Title

Essays on the Role of Supply Chains in the Propagation of Macroeconomic Shocks

Author

Chen, Junyuan

Publication Date

2024

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Essays on the Role of Supply Chains in the Propagation of Macroeconomic Shocks

A dissertation submitted in partial satisfaction of
the requirements for the degree Doctor of Philosophy

in

Economics

by

Junyuan Chen

Committee in charge:

Professor Marc-Andreas Muendler, Co-Chair
Professor Valerie A. Ramey, Co-Chair
Professor Munseob Lee
Professor Kanishka Misra
Professor Johannes F. Wieland

2024

The dissertation of Junyuan Chen is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2024

DEDICATION

To my family

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

I would like to acknowledge Professor Valerie A. Ramey and Professor Marc-Andreas Muendler for their support as the co-chairs of my committee. My progress over the years would not be possible without their patience and trust. I would like to acknowledge Professor Johannes F. Wieland, Professor Juan Herreño, and Professor Fabian Trottner for their invaluable feedback as I develop this dissertation. I would like to acknowledge Professor Munseob Lee for his support during the early stage of my graduate study.

Chapter 1, in part, is currently being prepared for submission for publication of the material. The dissertation author was the primary author of this paper.

Chapter 2, in part, is currently being prepared for submission for publication of the material. The dissertation author was the primary author of this paper.

Chapter 3, in full, is currently being prepared for submission for publication of the material. Chapter 3 is coauthored with Carlos Góes, Marc-Andreas Muendler, and Fabian Trottner. The dissertation author was one of the primary authors of this paper.

I would like to acknowledge the partial financial support received as a Graduate Student Researcher at the Globalization and Prosperity Lab of University of California San Diego as I made progress on this dissertation.

VITA

EDUCATION

University of California San Diego

Doctor of Philosophy in Economics	2024
Candidate in Philosophy in Economics	2021

University of Minnesota Twin Cities

Bachelor of Science in Economics, *summa cum laude*, high distinction	2017
Bachelor of Arts in Mathematics, high distinction	2017

EXPERIENCE

Research assistant, University of California San Diego	2021–2024
Associate instructor, University of California San Diego	2023 Summer
Teaching assistant, University of California San Diego	2018–2023

ABSTRACT OF THE DISSERTATION

Essays on the Role of Supply Chains in the Propagation of Macroeconomic Shocks

by

Junyuan Chen

Doctor of Philosophy in Economics

University of California San Diego, 2024

Professor Marc-Andreas Muendler, Co-Chair
Professor Valerie A. Ramey, Co-Chair

Chapter 1 studies the dynamic behavior of durable input inventories. The positive contemporaneous comovement between aggregate inventories and sales is a well-known stylized fact. I highlight an overlooked feature that for durable input inventories, which constitute a volatile component of the aggregate inventories, the inventory movements lag sales movements by around three quarters. This lagged comovement is discernible both in the unconditional cyclical components of data and in the impulse responses to identified aggregate shocks. I develop a tractable supply chain production problem that is capable of reproducing the lagged comovement. In this model, producers are required to order critical inputs from suppliers one quarter in advance and they occasionally adjust their optimal order sizes based on forecasts of their own future sales subject to information frictions.

Fitting the model impulse responses with the empirical counterparts suggest the frictions need to be strong.

Chapter 2 studies the aggregate implications of the lagged inventory-sales comovement. I embed the production problem developed in Chapter 1 into a multisector New Keynesian framework featuring input-output relations between sectors and conventional real and nominal frictions. Following a monetary shock, relative to a counterfactual scenario in which the inventory-sales comovement is fully synchronized, the estimated model demonstrates dampened responses of aggregate output over the first year but more gradual recovery over later horizons due to the reduced sensitivity of user cost of capital with respect to real interest rate changes.

Chapter 3 studies the dynamic welfare implications following tariff shocks. We develop a tractable general equilibrium framework that features staggered sourcing decisions, nests the Eaton-Kortum model as the limiting long-run scenario, and accounts for the horizon-specific trade elasticity. We calibrate the model and use it to quantify the welfare impact of the 2018 US-China trade war. Model outcomes imply that applying the well-known static welfare formula on observed domestic trade shares miss important distortions. The short-run welfare impact can be smaller than the long-run impact for the US but larger for China despite the same low short-run trade elasticity. Third countries such as Mexico and Vietnam may experience welfare losses in the short run but welfare gains in the long term.

Chapter 1. Supply Chain Frictions and the Dynamic Behavior of Durable Input Inventories

Macroeconomists have long been interested in how inventory behavior has affected aggregate outcomes because inventories are procyclical and volatile (Ramey and West 1999). In particular, the capability of generating a positive contemporaneous cross-correlation between inventories and sales is widely considered to be a requirement for an inventory model to be useful for shedding light on aggregate implications. By introducing inventories into a macroeconomic model, a majority of the existing studies have examined the implications on the volatility of aggregate output, another contemporaneous second moment.

Going beyond the contemporaneous inventory-sales comovement, this paper establishes an overlooked fact that, especially for durable input inventories, the cross-correlation is much stronger between inventories and lagged sales.[1] To that end, I first document that the movements of durable input inventories lag movements of durable sales by about three quarters in the unconditional cyclical components of aggregate time series data. This is reflected by the cross-correlation coefficients between the cyclical components of inventory and different leads or lags of sales. I then show that the lagged comovement is also evident in the estimated impulse responses with respect to identified aggregate shocks. In particular, I focus on the responses to monetary shocks identified based on the Romer and Romer (2004) shocks extended by Wieland and Yang (2020). After a contractionary monetary shock that raises the nominal interest rates, durable sales experience a decline, reaching a trough by the end of the first year. In contrast, durable input inventories remain mostly unchanged in the first year, only start to decline in the second year and reach a trough in the third year. Although both sales and inventories decline following monetary shocks, the durable input inventories comove with a substantial lag. This empirical pattern is robust for the durable input inventories but seems to hold less well for the other types of inventories, especially when considering the changes induced by the monetary shocks. I therefore deem the lagged inventory-sales comovement as a distinct feature of the durable input inventories and focus on this specific type of inventories.

[1] Input inventories refer to the stock of goods held by manufacturers for their production that are not yet part of the output. Based on the classification by the Bureau of Economic Analysis (BEA), they include inventories in two of the three stages of fabrication: "materials and supplies" and work-in-process.

Why do the movements of durable input inventories lag the movements of sales? How does knowing this fact change the way how we understand the aggregate implications of inventory behavior? To address such questions, I need a structural model that is capable of reproducing the lagged inventory-sales comovement observed in data. Prior studies have considered embedding inventory problems based on different rationales into macroeconomic environment (e.g., Khan and Thomas 2007; Wen 2011). However, to the best knowledge of the author, they are not readily applicable for capturing the lagged comovement observed for durable input inventories. Intuitively, in anticipation of a future decline in sales due to a contractionary monetary shock, producers would set a lower inventory target on average because of the lowered expected value of holding inventories.[2] Such a model would make a counterfactual prediction that inventory movements should lead sales fluctuations.

To understand this limitation of existing models in reproducing the lagged inventory-sales comovement, I draw on knowledge outside of the economics literature for inspiration. In the field of supply chain management, a substantial amount of effort has been dedicated to understanding the so-called bullwhip effect, a phenomenon that involves the increasing amplification of demand variability along a supply chain (Giard and Sali 2013). This literature often considers the lack of shared information on the final demand, which results in the use of order data from the immediate downstream customer for forecasting future demand, to be a cause of the bullwhip effect (Lee, Padmanabhan and Whang 1997). Numerous studies conclude that improving information shared among supply chain members could substantially improve the profitability (e.g., Metters 1997).

What can we learn from management scientists' effort on mitigating the bullwhip effect? I identify two closely related real-world characteristics of manufacturing production that are not stressed enough by the inventory problems examined by macroeconomists. First, in reality, production for a wide range of products happens in multiple stages across different producers along a supply chain that do not share sufficient information among themselves. For this reason, the information relevant to production decisions may well come from past order sizes that do not necessarily align with the current aggregate state of the economy. Second, since a supply chain member may rely on the order information from its direct customers for making production arrangements, delay in

[2]The decline in the expected value can arise from both the lowered expected future sales and the increase in financial cost.

responding to the changes in demand from customer producers may be further accumulated as the supply chain member itself makes adjustment with its own suppliers. As long as the adjustment for production arrangements at each stage is not flexible enough, a change in demand from the most downstream customers may not reach the upstream producers instantly as it passes through the production stages. With these two characteristics in mind, I develop a supply chain production model that aims to capture the essence of these missing details.

The model consists of a serial supply chain in which varieties from one production stage are delivered to producers in the next stage as critical inputs for their production. The production technology takes the same functional form that combines the critical inputs with a bundle of other input factors. Three main assumptions govern the model outcomes. First, since production takes time, the critical inputs passed along the supply chain must be ordered one quarter in advance when the producers who use these inputs for production have not observed their own sales. Second, producers observe their own historical sales perfectly but infer aggregate sales only from noisy private signals and their own historical sales. Third, producers receive critical inputs in each period but only adjust the size of the order placed for their critical inputs occasionally in a Calvo (1983*a*) fashion.

The first assumption specifies the timing for decision making and justifies the presence of inventories in this model. The lead time for acquiring the critical inputs is the sole reason for input inventories to exist in this model. In particular, the level of input inventories at the end of a period is exactly the size of the order for the critical inputs placed at the beginning of the period.[3] The second assumption specifies the information frictions faced by the producers when they choose the order size for their critical inputs. I consider an environment where the demand fluctuations from the immediate downstream customers arise from both idiosyncratic taste shocks and aggregate shocks. The result is that producers make biased predictions on future order sizes due to the difficulty in distinguishing fluctuations in order sizes induced by aggregate shocks from those induced by idiosyncratic taste shocks. Lastly, the third assumption affects the relevant time horizons under consideration when producers choose their order size. Intuitively, when the frequency of adjustment is low, producers make production arrangements based on their forecasts over a longer span of time, and hence any bias in their forecasts results in more sustained deviations from the optimal choices in the absence of frictions. This third assumption adds more flexibility to the model for reproducing

[3]Throughout the paper, I assume that a period is a quarter of a year.

the impulse responses observed in data and also smooths out the inventory responses.[4]

The model is parameterized to reproduce the impulse responses of durable input inventories following a monetary shock when paths of changes in sales, nominal interest rates and factor prices induced by a monetary shock are fed into the model. The model is able to fit the empirical target remarkably well. To achieve the fitness, the critical inputs and the bundle of other input factors need to demonstrate relatively strong complementarity; the signals on aggregate shocks need to be uninformative; and the persistance of idiosyncratic taste shocks need to be high. These requirements seem to be plausible and are broadly in line with the popular beliefs among management scientists. In particular, the latter two requirements suggest that, from the lens of the model, producers are ordering critical inputs as if they perceive the past sales changes will persist over time and completely ignore the dynamic effects of monetary shocks. Furthermore, introducing Calvo-style sluggish adjustment of order size does improve the fitness of the model. It effectively strengthens the impact of information frictions.[5]

This paper is related to several strands of literature. First, the empirical finding on the inventory behavior joins a long-lasting literature on inventory facts that are related to business cycle fluctuations. Blinder (1981) and Blinder and Holtz-Eakin (1986) document the surprising extent to which business cycle fluctuations can be accounted by inventory changes. Ramey (1989) raises the importance of manufacturing input inventories which are more volatile than the other types of inventories. Several papers including Mcconnell and Perez-Quiros (2000), Stock and Watson (2002) and Davis and Kahn (2008) consider the possibility for changing inventory behavior to explain the Great Moderation started around 1984. Wen (2005) examines the cyclicality of inventories over different frequencies using a band-pass filter. Relative to these earlier works which focus on the contemporaneous correlations involving aggregate inventories, I examine the correlations of inventories with the leads and lags of sales and highlight the lagged inventory-sales comovement that holds particularly well for durable input inventories.

Second, the model ingredients for the production problem developed in this paper is related to, yet distinct from, several ideas in prior works. To begin with, several studies have considered

[4]The Calvo-style adjustment does not necessarily slow down the changes in aggregate inventory level. Its effect depends on how producers perceive future changes. This will be discussed in Section III and revisited in Section V.

[5]Another reason for introducing this Calvo-style friction is that, by adjusting its strength, it absorbs the impact of other aspects of model calibration, such as the number of production stages, which are fixed unchanged.

a direct connection between inventory levels and the cost of production. For example, Christiano (1988) gives inventory stocks a direct role in the production function and views inventory volatility as resulting from buffering unexpected shocks. Eichenbaum (1983, 1984, 1989) views inventories as a device that allows producers to smooth production cost, instead of production level as in conventional wisdom. In my model, input inventories are the output collected from suppliers for the production in the proceeding period. The abundance of them affects the production rate and hence producers determine their levels in an attempt to lower the present value of production cost.

In the macroeconomics literature, the notion of multistage production or time to build is first formally considered by Kydland and Prescott (1982) for the production of capital. More recently, Sarte, Schwartzman and Lubik (2015) formulate an extended version of multistage technology that allows freely shifting resources across the production of input inventories in different stages. There is a key distinction between the multistage production considered in this paper and theirs. Namely, I view production happening in different stages as being conducted by different producers along a supply chain that only interact among each other via the information on order sizes. This is very different from the problems considered in prior studies where the entire production process is conducted as if there is a single producer that coordinates resources across different stages. The distinction arises from the intention to capture the missing feature emphasized in the supply chain management studies that supply chain members do not coordinate with each other in a manner that results in optimal behavior for the entire supply chain (Lee, Padmanabhan and Whang 1997; Metters 1997; Chatfield et al. 2004; Ouyang 2007). Furthermore, information frictions are introduced as a Bayesian learning process in a manner similar to models developed in other contexts including the price-setting problems (Morris and Shin 2002; Nimark 2008; Angeletos and Huo 2021). Unlike the price-setting problems, production planning by an individual producer does not involve forming anticipation on how other producers make choices. I therefore do not deal with the complications of higher-order beliefs. For the model to maintain its prediction on lagged inventory-sales comovement, what matters the most is the forecasting bias arising from the frictions.

