rho^2 + \frac{\frac{2\rho^2}{(1-\rho^2)^2}(\sigma_i^u/\sigma_i^\epsilon)^2}{\frac{1+\rho^2}{(1-\rho^2)^3}(\sigma_i^u/\sigma_i^\epsilon)^2 + \frac{1}{(1-\rho^2)}} - \phi_{i,1}^2 \right) / (1 - \phi_{i,1}^2)$.

stream sectors in quarter t on their own value-added share in quarter $t+h$ for a 2-year (8 quarters) horizon. Panel (a) in Figure 6 shows hump-shaped, mean-reverting impulse response functions, especially those sectors with relatively large AR(2) components. This is consistent with our model specification of AR(2) shocks, which implies hump-shaped response in downstream value-added following its own innovation. In terms of magnitude, for those sectors with large AR(2) components, an one standard deviation innovation in value-added to a downstream sector raises its own concurrent value-added by one unit by our construction. The expected value added increases over time, peaking in the third quarter at 1.6 standard deviation above the mean. The impact starts to decline after three quarters, but the impact is persistent and only decays to about half of the initial impact after two years. For those sectors with small AR(2) components, the hump-shaped response is present but significantly less pronounced. The impact of the shock peaks in one quarter but is nevertheless persistent with a half-life of about two years.

Appendix Figure A.4 further confirms the presence of a pronounced hump-shaped impulse response function when all downstream sectors are grouped together.

Figure 6. Cumulative impulse response along the supply chain



Impulse-Response Along the Supply Chain We now use the innovations in downstream value-added to estimate the impulse-responses of upstream sectors along each supply chain. Since each supply chain involves many sectors, we group suppliers based on their supply chain positions. Specifically, for each downstream consumer good j , we classify input suppliers (indexed by i) into three groups according to terciles of our upstreamness/production delay measure ξ_{ij} . To maintain consistency with our four-sector vertical supply chain example in Section 3.2.2, we number groups in decreasing order of upstreamness: the most upstream suppliers are assigned to group $g = 1$, while suppliers closest to the downstream sector belong to group $g = 3$. We refer to the 25 downstream sectors themselves as group $g = 4$. Note that since the same sector can supply multiple downstream goods, it may occupy different positions and thus belong to different groups across various supply chains.

We estimate the impulse-response of sectors along the supply chain to innovations in downstream value-added using the following specification:

$$VA_{i,t+h}^{up,g} - VA_{i,t-1}^{up,g} = \beta^{h,g} \sum_{i \in \mathcal{U}_j} w_{ij} r_{jt} + \delta^{h,g} \sum_{i \in \mathcal{U}_j} w_{ij} r_{jt} \times P_j + Z_{it}^g + \epsilon_{it}, \quad (32)$$

$$\text{with controls } Z_{it}^g \in \left\{ \left\{ \sum_{i \in \mathcal{U}_j} w_{ij} VA_{j,t-s}^{down}, VA_{i,t-s}^{up,g} \right\}_{s=1}^p, \mu_i, \mu_t \right\}.$$

That is, we project the revenue of each sector along the supply chain of downstream consumption good j onto value-added innovations r_{jt} , estimating group-specific coefficients $\beta^{h,g}$ and $\delta^{h,g}$ for each time h after the shock.¹⁷ Because an upstream sector i may supply multiple downstream goods j , we aggregate the downstream innovations using weights proportional to each sector's input usage intensity: $w_{ij} = L_{ij} \times \frac{\text{std}(VA_{j,t})/\alpha_j}{\text{std}(VA_{i,t})/\alpha_i}$. The term L_{ij} is the ij -th entry of the IO table's Leontief-inverse, representing the importance of supplier i for producing good j . The ratio of standard deviations and value-added intensities adjust for the fact that both the downstream innovations r_{jt} and the upstream outcome variables are standardized in the regression. We exclude four commodity sectors (c.f. footnote 15) because they serve as broadly upstream suppliers across many different supply chains.

We plot the coefficients $\beta^{h,g}$ and $\beta^{h,g} + \delta^{h,g}$ across time h for different supply chain groups $g \in \{3, 2, 1\}$ in the remaining panels (b)–(d) of Figure 6. The estimate $\beta^{h,g}$, shown in the solid lines, represents the impulse response functions at horizon h for upstream group g 's value-added when they situate along the supply chain for a downstream sector j with a relatively small hump-shaped AR(2) component (i.e., with the dummy P_j turned off). The estimate $\beta^{h,g} + \delta^{h,g}$, shown in the dashed lines, represents the impulse response functions at horizon h for upstream group g 's

¹⁷ As a robustness check, in Appendix D.7 we replace the two lags of upstream value-added in the control variables by $\Delta VA_{i,t-1}^{up,g}$. This is a long-difference specification that could suppress small-sample bias (Jordà and Taylor, 2025).

value-added when they situate along the supply chain for a downstream sector j with a relatively large hump-shaped AR(2) component (i.e., with the dummy P_j turned on).

The heterogeneity in the impulse response functions provide a sharp empirical test of our theory. Our estimated impulse response functions are consistent with all four of our theory's key predictions under incomplete information. First, downstream sectors (group $g = 4$) where AR(2) shocks dominate (dashed line in panel a) exhibit a distinctly hump-shaped response in their own value-added following demand shocks. Second, the value-added response along the supply chain is also hump-shaped and initially amplified as shocks propagate upstream, generating pronounced volatility: the response is amplified from the downstream sectors in Figure 6 to group $g = 3$ in Figure 6, and is further amplified in group $g = 2$. Third, this amplification diminishes in sectors located further upstream (group $g = 1$). Fourth—and most importantly—the downstream hump-shaped responses and upstream amplification effects are markedly stronger in sectors where AR(2) shocks dominate, while such patterns are weak or absent in sectors characterized primarily by AR(1) shocks. Specifically, we observe a pronounced gap between the impulse responses for AR(2)-dominant supply chains (dashed lines) and AR(1)-dominant supply chains (solid lines). This gap initially widens and then gradually narrows over time; similarly, moving upstream from downstream sectors (group $g = 4$) towards the most upstream suppliers (group $g = 1$), the gap initially widens before eventually closing. These empirical results closely align with our theoretical predictions.

4.4 Quantitative Implications

In this section we estimate the demand shock process from the data, and we demonstrate the quantitative importance of hump-shaped shocks in a time-to-build environment for explaining sectoral fluctuations along the supply chain.

Estimation We assume (16) and (17) are the true data-generating process, and we estimate the covariance matrices Σ^u and Σ^e for the two types of demand shocks as well as the persistence parameter ρ .¹⁸ As we do not separately observe the two types of shocks, we estimate these three objects $\eta \equiv \{\Sigma^u, \Sigma^e, \rho\}$ following our derivation of the incomplete information case in Section 3.3. Specifically, based on the initial guess η_0 , we first run a Kalman filter on the panel data of realized downstream value-added. We then estimate these three objects using maximum likelihood, update the Kalman filter, and iterate until convergence. To ensure the covariance matrices are positive semi-definite, we estimate them by performing Cholesky decompositions.¹⁹


We find $\rho = 0.818$. Figure 7 plots the correlation matrices—which are easier to visualize than

¹⁸In Appendix D.10 we re-do the exercises in this section allowing for separate persistent parameters for the hump-shaped and monotone-decay shocks, following the derivations in Appendix B.2.

¹⁹In Appendix C, we document the details of MLE estimation.

covariances—for hump-shaped shocks in panel (a) and for monotone-decay shocks in panel (b). The figure shows that the demand shocks are not too correlated across sectors.²⁰ Accordingly, the Kalman gain matrix also has large diagonal entries. Table 2 reports the diagonal entries of the Kalman gain matrix. A higher Kalman gain in sector i indicates that hump-shaped AR(2) shocks are the more important driver for a downstream sector’s demand variation. We find that “Aerospace Products and Parts”, “Soft Drinks and Ice”, “Apparel”, “Medical Equipment and Supplies” and “Navigational/Measuring/Electromedical Instruments” are downstream sectors most exposed to hump-shaped shocks.

Figure 7. Correlation matrices for hump-shaped and monotone-decay shocks



Validation of the Structural Estimates against the Reduced-Form PACF Coefficients As footnote 16 shows, there is a monotone mapping from the PACF lag 2 coefficient ϕ_{i2} to the relative importance of AR(2) shocks in a sector, $(\sigma_i^u/\sigma_i^\epsilon)^2$. In the theoretical scenario where demand shocks are independent across sectors, this relationship implies a monotonic mapping from ϕ_{i2} to each sector's Kalman gain κ_i (see Appendix A.6 for derivation).

Figure 8 plots both the theoretical relationship (blue curve, assuming independence of demand shocks across sectors) and our empirical estimates of ϕ_{i2} (red scatter points) against the corresponding Kalman gains κ_i . Recall we empirically estimate ϕ_{i2} from a reduced-form regression of VA_{it} on VA_{it-2} , controlling for VA_{it-1} .

Although demand shocks across sectors are not strictly independent, the cross-sector correlations are relatively weak (c.f. Figure 7), with most variation coming from within-sector shocks.

²⁰For the covariance matrices Σ^u and Σ^ϵ , their diagonals explain 91% and 92% of the variation for hump-shaped and monotone-decay shocks respectively, according to the Frobenius norm.

Table 2. Estimation results: Kalman gains across downstream sectors


$\rho = 0.818$		κ
Transportation	Ship and Boat Building	0.18
	Heavy Duty Trucks	0.22
	Automobiles and Light Duty Motor Vehicles	0.26
	Motor Vehicle Bodies and Trailers	0.57
	Other Transportation Equipment	0.70
	Aerospace Products and Parts	0.76
Food Products	Fruit and Vegetable Preserving and Specialty Foods	0.07
	Coffee and Tea	0.08
	Sugar and Confectionery Products	0.30
	Breweries	0.42
	Other Food Except Coffee and Tea	0.46
	Animal Slaughtering and Processing	0.59
	Bakeries and Tortilla	0.68
	Soft Drinks and Ice	0.76
Other Final Goods	Major Electrical Household Appliances	0.12
	Office And Other Furniture	0.22
	Carpet and Rug Mills	0.31
	Tobacco	0.39
	Newspaper Publishers	0.42
	Soap, Cleaning Compounds, and Toilet Preparation	0.50
	Periodical, Book, and Other Publishers	0.50
	Pharmaceuticals and Medicines	0.50
	Apparel	0.77
	Medical Equipment and Supplies	0.83
	Navigational/Measuring/Electromedical Instruments	0.98

The empirical scatter points generally align with the theoretical curve, validating our model specification. The small deviations observed may be attributed to these mild cross-sector dependencies in demand shocks. Appendix Figure A.1 further confirms that the relationship between our empirical estimates of the PACF lag 1 coefficients ϕ_{i1} and the Kalman gains also aligns with the theoretical prediction.

The impulse-response exercises in Section 4.3 are done by sorting downstream sectors into those with above- and below-median PACF lag 2 coefficients. In Appendix D.3, we show the estimated impulse response functions are robust to sorting sectors by their Kalman gains.

Variance Decompositions The source of the bullwhip effect—that downstream demand shocks create amplified volatility in upstream sectors—are the hump-shaped AR(2) shocks to consumer demand. How quantitatively important are AR(2) shocks for supply chain volatility as a whole? To answer this, we conduct variance decompositions following equation (28) in Proposition 4. That equation implies that the volatility of demand for downstream goods, $Var_t(\tilde{\theta}_{t+1})$, consists of three components: the hump-shaped AR(2) demand shocks, the monotone-decay AR(1) shocks, and forecast error of the hidden states. Accordingly, for each upstream sector j along the supply chain of a downstream consumer good i , we compute sector j 's total volatility attributable to each of the three components of the demand volatility in good i . Then, for each downstream

Figure 8. Relationship between Kalman gain and partial autocorrelation function



sector i , we calculate the average share of volatility across sectors along i 's supply chain that is attributable to each of the three components.


Figure 9 shows the decomposition results. Each bar represents a downstream sector and has three portions that add up to one. The blue portion captures the supply chain volatility associated with the downstream sector's hump-shaped shocks. The red portion captures the supply chain volatility associated with the downstream sector's monotone-decay shocks. The green portion captures the supply chain volatility associated with the forecast error associated with not observing the hidden state. The downstream sectors are arranged in increasing order of their Kalman gains (i.e., the diagonals of the Kalman gain matrix), meaning that hump-shaped shocks tend to play a larger role in sectors towards the right-hand side.

Across supply chains in the US economy, hump-shaped AR(2) shocks account for 47.6% of volatility on average. The fraction ranges from 9.2% on the lower end (for the supply chain associated with "Coffee and Tea") to 92.2% on the higher end (for the supply chain associated with "Navigational/Measuring/Electromedical/Control Instruments"). The volatility due to incomplete information—the green portions—is small across all supply chains.²¹

Counterfactuals We now compute several counterfactuals based on the estimated model to demonstrate the importance of hump-shaped shocks and the time-to-build dynamics in generating supply chain volatility.

²¹Going from sectors on the left-hand side to those on the right, the green portions first increase and then decrease. This is because the forecast errors are greatest for intermediate levels of Kalman gain, i.e., when both types of demand shocks are important.

Figure 9. Variance decompositions along supply chains



In the first set of counterfactuals, we hold constant the downstream demand volatility $Var_t(\tilde{\theta}_{t+1})$, change the composition of demand shocks, and investigate the corresponding changes in the volatility of supply chains associated with broad groups of downstream sectors.

Column (1) of Table 3 shows the change in supply chain volatility when the volatility of a group of downstream sectors arises entirely from hump-shaped shocks. Specifically, consider row (a) in Table 3 labeled “Transportation”, which represents a group of downstream industries that belong to the broader transportation sector. For each industry j along the supply chains for the transportation sector, we compute $\widehat{Var}_t(\tilde{\gamma}_{j,t+1})$ by first setting the non-transportation entries of $Var_t(\tilde{\theta}_{t+1})$ to zero and then plugging it into equation (29). The resulting measure $\widehat{Var}_t(\tilde{\gamma}_{j,t+1})$ is the volatility of sector j attributable to demand volatility in the transportation sector given the Kalman gains based on our empirical estimates. We then compute the alternative measure $\widehat{Var}_t^{AR(2)}(\tilde{\gamma}_{j,t+1})$ by setting $K = I$ when using equation (29). This measure $\widehat{Var}_t^{AR(2)}(\tilde{\gamma}_{j,t+1})$ captures the counterfactual volatility in sector j that arises from AR(2) demand shocks to the transportation sector, while holding constant the levels of demand volatility in transportation. Column (1) of Table 3 then reports the ratio $\widehat{Var}_t^{AR(2)}(\tilde{\gamma}_{j,t+1}) / \widehat{Var}_t(\tilde{\gamma}_{j,t+1})$ averaged across all sectors j along the supply chains of the corresponding downstream groups.

We find that hump-shaped AR(2) shocks significantly amplifies sectoral volatility: for supply chains associated with transportation sectors, the volatility increases by 145.7%. The amplification is notably larger for transportation-related sectors compared to food products or other final goods, partly because transportation sectors have longer supply chains delays, as demonstrated

in columns (5) and (6). These columns report the average delay and the average delay among the top five longest delays along supply chains corresponding to each downstream sector. When demand shocks in all downstream sectors become hump-shaped, the overall supply chain volatility in the economy increases by 59.8% (row d, column 2).

Column (2) of Table 3 conducts an analogous exercise: we turn off hump-shaped shocks—setting the Kalman gain matrix $\mathbf{K}$ to zero, thereby ensuring that demand volatility is driven entirely by monotone-decay shocks—while holding downstream demand volatility constant. We then examine the implied changes in supply chain volatility associated with particular downstream industry groups. In this scenario, supply chain volatility declines by more than 69% across all downstream groups compared to the baseline.

Column (3) of Table 3 computes the implied supply chain volatility in a static model relative to our dynamic model, again holding demand volatility constant. Specifically, we use the formula $Var_t^{Static}(\tilde{\gamma}_{t+1}) = (\mathbf{I} - \mathbf{\Omega})^{-1} Var_t(\tilde{\theta}_{t+1})(\mathbf{I} - \mathbf{\Omega}')^{-1}$, which follows from relationship between sectoral revenue γ and demand θ in Acemoglu et al. (2012), given by $\gamma = (\mathbf{I} - \mathbf{\Omega})^{-1} \theta$. We find that a static model without time-to-build dynamics significantly understates the supply chain volatility attributable to downstream demand shocks—by more than 25.4% overall. The understatement is particularly pronounced (exceeding 50%) among the downstream group “Other Final Goods”, which includes “Apparel”, “Carpet and Rug Mills”, “Office And Other Furniture”, “Office And Other Furniture”, “Tobacco”, “Medical Equipment and Supplies” and “Navigational/Measuring/Electromedical Instruments”, “Newspaper Publishers”, “Soap, Cleaning Compounds, and Toilet Preparation”, “Periodical, Book, and Other Publishers” and “Pharmaceuticals and Medicines”.

In Appendix Table A.3, we report the outcomes for the same set counterfactual exercises, except that we examine each downstream industry individually rather than in groups.

Table 3. Counterfactual change in supply chain volatility based on the nature of downstream demand shocks

		% Change in supply chain volatility			Delays along supply chain	
		Only AR(2) shocks (1)	Only AR(1) shocks (2)	Static model (3)	Average (5)	Top 5 (6)
(a)	Transportation	145.73	-71.05	-25.38	2.32	3.84
(b)	Food Products	83.06	-69.09	-37.66	2.14	3.31
(c)	Other Final Goods	42.42	-77.12	-53.06	2.04	3.03
(d)	All Sectors	59.82	-69.08	-42.67	2.14	3.31

5 Conclusion

We develop a framework to analyze dynamic interactions over production networks, with two main ingredients: sectoral demand shocks and heterogeneous time-to-build delays. Each producer’s input choices depend on expectations of future demand all time horizons, arising from both exogenous consumer demand shocks and endogenous responses from other producers along the supply chain, each potentially facing distinct time-to-build delays. The model highlights how the bullwhip effect emerges when current demand shocks generate amplified expectations of future demand. We characterize the solution in closed form over primitives and show how equilibrium input decisions depend on the entire temporal structure of shocks and supply-chain linkages. Empirically and quantitatively, we show that the bullwhip is significant across downstream sectors that are important for final consumption.

References

- Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi**, “The Network Origins of Aggregate Fluctuations,” *Econometrica*, 2012, 80 (5), 1977–2016.
- Alfaro, Laura, Davin Chor, Pol Antras, and Paola Conconi**, “Internalizing Global Value Chains: A Firm-Level Analysis,” *Journal of Political Economy*, 2019, 127 (2), 508–559.
- Antràs, Pol and Vitalii Tubdenov**, “Measuring the Average Period of Production,” *Working Paper*, 2025.
- , **Davin Chor, Thibault Fally, and Russell Hillberry**, “Measuring the Upstreamness of Production and Trade Flows,” *American Economic Review*, 2012, 102 (3), 412–416.
- Baqaei, David Rezza and Emmanuel Farhi**, “The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten’s Theorem,” *Econometrica*, 2019, 87.
- and —, “Productivity and Misallocation in General Equilibrium,” *The Quarterly Journal of Economics*, February 2020, 135 (1), 105–163.
- Bayraktar, Erhan and Christian Keller**, “Path-dependent Hamilton–Jacobi equations in infinite dimensions,” *Journal of Functional Analysis*, 2018, 275 (8), 2096–2161.
- Beaudry, Paul and Franck Portier**, “Stock Prices, News, and Economic Fluctuations,” *American Economic Review*, 2006.
- Bigio, Saki and Jennifer La’O**, “Distortions in Production Networks,” *The Quarterly Journal of Economics*, November 2020, 135 (4), 2187–2253.
- Bizzarri, Matteo, Marco Pangallo, and Francisco Queiros**, “Production Networks, Time to Build, and Endogenous Oscillations,” *Working paper*, 2025.
- Blanchard, Olivier J., Jean-Paul L’Huillier, and Guido Lorenzoni**, “News, Noise, and Fluctuations: An Empirical Exploration,” *American Economic Review*, 2013.
- Boerma, Job, Georgii Riabov, and Aleh Tsyvinski**, “Making Memories: Memory Effects and Their Dynamics,” *Working paper*, 2024.

- Box, George EP, Gwilym M Jenkins, Gregory C Reinsel, and Greta M Ljung**, *Time Series Analysis: Forecasting and Control*, John Wiley & Sons, 2015.
- Carvalho, Vasco**, “Aggregate Fluctuations and the Network Structure of Intersectoral Trade,” *Working paper*, 2010.
- Elliott, Matthew, Benjamin Golub, and Matthew V Leduc**, “Supply Network Formation and Fragility,” *American Economic Review*, 2022, 112 (8), 2701–2747.
- Farmer, Leland E**, “The Discretization Filter: A Simple Way to Estimate Nonlinear State Space Models,” *Quantitative Economics*, 2021, 12 (1), 41–76.
- Ferrari, Alessandro**, “Inventories, Demand Shocks Propagation, and Amplification in Supply Chains,” *Working Paper*, 2025.
- Foerster, Andrew T, Pierre-Daniel G Sarte, and Mark W Watson**, “Sectoral Versus Aggregate Shocks: A Structural Factor Analysis of Industrial Production,” *Journal of Political Economy*, 2011, 119 (1), 1–38.
- Forrester, Jay W**, “Industrial Dynamics: A Major Breakthrough for Decision Makers,” *Harvard business review*, 1958, 36 (4), 37–66.
- Gabaix, Xavier**, “The Granular Origins of Aggregate Fluctuations,” *Econometrica*, 2011.
- Hamilton, James D**, “A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle,” *Econometrica: Journal of the econometric society*, 1989, pp. 357–384.
- Humpherys, Jeffrey, Preston Redd, and Jeremy West**, “A Fresh Look at the Kalman Filter,” *SIAM review*, 2012, 54 (4), 801–823.
- Ito, Kazufumi and Kaiqi Xiong**, “Gaussian Filters for Nonlinear Filtering Problems,” *IEEE transactions on automatic control*, 2000, 45 (5), 910–927.
- Jones, Charles I.**, “Intermediate Goods and Weak Links in the Theory of Economic Development,” *American Economic Journal: Macroeconomics*, April 2011, 3 (2), 1–28.
- , “Misallocation, Economic Growth, and Input-Output Economics,” *Advances in Economics and Econometrics*, 2013, 2.
- Jordà, Òscar and Alan M Taylor**, “Local Projections,” *Journal of Economic Literature*, 2025, 63 (1), 59–110.
- Kalman, Rudolph E and Richard S Bucy**, “New Results in Linear Filtering and Prediction Theory,” 1961.
- Kalman, Rudolph Emil**, “A New Approach to Linear Filtering and Prediction Problems,” 1960.
- Khan, Aubhik and Julia K Thomas**, “Inventories and the Business Cycle: An Equilibrium Analysis of (S, s) Policies,” *American Economic Review*, 2007, 97 (4), 1165–1188.
- **and** —, “Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment dynamics,” *Econometrica*, 2008, 76 (2), 395–436.
- **and** —, “Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity,” *Journal of Political Economy*, 2013, 121 (6), 1055–1107.