The paper proceeds as follows. Section I documents the lagged inventory-sales comovement for the durable input inventories. Section II introduces the supply chain production problem. Section III examines the capability of the model to reproduce the lagged inventory-sales comovement. Section IV discusses the relations to other models. Section V concludes.

I. The Lag in Inventory-Sales Comovement

Documenting the lagged comovement between durable input inventories and sales is an important empirical contribution of this paper. In this section, I explain in detail how this empirical finding is established. While the focus is on the durable input inventories, other types of inventories are considered in Appendix 1.A.

A. Unconditional Cross-Correlations

The monthly inventory and sales data for the empirical finding are based on the underlying detail tables for National Income and Product Account (NIPA) from the Bureau of Economic Analysis.[6] Unlike the inventory changes and final sales in the NIPA tables that are measured in terms of value added as components of GDP, inventories and sales in the detail tables are monthly inventory stocks and gross sales.[7] For the empirical estimation, I focus on the sample over 1967:1–2007:12 before the Great Recession.

Figure 1.1 plots the transformed data for input inventories and sales in the durable manufacturing sector. To deal with the nonstationarity of the time series, I have computed both the annual growth rate and the cyclical component of the original data. Two features of the data are noticeable by observing Figure 1.1. First, echoing what earlier studies back in 1980s have emphasized, the fluctuations in durable input inventories closely keep track of those of the sales, confirming the well-known inventory-sales comovement.[8] Second, by paying attention to the peaks and troughs, it is not hard to see that input inventory movements tend to lag sales movements. That is, for the durable manufacturing sector, the comovement is better described as lagged comovement rather than contemporaneous comovement. Although this second feature is discernible from the raw data before any transformation as well, a formal discussion of it is surprisingly absent in the literature.[9]

[6] The NIPA underlying detail tables are considered to be the standard data sources among macroeconomic inventory studies and have been employed for numerous studies over decades.

[7] Gross sales and inventory stocks are more natural measurement for considering inventory problems, as they are closer to what economic agents observe for making decisions. Notice that BEA data have been adjusted to take into account the discrepancy between book values and market values. Like earlier studies, I take the BEA data construction procedures as given and view them as the best available aggregate data for US economy.

[8] When Blinder (1981) commented that "... to a great extent, business cycles are inventory fluctuations", he was comparing the peak-to-trough real GNP changes with the contemporaneous inventory changes, which are in roughly similar magnitude. Similar peak-to-trough comparisons were used by Ramey and West (1999) for motivating the importance of inventories.

[9] The lagged comovement is consistent with another well-known feature that the inventory-to-sales ratio is counter-

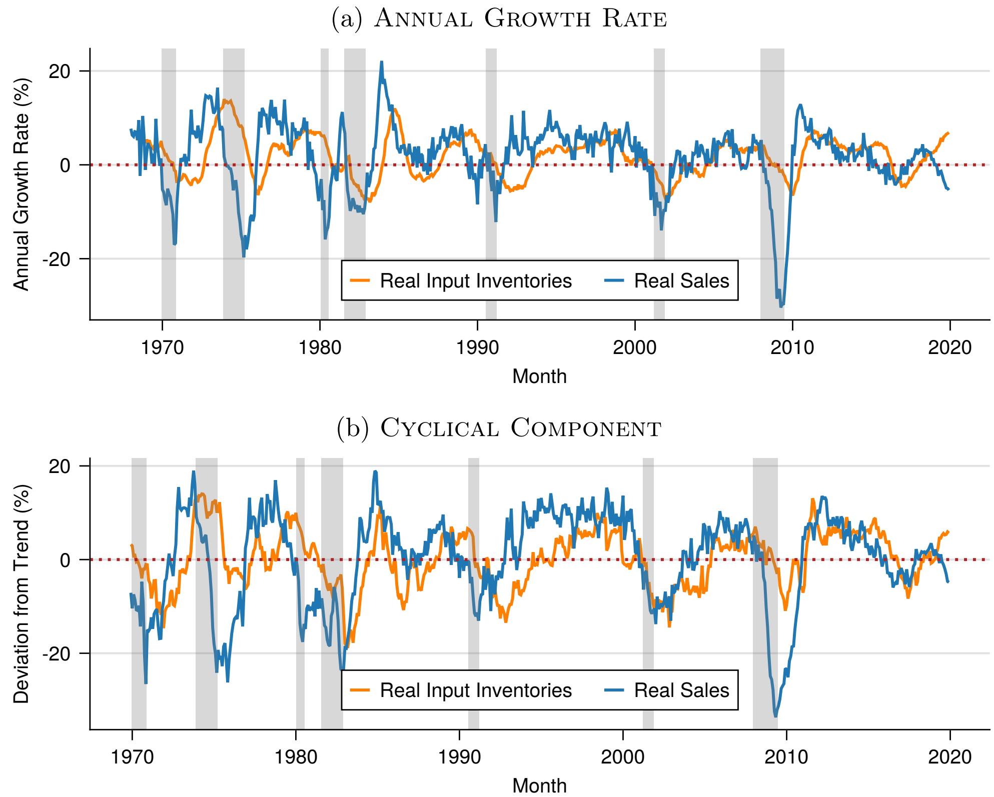

FIGURE 1.1: Real Input Inventories and Sales in Durable Manufacturing Sector

Note: Annual growth rates are computed as log differences with the one-year lag. Cyclical components are two-year horizon forecast errors following Hamilton (2018).

Loosely speaking, these two features combined together is what I call the "lagged inventory-sales comovement".

To highlight the timing of how durable input inventories comove with sales, I compute cross-correlation coefficients between inventories and different leads or lags of sales, $\rho(\text{Inventory}_t, \text{Sales}_{t-l})$ with different l. Figure 1.2 plots the results based on both the annual growth rates and cyclical components. The two sets of results look very similar, both suggesting that the correlation coefficients are larger when comparing inventory movements with the lagged sales. In particular, when we consider only the contemporaneous correlation coefficients, the coefficients of 0.2 and 0.34 only indicate some modest strength of comovement. However, when we look at the greatest correlation coefficients attained when lagging sales by around three quarters, the correlation coefficients of

cyclical. However, they are two different characteristics because countercyclical inventory-to-sales ratio can arise from contemporaneous inventory-sales comovement with inventories move by a smaller dollar amount relative to sales. A countercyclical inventory-to-sales ratio is not informative on the specific timing relation emphasized in this paper.

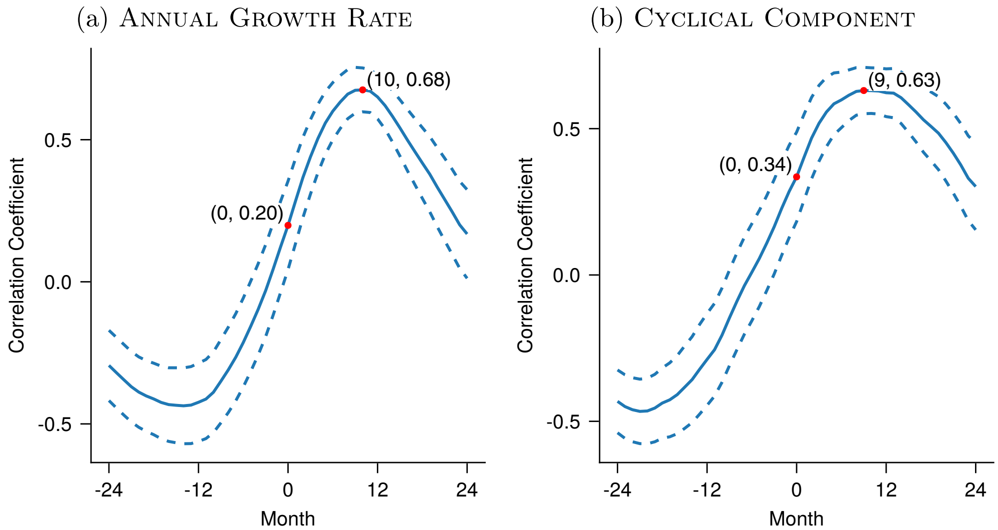

FIGURE 1.2: Cross-Correlations between Durable Input Inventories and Sales

Note: The correlation coefficients are for $\rho(\text{Inventory}_t, \text{Sales}_{t-l})$ across different values of l. Annual growth rates are computed as log differences with the one-year lag. Cyclical components are two-year horizon forecast errors following Hamilton (2018). Dashed lines represent 90% confidence intervals based on robust standard errors obtained from a GMM Newey and West (1987) procedure.

0.68 and 0.63 suggest rather strong comovement relative to what the contemporaneous correlation coefficients would suggest. In a nutshell, durable input inventory fluctuations lag behind sales by a non-negligible amount of time.

How about the other types of inventories? In Appendix 1.A, I repeat the estimation for other types of inventories available from the same data source. The lag in the inventory-sales comovement seems to be less important for the nondurable manufacturing sector. Yet, a somewhat similar pattern of cross-correlation coefficients shows up for the output inventories in the durable wholesale sector and the retail sector. If we only consider the unconditional comovement based on these cross-correlation coefficients, the input inventories in the durable manufacturing sector does not seem to be the only type of inventories that demonstrates the lagged inventory-sales comovement. However, once we start to consider conditional comovements induced by specific aggregate shocks, the answer is going to be different. In particular, as we move on to the next set of results, we see that for the durable manufacturing sector, the lagged inventory-sales comovement also holds well when only considering the changes induced by monetary shocks. Since the durable sector is more responsive to monetary shocks, I consider it reasonable to focus on this sector for the main text.

B. Cross-Correlations Conditioning on Shocks

One way to think about the fluctuations of an aggregate variable over business cycles is to imagine it being hit by various shocks in each period, with the observed outcomes as a combination of responses to all shocks over the entire history. The structural shocks with concrete economic interpretations are typically not observed directly. For this reason, cross-correlation coefficients directly computed from the annual growth rates or cyclical components depict the overall relations between inventories and sales without taking a stance on the driving forces behind the fluctuations. In consideration of later structural analysis, I examine inventory-sales comovement conditioning on monetary shocks, for which arguably plausible instruments are readily available.[10]

1. Methodology

Formally, let $\{\varepsilon_t^i\}_t$ denote a sequence of i.i.d. innovations to a Taylor rule of monetary policy. A path of an aggregate variable $\{x_t\}_t$ can then be written as

$$x_t = \sum_{l=0}^{\infty} \Psi_l^x \varepsilon_{t-l}^i + \tilde{x}_t$$

where the first term is an infinite-order moving-average (MA) process driven by the monetary shocks and the second term is a residual. By conditioning on a specific structural shock, I ignore $\tilde{x}_t$ and only consider the first MA(∞) term for determining the dynamic cross-correlations between two paths.[11] It is not hard to see that for any two MA processes $\{x_t\}_t$ and $\{y_t\}_t$ with uncorrelated innovations, their MA coefficients fully determine the cross-correlations between $\{x_t\}_t$ and $\{y_{t-l}\}_t$ for any l. Therefore, for an empirical investigation, the primary task is to estimate the structural MA coefficients $\{\Psi_l^x\}_{l\geqslant 0}$ and $\{\Psi_l^y\}_{l\geqslant 0}$.

To that end, I estimate the impulse responses of various types of inventories and sales to monetary shocks. With the real quantities transformed to the log scale, I estimate structural vector autoregressions (SVARs) with the instrumental variable (IV) ordered the first.[12] Each

[10] The results presented here are not special to monetary shocks. Similar results are obtained when considering aggregate TFP shocks.

[11] As innovations to a structural shock that represents a primitive force, $\{\varepsilon_t^i\}_t$ are by definition orthogonal to any other structural force that drives the fluctuations.

[12] This is the VAR estimation with "internal instrument". As shown by Plagborg-Møller and Wolf (2021), the same estimand could be estimated with local projection (LP). The empirical findings hold true no matter whether estimation

SVAR estimation involves eight variables with the lag length being 12 months: the IV, inventory, sales, industrial production index, unemployment rate, federal funds rate, deflator for the personal consumption expenditures (PCE) and producer price index for commodities.[13] Since only the responses with respect to monetary shocks are of interest, only the ordering for the IV is relevant. For the IV, I use the shock series constructed by Romer and Romer (2004) and extended by Wieland and Yang (2020). Given the coefficient estimates from SVAR, I compute impulse response estimates for the inventory and sales with respect to the Romer-Romer shocks over the first 49 monthly horizons.[14]

The impulse response estimates derived from the SVAR estimation are interpreted as estimates for the structural MA coefficients with respect to monetary shocks for the first 49 horizons. The dynamic cross-correlations between inventory and sales are computed based on these MA estimates. Specifically, for a given time gap $s \geqslant 0$, the dynamic cross-correlation coefficients between two MA(q) processes $\{x_t\}_t$ and $\{y_{t+s}\}_t$ are computed as

$$\rho(x_t, y_{t+s}) = \frac{\sum_{l=0}^{q-s-1} \hat{\Psi}_l^x \hat{\Psi}_{l+s}^y}{\sqrt{\sum_{l=0}^{q-1} (\hat{\Psi}_l^x)^2} \sqrt{\sum_{l=0}^{q-1} (\hat{\Psi}_l^y)^2}}.$$

The order of the true structural MA processes induced by monetary shocks may be of infinite order in principle. However, for practical purposes, it makes no meaningful difference to truncate the higher order terms when computing the cross-correlation coefficients as the effects of monetary shocks diminish over time. Below, I always use the 49 MA coefficients for computing the cross-correlations and disregard any potential (negligibly small) effects from monetary shocks after the 49th quarter.

For statistical inference, I implement the recursive-design wild bootstrap following Gonçalves and Kilian (2004). The confidence bands appear in the plots of estimation results are all pointwise 90% equal-tailed percentile confidence intervals based on bootstrap samples of 10,000 draws.

2. Results

The empirical finding established here is that, following a monetary shock, the responses of input inventory lag responses of sales by about 3 quarters among the durable manufacturing

is based on SVAR or LP. However, SVAR is preferred due to the simulation results by Li, Plagborg-Møller and Wolf (2021). They find that estimators based on VAR demonstrate advantages over the bias-variance tradeoff.