- Kopytov, Alexandr, Bineet Mishra, Kristoffer Nimark, and Mathieu Taschereau-Dumouchel**, “Endogenous Production Networks under Supply Chain Uncertainty,” *Econometrica*, 2024, 92 (5), 1621–1659.
- Kydland, Finn E and Edward C Prescott**, “Time to Build and Aggregate Fluctuations,” *Econometrica: Journal of the Econometric Society*, 1982, pp. 1345–1370.
- Lee, Hau L, Venkata Padmanabhan, and Seungjin Whang**, “Information Distortion in A Supply Chain: The Bullwhip Effect,” *Management science*, 1997, 43 (4), 546–558.
- , **Vineet Padmanabhan, and Seungjin Whang**, “Comments on Information distortion in a supply chain: The bullwhip effect;,” *Management science*, 2004, 50 (12_supplement), 1887–1893.
- Lehn, Christian Vom and Thomas Winberry**, “The Investment Network, Sectoral Comovement, and the Changing US Business Cycle,” *The Quarterly Journal of Economics*, 2022, 137 (1), 387–433.
- Liu, Ernest**, “Industrial Policies in Production Networks*,” *The Quarterly Journal of Economics*, November 2019, 134 (4), 1883–1948.
- **and Aleh Tsyvinski**, “A Dynamic Model of Input–Output Networks,” *Review of Economic Studies*, 2024, p. rdae012.
- **and Song Ma**, “Innovation Networks and R&D Allocation,” *Working Paper*, 2024.
- Long, John B and Charles I Plosser**, “Real Business Cycles,” *Journal of Political Economy*, 1983, 91 (1), 39–69.
- Long, John B. Jr and Charles I. Plosser**, “Real Business Cycles,” *Journal of Political Economy*, 1983.
- Lorenzoni, Guido**, “A Theory of Demand Shocks,” *American Economic Review*, 2009.
- , “News and Aggregate Demand Shocks,” *Annual Review of Economics*, 2011.
- Nikolakoudis, George**, “Incomplete Information in Production Networks,” *Working Paper*, 2024.
- Oberfield, Ezra**, “A Theory of Input-Output Architecture,” *Econometrica*, March 2018, 86.
- Pellet, Thomas and Alireza Tahbaz-Salehi**, “Rigid Production Networks,” *Journal of Monetary Economics*, 2023, 137, 86–102.
- Schaal, Edouard and Mathieu Taschereau-Dumouchel**, “Echoes and Delays: Time-to-Build in Production Networks,” *Working Paper*, 2025.
- Taschereau-Dumouchel, Mathieu**, “Cascades and Fluctuations in an Economy with an Endogenous Production Network,” *Working Paper*, 2020.

Appendix

Table of Contents

Section A : Proofs	
A.1	Proof of Theorem 1
A.2	Proof of Corollary 1
A.3	Proof of Proposition 2
A.4	Proof of Theorem 2
A.5	Proof of Proposition 3
A.6	Derivation of the Kalman gain matrix and the PACF coefficients
A.7	Proof of Theorem 3
A.8	Proof of Proposition 4

Section B : Theoretical Results not in the Main Text	
B.1	Decentralizing the Planner's Solution as the Competitive Equilibrium
B.2	Heterogeneous Rates of Convergence for Demand Shock Components

Section: C : MLE Algorithm Details	
--	--

Section D : Alternative Empirical Specifications and Robustness Checks	
D.1	Higher Volatility in Upstream Sectors Across Supply Chains
D.2	Downstream Sector's Own Impulse-Response: Pooled Specification
D.3	Classifying Downstream Sectors by Kalman Gains
D.4	Alternative Measures of Time-to-Build
D.5	Using Alternative Parameters to Construct Supply Chain Distance
D.6	Using Alternative Numbers of Lags When Estimating Impulse-Response Functions
D.7	Impulse-Response Using Long Difference Specification
D.8	Impulse-Response with Monthly Data
D.9	Controlling Supply Shock
D.10	Quantitative Implications Given Heterogeneous Rates of Convergence
D.11	Counterfactuals Each Downstream Industry Individually

A Proofs

A.1 Proof of Theorem 1

Proof. We iteratively substitute the left-hand side of (9) into the right-hand side to obtain

$$\begin{aligned}
\gamma_t &= \theta_t + \sum_{d=1}^{\infty} \beta^d \Omega_d \mathbb{E}_t [\gamma_{t+d}] \\
&= \theta_t + \mathbb{E}_t \left[\sum_{d=1}^{\infty} \beta^d \Omega_d \left(\theta_{t+d} + \sum_{d'=1}^{\infty} \beta^{d'} \Omega_{d'} \gamma_{t+d+d'} \right) \right] \\
&\vdots \\
&= \theta_t + \beta \Omega_1 \mathbb{E}_t [\theta_{t+1}] + \beta^2 (\Omega_1^2 + \Omega_2) \mathbb{E}_t [\theta_{t+2}] \\
&\quad + \beta^3 (\Omega_1^3 + \Omega_2 \Omega_1 + \Omega_1 \Omega_2 + \Omega_3) \mathbb{E}_t [\theta_{t+3}] + \dots \\
&= \theta_t + \sum_{s=1}^{\infty} \mathbf{G}_s \mathbb{E}_t [\theta_{t+s}], \tag{A1}
\end{aligned}$$

where $\mathbf{G}_s \equiv \beta^s \sum_{\phi \in \Phi_s} \prod_{\phi_j \in \phi} \Omega_{\phi_j}$. Note that the proof of Corollary 1 below implies that $\mathbf{I} + \sum_{s=1}^{\infty} \mathbf{G}_s = (\mathbf{I} - \sum_{d=1}^{\infty} \beta^d \Omega_d)^{-1}$, where invertibility is implied by $[\sum_{d=0}^{\infty} \beta^d \Omega_d]_{ji} \leq \Omega_{ji}$, $\sum_j \Omega_{ji} \leq 1$, and that $\mathbf{I} - \Omega$ is invertible (see footnote 2). This and the fact that demand is mean-reverting in the long-run (i.e., $\lim_{s \rightarrow \infty} \mathbb{E}_t [\theta_{t+s}] = \bar{\theta}$ exists) directly imply that the infinite series in (A1) always converge. $\square$

A.2 Proof of Corollary 1

Proof. Let $\mathbf{G}_{\infty} \equiv \mathbf{I} + \sum_{s=1}^{\infty} \mathbf{G}_s$. To show $\mathbf{G}_{\infty} = (\mathbf{I} - \sum_{d=1}^{\infty} \beta^d \Omega_d)^{-1}$, note

$$\mathbf{G}_{\infty} = \mathbf{I} + \beta \Omega_1 \mathbf{G}_{\infty} + \beta^2 \Omega_2 \mathbf{G}_{\infty} + \dots = \mathbf{I} + \left(\sum_{d=0}^{\infty} \beta^d \Omega_d \right) \mathbf{G}_{\infty}.$$

Rearranging yields the result. Note that the invertibility of $(\mathbf{I} - \sum_{d=0}^{\infty} \beta^d \Omega_d)$ follows from the invertibility of $\mathbf{I} - \Omega$ and that for all entries, $[\sum_{d=0}^{\infty} \beta^d \Omega_d]_{ji} \leq \Omega_{ji}$. $\square$

A.3 Proof of Proposition 2

Proof. We start from taking logs of the production function (2):

$$\ln y_{it} = \ln z_{it} + \alpha_i \ln \ell_{it} + \sum_j \omega_{ij} \ln m_{ij,t-d_{ij}}$$

We obtain (15) by substituting out $\frac{m_{ijt}}{y_{jt}}$ and $\frac{\ell_{jt}}{\ell_t}$ using (8), and we obtain (14) by summing (5) across j and substitute out w_t using (6). $\square$

A.4 Proof of Theorem 2

Proof. Equation (19) and Theorem 1 jointly imply that

$$\tilde{\gamma}_t = \tilde{\theta}_t + \sum_{s=1}^{\infty} \mathbf{G}_s \left(\rho^s (\tilde{\theta}_t - \tilde{\mathbf{x}}_t) + (s+1) \rho^s \tilde{\mathbf{x}}_t \right).$$

Let $\tilde{\gamma}_t = \mathbf{G}_{\infty}^{\epsilon} (\tilde{\theta}_t - \tilde{\mathbf{x}}_t) + \mathbf{G}_{\infty}^x \tilde{\mathbf{x}}_t$. The matrix $\mathbf{G}_{\infty}^{\epsilon}$ summarizes the impact of demand's monotone-decay component on revenue, and $\mathbf{G}_{\infty}^x$ captures the impact of the persistent component. We derive these two components separately.

To derive $\mathbf{G}_{\infty}^{\epsilon}$, we can follow the same logic as in the proof for Proposition 1 and decompose the impact of demand's monotone-decay component recursively as the direct and indirect network effects, where each round of the network effect is summarized by $\mathbf{G}_{\infty}^{\epsilon}$, discounted to the presented by both the discount rate (β) and also the rate at which the shock dissipates (ρ). Specifically, note

$$\begin{aligned} \mathbf{G}_{\infty}^{\epsilon} &= \mathbf{I} + \rho\beta\mathbf{\Omega}_1 + (\rho\beta)^2 (\mathbf{\Omega}_1^2 + \mathbf{\Omega}_2) + (\rho\beta)^3 (\mathbf{\Omega}_1^3 + \mathbf{\Omega}_1\mathbf{\Omega}_2 + \mathbf{\Omega}_2\mathbf{\Omega}_1 + \mathbf{\Omega}_3) + \dots \\ &= \mathbf{I} + \rho\beta\mathbf{\Omega}_1\mathbf{G}_{\infty}^{\epsilon} + (\rho\beta)^2 \mathbf{\Omega}_2\mathbf{G}_{\infty}^{\epsilon} + \dots \\ &= \mathbf{I} + \left(\sum_{d=1}^{\infty} (\rho\beta)^d \mathbf{\Omega}_d \right) \mathbf{G}_{\infty}^{\epsilon}, \end{aligned}$$

and re-arranging yields the that $\mathbf{G}_{\infty}^{\epsilon} = \left(\mathbf{I} - \sum_{d=1}^{\infty} (\rho\beta)^d \mathbf{\Omega}_d \right)^{-1}$.

It is slightly more involved to derive $\mathbf{G}_{\infty}^x$ because the impact of a persistent shock does not decay monotonically; instead, the impact has an arithmetically increasing component as well as an exponentially decaying component. We can write $\mathbf{G}_{\infty}^x$ as

$$\mathbf{G}_{\infty}^x = \mathbf{I} + 2\rho\beta\mathbf{\Omega}_1 + 3(\rho\beta)^2 (\mathbf{\Omega}_1^2 + \mathbf{\Omega}_2) + 4(\rho\beta)^3 (\mathbf{\Omega}_1^3 + \mathbf{\Omega}_1\mathbf{\Omega}_2 + \mathbf{\Omega}_2\mathbf{\Omega}_1 + \mathbf{\Omega}_3) + \dots$$

1. The first term $\mathbf{I}$ captures the impact of the persistent component $\tilde{\mathbf{x}}$ on time- t revenue through concurrent demand.
2. The second term $2\rho\beta\mathbf{\Omega}_1$ captures the impact of $\tilde{\mathbf{x}}_t$ on time- t revenue through expected demand at $t+1$. The impact of $\tilde{\mathbf{x}}$ on $t+1$ demand is captured by 2ρ . The impact of $t+1$ demand on time- t revenue goes through input-output linkages with one-period delay and is captured by $\mathbf{\Omega}_1$.
3. The third term $3(\rho\beta)^2 (\mathbf{\Omega}_1^2 + \mathbf{\Omega}_2)$ captures the impact of $\tilde{\mathbf{x}}_t$ on time- t revenue through expected demand at $t+2$. The impact of $\tilde{\mathbf{x}}$ on $t+2$ demand is captured by $3\rho^2$. The $t+2$ demand on time- t revenue goes through input-output linkages with two periods delay and is captured by $\mathbf{\Omega}_1^2 + \mathbf{\Omega}_2$. The term $\mathbf{\Omega}_1^2$ captures the linkages from inputs supplied by the focal sector to the sector experiencing the demand shock through another intermediate

producer, with two walks in total and one-period delay along each walk. The term Ω_2 captures the linkages from the focal sector directly to the sector experiencing the demand shock, with one walk of two periods delay.

Each successive term in $\mathbf{G}_\infty^x$ summarizes the impact of $\tilde{\mathbf{x}}_t$ on time- t revenue through demand at a future time $t + s$, accounting for all possible direct and indirect network linkages with a delay of s periods in total. To derive $\mathbf{G}_\infty^x$ in closed-form, we re-write it recursively.

Define $\mathbf{Z} \equiv \sum_{d=1}^{\infty} d(\rho\beta)^d \Omega_d$. We can then write $\mathbf{G}_\infty^x$ as

$$\begin{aligned}
\mathbf{G}_\infty^x &= \mathbf{I} + \rho\beta\Omega_1 (2\mathbf{I} + 3\rho\beta\Omega_1 + 4(\rho\beta)^2 (\Omega_1^2 + \Omega_2) + \dots) \\
&\quad + (\rho\beta)^2 \Omega_2 (3\mathbf{I} + 4\rho\beta\Omega_1 + 5(\rho\beta)^2 (\Omega_1^2 + \Omega_2) + \dots) \\
&\quad + (\rho\beta)^3 \Omega_3 (4\mathbf{I} + 5\rho\beta\Omega_1 + 6(\rho\beta)^2 (\Omega_1^2 + \Omega_2) + \dots) \\
&= \mathbf{I} + \rho\beta\Omega_1 (\mathbf{G}_\infty^\epsilon + \mathbf{G}_\infty^x) + (\rho\beta)^2 \Omega_2 (2\mathbf{G}_\infty^\epsilon + \mathbf{G}_\infty^x) + (\rho\beta)^3 \Omega_3 (3\mathbf{G}_\infty^\epsilon + \mathbf{G}_\infty^x) + \dots \\
&= \mathbf{I} + \mathbf{Z}\mathbf{G}_\infty^\epsilon + \sum_{d=1}^{\infty} (\rho\beta)^d \Omega_d \mathbf{G}_\infty^x \\
&= \left(\mathbf{I} - \sum_{d=1}^{\infty} (\rho\beta)^d \Omega_d \right)^{-1} (\mathbf{I} + \mathbf{Z}\mathbf{G}_\infty^\epsilon) \\
&= \mathbf{G}_\infty^\epsilon + \mathbf{G}_\infty^\epsilon \mathbf{Z} \mathbf{G}_\infty^\epsilon
\end{aligned}$$

To derive the impulse-response, equation (18) implies that

$$\begin{aligned}
\partial \tilde{\theta}_{t+s} / \partial \epsilon_t &= \rho^s, \quad \partial \tilde{\mathbf{x}}_{t+s} / \partial \epsilon_t = 0, \\
\partial \tilde{\theta}_{t+s} / \partial \mathbf{u}_t &= (s+1) \rho^s, \quad \text{and} \quad \partial \tilde{\mathbf{x}}_{t+s} / \partial \epsilon_t = \rho^s.
\end{aligned}$$

Hence

$$\begin{aligned}
\frac{\partial \tilde{\gamma}_{t+s}}{\partial \epsilon_t} &= \mathbf{G}_\infty^\epsilon \frac{\partial \tilde{\theta}_{t+s}}{\partial \epsilon_t} = \rho^s \mathbf{G}_\infty^\epsilon, \\
\frac{\partial \tilde{\gamma}_{t+s}}{\partial \mathbf{u}_t} &= \mathbf{G}_\infty^\epsilon \frac{\partial \tilde{\theta}_{t+s}}{\partial \mathbf{u}_t} + (\mathbf{G}_\infty^x - \mathbf{G}_\infty^\epsilon) \frac{\partial \tilde{\mathbf{x}}_{t+s}}{\partial \mathbf{u}_t} \\
&= (s+1) \rho^s \mathbf{G}_\infty^\epsilon + \rho^s \mathbf{G}_\infty^\epsilon \left(d \sum_{d=1}^{\infty} (\rho\beta)^d \Omega_d \right) \mathbf{G}_\infty^\epsilon, \quad \text{as desired.}
\end{aligned}$$

□

A.5 Proof of Proposition 3

Proof. Equation (20) implies that

$$\begin{aligned}
Var_t(\tilde{\gamma}_{t+1}) &= Var_t\left(G_\infty^\epsilon\left(\tilde{\theta}_{t+1} - \mathbf{x}_{t+1}\right) + G_\infty^x \mathbf{x}_{t+1}\right) \\
&= Var_t\left(G_\infty^\epsilon\left(\rho\tilde{\theta}_t + \epsilon_{t+1}\right) + G_\infty^x \mathbf{u}_{t+1}\right) \\
&= G_\infty^\epsilon \Sigma_t^\epsilon G_\infty^{\epsilon'} + G_\infty^x \Sigma_t^u G_\infty^{x'}, \quad \text{as desired.}
\end{aligned}$$

□

A.6 Derivation of Kalman Gains and PACF Coefficients

We observe $\tilde{\theta}_t$, and the state variables are $\tilde{\theta}_t$ and $\mathbf{x}_t$. We write down the measurement equation and state equation:

$$\begin{aligned}
S_t \equiv \begin{bmatrix} \tilde{\theta}_t \\ \mathbf{x}_t \end{bmatrix} &= \underbrace{\begin{bmatrix} \rho I & \rho I \\ 0 & \rho I \end{bmatrix}}_{\equiv M} S_{t-1} + \underbrace{\begin{bmatrix} I & I \\ 0 & I \end{bmatrix}}_{\equiv G} \begin{bmatrix} \epsilon_t \\ \mathbf{u}_t \end{bmatrix} \\
\tilde{\theta}_t &= \underbrace{\begin{bmatrix} I & 0 \end{bmatrix}}_{\equiv H'} S_t
\end{aligned}$$

Updating equation:

$$S_{t|t} = M S_{t-1|t-1} + P_t \left(\tilde{\theta}_t - H' M S_{t-1|t-1} \right),$$

where

$$P_t = \Sigma_{t|t-1} H \left(H' \Sigma_{t|t-1} H \right)^{-1}$$

and the conditional variance terms evolve according to

$$\Sigma_{t|t} = (I - P_t H') \Sigma_{t|t-1}$$

$$\Sigma_{t+1|t} = M \Sigma_{t|t} M' + G Q G'$$

where $Q = \begin{bmatrix} \Sigma^\epsilon & 0 \\ 0 & \Sigma^u \end{bmatrix}$. As $t \rightarrow \infty$, P_t , $\Sigma_{t|t}$, and $\Sigma_{t+1|t}$ converge to:

$$P \equiv \lim_{t \rightarrow \infty} P_t = \hat{\Sigma} H \left(H' \hat{\Sigma} H \right)^{-1},$$

$$\Sigma \equiv \lim_{t \rightarrow \infty} \Sigma_{t|t} = (I - P H') \hat{\Sigma},$$

$$\hat{\Sigma} \equiv \lim_{t \rightarrow \infty} \Sigma_{t+1|t} = M \hat{\Sigma} M' + G Q G'.$$

We know that matrices Σ and P must take the form $\Sigma = \begin{bmatrix} 0 & 0 \\ 0 & F \end{bmatrix}$ and $P = \begin{bmatrix} I \\ K \end{bmatrix}$, where F is the variance for the nowcast and K is the Kalman gain. Substituting these forms into P , Σ ,

and $\hat{\Sigma}$ enables us to derive $\mathbf{F}$ and $\mathbf{K}$:

$$\begin{aligned} \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{F} \end{bmatrix} &= \Sigma = (\mathbf{I} - \mathbf{P}\mathbf{H}') \hat{\Sigma} = \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ -\mathbf{K} & \mathbf{I} \end{bmatrix} \hat{\Sigma} \\ \hat{\Sigma} &= \rho^2 \begin{bmatrix} \mathbf{I} & \mathbf{I} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{F} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{I} & \mathbf{I} \end{bmatrix} + \begin{bmatrix} \mathbf{I} & \mathbf{I} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \Sigma^\epsilon & \mathbf{0} \\ \mathbf{0} & \Sigma^u \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{I} & \mathbf{I} \end{bmatrix} \\ &= \rho^2 \begin{bmatrix} \mathbf{F} & \mathbf{F} \\ \mathbf{F} & \mathbf{F} \end{bmatrix} + \begin{bmatrix} \Sigma^\epsilon + \Sigma^u & \Sigma^u \\ \Sigma^u & \Sigma^u \end{bmatrix} \\ &= \begin{bmatrix} \Sigma^\epsilon + \rho^2 \mathbf{F} + \Sigma^u & \rho^2 \mathbf{F} + \Sigma^u \\ \rho^2 \mathbf{F} + \Sigma^u & \rho^2 \mathbf{F} + \Sigma^u \end{bmatrix} \end{aligned}$$

Substitute out $\hat{\Sigma}$ from the right-hand side of Σ to get

$$\begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{F} \end{bmatrix} = \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ -\mathbf{K} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \Sigma^\epsilon + \rho^2 \mathbf{F} + \Sigma^u & \rho^2 \mathbf{F} + \Sigma^u \\ \rho^2 \mathbf{F} + \Sigma^u & \rho^2 \mathbf{F} + \Sigma^u \end{bmatrix}$$

which implies

$$\mathbf{F} = (\mathbf{I} - \mathbf{K}) (\rho^2 \mathbf{F} + \Sigma^u) = \Sigma^\epsilon (\Sigma^\epsilon + \rho^2 \mathbf{F} + \Sigma^u)^{-1} (\rho^2 \mathbf{F} + \Sigma^u).$$

The form of $\mathbf{P}$ enables us to obtain the Kalman gain matrix:

$$\begin{bmatrix} \mathbf{I} \\ \mathbf{K} \end{bmatrix} = \mathbf{P} = \hat{\Sigma} \mathbf{H} (\mathbf{H}' \hat{\Sigma} \mathbf{H})^{-1} = \begin{bmatrix} \mathbf{I} \\ (\rho^2 \mathbf{F} + \Sigma^u) (\Sigma^\epsilon + \rho^2 \mathbf{F} + \Sigma^u)^{-1} \end{bmatrix}$$

which implies

$$\mathbf{K} = (\rho^2 \mathbf{F} + \Sigma^u) (\Sigma^\epsilon + \rho^2 \mathbf{F} + \Sigma^u)^{-1},$$

as desired.