[13] Selection of the additional variables included in the SVAR estimation follows Ramey (2016).

[14] The estimates are normalized so that the peak impact of Romer-Romer shocks on federal funds rate is 1%.

industries. Figure 1.3a shows the estimated impulse responses of durable input inventory and durable sales with respect to monetary shocks. As expected, following an exogenous increase in federal funds rate, the sales from the durable manufacturing sector demonstrates hump-shaped responses that reach a trough at the end of the first year. The durable input inventory, in contrast, first remains mostly unchanged during the first year and only starts to decline over the second year, reaching a trough in the third year. The responses of input inventory are therefore considered to be lagging behind sales by a substantial amount of time. To highlight this timing relations, Figure 1.3b shows the dynamic cross-correlation coefficients between input inventory and sales across different leads and lags. Here, a positive value on the horizontal axis represents the number of months between the horizons of responses from input inventory and sales. The peak correlation coefficient of 0.88 is reached when comparing the responses of sales with the responses of inventory 9 months later (excluding the contemporaneous one).[15] The contemporaneous correlation coefficient $\rho(h_t, y_t)$ takes a smaller value of 0.65. This suggests that the lagged inventory-sales comovement observed in the cyclical components of data is still relevant when only considering the fluctuations induced by monetary shocks. For the other types of inventories, the lagged comovement following monetary shocks seems to be less important (see Appendix 1.A).

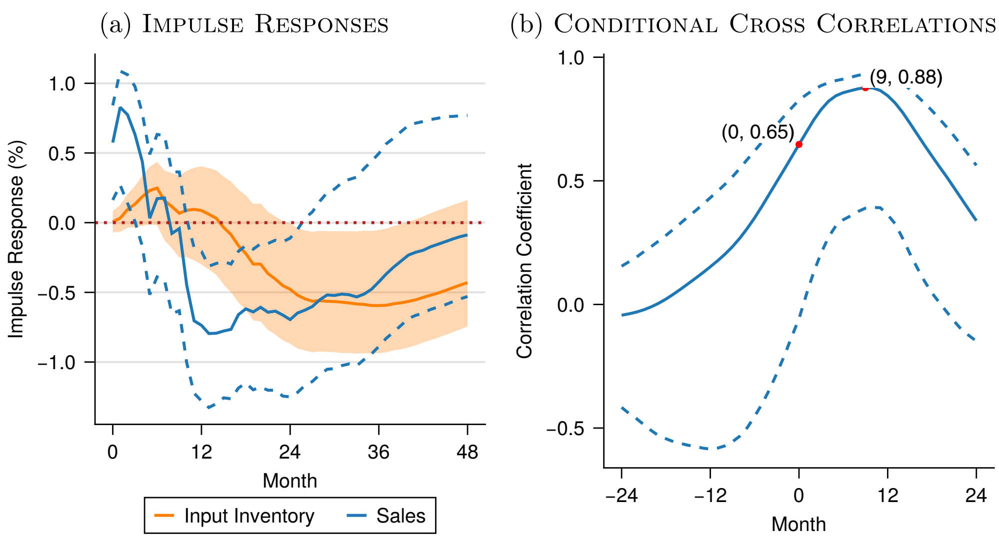

FIGURE 1.3: Lagged Comovement Following Monetary Shocks in the Durable Sector

[15] To clarify the interpretation of such estimates, consider two structural MA processes that represent the fluctuations of input inventory and sales induced by monetary shocks denoted as $\{h_t\}_t$ and $\{y_t\}_t$ respectively. The correlation coefficient at the peak refers to $\rho(h_{t+9}, y_t)$ being 0.88.

C. Comovement in Existing Models

Do existing models capture the lagged inventory-sales comovement observed in data? This is unlikely for two reasons. First, the empirical regularity documented here is not explicitly taken into account in prior work. It would be a coincidence if a model happens to demonstrate this feature. Second, the presence of the lag is actually contradicting the anticipatory behavior that is often thought to be captured by inventory movements. In particular, conventional wisdom suggests that if producers anticipate rising sales in the near future, they should increase the inventories early to smooth the production. This would suggest the opposite scenario that inventory movements lead sales movements.

Given the vast amount of prior works on inventory behavior, I solve two representative models under a partial equilibrium environment to verify that they do not capture the lagged inventory-sales comovement. This is apparently not an exhaustive coverage of all variants of inventory models in the literature. However, this process helps illuminate the reason why one should not view the lagged inventory-sales comovement as a feature that can be readily reproduced in a standard inventory model. Below, I consider a model based on the precautionary stock-out avoidance motive and a model featuring multistage production (time to build). The former is adapted from Wen (2011); while the latter is based on Sarte, Schwartzman and Lubik (2015). Essential model details are provided in Appendix 1.B. For both of them, I feed the same path of sales changes and real interest rate path into the models and compute the responses in the input inventories. I then compute cross-correlations based on the impulse responses in a way analogous to how they are obtained from the empirical counterparts.

The results are shown in Figures 1.4 and 1.5. For both models, they generate the counterfactual outcomes that input inventory movements lead sales movements. One reason for inventories to decline is the rise of real interest rates that affects the intertemporal allocation of resources for producing the inventories. However, even without feeding the real interest rate path into the models, there is no sign of lagged comovement. The optimality conditions from these models imply immediate responses in inventories that are in sharp contrast to what the data suggest. This suggests that these models lack important ingredients that are in operation in reality. In Section III**??**, I provide further discussion on why these models do not generate the lagged inventory-sales comovement after

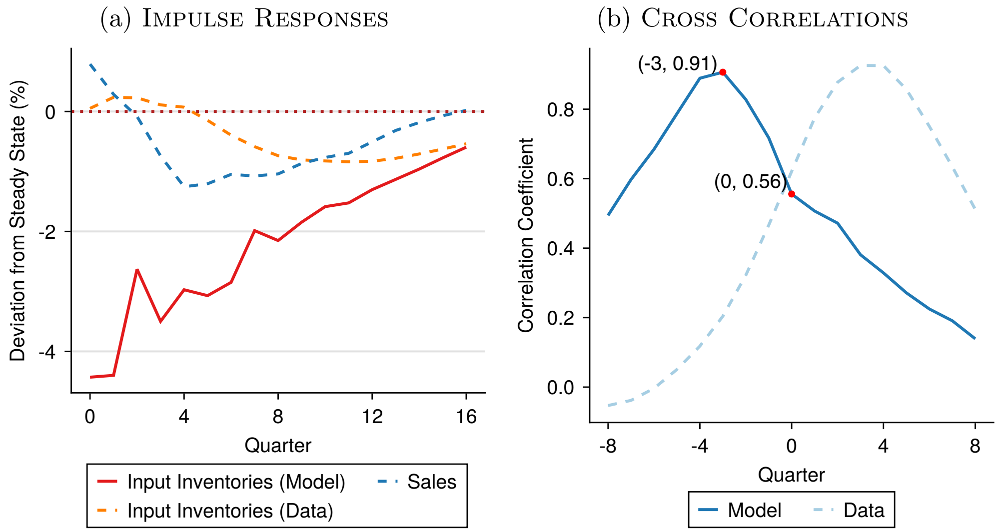

FIGURE 1.4: Stockout-Avoidance Model: Partial Equilibrium Outcomes

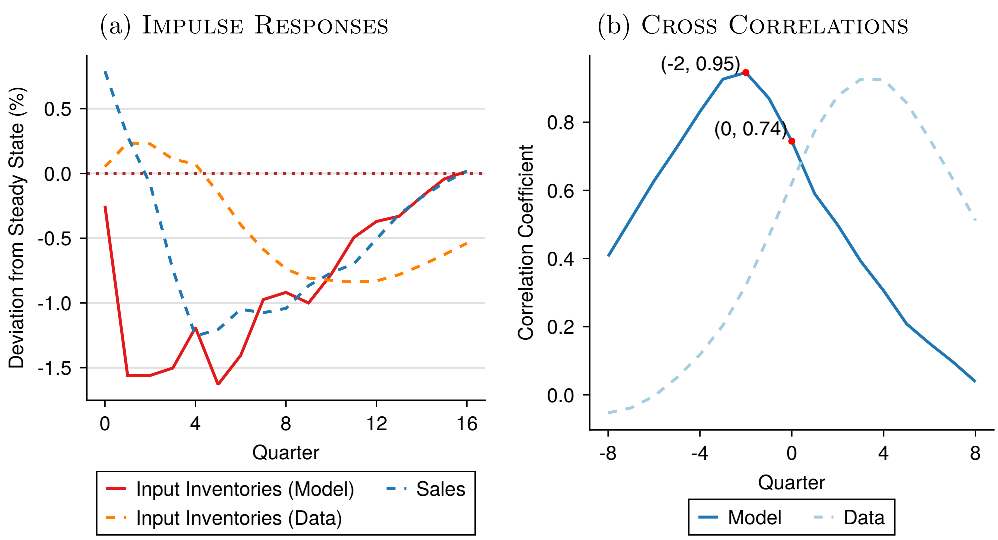

FIGURE 1.5: Multistage Production: Partial Equilibrium Outcomes

presenting the results for my model introduced in Section II.

II. A Supply Chain Problem with Frictions

Is the lag in the inventory-sales comovement economically meaningful? The answer to this question inevitably relies on some structure of the economy and the lag would otherwise be hard to interpret. In this section, inspired by the real-world challenges encountered in supply chain management, I develop a supply chain production problem that is capable of generating the lagged inventory-sales comovement documented in Section I.

A. Model

Time is discrete and each period is a quarter denoted as t. There is a set of products that are only produced along a supply chain. The supply chain consists of S stages with each stage operated by a continuum of producers each producing a distinct variety. Producers purchase a bundle of stage-specific varieties only from their immediate upstream suppliers as critical inputs and they sell their own products to only the immediate downstream customers. Production in a stage involves combining the critical inputs that must be ordered one period in advance with other input factors (labor, capital and intermediate inputs). The problem faced by producers in all stages is identical except that producers in the most upstream first stage do not require obtaining any critical input in advance.

Timing At the beginning of each period, producer i in stage s with $s > 1$ places orders for a bundle of stage-$(s-1)$ varieties before observing new orders for its own products. After learning the order size, it combines the existing input inventories ordered in the previous period with other factors to deliver the order received for the current period. In this model, input inventories exist solely because of the one-period gap between the time when orders are placed and the time when goods are used for production. As the critical inputs ordered at the beginning of period t arrive, they are counted as end-of-period input inventories of period t.[16] Input inventories are therefore goods that are temporarily held in a production stage, to be passed from one stage to another along the supply chain. They would not exist if the production process completes instantaneously.[17]

Production Let $Y_{it|s}$ be the stage-s output produced by producer i in period t. $X_{it|s-1}$ is a bundle of stage-$(s-1)$ varieties ordered by a stage-s producer in period t. $Z_{it|s}$ is a bundle of other input

[16] I have assumed that critical inputs ordered in each period are only used for a single period after they arrive. Therefore, there is no distinction between order size and the end-of-period input inventories. This also rules out the possibility of accumulating input inventories for periods beyond the next quarter, which would complicate the model.

[17] This justification for the presence of inventories is similar to the time-to-build models including Kydland and Prescott (1982) and Sarte, Schwartzman and Lubik (2015).

factors used for production that can be freely allocated across producers in all stages and satisfies

$$\sum_{s=1}^{S}\int_0^1 Z_{it|s}\mathrm{d}i = A_{jt}((K_{jt}^s)^{\alpha_j}L_{jt}^{1-\alpha_j} - \phi_p)^{\theta_j}M_{jt}^{1-\theta_j}$$

where A_{jt} is total factor productivity of sector j; K_{jt}^s is the aggregate capital utilization of sector-j capital; L_{jt} is the aggregate labor input in sector j; ϕ_p is a fixed cost that lowers the effective value added component in production; and M_{jt} is the aggregate intermediate inputs excluding the stage-specific critical inputs.[18] The production technology takes the CES form and can be written as

$$Y_{it|s} = A_s\left[\upsilon^{\frac{1}{\kappa}}(X_{it-1|s-1})^{\frac{\kappa-1}{\kappa}} + (1-\upsilon)^{\frac{1}{\kappa}}Z_{it|s}^{\frac{\kappa-1}{\kappa}}\right]^{\frac{\kappa}{\kappa-1}}.$$

Notice that $X_{it-1|s-1}$ is predetermined in period t and the producer adjusts only $Z_{it|s}$ for achieving different output levels.

1. Information Frictions

A key ingredient of the model is the information frictions on aggregate sales. In this model, sales of output varieties fluctuate for two reasons. First, there are various aggregate shocks that disturb economic activities. Sales are affected by these aggregate shocks. Second, producers face idiosyncratic taste shocks that affect market shares of the output varieties. These taste shocks are assumed to be independent from any aggregate shock and persist over time. Whenever producers choose their order size for the critical inputs, they make forecasts on the paths of future sales based on their belief formed over existing information. Information frictions come into play when producers predict future sales.

Noisy Signals In contrast to a model with complete information, producer i only observes its own history of sales $\{Y_{it-l|s}\}_{l>0}$ perfectly at the beginning of period t. The volume of aggregate sales $Y_{t|s}$ in each period is not a common knowledge due to the noisy information on the magnitude of each innovation to aggregate shocks.[19] For each innovation ε_t to a specific structural aggregate shock

[18] Production of this other input bundle takes the same form of technology as the one for sector-specific varieties to be introduced in **??**.