In the theoretical scenario where demand shocks are independent across sectors, each sector's Kalman gain κ_i simplifies to:

$$\kappa_i = \frac{-\left(1 + (\sigma_i^u / \sigma_i^\epsilon)^2 - \rho^2\right) + \sqrt{\left(1 + (\sigma_i^u / \sigma_i^\epsilon)^2 - \rho^2\right)^2 + 4\rho^2 (\sigma_i^u / \sigma_i^\epsilon)^2}}{2\rho^2}$$

Hence, there is a monotone mapping from Kalman gain κ_i to the relative importance of AR(2) shocks in each sector, $(\sigma_i^u / \sigma_i^\epsilon)^2$.

We next derive the PACF coefficients at lags 1 and 2. According to the Durbin–Levinson Algorithm,

$$\phi_{i,1} = \text{Corr}(\theta_{it}, \theta_{it-1}) = \rho + \frac{\frac{\rho}{(1-\rho^2)^2} (\sigma_i^u / \sigma_i^\epsilon)^2}{\frac{1+\rho^2}{(1-\rho^2)^3} (\sigma_i^u / \sigma_i^\epsilon)^2 + \frac{1}{1-\rho^2}},$$


$$\phi_{i,2} = \text{Corr}(\theta_{it}, \theta_{it-2} | \theta_{it-1}) \quad (\text{A2})$$

$$= \frac{\text{Corr}(\theta_{it}, \theta_{it-2}) - \phi_{i,1}^2}{1 - \phi_{i,1}^2} \quad (\text{A3})$$

$$= \left(\rho^2 + \frac{\frac{2\rho^2}{(1-\rho^2)^2} (\sigma_i^u / \sigma_i^\epsilon)^2}{\frac{1+\rho^2}{(1-\rho^2)^3} (\sigma_i^u / \sigma_i^\epsilon)^2 + \frac{1}{(1-\rho^2)}} - \phi_{i,1}^2 \right) / (1 - \phi_{i,1}^2). \quad (\text{A4})$$

Figure A.1 plots both the theoretical relationship (blue curve, assuming independence of demand shocks across sectors) and our empirical estimates of PACF at lag 1 (ϕ_{i1} , red scatter points) against the corresponding estimated Kalman gains κ_i .

Figure A.1. PACF and Kalman gain



A.7 Proof of Theorem 3

Proof. Equation (25) states that

$$\hat{\mathbf{x}}_t = \rho \hat{\mathbf{x}}_{t-1} + \mathbf{K} \left(\tilde{\boldsymbol{\theta}}_t - \rho \left(\tilde{\boldsymbol{\theta}}_{t-1} + \hat{\mathbf{x}}_{t-1} \right) \right)$$

Hence,

$$\mathbb{E}_t \left[\tilde{\boldsymbol{\theta}}_{t+s} \right] = \rho^s \tilde{\boldsymbol{\theta}}_t + s \rho^s \hat{\mathbf{x}}_t$$

The demand innovation at time t is defined as

$$\mathbf{v}_t \equiv \tilde{\boldsymbol{\theta}}_t - \mathbb{E}_{t-1} [\tilde{\boldsymbol{\theta}}_t] = \tilde{\boldsymbol{\theta}}_t - \rho (\tilde{\boldsymbol{\theta}}_{t-1} + \hat{\mathbf{x}}_{t-1})$$

We have

$$\begin{aligned} \mathbb{E}_{t-1} [\tilde{\boldsymbol{\theta}}_{t+s}] &= \rho^{s+1} (\tilde{\boldsymbol{\theta}}_{t-1} + \hat{\mathbf{x}}_{t-1}) + s\rho^{s+1} \hat{\mathbf{x}}_{t-1} \\ \mathbb{E}_t [\tilde{\boldsymbol{\theta}}_{t+s}] &= \rho^s \tilde{\boldsymbol{\theta}}_t + s\rho^s \hat{\mathbf{x}}_t \\ &= \rho^s (\mathbb{E}_{t-1} [\tilde{\boldsymbol{\theta}}_t] + \mathbf{v}_t) + s\rho^s (\rho \hat{\mathbf{x}}_{t-1} + \mathbf{K} \mathbf{v}_t) \\ \mathbb{E}_t [\mathbf{x}_{t+s}] &= \rho^s \hat{\mathbf{x}}_t = \rho^{s+1} \hat{\mathbf{x}}_{t-1} + \rho^s \mathbf{K} \mathbf{v}_t \end{aligned}$$

Hence

$$\frac{\partial \mathbb{E}_t [\tilde{\boldsymbol{\theta}}_{t+s}]}{\partial \mathbf{v}_t} = \rho^s \mathbf{I} + s\rho^s \mathbf{K}, \quad \frac{\partial \mathbb{E}_t [\mathbf{x}_{t+s}]}{\partial \mathbf{v}_t} = \rho^s \mathbf{K}.$$

Since

$$\mathbb{E}_t [\tilde{\boldsymbol{\gamma}}_{t+s}] = \mathbf{G}_\infty^\epsilon (\mathbb{E}_t [\tilde{\boldsymbol{\theta}}_{t+s}] - \mathbb{E}_t [\tilde{\mathbf{x}}_{t+s}]) + \mathbf{G}_\infty^x \mathbb{E}_t [\tilde{\mathbf{x}}_{t+s}]$$

We have

$$\begin{aligned} \frac{\partial \mathbb{E}_t [\tilde{\boldsymbol{\gamma}}_{t+s}]}{\partial \mathbf{v}_t} &= \mathbf{G}_\infty^\epsilon \left(\frac{\partial \mathbb{E}_t [\tilde{\boldsymbol{\theta}}_{t+s}]}{\partial \mathbf{v}_t} - \frac{\partial \mathbb{E}_t [\tilde{\mathbf{x}}_{t+s}]}{\partial \mathbf{v}_t} \right) + \mathbf{G}_\infty^x \frac{\partial \mathbb{E}_t [\tilde{\mathbf{x}}_{t+s}]}{\partial \mathbf{v}_t} \\ &= \mathbf{G}_\infty^\epsilon (\rho^s \mathbf{I} + s\rho^s \mathbf{K} - \rho^s \mathbf{K}) + \mathbf{G}_\infty^x \rho^s \mathbf{K} \\ &= \rho^s \mathbf{G}_\infty^\epsilon (\mathbf{I} + (s-1) \mathbf{K}) + \mathbf{G}_\infty^x \rho^s \mathbf{K}, \end{aligned}$$

as desired. □

A.8 Proof of Proposition 4

Proof. Under incomplete information,

$$\mathbb{E}_t [\tilde{\boldsymbol{\theta}}_{t+s}] = \rho^s \tilde{\boldsymbol{\theta}}_t + s\rho^s \hat{\mathbf{x}}_t$$

Substitute this into equation (A1) and follow the same derivation as in the proof of Theorem 2, we get that under incomplete information,

$$\tilde{\boldsymbol{\gamma}}_t = \mathbf{G}_\infty^\epsilon (\tilde{\boldsymbol{\theta}}_t - \hat{\mathbf{x}}_t) + \mathbf{G}_\infty^x \hat{\mathbf{x}}_t.$$

Hence,

$$\text{Var}_t (\tilde{\boldsymbol{\gamma}}_{t+1}) = \text{Var}_t (\mathbf{G}_\infty^\epsilon (\tilde{\boldsymbol{\theta}}_{t+1} - \hat{\mathbf{x}}_{t+1}) + \mathbf{G}_\infty^x \hat{\mathbf{x}}_{t+1})$$

Since $\tilde{\theta}_{t+1} = \rho\tilde{\theta}_t + \rho\mathbf{x}_t + \mathbf{u}_{t+1} + \epsilon_{t+1}$, $\hat{\mathbf{x}}_{t+1} = \rho\hat{\mathbf{x}}_t + \mathbf{K}(\tilde{\theta}_{t+1} - \rho(\tilde{\theta}_t + \hat{\mathbf{x}}_t))$ and $Var_t(\mathbf{x}_t) = \rho^2\mathbf{F}$, we have:

$$\begin{aligned} Var_t(\tilde{\gamma}_{t+1}) &= Var_t\left(\mathbf{G}_\infty^\epsilon(\rho\mathbf{x}_t + \mathbf{u}_{t+1} + \epsilon_{t+1} - \mathbf{K}\rho\mathbf{x}_t - \mathbf{K}\mathbf{u}_{t+1} - \mathbf{K}\epsilon_{t+1})\right. \\ &\quad \left.+ \mathbf{G}_\infty^x(\mathbf{K}\rho\mathbf{x}_t + \mathbf{K}\mathbf{u}_{t+1} + \mathbf{K}\epsilon_{t+1})\right) \\ &= (\mathbf{G}_\infty^\epsilon - \mathbf{G}_\infty^\epsilon\mathbf{K} + \mathbf{G}_\infty^x\mathbf{K})(\rho^2\mathbf{F} + \Sigma^\epsilon + \Sigma^u)(\mathbf{G}_\infty^\epsilon - \mathbf{G}_\infty^\epsilon\mathbf{K} + \mathbf{G}_\infty^x\mathbf{K})', \quad \text{as desired.} \end{aligned}$$

□

B Theoretical Results not in the Main Text

B.1 Decentralizing the Planner's Solution as the Competitive Equilibrium

We now set up the competitive equilibrium to decentralize the planner's solution, and we show that the Lagrange multipliers in the planner's solution indeed correspond to prices.

Consumer Problem At each time t , the consumer solves

$$W_t(a_t) = \max_{\{c_{it}, \bar{\ell}_t\}} \sum_i \theta_{it} \ln c_{it} - v(\bar{\ell}_t) + \beta \mathbb{E}_t[W_{t+1}(a_{t+1})]$$

subject to the budget constraint

$$\sum_i p_{it}c_{it} + a_{t+1}/R_t = w_t\bar{\ell}_t + a_t + \Pi_t,$$

where R_t is the interest rate, a_t is savings in terms of a numeraire bond in zero net supply, p_{it} is the price of good i , w_t is the wage rate, and Π_t is the firm profits. We choose the numeraire good so that the marginal utility of income in any period t (i.e., the Lagrange multiplier on the budget constraint in each period) is always equal to one. Given the normalization, the intra-temporal Euler equation is

$$\theta_{it} = p_{it}c_{it}, \tag{A5}$$

and the inter-temporal Euler equation is

$$1 = \beta R_t.$$

Producer Problem The representative producer in sector i solves

$$V_{it} \left(\{m_{ijt-q}\}_q \right) = \max_{\ell_{it}, \{m_{ijt}\}} p_{it}y_{it} - w_t\ell_{it} - \sum_j p_{jt}m_{ijt} + \mathbb{E}_t \left[\Lambda_{t,t+1} V_{it+1} \left(\{m_{ijt+1-q}\}_q \right) \right]$$

$$\max_{m_{ijt}} \mathbb{E}_t \left[\sum_s \Lambda_{t,t+s} \max_{\ell_{it+s}} (p_{it+s}y_{it+s} - w_{t+s}\ell_{it+s}) - p_{jt}m_{ijt} \right]$$

where $\Lambda_{t,t+s}$ is the stochastic discount factor. Note that given log-linear utility, the stochastic discount factor is equal to $\Lambda_{t,t+s} = \beta^s$ in all states—across all realizations of the productivity or demand shocks. Specifically, $\Lambda_{t,t+s} = \beta^s \frac{\theta_{i,t+s}c_{it}p_{it}}{\theta_{it}c_{i,t+s}p_{i,t+s}}$ for all i , which simplifies to β^s given (A5).