[19] Producers understand the dynamic effects of all aggregate shocks on aggregate sales perfectly. The missing information is the magnitude of each innovation to aggregate shocks but not how they are transmitted and propagated.

that affects $Y_{t|s}$, each producer receives a private signal z_{it} of precision τ_v in every ensuing period $t+l$ with $l \geqslant 0$ such that

$$z_{it} = \varepsilon_t + v_{it} \quad \text{with} \quad v_{it} \sim \text{Normal}(0, 1/\tau_v).$$

Suppose that each producer forms the same prior distribution on ε_t that is Normal$(0, 1/\tau_\varepsilon)$. With Bayesian updating in each period $t+l$, the average of the perceived magnitude of ε_t among all producers satisfies

$$\overline{\mathbb{E}}_{it+l}\varepsilon_t \equiv \int_0^1 \mathbb{E}_{it+l}\varepsilon_t \mathrm{d}i = \frac{(l+1)\tau_v}{\tau_\varepsilon + (l+1)\tau_v}\varepsilon_t \quad \text{for } l \geqslant 0. \tag{1.1}$$

Taste Shocks In addition to the noisy signals, producers also make use of the information on their individual sales that are perfectly observed when they forecast future sales. Let $\{\omega_{it|s}\}_{t\geqslant 0}$ denote a path of market share intensities of output variety i in stage s and assume that it follows the same mean-zero AR(1) process across all varieties after log transformation

$$\log \omega_{it|s} = \rho_\omega \log \omega_{it-1|s} + \varepsilon^\omega_{it} \quad \text{with} \quad 0 < \rho_\omega < 1. \tag{1.2}$$

and that $\int_0^1 \omega_{it|s}\mathrm{d}i = 1$. Changes in $\omega_{it|s}$ can be interpreted as being induced by some form of idiosyncratic taste shocks. They result in changes in $Y_{it|s}$ that are independent from changes in aggregate sales. It is assumed that the structure of the economy is perfectly understood by the producers. However, since the aggregate sales $\{Y_{t|s}\}_{t\geqslant 0}$ is not well observed, producer i does not observe the true values of $\{\omega_{it|s}\}_{t\geqslant 0}$ either.

2. Adjustment Frequency

Related to the specification on information frictions is the frequency at which producers adjust the order size for critical inputs. Suppose that producers choose the optimal order size in each period. This is effectively assuming that producers update their information in every period. In this scenario, there is no need to forecast changes beyond the next period. However, if producers adjust their order size only occasionally, the same information will be utilized for forecasts over multiple horizons and the choice made today will stay relevant for multiple periods. Any belief held by a producer at a moment will therefore be more influential for its future outcomes.

In an attempt to introduce further flexibility into the model, I allow the choice on order size to be adjusted only occasionally after a random amount of time in a fashion as in Calvo (1983*a*). Specifically, let λ be the probability that a producer readjust its order size for critical inputs in a period. The event that a readjustment happens is assumed to be independent across producers. Then, at the time when a producer adjusts its order size, the choice is made in anticipation of all future changes discounted by the chance that they will stay relevant. More precisely, the cost minimization problem at the time of adjustment can be written as

$$\min_{\substack{X^*_{it|s-1},\\ \{Z_{it+l+1|s}\}_{l\geqslant 0}}} \quad \mathbb{E}_{it} \sum_{l=0}^{\infty} \frac{(1-\lambda)^l}{\prod_{l'=0}^{l}(1+i_{t+l'})} \left[(1+i_{t+l})P_{jt+l|s-1}X^*_{it|s-1} + P^Z_{jt+l+1}Z_{it+l+1|s}\right]$$

$$\text{s.t.} \qquad Y_{it+l+1|s} = A_s \left[v^{\frac{1}{\kappa}}(X^*_{it|s-1})^{\frac{\kappa-1}{\kappa}} + (1-v)^{\frac{1}{\kappa}} Z^{\frac{\kappa-1}{\kappa}}_{it+l+1|s}\right]^{\frac{\kappa}{\kappa-1}} \qquad \text{for all } l \geqslant 0$$

where $X^*_{it|s-1}$ is the order size set in period t for each period until the producer is able to adjust it again; i_t is the nominal interest rate.[20] The notation $\mathbb{E}_{it}$ stresses that the expectation is taken over the subjective belief of the individual producer. Notice that if λ takes 1, then the problem degenerates to a static version where adjustment is made in every period.

Depending on the perception on future changes, $X^*_{it|s-1}$ may vary across i. Yet, for aggregate outcomes, it suffices to know the average across producers, which is denoted as $\overline{\mathbb{E}}_{it}X^*_{it|s-1}$. The aggregate input inventories for all producers in stage s depends on the level of $\overline{\mathbb{E}}_{it}X^*_{it|s-1}$ over the history and the distribution of producers who have last adjusted in different periods. The law of motion for the aggregate stage-$(s-1)$ input inventories can be written as

$$X_{t|s-1} = \lambda \sum_{l\geqslant 0}(1-\lambda)^l \overline{\mathbb{E}}_{it-l}X^*_{it-l|s-1} \tag{1.3}$$

where $\lambda(1-\lambda)^l$ is the share of producers who have not adjusted again for l periods since their last adjustment. The lack of flexibility in adjusting the order size strengthens the importance of producers' perception on future sales at the time of adjustment on shaping the aggregate behavior.

[20] The discount factor for the lth period is $\frac{\beta^l \Xi_{t+l} P_t}{\Xi_t P_{t+l}}$ where Ξ_t is the expected marginal value of a dollar in period t and P_t is the price index for the final consumption bundle to be introduced in **??**. With additional model details to be introduced, the discount rate is equal to $\prod_{l'=0}^{l}(1+i_{t+l'})^{-1}$ due to the consumption Euler equation.

3. Pricing

For simplicity, I do not consider strategic behavior that may arise between supplier-customer pairs. Instead, I assume that outputs from all producers along the supply chain are bundled together before being distributed to the downstream customers in a market environment. Each individual producer charges a constant markup over its marginal cost in each period. The downstream customers need only consider the aggregate price for the critical inputs from the suppliers that are bundled together using a CES aggregator. For the linearized model, the level of markup charged by the producers is not required for solving the impulse responses. It suffices to know the log deviations in nominal marginal cost over time. Since the most upstream stage-1 producers use only the factor bundle $Z_{it|1}$, their nominal marginal cost is always equal to P_t^Z, the price index for this input factor bundle. Since $Z_{it|s}$ is aggregated using the same technology for all s, all supply chain members face a common price P_t^Z for this bundle. For stage-s producers with $s > 1$, their marginal costs vary depending on their input inventories. Let $\mu_{it|s}$ denote the marginal cost of producer i in stage s. The individual marginal cost satisfies

$$\mu_{it|s} = P_t^Z \left[\frac{(1-\upsilon)Y_{it|s}}{Z_{it|s}} \right]^{-\frac{1}{\kappa}}. \tag{1.4}$$

B. Optimal Partial Equilibrium Responses to Aggregate Shocks

To characterize responses to aggregate shocks, I consider a linearized economy described in terms of log deviations around a steady state that are linear transformations of aggregate shocks. To get into the essence without dealing with auxiliary complexity, here I only consider a partial equilibrium where the sales of the most downstream stage-3 producers are exogenously given. Let variables with a hat on top denote their log deviations and variables with a star on subscripts denote their steady-state levels. Results in Appendix 1.C imply that, when a stage-s producer adjusts at time t, the optimal order size of the stage-$(s-1)$ output satisfies

$$\hat{X}^*_{it|s-1} = \frac{1}{m} \sum_{l=0}^{\infty} \beta^l (1-\lambda)^l \left[\kappa \chi_s (\hat{P}^Z_{jt+l+1} - \hat{\imath}_{t+l} - \hat{P}_{jt+l|s-1}) + \mathbb{E}_{it} \hat{Y}_{it+l+1|s} \right] \tag{1.5}$$

where $\overline{m}$ is a constant only depending on parameters and

$$\chi_s \equiv \frac{(1-\upsilon)^{\frac{1}{\kappa}} Z_{*|s}^{\frac{\kappa-1}{\kappa}}}{\upsilon^{\frac{1}{\kappa}} X_{*|s-1}^{\frac{\kappa-1}{\kappa}} + (1-\upsilon)^{\frac{1}{\kappa}} Z_{*|s}^{\frac{\kappa-1}{\kappa}}} \tag{1.6}$$

is a value arising from log-linearization that can be treated as a parameter.

Equation (1.5) says that given anticipated future price changes, the producer chooses the order size based on its forecasted changes in its own future order sizes while it is not adjusting again. The optimal order size chosen at time t therefore hinges on how producers perceive future sales. Taking the average across all producers yields

$$\hat{X}_{t|s-1} \equiv \int_0^1 \hat{X}^*_{it|s-1} \mathrm{d}i = \frac{1}{\overline{m}} \sum_{l=0}^{\infty} \beta^l (1-\lambda)^l \left[\kappa \chi_s (\hat{P}^Z_{jt+l+1} - \hat{i}_{t+l} - \hat{P}_{jt+l|s-1}) + \overline{\mathbb{E}}_{it} \hat{Y}_{it+l+1|s} \right] \tag{1.7}$$

where the individual index i is dropped from $\hat{X}_{t|s-1}$ as the variable represents the aggregate changes. With a CES aggregator for stage-s varieties, it must be that $\hat{X}_{t|s-1} = \hat{Y}_{t|s-1}$. Hence, Equation (1.7) also characterizes how changes in one production stage affect the sales of the immediate upstream suppliers. For characterizing aggregate responses, the remaining task is to characterize the average belief on future order size $\overline{\mathbb{E}}_{it} \hat{Y}_{it+l+1|s}$.

Suppose ε_0 is the magnitude of a one-time shock to an aggregate variable that induces all fluctuations since time 0.[21] With $\hat{Y}_{it|s} = \hat{\omega}_{it|s} + \hat{Y}_{t|s}$ for all t, the observed changes in individual order sizes are either due to idiosyncratic taste shocks or aggregate shocks. Since the aggregate change $\hat{Y}_{t|s}$ is linear with respect to the magnitude of ε_0, Equation (1.1) implies that at any given time $t > 0$,

$$\overline{\mathbb{E}}_{it} \hat{Y}_{t|s} = \frac{(t+1)\tau_v}{\tau_\varepsilon + (t+1)\tau_v} \hat{Y}_{t|s}.$$

With the observed $\hat{Y}_{it-1|s}$, the producer infers the component induced by the taste shocks. Since $\overline{\mathbb{E}}_{it} \hat{Y}_{it-1|s} = \hat{Y}_{t-1|s}$, the average perception of the sales changes due to market share fluctuations satisfies

$$\overline{\mathbb{E}}_{it} \hat{\omega}_{it-1|s} = \overline{\mathbb{E}}_{it} \hat{Y}_{it-1|s} - \overline{\mathbb{E}}_{it} \hat{Y}_{t-1|s} = \frac{\tau_\varepsilon}{\tau_\varepsilon + (t+1)\tau_v} \hat{Y}_{t-1|s}$$

[21] The characterization here for impulse responses with respect to a single shock is the building block for the realistic case in which multiple shocks are hitting the economy in each period.

which would be zero under full information. That is, with noisy signals on aggregate shocks, producers discount the information on aggregate shocks on average and "mistakenly" attribute changes induced by aggregate shocks to those induced by idiosyncratic shocks. Since the idiosyncratic shocks are persistent as in Equation (1.2), the average prediction of future sales satisfies

$$\overline{\mathbb{E}}_{it}\hat{Y}_{it+l+1|s} = \frac{\tau_\varepsilon}{\tau_\varepsilon + (t+1)\tau_v}\rho_\omega^{l+2}\hat{Y}_{t-1|s} + \frac{(t+1)\tau_v}{\tau_\varepsilon + (t+1)\tau_v}\hat{Y}_{t+l+1|s} \quad \text{for} \quad l \geqslant 0. \tag{1.8}$$

Interestingly, Equation (1.8) combined with Equation (1.7) implies that the noisy information on ε_0 induces backward-looking behavior among forward-looking individuals.[22] Given a path of aggregate sales $\{\hat{Y}_{t|s}\}_{t\geqslant 0}$ and paths of price changes $\{\hat{P}^Z_{jt}\}_{t\geqslant 0}$, $\{\hat{P}_{jt|s-1}\}_{t\geqslant 0}$ and $\{\hat{i}_t\}_{t\geqslant 0}$, Equations (1.7) and (1.8) provide the average optimal order size set by stage-s producers when they adjust at time t. The aggregate order size placed in each period by all stage-s producers including those who have not adjusted again at time t can then be obtained with a log-linearized version of Equation (1.3). The model counterpart of the input inventories in data is the sum of input inventories held by producers in all stages. Let $\hat{X}_t$ denote the log deviations in the aggregate input inventories. Assuming that all producers have the same inventory-to-sales ratio h_* in the steady state,[23] we have

$$\hat{X}_t = \sum_{s=2}^{S} \frac{X_{*|s}}{\sum_{s=2}^{S} X_{*|s}}\hat{X}_{t|s} = \frac{1}{\overline{h}}\sum_{s=2}^{S} h_*^{S-s}\hat{X}_{t|s} \tag{1.9}$$

where $\overline{h} \equiv \sum_{s=2}^{S} h_*^{S-s}$.

Lastly, in Appendix 1.C, I provide the characterization on the last piece of component required for characterizing the partial equilibrium responses. Changes in the aggregate price index for stage-s output after linearizing the model can be written as

$$\hat{P}_{t|s} = \hat{P}^Z_t + \frac{1-\chi_s}{\kappa\chi_s}(\hat{Y}_{t|s} - \hat{X}_{t-1|s-1}) \quad \text{for } s > 1. \tag{1.10}$$

The dynamic behavior of input inventories is jointly characterized by Equations (1.7)–(1.10).

[22] Although the context is different, results sharing similar intuition may arise when considering a pricing problem. For example, see Nimark (2008).

[23] The inventory-to-sales ratio here refers to $h_* \equiv \frac{P_{*|s-1}X_{*|s-1}}{P_{*|s}Y_{*|s}}$.

III. Input Inventory Responses in Model

In this section, I assess the capability of the model described in Section II for reproducing the lagged inventory-sales comovement documented in Section I. To highlight the main insight, I examine the production problem in isolation, taking as exogenous the path of sales of the most-downstream producers. Towards the end, I discuss the distinction from some related models.

A. Bringing the Model to Data

To assess the capability of the model to reproduce the observed impulse responses shown in Section I, I estimate the key model parameters by minimizing the Euclidean distance between model impulse responses and their empirical counterparts following a monetary shock. Equations (1.7)–(1.10) imply that given a set of parameters, the aggregate input inventory responses are fully determined if we feed in the model a path of changes in sales in the most downstream stage S, a path of interest rate changes and a path of input factor price changes. I therefore first obtain the impulse responses of the empirical counterparts of these variables.

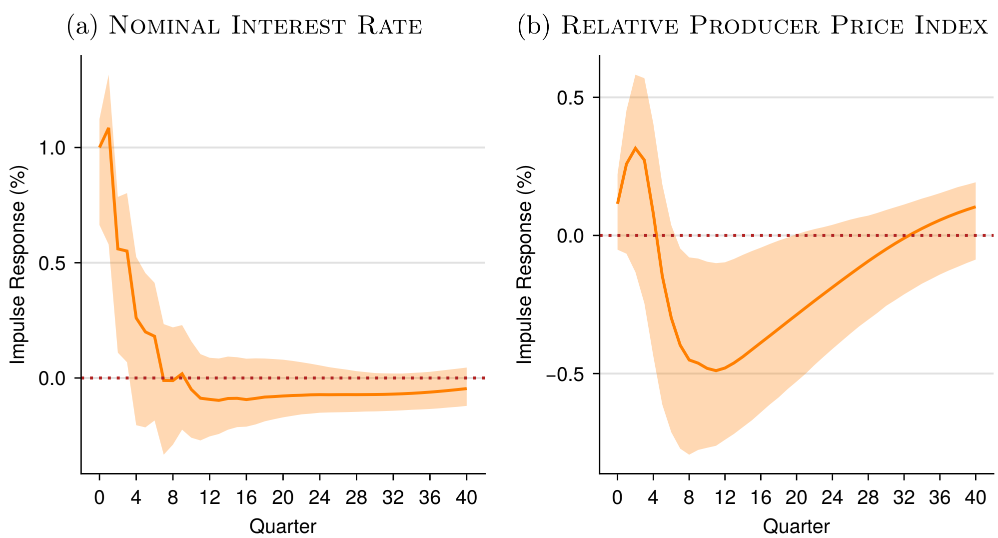

FIGURE 1.6: Impulse Responses to a Romer-Romer Shock

The impulse responses of sales are obtained in the same way as before and have been shown in Figure 1.3a. These estimates are interpreted as the empirical counterpart of impulse responses of sales among the most downstream producers $\{\hat{Y}_{t|S}\}_{t\geqslant 0}$ following an innovation to the nominal interest rate i_t. In the model, sales of producers that are not in the last production stage are counted as input inventories. Impulse response estimates of the federal funds rate with respect to the Romer-Romer shocks are treated as the empirical counterpart of the impulse responses of nominal interest rate and are shown in Figure 1.6a. Selection for the empirical counterpart of the changes in the price of the other input factor bundle $\{\hat{P}_t^Z\}_{t\geqslant 0}$ is less obvious. If we assume that input prices faced by producers capture the changes in factor prices, then a producer price index (PPI) could be a reasonable candidate. Here, I use the PPI of materials and components for manufacturing as a proxy for factor prices. Since the price index of a final consumption bundle will be used for measuring the overall price level of the model economy in the complete model to be introduced in **??**, I consider the ratio between the PPI and the measure of consumption price as a measurement for P_t^Z/P_t. The impulse responses of that are shown in Figure 1.6b.

I now turn to the relevant parameters for generating the impulse responses of input inventories in the model. The distribution of the stages of production is not observed and clearly varies across the products in the reality. Nonetheless, throughout the main text, I assume that there are three production stages ($S = 3$). Recall that producers in the most upstream stage 1 does not require ordering critical inputs one period in advance and hence they do not hold input inventories. Only the producers in the stages 2 and 3 hold input inventories. In addition, I assume that the idiosyncratic taste shocks induce market share fluctuations that follow an AR(1) process. For the production technology that combines the critical input with a bundle of other input factors, the elasticity of substitution κ remains the same across production stages and χ_s defined in Equation (1.6) takes the same value χ. With these simplifications, the parameters that need to be specified are the precision ratio of signal vs prior τ_v/τ_ε, the persistance of taste shocks ρ_ω, the arrival rate of the Calvo adjustment shock λ, the technological parameters κ and χ.

In principle, the parameters can be chosen by solving the nonlinear least squares

$$\min_{\Theta}[\ J(\Theta) - \hat{J}\ ]^{\mathsf{T}} W^{-1} [\ J(\Theta) - \hat{J}\ ]$$

with $J(\Theta)$ being the model outcomes associated with a parameter vector Θ and $\hat{J}$ being the empirical counterpart used as targets. W is a diagonal weight matrix with weights proportional to the length of each bootstrap percentile interval of each impulse response estimate. Here, only the impulse responses of input inventory are involved in $J(\Theta)$. Model parameters would therefore be selected by making the model predictions on input inventory responses as close as possible to the empirical counterpart. However, with the impulse responses of input inventory alone as the empirical target, not all parameters can be reliably determined by solely relying on the process described above.[24] For the technological parameters κ and χ, I directly impose their values to be 0.1 and 0.95 respectively. The former is chosen to allow complementarity between the critical input and the other input factor bundle. The latter affects the sensitivity of the adjustment in the other input factor bundle for fulfilling the sales when the predetermined input inventory level deviates from the optimal level. A greater value of χ weakens the responses of $\hat{Z}_{t|s}$ with respect to $\hat{Y}_{t|s}$. This can affect the general equilibrium outcomes but are not influential in the partial equilibrium context considered in this section. The resulting parameters are collected in Table 1.1.

TABLE 1.1: Model Parameters for Partial Equilibrium

Estimated parameters	
Precision of signal relative to precision of prior (τ_v/τ_ε)	0.00
Persistence of idiosyncratic taste shocks (ρ_ω)	0.99
Probability of adjusting order size (λ)	0.32
Calibrated technological parameters	
Elasticity of substitution (κ)	0.10
Steady-state share (χ)	0.95

B. Evaluating the Model

Figure 1.7a shows that, under the estimated parameters, the model outcome fits the empirical target remarkably well. However, for the model outcome to fit this well, the private signals need to be uninformative. In other words, the model combined with the data does not suggest a meaningful learning process among the producers on the magnitude of aggregate shocks. The model structure implies that information frictions faced by the producers must have been substantial.[25] With λ being around 0.3, the model also suggests quite strong inertia in the adjustment of order sizes placed for the critical inputs.

To gain insights on how the model achieves the fit, I examine the underlying components of the model. Since the changes in historical sales are important for producers from different stages to

[24]Intuitively, identification requires additional information not captured in $\hat{J}$.

[25]An alternative interpretation could be that producers do not recognize the impulse responses induced by monetary shocks.

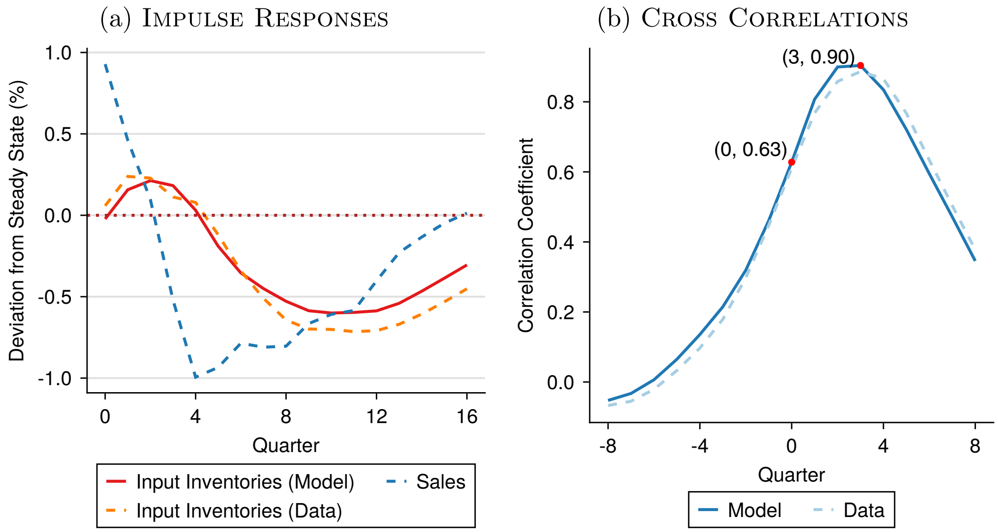

FIGURE 1.7: Partial Equilibrium Outcomes with Given Sales Path

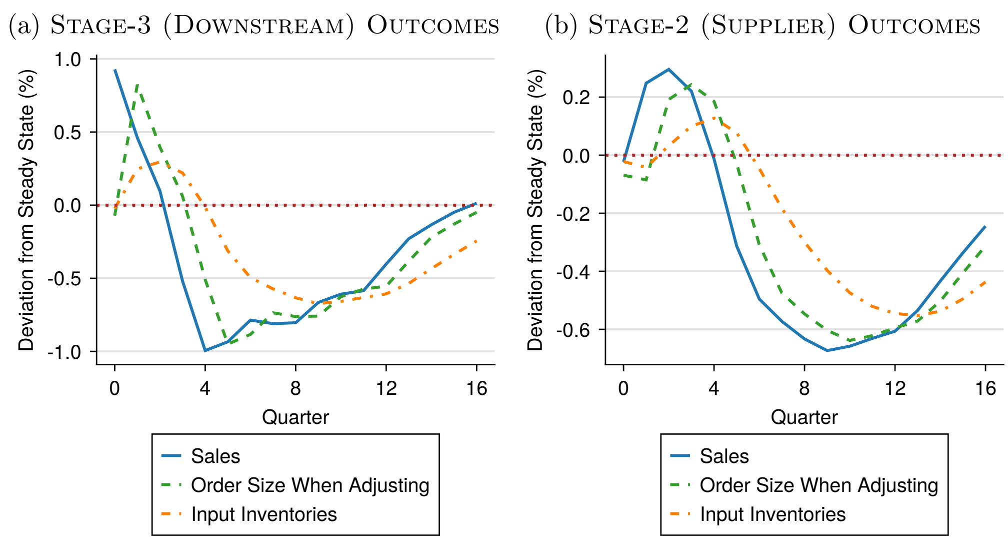

FIGURE 1.8: More Details on Partial Equilibrium Outcomes

determine their order sizes for critical inputs, it is useful to see how the past sales affect their choices when they have the chance to adjust order size. Figure 1.8a compares the average changes in order sizes chosen by producers in the most downstream stage 3 when they adjust with changes in their sales. Notice that because of the use of sales in the last quarter for forecasting future sales, changes in order sizes closely follow changes in their sales in the last quarter. Except some small effects due to price changes, the producers do not respond in anticipation of future changes induced by the aggregate shock. They perceive the sales changes as being induced by the persistent idiosyncratic taste shocks. Figure 1.8a additionally shows the input inventories held by the producers. Recall that in this model, the evolution of input inventories is driven by changes in order sizes. Due to the staggered Calvo-style adjustment in order size, the changes in input inventories demonstrate smoother gradual responses following the adjustment in order sizes. In Figure 1.8b, similar patterns hold true for the stage-2 producers that serve as the suppliers for the stage-3 producers.[26]

The model involves three types of frictions. First, the outputs from the suppliers have to be ordered one period in advance. Second, producers predict future sales without the information on aggregate shocks. Third, order sizes are only occasionally adjusted. The first friction is necessary for justifying the presence of input inventories. How about the other two? Figure 1.9a shows the best that the model can achieve for fitting the empirical target when the private signals on aggregate shocks are restricted to be precise. Since sales are declining over the first year, input inventories also decline from the beginning without any sign of delay existing in the empirical counterpart. Figure 1.9b shows the result when the Calvo-style adjustment friction is mostly removed with a high probability of adjustment ($\lambda = 0.99$). In this case, the model outcome still demonstrates some extent of the lagged inventory-sales comovement. However, the fitness with the empirical counterpart is noticeably undermined as the input inventories reach the trough too early. Combined together, the results suggest that the information friction is the most important ingredient for generating the lagged comovement; while the sluggish adjustment greatly improves the fitness of the model.

How do the parameters for the CES aggregator between the critical input and the bundle of other factors affect the model outcome? Figure 1.10a shows that as the elasticity goes from the baseline level of 0.1 to 0.9, the response of input inventories on impact becomes more negative. Intuitively, as interest rate increases, the opportunity cost of ordering the critical input rises due to the need of incurring the cost in advance. As a result the producers would rather reduce the order size to substitute more of the other input factors that do not require incurring cost in advance.[27] To fit the empirical target, the elasticity of substitution needs to be relatively small so that the impact of price changes on the order sizes chosen by producers are limited. The value of the other parameter χ is not well identified with the data utilized so far. Holding κ at the level of 0.1, varying

[26]Notice that sales of stage-2 producers match the input inventories of stage-3 producers because the entire sales arrive at the stage-3 producers as input inventories at the end of each period.

[27]A similar result holds if the suppliers provide credit to the customers. In that case the price of the critical input would still increase as the cost of credit increases.