Planner’s Lagrange Multipliers By comparing the equilibrium conditions with the planner’s solution, it is evident that the planner’s Lagrange multipliers coincide with equilibrium prices.

Remark on the Choice of Numeraire and Our Characterization The object γ_t in the planner’s solution coincides with the vector of sectoral revenues in the decentralized equilibrium with the marginal utility of income normalized to one in each period. This normalization is equivalent to choosing a numeraire in each period and is therefore inconsequential for analyzing the relative revenue across sectors at each time (i.e., γ_{it}/γ_{jt}). The nominal interest rate is $1/\beta$ under our choice of numeraire. Theorem 1 further implies that the relative sectoral revenue is invariant to the labor supply function $v(\cdot)$ and sectoral productivities. The labor supply function and sectoral productivities affect quantities—as analyzed in section 2.2—but they do not affect relative revenue across sectors. The property also holds in static production network models with log-linear preferences and production functions (Acemoglu et al. 2012) as well as in dynamic production network models with log-linear adjustment costs (Liu and Tsyvinski 2024). Our analysis here establishes this property in a dynamic production network model with heterogeneous time-to-build.

B.2 Heterogeneous Rates of Convergence for Demand Shock Components

Our analysis thus far has focused on the case where both the monotone-decay and hump-shaped components of demand shocks have the same parameter ρ capturing the rate of convergence. In this section, we generalize our results in section 3 to the case where the two demand shock components have heterogeneous rates of convergence.

Specifically, we extend the demand shock process described by equations (16) and (17) as follows:

$$\theta_{it} - \bar{\theta}_i = \rho_\epsilon (\theta_{it-1} - \bar{\theta}_i) + x_{it} + \epsilon_{it}, \quad (\text{A6})$$

$$x_{it} = \rho_x x_{i,t-1} + u_{it}. \quad (\text{A7})$$

where ρ_ϵ defines the rate at which demand reverts back to the steady-state under monotone-decay shocks and ρ_x defines the rate of convergence under hump-shaped shocks. The previous case is nested when $\rho_\epsilon = \rho_x = \rho$.

The expected future demand can be then determined as

$$\mathbb{E}_t [\tilde{\theta}_{t+s}] = \rho_\epsilon^s (\tilde{\theta}_t - \tilde{x}_t) + \frac{\rho_\epsilon^{s+1} - \rho_x^{s+1}}{\rho_\epsilon - \rho_x} \tilde{x}_t$$

which generalizes equation (19) of the baseline model to the case of heterogenous persistence. We now generalize Theorems 2 and Proposition 3 to this setting.

Proposition 5. *Under complete information, with demand characterized by equations (A6) and (A7), sectoral revenue at time t is given by*

$$\tilde{\gamma}_t = \mathbf{G}_\infty^\epsilon (\tilde{\theta}_t - \tilde{x}_t) + \mathbf{G}_\infty^x \tilde{x}_t, \quad (\text{A8})$$

where $\mathbf{G}_\infty^\epsilon$ determines the impact of demand's monotone-decay component on revenue, and $\mathbf{G}_\infty^x$ determines the impact of demand's hump-shaped component, and

$$\mathbf{G}_\infty^\epsilon \equiv \left(\mathbf{I} - \sum_{d=1}^{\infty} (\rho_\epsilon \beta)^d \boldsymbol{\Omega}_d \right)^{-1}, \quad \mathbf{G}_\infty^x \equiv \frac{\rho_\epsilon}{\rho_\epsilon - \rho_x} \mathbf{G}_\infty^\epsilon - \frac{\rho_x}{\rho_\epsilon - \rho_x} \left(\mathbf{I} - \sum_{d=1}^{\infty} (\rho_x \beta)^d \boldsymbol{\Omega}_d \right)^{-1}.$$

The impulse-response functions are:

$$\frac{\partial \mathbb{E}_t [\tilde{\gamma}_{t+s}]}{\partial \epsilon_t} = \rho_\epsilon^s \mathbf{G}_\infty^\epsilon, \quad \frac{\partial \mathbb{E}_t [\tilde{\gamma}_{t+s}]}{\partial u_t} = \mathbf{G}_\infty^\epsilon \frac{\rho_\epsilon^{s+1} - \rho_x^{s+1}}{\rho_\epsilon - \rho_x} + (\mathbf{G}_\infty^x - \mathbf{G}_\infty^\epsilon) \rho_x^s.$$

The revenue volatility under complete information is $\text{Var}_t (\tilde{\gamma}_{t+1}) = \mathbf{G}_\infty^\epsilon \boldsymbol{\Sigma}_t^\epsilon \mathbf{G}_\infty^{\epsilon'} + \mathbf{G}_\infty^x \boldsymbol{\Sigma}_t^u \mathbf{G}_\infty^{x'}$.

The revenue volatility under incomplete information is

$$\text{Var}_t (\tilde{\gamma}_{t+1}) = (\mathbf{G}_\infty^\epsilon - \mathbf{G}_\infty^\epsilon \mathbf{K} + \mathbf{G}_\infty^x \mathbf{K}) (\rho_x^2 \mathbf{F} + \boldsymbol{\Sigma}^\epsilon + \boldsymbol{\Sigma}^u) (\mathbf{G}_\infty^\epsilon - \mathbf{G}_\infty^\epsilon \mathbf{K} + \mathbf{G}_\infty^x \mathbf{K})'. \quad (\text{A9})$$

Proof. The expected future demand can be determined as

$$\mathbb{E}_t [\tilde{\theta}_{t+s}] = \rho_\epsilon^s (\tilde{\theta}_t - \tilde{x}_t) + \frac{\rho_\epsilon^{s+1} - \rho_x^{s+1}}{\rho_\epsilon - \rho_x} \tilde{x}_t \quad (\text{A10})$$

To derive $\mathbf{G}_\infty^\epsilon$, we can follow the same logic as in the proof for Proposition 1 and decompose the impact of demand's monotone-decay component recursively as the direct and indirect network effects, where each round of the network effect is summarized by $\mathbf{G}_\infty^\epsilon$, discounted to the presented by both the discount rate (β) and also the rate at which the shock dissipates (ρ_ϵ).

Specifically, note

$$\begin{aligned}
\mathbf{G}_\infty^\epsilon &= \mathbf{I} + \rho_\epsilon \beta \mathbf{\Omega}_1 + (\rho_\epsilon \beta)^2 (\mathbf{\Omega}_1^2 + \mathbf{\Omega}_2) + (\rho_\epsilon \beta)^3 (\mathbf{\Omega}_1^3 + \mathbf{\Omega}_1 \mathbf{\Omega}_2 + \mathbf{\Omega}_2 \mathbf{\Omega}_1 + \mathbf{\Omega}_3) + \dots \\
&= \mathbf{I} + \rho_\epsilon \beta \mathbf{\Omega}_1 \mathbf{G}_\infty^\epsilon + (\rho_\epsilon \beta)^2 \mathbf{\Omega}_2 \mathbf{G}_\infty^\epsilon + \dots \\
&= \mathbf{I} + \left(\sum_{d=1}^{\infty} (\rho_\epsilon \beta)^d \mathbf{\Omega}_d \right) \mathbf{G}_\infty^\epsilon,
\end{aligned}$$

To derive $\mathbf{G}_\infty^x$ in closed-form, we re-write it recursively.

$$\begin{aligned}
\mathbf{G}_\infty^x &= \mathbf{I} + \frac{\rho_\epsilon^2 - \rho_x^2}{\rho_\epsilon - \rho_x} \beta \mathbf{\Omega}_1 + \frac{\rho_\epsilon^3 - \rho_x^3}{\rho_\epsilon - \rho_x} \beta^2 (\mathbf{\Omega}_1^2 + \mathbf{\Omega}_2) + \frac{\rho_\epsilon^4 - \rho_x^4}{\rho_\epsilon - \rho_x} \beta^3 (\mathbf{\Omega}_1^3 + \mathbf{\Omega}_1 \mathbf{\Omega}_2 + \mathbf{\Omega}_2 \mathbf{\Omega}_1 + \mathbf{\Omega}_3) + \dots \\
&= \frac{\rho_\epsilon}{\rho_\epsilon - \rho_x} [\mathbf{I} + \rho_\epsilon \beta \mathbf{\Omega}_1 + (\rho_\epsilon \beta)^2 (\mathbf{\Omega}_1^2 + \mathbf{\Omega}_2) + (\rho_\epsilon \beta)^3 (\mathbf{\Omega}_1^3 + \mathbf{\Omega}_1 \mathbf{\Omega}_2 + \mathbf{\Omega}_2 \mathbf{\Omega}_1 + \mathbf{\Omega}_3) + \dots] \\
&\quad - \frac{\rho_x}{\rho_\epsilon - \rho_x} [\mathbf{I} + \rho_x \beta \mathbf{\Omega}_1 + (\rho_x \beta)^2 (\mathbf{\Omega}_1^2 + \mathbf{\Omega}_2) + (\rho_x \beta)^3 (\mathbf{\Omega}_1^3 + \mathbf{\Omega}_1 \mathbf{\Omega}_2 + \mathbf{\Omega}_2 \mathbf{\Omega}_1 + \mathbf{\Omega}_3) + \dots] \\
&= \frac{\rho_\epsilon}{\rho_\epsilon - \rho_x} \left(\mathbf{I} - \sum_{d=1}^{\infty} (\rho_\epsilon \beta)^d \mathbf{\Omega}_d \right)^{-1} - \frac{\rho_x}{\rho_\epsilon - \rho_x} \left(\mathbf{I} - \sum_{d=1}^{\infty} (\rho_x \beta)^d \mathbf{\Omega}_d \right)^{-1}
\end{aligned}$$

To derive the impulse-response, equation A10 implies that

$$\begin{aligned}
\frac{\partial \mathbb{E}_t [\theta_{i,t+s}]}{\partial \epsilon_{it}} &= \rho_\epsilon^s \\
\frac{\partial \mathbb{E}_t [\theta_{i,t+s}]}{\partial u_{it}} &= \frac{\rho_\epsilon^{s+1} - \rho_x^{s+1}}{\rho_\epsilon - \rho_x}
\end{aligned}$$