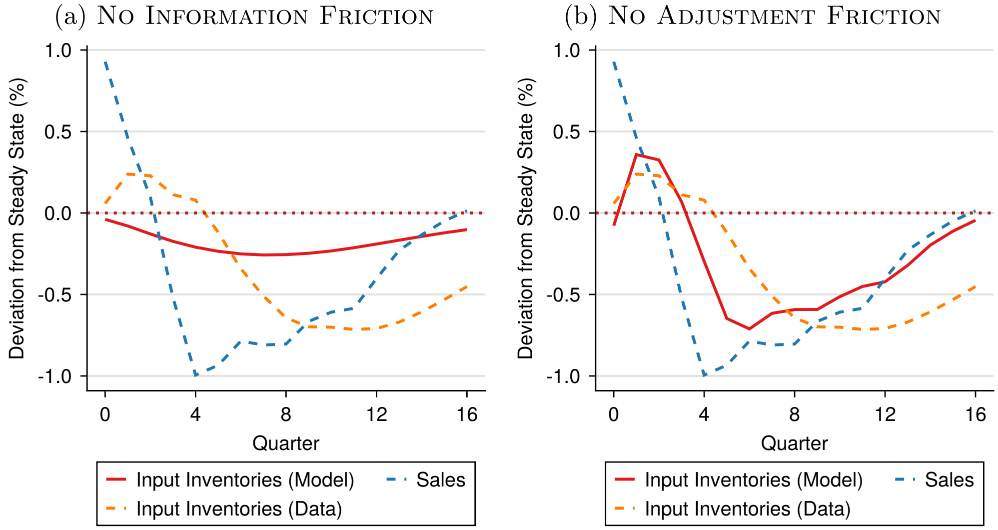

FIGURE 1.9: Outcomes from Restricted Models

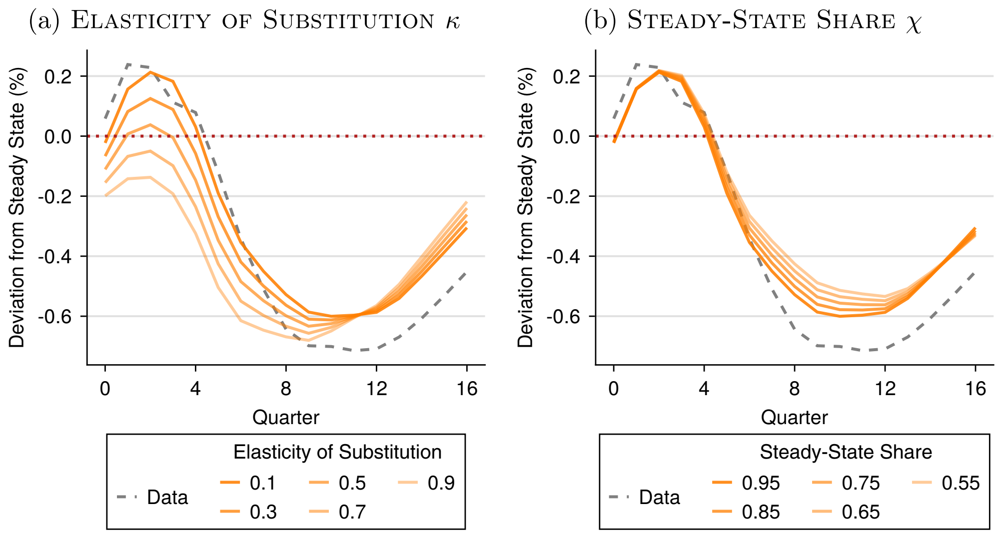

FIGURE 1.10: Sensitivity to Technological Parameters

the value of χ does not affect the input inventory responses much. This parameter, however, will be more important when considering the general equilibrium outcomes as it will affect the sensitivity of marginal cost of production when the input inventory deviates from the optimal level that would be attained in the absence of frictions. The value of χ needs to be large to prevent the marginal cost from responding too much to aggregate shocks.

C. Plausibility

The capability of the model for reproducing the lagged inventory-sales comovement hinges on a combination of assumptions. Here, I discuss the plausibility of the crucial ones based on suggestive evidence.

The Lag in Order Fulfillment The model requires a one-quarter lag for receiving the critical inputs from suppliers. For durable manufacturing industries, this assumption seems acceptable based on data from the US Census Manufacturers' Shipments, Inventories, and Orders (M3) survey. Figure 1.11 plots the magnitude of unfilled order over time as a ratio over the monthly shipment among producers that hold unfilled orders.[28] The solid blue line is for the data across all durable manufacturing industries; while the thinner grey lines in the background are for the 2-digit SIC level data among 7 durable manufacturing industries. Assuming that each dollar value of unfilled orders is being fulfilled at the same rate, the data suggest that it typically takes about 4 months for orders to be fulfilled. Since a substantial fraction of intermediate inputs for the durable sector come from suppliers within the durable sector, the data from the M3 survey suggests that assuming a one-quarter lag for receiving the orders is realistic.

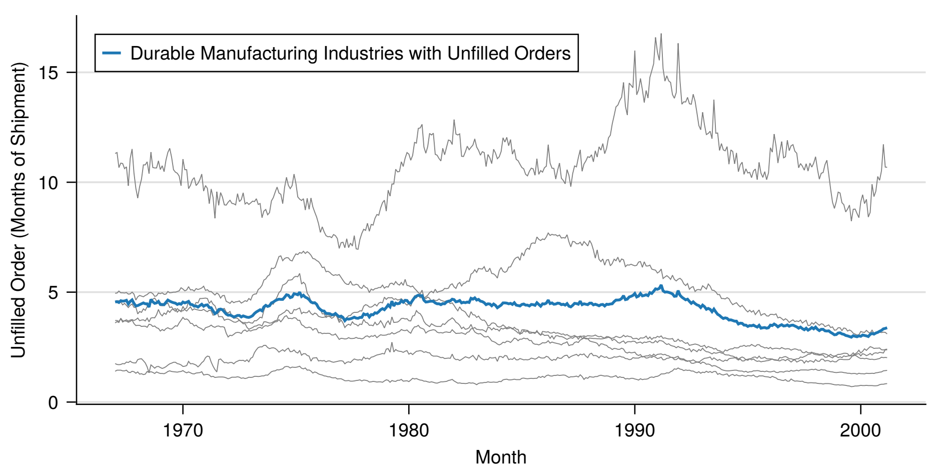

FIGURE 1.11: Unfilled Orders in Durable Manufacturing Sector

Note: Data are based on the US Census Manufacturers' Shipments, Inventories, and Orders (M3) survey. The grey lines represent data for seven 2-digit SIC level industries within the durable manufacturing sector.

[28]The ratios are readily available from the dataset and only cover manufacturers that report both unfilled order and shipment. Although not all producers follow a "ship-to-order" mode, this ratio is still representative as the bulk of shipment comes from firms that hold unfilled orders.

Information Frictions As an important model ingredient for generating the lagged inventory-sales comovement, producers forecast future sales changes based on their own historical sales. Due to the nature of information frictions, direct observation on how producers make forecasts are not generally available. However, the presence of information frictions is a critical consideration in supply chain management and their real-world significance can hardly be denied. In fact, management scientists widely accept the presumption that orders are the only information firms exchange under conventional supply chain management. For example, in an influential work, Lee, Padmanabhan and Whang (1997) write that "Typically, an upstream supplier relies only on the order data from the downstream retailer". They deem the lack of information other than the orders received by producers as a source of distortion, as the orders from the downstream customers do not necessarily reflect the current final demand. Furthermore, a substantial amount of analysis in the supply chain management literature concerns the performance of supply chain under different schemes of information sharing (e.g., Cachon and Fisher 2000; Aviv 2001; Simchi-Levi and Zhao 2003; Ouyang 2007). Considering the sole amount of efforts invested in understanding the impact of information sharing, I consider the model ingredient on information frictions as a reasonable approximation to the reality.

IV. Relations to Other Models

The inventory problem considered in the last section is related but different from models considered in prior work. The distinction is motivated by the goal of reproducing the lagged inventory-sales comovement that the prior work does not consider. Here, I highlight the important distinctions from some representative models.

The Precautionary Stockout-Avoidance Motive The precautionary stockout-avoidance motive is a well-known justification for the presence of inventories (Kahn 1987). In Section IC, I have shown results for a version of such models adapted from Wen (2011), where the same precautionary motive is used to model input inventories.[29] The core ingredient of the model is that the downstream producers face idiosyncratic demand shocks in each period that are not observed at the time when the producers obtain their inputs at the beginning of each period. Under this assumption, the inventories

[29] Under this framework, the input inventories are simply output inventories from the upstream producers that are held by the downstream producers.

held by each individual producer either runs down to zero because of demand that is sufficiently strong or otherwise remains positive. Notice that, although the idiosyncratic demand shocks are not observed, producers anticipate the (correct) aggregate variables in this model economy. In particular, producers choose the same inventory level at the beginning of each period that is proportional to the contemporaneous aggregate demand. There are only two possible ways for aggregate shocks to affect the inventory level. First, aggregate shocks may alter the inventory-to-sales ratio because of their impact on the expected value of inventories that are left to the next period. For this channel, the effects are completely forward-looking because only the future changes matter. Second, aggregate shocks may directly affect the aggregate sales. For this channel, since the optimal inventory level is only proportional to the contemporaneous aggregate sales, there is no way for the inventory level to track the lagged sales. Combined together, it is impossible for the optimal inventory stock to reproduce the lagged inventory-sales comovement demonstrated in Section I in such a model.

Time to Build Kydland and Prescott (1982) consider a model in which investment goods take multiple periods to build. Sarte, Schwartzman and Lubik (2015) develop a model that captures this idea with a more flexible multistage production technology and view the goods involved in each stage of the production as input inventories. In Section IC, I have shown that this model does not generate the lagged inventory-sales comovement. To see why their model does not have such a capability but mine does, it is important to recognize the distinctions on how production coordinates across different stages. In particular, allocations of input factors across the stage-specific goods in Sarte, Schwartzman and Lubik (2015) are determined via a set of production Euler equations that equalize the discounted marginal product across goods in different stages within each period. The coordination across different production stages is too efficient in the sense that there is no suboptimal allocation of input factors from the perspective of a social planner. In contrast, in my model, because producers order critical inputs in advance based on their biased forecasts, allocations of input factors do not necessarily reflect the optimal allocations in a frictionless environment. In particular, receiving unexpectedly high order sizes, only the most downstream producers are directly exposed to the sales changes on impact. Suppliers of the directly affected producers only start to adjust their order sizes for critical inputs when their own sales get affected with a delay. Such delayed responses induce sales fluctuations going upward along the supply chain that do not necessarily align with contemporaneous sales fluctuations among the most downstream producers. Such frictional transmission of sales fluctuations across stages of supply chain is a unique feature of my model that is distinct from the conventional time-to-build models.

(S, s) **Models** A popular class of inventory models is based on the (S, s) policy induced by a fixed cost incurred when making orders (Caplin 1985; Khan and Thomas 2007). In the international trade literature, Alessandria, Kaboski and Midrigan (2010) highlight the capability of an (S, s) inventory problem for capturing the shipping lags and fixed costs involved in importing goods. More recently, Alessandria et al. (2023) introduce shocks of shipping delays into the (S, s) inventory problem to quantify the aggregate impact of increased shipping time observed during the pandemic of COVID-19. Although the (S, s) models are powerful modeling devices, the lagged inventory-sales comovement highlighted in this paper is not what such models intend to capture when they are developed. For reasons that are similar to the other models, the (S, s) models by themselves are not readily capable of generating outcomes considered in this section. In particular, although Alessandria et al. (2023) highlight shipping delays involved in the orders for intermediate inputs, both the shocks of interest and model mechanism considered are different. They are interested in the shocks to shipping time observed during the unusual supply-chain disruptions following the pandemic; while adjustment of orders for critical inputs is treated as a technological assumption that is held unchanged under my model environment. In (S, s) models, idiosyncratic demand shocks are needed for smoothing out the time when producers place orders; while in the model of this paper, the idiosyncratic taste shocks are introduced to justify the bias in the sales predictions made by producers. The possibility of stocking out is important for shaping how producers order and use the inputs under an (S, s) model; while in my model, perception of future changes and frictions jointly govern the ordering behavior.

V. Conclusions

In this paper, I document an overlooked feature of the comovement between inventories and sales. Namely, durable input inventory movements lag sales movements by about three quarters. I show that this feature is discernible in both the unconditional cyclical variation and in impulse responses with respect to identified aggregate shocks. Existing inventory models are not readily capable of reproducing this feature observed in data.

Motivated by the empirical finding, I develop a tractable supply chain production problem that is capable of reproducing the lagged inventory-sales comovement. The model features multiple production stages across a serial supply chain where producers have to order critical inputs from suppliers in advance. In this model, producers determine the order sizes based on their biased forecasts on future sales due to information frictions and only adjust their order sizes occasionally in a staggered fashion. With these two types of frictions, the model is capable of fitting the empirical impulse responses following a monetary shock very well.

Acknowledgements

Chapter 1, in part, is currently being prepared for submission for publication of the material. The dissertation author was the primary author of this paper.

Appendix 1.A. Additional Details on Inventory-Sales Comovement

This section provides additional details on the inventory-sales comovement.

A. Data

Real inventories and sales data from NIPA underlying detail tables are only based on the Standard Industrial Classification (SIC) before 1997 and only based on the North American Industrial Classification System (NAICS) since then. To obtain continued measurement at the broad level of durable manufacturing industry, I splice the corresponding time series and adjust the log scale of the SIC series so that the one-year growth rates over months that are affected by the transition point are comparable to the adjacent ones in 1997 that are based on NAICS.[30] To construct data for real input inventories, I use the real inventory data by stage of fabrication and sum up the materials and supplies with work-in-process.[31]

B. Summary Statistics

Table 1.2 collects summary statistics for different types of inventories across two sample periods. Data for the earlier sample are based on the SIC industries; while those for the more recent sample are based on the NAICS industries. Notice that input inventory (sum of materials and supplies and works-in-process) from the durable manufacturing sector constitutes about a quarter of the total inventory over the main sample period (1970–1996). It has also been the most volatile based on the standard deviation of the annual growth rates over this earlier sample period. However, the volatility of the finished-goods inventory in the durable manufacturing sector seems to become larger over the second sample period. This is largely due to the extraordinarily fast growth in the finished-goods inventory towards the end of 1990s. The share of input inventory also seems to be declining. Nonetheless, it still remains an important component of aggregate inventory based on the magnitude and volatility.

[30] Inventory data in 1997 are available under both SIC and NAICS. I take advantage of the overlapped data for rescaling the SIC data. For sales data, there is no overlap and I instead use the 1996 and 1997 data.

[31] In general, summation of time series data should be based on nominal data for maintaining appropriate weights across individual components over time. However, there are no separate price indices for inventories in different stages. Since the summation of real inventories across all three stages is very similar to the provided total real inventories, I consider the summation across real inventories as innocuous.