Hence

$$\begin{aligned}
\frac{\partial \mathbb{E}_t [\tilde{\gamma}_{t+s}]}{\partial \epsilon_t} &= \rho_\epsilon^s \left(\mathbf{I} - \sum_{d=1}^{\infty} (\rho_\epsilon \beta)^d \mathbf{\Omega}_d \right)^{-1} \\
\frac{\partial \mathbb{E}_t [\tilde{\gamma}_{t+s}]}{\partial \mathbf{u}_t} &= \left(\mathbf{I} - \sum_{d=1}^{\infty} (\rho_\epsilon \beta)^d \mathbf{\Omega}_d \right)^{-1} \frac{\rho_\epsilon^{s+1} - \rho_x^{s+1}}{\rho_\epsilon - \rho_x} \\
&\quad + \frac{\rho_x}{\rho_\epsilon - \rho_x} \left(\left(\mathbf{I} - \sum_{d=1}^{\infty} (\rho_\epsilon \beta)^d \mathbf{\Omega}_d \right)^{-1} - \left(\mathbf{I} - \sum_{d=1}^{\infty} (\rho_x \beta)^d \mathbf{\Omega}_d \right)^{-1} \right) \rho_x^s \quad \text{as desired.}
\end{aligned}$$

Hence, under complete information,

$$\text{Var}_t (\tilde{\gamma}_{t+1}) = \mathbf{G}_\infty^\epsilon \Sigma_t^\epsilon \mathbf{G}_\infty^{\epsilon'} + \mathbf{G}_\infty^x \Sigma_t^u \mathbf{G}_\infty^{x'}$$

Under incomplete information,

$$\text{Var}_t (\tilde{\gamma}_{t+1}) = (\mathbf{G}_\infty^\epsilon - \mathbf{G}_\infty^\epsilon \mathbf{K} + \mathbf{G}_\infty^x \mathbf{K}) (\rho_x^2 \mathbf{F} + \Sigma^\epsilon + \Sigma^u) (\mathbf{G}_\infty^\epsilon - \mathbf{G}_\infty^\epsilon \mathbf{K} + \mathbf{G}_\infty^x \mathbf{K})', \quad \text{as desired.}$$

□

C MLE Algorithm Details

To estimate the covariance matrices Σ^ϵ and Σ^u , and the persistent parameter ρ , we integrate a Kalman filter into a maximum likelihood estimator (MLE). First, we specify two lower triangular matrices L^ϵ and L^u with positive diagonal entries and parametrize the covariance matrices as $\Sigma^\epsilon = L^\epsilon (L^\epsilon)^\top$, $\Sigma^u = L^u (L^u)^\top$. This Cholesky decomposition ensures that the covariance matrices are positive semi-definite.

First, based on the initial guess of parameters, we run the Kalman filter to compute the log-likelihood of the observed data, $\tilde{\theta}_t$, given the model parameters. We next maximize the log-likelihood with respect to Σ^ϵ , Σ^u and ρ . We then use the parameter estimates to update the Kalman filter, and iterate until convergence.

The log-likelihood function, assuming a joint normal distribution, is given by:

$$f = \frac{1}{2\pi\sqrt{|\mathbf{W}_t|}} \exp\left(-\frac{1}{2}\mathbf{v}_t'\mathbf{W}_t^{-1}\mathbf{v}_t\right),$$

and the log-likelihood function is given by:

$$\mathcal{L}(\{\theta_t\}_{t=1}^T; \Sigma^\epsilon, \Sigma^u, \rho) = -\frac{1}{2} \sum_{t=1}^T (\log |\mathbf{W}_t| + \mathbf{v}_t'\mathbf{W}_t^{-1}\mathbf{v}_t),$$

where $\tilde{\theta}_{t|t-1} = \mathbf{H}'\mathbf{M}\mathbf{S}_{t-1|t-1}$ is the forecast, $\mathbf{v}_t = \theta_t - \theta_{t|t-1}$ is the forecast error, $\mathbf{W}_t = \mathbf{H}'\Sigma_{t|t-1}\mathbf{H}$ is the variance of forecast error, with $\mathbf{H}$, $\mathbf{M}$, $\mathbf{S}_{t-1|t-1}$, and $\Sigma_{t|t-1}$ as defined in Appendix Section A.6.

We initialize the Kalman filter with a diffused prior, which places zero weight (up to machine precision) on the initial forecast ($\theta_{0|-1}$). Our results are also robust to initializing the Kalman with the long-run variance of the forecast error.

Practically, we implement this algorithm by using the python library “statsmodels.api.tsa.statespace.MLEModel”, with the initialization option set to “approximate_diffuse”, and we minimize the negative log likelihood function using ‘L-BFGS-B’ method. We set both the function tolerance and the gradient tolerance to 1e-6.

D Alternative Empirical Specifications and Robustness Checks

In this section, we conduct several alternative specifications and robustness checks to complement our empirical results in the main text.

D.1 Higher Volatility in Upstream Sectors Across Supply Chains

Figure 5 in the main text shows the relationship between upstreamness, as measured by supply chain delays, and volatility, along the sector-specific supply chain for 9 of the 25 downstream sectors. Appendix Figure A.2 shows the relationship for the remaining 16 downstream sectors. Appendix Figure A.3 reproduces the main text figure after dropping the four commodity sectors (“Oil and Gas extraction”, “Coal mining”, “Copper, Nickel, Lead, and Zinc Mining”, and “Iron, Gold, Silver and Other Metal Ore Mining”).

Figure A.2. Along supply chains, upstream sectors exhibit higher volatility
(supply chains for downstream sectors not shown in the main text)





Figure A.3. Along supply chains, upstream sectors exhibit higher volatility (dropping commodity sectors)




D.2 Downstream Sector's Own Impulse-Response: Pooled Specification

Figure 6 panel (a) in the main text shows the impulse-response of downstream value-added to its own innovations. Downstream sectors in that figure are classified into two groups according to whether a sector has a dominant AR(2) component. Appendix Figure A.4 reproduces the impulse-response by pooling all 25 downstream sectors together. The figure demonstrates a hump-shaped

pattern.

Figure A.4. Downstream sector's own impulse-response: pooled specification



D.3 Classifying Downstream Sectors by Kalman Gains

Figure 6 in the main text shows the estimated impulse response functions when downstream sectors (and the associated supply chains) are classified by their PACF lag 2 coefficients. Appendix Figure A.5 shows that our results are robust to classifying downstream sectors according to the diagonals of the estimated Kalman gains. Specifically, we estimate the following local projection regressions:

$$VA_{j,t+h}^{down} - VA_{j,t-1}^{down} = \alpha_0^h r_{jt} + \alpha_1^h r_{jt} \times K_j + X_{jt} + \zeta_{jt}, \quad (\text{A11})$$


$$\text{with controls } X_{jt} \in \left\{ \left\{ VA_{j,t-s}^{down} \right\}_{s=1}^p, \mu_j, \mu_t \right\}.$$

$$VA_{i,t+h}^{up,g} - VA_{i,t-1}^{up,g} = \beta^{h,g} \sum_{i \in \mathcal{U}_j} w_{ij} r_{jt} + \delta^{h,g} \sum_{i \in \mathcal{U}_j} w_{ij} r_{jt} \times K_j + Z_{it}^g + \epsilon_{it}, \quad (\text{A12})$$

$$\text{with controls } Z_{it}^g = \left\{ \left\{ \sum_{i \in \mathcal{U}_j} w_{ij} VA_{j,t-s}^{down}, VA_{i,t-s}^{up,g} \right\}_{s=1}^p, \mu_i, \mu_t \right\}.$$

where K_j is a dummy variable if j -th diagonal Kalman gain is above the median among 25 downstream sectors.

Figure A.5. Cumulative impulse response along the supply chain
(classifying downstream sectors by their Kalman gains)



D.4 Alternative Measures of Time-to-Build

Figure 6 in the main text show the estimated impulse response functions when supply chain delays are based on the backlog-ratio for each input, i.e., the ratio between the stock value of unfilled orders and the flow value of goods delivered (Liu and Tsyvinski, 2024). The strategy yields a supplier-specific measure of delay.

In this Appendix we re-estimate the impulse-response functions using two alternative strategies to measure delays. The first alternative strategy follows Antràs and Tubdenov (2025) and infers an input-buyer-specific measure of time-to-build based on the ratio between the value of inventories and cost of goods sold (COGS) using the COMPUSTAT data. Because this measure captures the average production period but not shipping delays, we round up the delay measures when converting from months to quarters. Appendix Figure A.6 shows the results.

The second alternative strategy specifies the total time-to-build delay from an input producer j to input user i is the sum of the backlog ratio of good j and the inventory-to-COGS ratio of good i . Appendix Figure A.7 shows the results.

Our main results remain robust across these alternative measures of time-to-build.

Figure A.6. Cumulative impulse response along the supply chain
(average production period as the time-to-build measure)

95% Bootstrapped CI — PACF at lag 2 below median - - - PACF at lag 2 above median





Figure A.7. Cumulative impulse response along the supply chain
(sum of the backlog-ratio and average production period as the time-to-build measure)



D.5 Using Alternative Parameters to Construct Supply Chain Distance

Figure 6 in the main text show the estimated impulse response functions when the supply chain distance ξ_{ij} are calculated according to (23) while setting $\rho = \beta = 1$. Appendix Figure A.8 reproduces the result using the realistic quarterly discount factor $\beta = 0.98$ and the rate of decay parameter $\rho = 0.818$ estimated from the data to construct our delay measure. Our results are robust to this change.

Appendix Figure A.9 further shows the correlation (Y-axis) between ξ_{ij} calculated using specific values of $\rho\beta$ (shown on the X-axis) and that calculated using $\rho\beta = 1$. The figure shows that the bilateral delay measure ξ_{ij} remains highly correlated across a wide range of $\rho\beta$.

Figure A.8. Cumulative impulse response along the supply chain ($\beta = 0.98, \rho = 0.818$)





Figure A.9. Correlation between ξ_{ij} under different values of $\rho\beta$ and the baseline measure ($\rho\beta = 1$)



D.6 Using Alternative Numbers of Lags When Estimating Impulse-Response Functions

Figure 6 in the main text show the estimated impulse response functions when we use $p = 2$ lags in the regression that extracts downstream value-added innovations (equation 30). In this Appendix we reproduce the impulse-response functions when we use alternative numbers of lags. For each specification, the number of lags p in the control variables for regression equations (31) and (32) are adjusted accordingly.

Appendix Figure A.10 shows the results for $p = 1$. Appendix Figure A.11 shows the results for $p = 3$.

Figure A.10. Cumulative impulse response along the supply chain ($p = 1$ lags)





Figure A.11. Cumulative impulse response along the supply chain ($p = 3$ lags)



D.7 Impulse-Response Using Long Difference Specification


Appendix Figure A.12 reproduces the estimated impulse-response functions when using the following specification:

$$VA_{i,t+h}^{up,g} - VA_{i,t-1}^{up,g} = \beta^{h,g} \sum_{i \in \mathcal{U}_j} w_{ij} r_{jt} + \delta^{h,g} \sum_{i \in \mathcal{U}_j} w_{ij} r_{jt} \times P_j + Z_{it}^g + \epsilon_{it}, \quad (\text{A13})$$

$$\text{with controls } Z_{it}^g \in \left\{ \left\{ \sum_{i \in \mathcal{U}_j} w_{ij} VA_{j,t-s}^{down} \right\}_{s=1}^p, \Delta VA_{i,t-1}^{up,g}, \mu_i, \mu_t \right\}.$$

We include $\Delta VA_{i,t-1}^{up,g}$ in our control instead of two lags of upstream value added. This is long difference specification could suppress small-sample bias (Jordà and Taylor, 2025).