TABLE 1.2: SUMMARY STATISTICS FOR INVENTORIES

	1970–1996		1997–2019	
	Share (%)	Volatility (%)	Share (%)	Volatility (%)
Manufacturing	48.31	3.07	39.93	2.84
Durable goods	31.49	4.21	24.40	3.61
Materials and supplies	9.56	5.59	8.64	3.81
Works-in-process	14.04	5.35	9.04	4.40
Finished goods	7.88	3.64	6.74	4.93
Nondurable goods	17.36	2.14	16.32	2.68
Materials and supplies	6.85	2.79	5.86	2.92
Works-in-process	2.76	2.99	2.91	5.63
Finished goods	7.74	3.24	7.56	2.26
Wholesale	25.07	3.47	27.82	3.74
Durable goods	17.56	4.14	17.59	5.50
Nondurable goods	7.88	4.36	11.08	2.76
Retail	26.62	3.96	32.25	3.87

Notes: "Share" refers to the average share of the (nominal) stock of each type of inventory in the total level. "Volatility" is computed as the standard error of the annual growth rate of each type of inventory. Data are based on the monthly time series from the NIPA underlying detail tables.

C. Other Inventory-Sales Comovement

The main text has focused on the lagged inventory-sales comovement between input inventory and sales in the durable manufacturing sector. Here, for the sake of completeness, I show the estimation results for the other types of inventories that are available in the NIPA underlying detail tables.

1. Unconditional Comovement

Figures 1.12–1.15 plot the coefficients for cross-correlations between inventories and sales over different leads/lags in different sectors. The coefficients for the durable wholesale sector and the retail sector also demonstrate a pattern that looks similar to the input inventories in the durable manufacturing sector. However, as we shall see shortly, this is no longer the case for the conditional comovement.

2. Comovement Following Monetary Shocks

Figure 1.16 shows the results between input inventory and sales for the nondurable manufacturing sector instead of the durable one. Although the nondurable sector also demonstrates some

(a) Annual Growth Rate (b) Cyclical Component

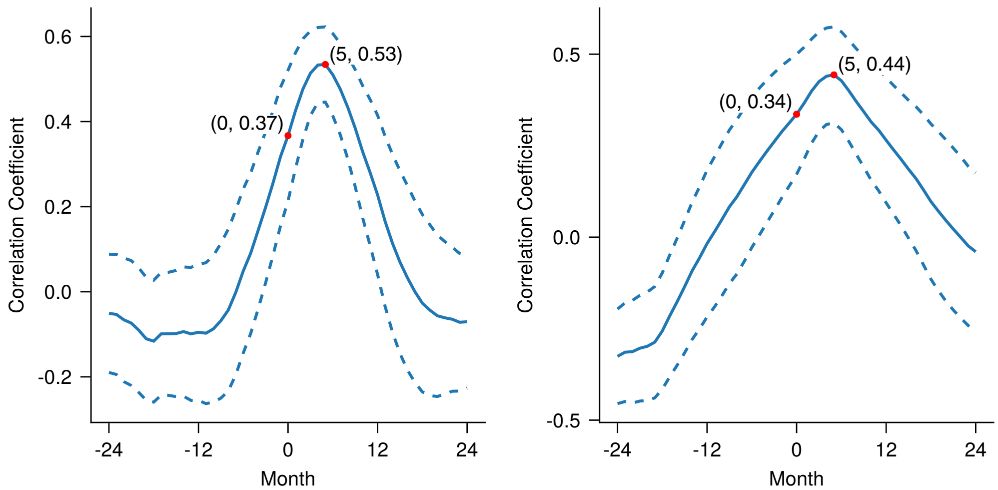

FIGURE 1.12: Cross-Correlations between Nondurable Input Inventories and Sales

Note: The correlation coefficients are for $\rho(\text{Inventory}_t, \text{Sales}_{t-l})$ across different values of l. Annual growth rates are computed as log differences with the 1-year lag. Cyclical components are two-year horizon forecast errors following Hamilton (2018). Dashed lines represent 90% confidence intervals based on robust standard errors obtained from a GMM Newey and West (1987) procedure.

(a) Annual Growth Rate (b) Cyclical Component

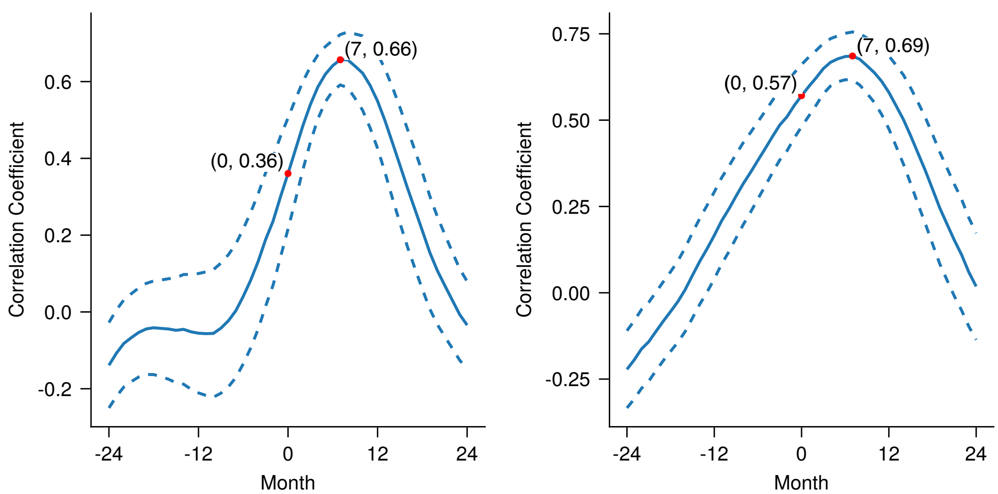

FIGURE 1.13: Cross-Correlations between Durable Wholesale Inventories and Sales

Note: The same note for Figure 1.12 applies.

lagged comovement, the contemporaneous correlation is much stronger and the lag is shorter. This is a reason for why the main text only focuses on the durable sector.

The remaining types of inventories are all output inventories in the sense that they are not

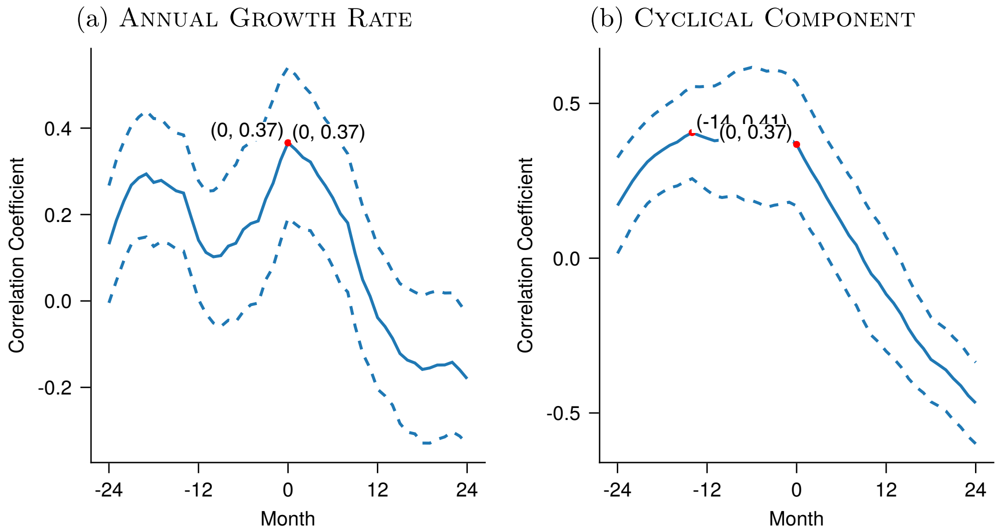

FIGURE 1.14: Cross-Correlations between Nondurable Wholesale Inventories and Sales

Note: The same note for Figure 1.12 applies.

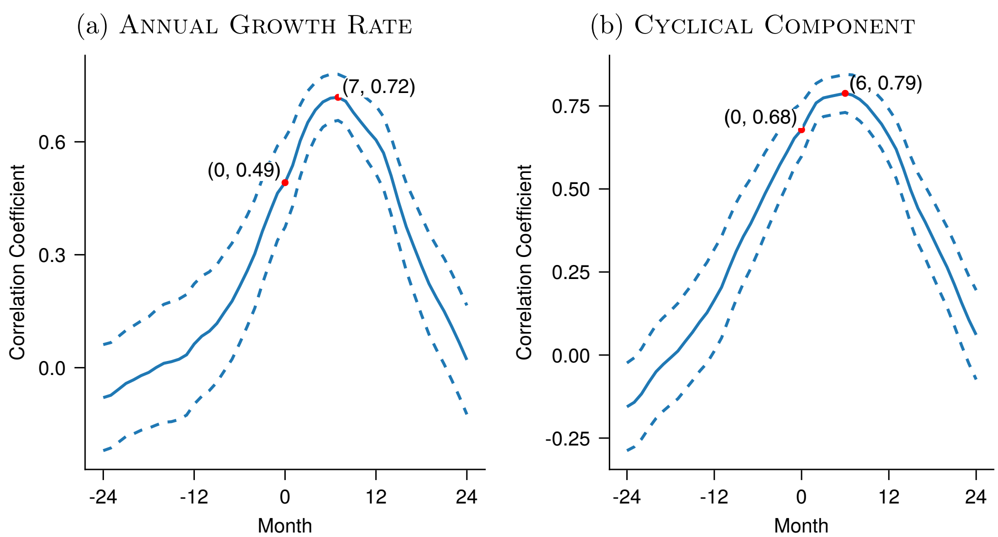

FIGURE 1.15: Cross-Correlations between Retail Inventories and Sales

Note: The same note for Figure 1.12 applies.

used as inputs for the production of the firms that hold these inventories. Notice that, as shown in Table 1.2, the majority of the output inventories are held by wholesalers and retailers. Output inventories held by manufacturers are often thought of as goods that are to be shipped to these wholesalers and retailers. Figures 1.17 and 1.18 show that the output inventory in the manufacturing

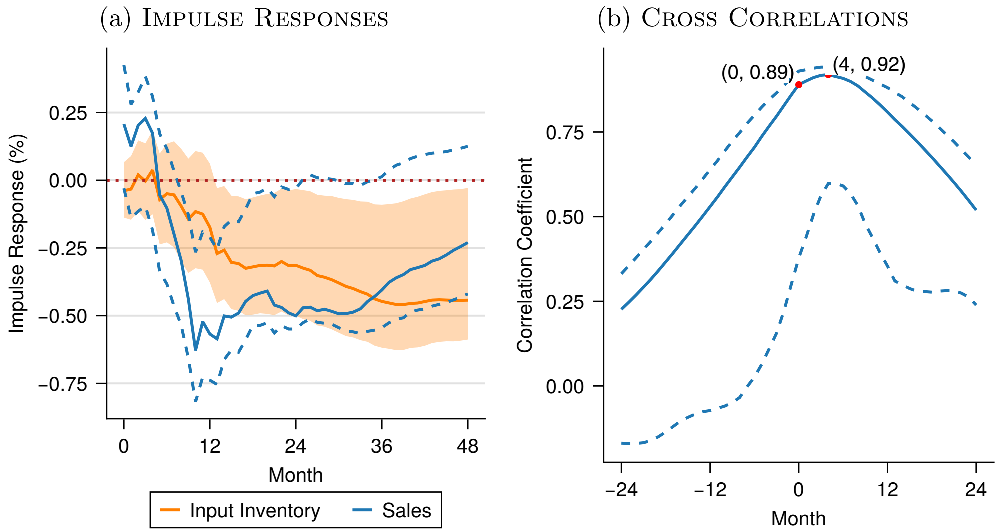

FIGURE 1.16: Input-Inventory-Sales Comovement in the Nondurable Manufacturing Sector

sector seems to act as a buffer for their sales. In particular, the contemporaneous correlation coefficients are both negative, suggesting a role that is consistent with the production smoothing motive instead of the (positive) comovement that is emphasized at the more aggregate level. For the wholesale and retail sectors, the results are shown in Figures 1.19–1.21. Unlike the output inventories held by the manufacturers that look like buffer, these inventories comove with the sales with rather high contemporaneous correlation coefficients. There does not seem to be meaningful lag with the inventory-sales comovement in these sectors. Combining all results together, it seems that the lagged comovement best describes the input inventories within the durable manufacturing sector but to a less extent for the other types of inventories.