Figure A.12. Cumulative impulse response along the supply chain (long difference specification)



D.8 Impulse-Response with Monthly Data

Figure 6 in the main text show the estimated impulse response functions using quarterly data. Appendix Figure A.13 reproduces the results using monthly data.

Figure A.13. Cumulative impulse response along the supply chain (monthly data)



D.9 Controlling for Productivity Shocks

The impulse-response functions estimated in the main text are based on extracted innovations to downstream value-added. Although in our theoretical framework, fluctuations in value-added (i.e., revenue multiplied by sectoral value-added intensity) originate exclusively from demand shocks—since productivity shocks have offsetting effects on prices and quantities—our empirically testable predictions directly concern fluctuations in sectoral value-added.

In this Appendix section, we conduct the following robustness check to control for sector-specific cost shocks. Specifically, we solve a generalized model with constant-elasticity-of-substitution (CES) production functions with one-period delays, linearize it around a steady-state, and derive separate implications of cost shocks and demand shocks.

Consider consumer preferences

$$V(\mathcal{S}_t) = \mathbb{E} \left[\sum_{s=t}^{\infty} \beta^{s-t} \left(\ln \left(\sum_i \theta_{it}^{1/\eta} c_{it}^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}} - v(\bar{\ell}_s) \right) | \mathcal{S}_t \right]$$

with production functions

$$y_{it} = z_{it} \ell_{it}^{\alpha_i} \left(\sum_{k=1}^N \lambda_{jk}^{\frac{1}{\xi}} m_{jk,t-d_{jk}}^{\frac{\xi-1}{\xi}} \right)^{\frac{\xi}{\xi-1} (1-\alpha_i)}.$$

Here, η is the elasticity of substitution across goods for the consumer, ξ is the elasticity of substitution across intermediate inputs for producers, λ_{jk} parametrizes input-output relationships

(with $\sum_k \lambda_{jk} = 1$), and θ_{it} again parametrizes consumer demand.

For simplicity, we solve this CES model assuming all inputs take one period to build, $d_{jk} = 1$ for all j, k .

Following our derivations in the main text, we can show that

$$p_{kt}m_{jkt} = \beta (1 - \alpha_j) \mathbb{E}_t [\gamma_{jt+1}] \frac{\lambda_{jk} p_{kt}^{1-\xi}}{\sum_i \lambda_{ji} p_{it}^{1-\xi}},$$

where $\gamma_{jt} \equiv p_{jt}y_{jt}$ is sectoral revenue as in the main text. That is, $(1 - \alpha_j) \mathbb{E}_t [\gamma_{jt+1}]$ is the total expenditure by producer j on intermediate inputs at time t , and $\frac{\lambda_{jk} p_{kt}^{1-\xi}}{\sum_i \lambda_{ji} p_{it}^{1-\xi}}$ is the fraction of intermediate input expenditure spent on input k .

Substitute m_{jkt} into the production function and then log-linearize around the steady-state (using $\tilde{\cdot}$ to denote log-deviation from the steady-state) we obtain:

$$\tilde{p}_{jt} = (1 - \alpha_j) (\tilde{\gamma}_{jt} - \mathbb{E}_{t-1} [\tilde{\gamma}_{jt}]) - \tilde{z}_{jt} + \alpha_j \tilde{w}_t + \sum_j \omega_{jk} [\xi \tilde{p}_{kt-1} + (1 - \xi) \tilde{p}_{jt-1}], \quad (\text{A14})$$

where ω_{jk} is the steady-state share cost expenditure by sector j on intermediate input k .

Equation (A14) enables us to extract productivity shocks from the data on prices. To get data on prices, we first divide value-added (IP data) by value-added intensity (IO data) to recover revenue, which is then divided by sectoral quantities (IP data).

We then use the following procedure based on equation (A14) to extract productivity shocks.

1. We residualize $\ln \gamma_{jt}$ with respect information available up to time $t - 1$. Specifically, we regress $\ln \gamma_{jt}$ on the first two principle components of all sectoral value-added at time $\{t - s\}_{s=1,\dots,4}$ and the first two principle components of all sectoral prices at time $\{t - s\}_{s=1,\dots,4}$. The residual from this regression is our measure of the innovation, $\tilde{\gamma}_{jt} - \mathbb{E}_{t-1} [\tilde{\gamma}_{jt}]$.
2. We multiply the innovation obtained in the previous step by $(1 - \alpha_j)$ and subtract the resulting product from $\ln p_{jt}$.
3. Finally, we residualize the outcome of the previous step with respect to information available up to time $t - 1$, to obtain residual $S_{jt} = -\tilde{z}_{jt} + \alpha_j \tilde{w}_t - \mathbb{E}_{t-1} [-\tilde{z}_{jt} + \alpha_j \tilde{w}_t]$.

S_{jt} is our measure of cost shocks; it captures the innovation in sectoral unit costs arising from productivity and labor—specifically, variations in unit costs not driven by changes in intermediate input prices.

We re-estimate our local projections as follows:

$$VA_{i,t+h}^{down} - VA_{i,t-1}^{down} = \alpha_0^h r_{it} + \alpha_1^h r_{it} \times P_i + X_{it} + \zeta_{it}, \quad (\text{A15})$$


with controls $X_{it} \in \left\{ \left\{ VA_{i,t-s}^{down} \right\}_{s=1}^p, S_{it}, \mu_i, \mu_t \right\}.$

$$VA_{i,t+h}^{up,g} - VA_{i,t-1}^{up,g} = \beta^{h,g} \sum_{i \in \mathcal{U}_j} w_{ij} r_{jt} + \delta^{h,g} \sum_{i \in \mathcal{U}_j} w_{ij} r_{jt} \times P_j + Z_{it}^g + \epsilon_{it}, \quad (\text{A16})$$

$$\text{with controls } Z_{it}^g \in \left\{ \left\{ \sum_{i \in \mathcal{U}_j} w_{ij} VA_{j,t-s}^{down}, VA_{i,t-s}^{up,g} \right\}_{s=1}^p, \sum_{i \in \mathcal{U}_j} w_{ij} S_{jt}, S_{it}, \mu_i, \mu_t \right\}.$$

Figure A.14 shows the estimated impulse response functions under this specification.

Figure A.14. Cumulative impulse response along the supply chain (controlling for cost shocks)



D.10 Quantitative Implications Given Heterogeneous Rates of Convergence

Section 4.4 in the main text conducts estimation and quantification based on the baseline model where both hump-shaped and monotone-decay shocks share a common persistence parameter ρ . In this Appendix we reproduce the analysis but allow for the two types of shocks to have different persistence parameters, based on our derivations in Section B.2. Following the MLE procedure described in Appendix Section C, we find $\rho_\epsilon = 0.667$ and $\rho_x = 0.928$. Table A.1 reports the diagonal entries of the Kalman gain matrix. This set of Kalman gains exhibits a correlation of 0.99 with the ones presented in Table 2 in the main text. Figure A.15 reproduces the variance decomposition results. Table A.2 reproduces the counterfactual exercises.

Table A.1. Estimation results: Kalman gains across downstream sectors (shock-specific persistence)

		κ
Transportation	Ship and Boat Building	0.27
	Heavy Duty Trucks	0.28
	Automobiles and Light Duty Motor Vehicles	0.29
	Motor Vehicle Bodies and Trailers	0.68
	Other Transportation Equipment	0.79
	Aerospace Products and Parts	0.79
Food Products	Fruit and Vegetable Preserving and Specialty Foods	0.09
	Coffee and Tea	0.11
	Sugar and Confectionery Products	0.35
	Breweries	0.51
	Other Food Except Coffee and Tea	0.53
	Animal Slaughtering and Processing	0.66
	Bakeries and Tortilla	0.77
	Soft Drinks and Ice	0.77
Other Final Goods	Major Electrical Household Appliances	0.11
	Office And Other Furniture	0.33
	Carpet and Rug Mills	0.34
	Newspaper Publishers	0.44
	Tobacco	0.49
	Periodical, Book, and Other Publishers	0.54
	Pharmaceuticals and Medicines	0.56
	Soap, Cleaning Compounds, and Toilet Preparation	0.59
	Apparel	0.83
	Medical Equipment and Supplies	0.86
	Navigational/Measuring/Electromedical Instruments	0.98

Figure A.15. Variance decomposition along supply chains (shock-specific persistence)




Table A.2. Counterfactual change in supply chain volatility based on the nature of downstream demand shocks (shock-specific persistence)

		% Change in supply chain volatility			Delays along supply chain	
		Only AR(2) shocks (1)	Only AR(1) shocks (2)	Static model (3)	Average (5)	Top 5 (6)
Downstream Group						
(a)	Transportation	154.50	-82.79	-17.66	4.30	8.90
(b)	Food Products	91.34	-77.97	-31.25	3.82	6.99
(c)	Other Final Goods	39.45	-85.95	-51.04	3.58	6.11
(d)	All Sectors	57.07	-77.58	-40.32	3.83	7.06

D.11 Counterfactuals Each Downstream Industry Individually

In Appendix Table A.3, we report the outcomes for the same set counterfactual exercises as in Table 3 of the main text, except that we examine each downstream industry individually rather than in groups.

Table A.3. Counterfactual change in supply chain volatility based on the nature of downstream demand shocks

Downstream Sector	% Change in supply chain volatility			Delays along supply chain	
	Only AR(2) shocks (1)	Only AR(1) shocks (2)	Static model (3)	Average (4)	Top 5 (5)
Apparel	40.02	-83.03	-57.47	2.18	3.35
Tobacco	181.66	-67.12	-16.94	2.06	2.60
Coffee and Tea	634.78	-25.18	108.95	2.12	3.12
Soft Drinks and Ice	41.25	-83.24	-57.74	2.14	3.47
Breweries	165.79	-69.79	-21.89	2.12	3.30
Carpet and Rug Mills	260.11	-61.93	6.61	2.21	3.41
Newspaper Publishers	161.96	-68.51	-21.43	2.04	2.94
Sugar and Confectionery Products	240.68	-58.50	2.60	1.95	2.96
Fruit and Vegetable Preserving	617.82	-22.62	109.32	2.02	3.27
Animal Slaughtering and Processing	102.51	-84.11	-44.28	2.77	3.89
Bakeries and Tortilla	59.16	-79.85	-51.49	2.02	3.31
Other Food Except Coffee and Tea	137.74	-69.77	-27.79	1.95	3.12
Pharmaceuticals and Medicines	120.85	-73.74	-33.01	1.99	2.31
Soap, Cleaning Compounds	117.40	-71.10	-33.02	1.83	2.58
Navigational Instruments	3.30	-86.38	-67.88	1.98	3.38
Major Electrical Appliances	480.26	-33.48	71.17	1.97	3.36
Automobiles and Light Duty Vehicles	326.57	-58.98	21.63	2.36	3.99
Heavy Duty Trucks	414.88	-57.88	48.24	2.61	4.15
Motor Vehicle Bodies and Trailers	94.46	-76.94	-41.11	2.15	4.03
Aerospace Products and Parts	43.78	-85.43	-52.95	2.33	3.68
Ship and Boat Building	409.20	-46.67	51.99	2.15	3.54
Other Transportation Equipment	60.26	-84.13	-50.11	2.34	3.64
Office And Other Furniture	303.88	-48.34	23.05	1.87	3.25
Medical Equipment and Supplies	26.35	-83.78	-60.38	1.96	2.96
Periodical, Book Publishers	131.93	-76.03	-32.50	2.35	3.16