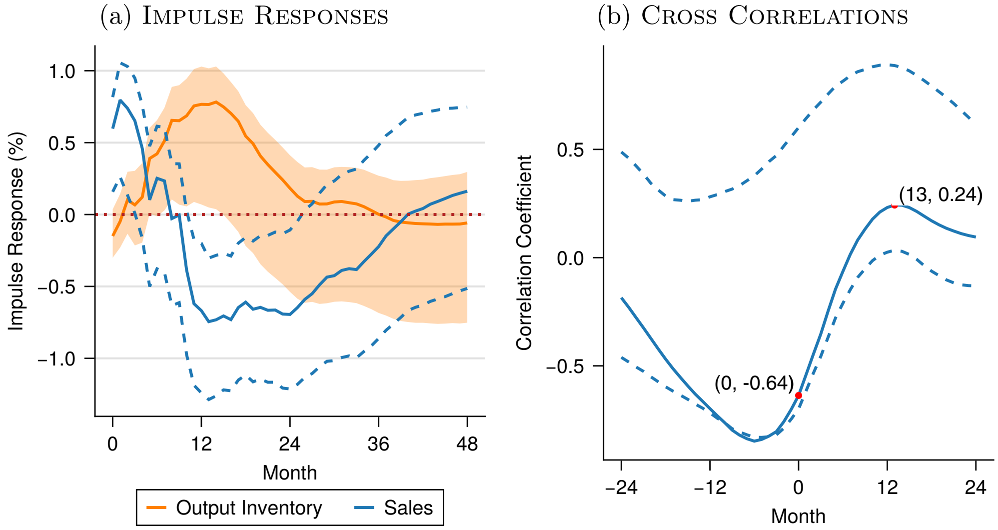

FIGURE 1.17: Output-Inventory-Sales Comovement in the Durable Manufacturing Sector

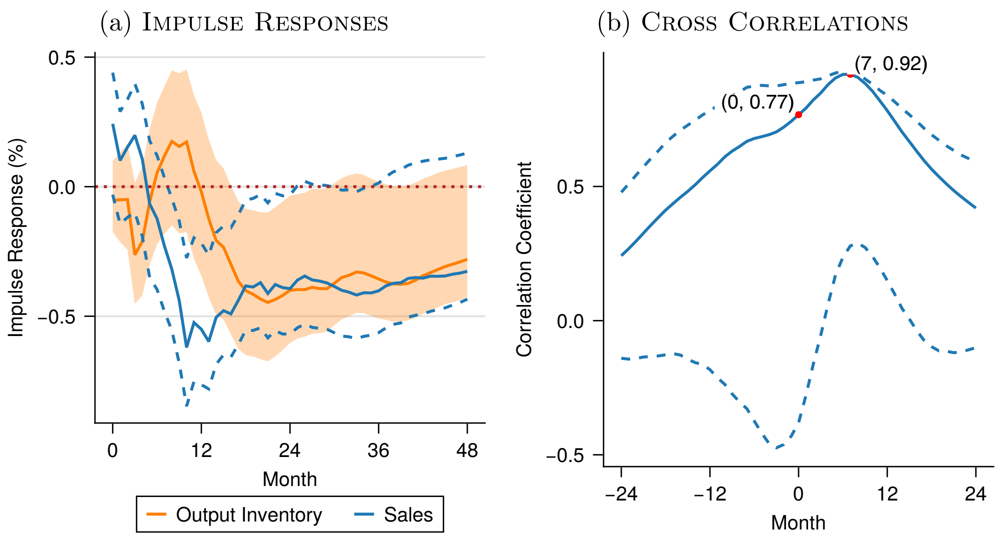

FIGURE 1.18: Output-Inventory-Sales Comovement in the Nondurable Manufacturing Sector

(a) Impulse Responses (b) Cross Correlations

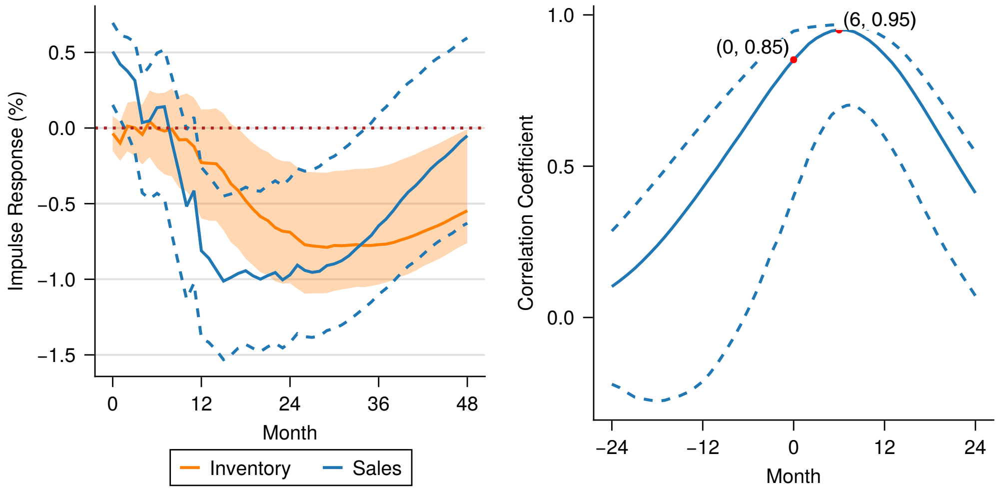

FIGURE 1.19: Inventory-Sales Comovement in the Durable Wholesale Sector

(a) Impulse Responses (b) Cross Correlations

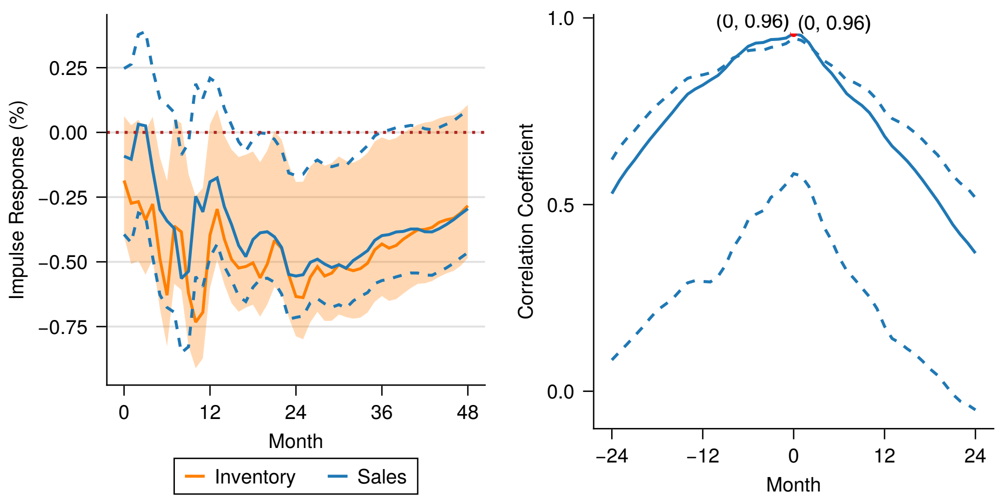

FIGURE 1.20: Inventory-Sales Comovement in the Nondurable Wholesale Sector

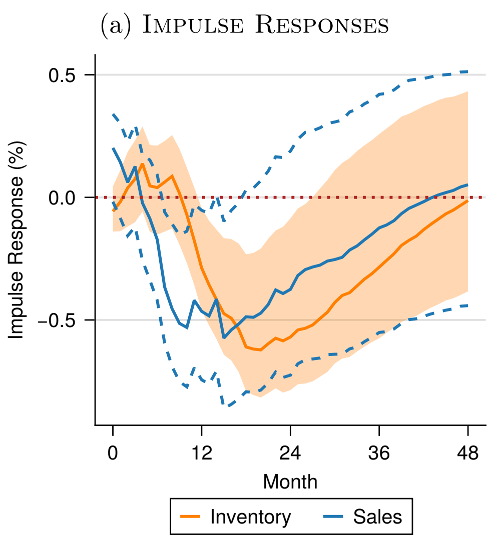

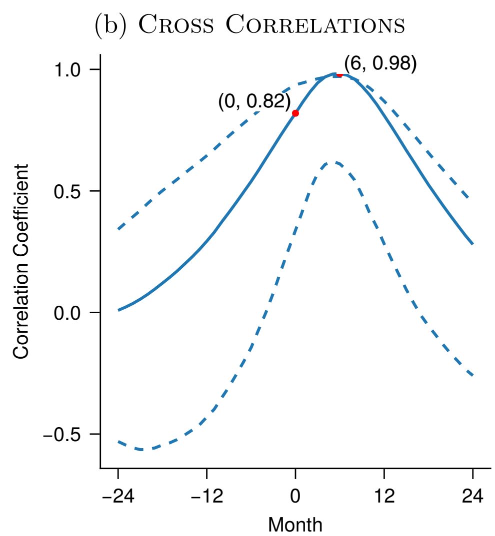

FIGURE 1.21: Inventory-Sales Comovement in the Retail Sector

D. Alternative Trend-Cycle Decompositions

The preferred detrending method is Hamilton (2018). Here, I explore the results when data are detrended with alternative methods. Figures 1.22 and 1.23 plots the cyclical fluctuations obtained with the band-pass filter proposed by Baxter and King (1999) and the Hodrick-Prescott filter. Despite the differences in how cyclical components are defined, we can discern the lagged inventory-sales comovement from all these plots.

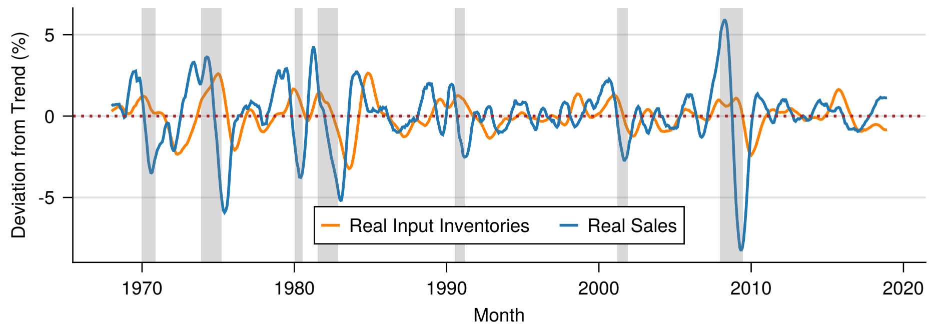

FIGURE 1.22: Cyclical Fluctuations Based on Baxter and King (1999)

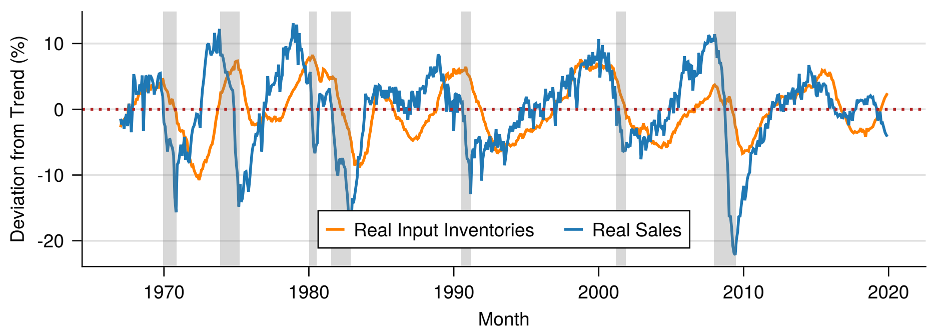

FIGURE 1.23: Cyclical Fluctuations Based on Hodrick-Prescott Filter

E. Results at the 2-Digit SIC Industry Level

I explore the robustness of the empirical results by reconducting the estimation at the 2-digit SIC industry level.[32] The purpose of doing so is to rule out the possibility that the lagged comovement observed is driven by a small subset of the industries with extraordinary properties. Figure 1.24 collects the cross correlations for each of the 11 durable 2-digit SIC industries. For conditional cross-correlations, the analogous estimation based on SVAR is repeated and the results are collected in Figure 1.25. It turns out that, at least at the 2-digit SIC level, the lagged inventory-sales comovement is a common feature across most of the durable industries.

(a) Annual Growth Rate (b) Cyclical Component

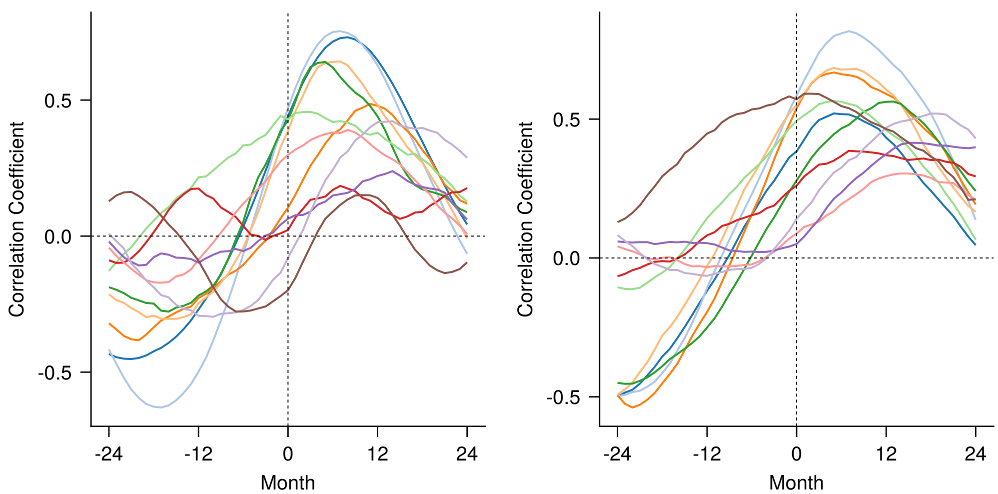

FIGURE 1.24: Unconditional Cross-Correlations Among 2-Digit SIC Industries

[32]For this purpose, the sample only covers 1967:1–1996:12 because data since 1997 are all based on the North American Industry Classification System.

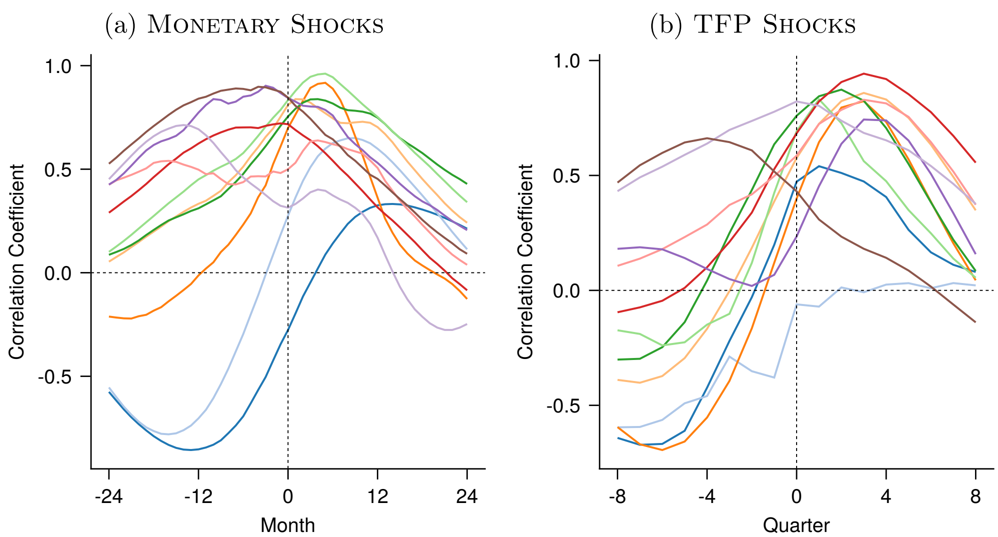

FIGURE 1.25: Conditional Cross-Correlations Among 2-Digit SIC Industries

Appendix 1.B. Details on Existing Inventory Models

This section provides some essential details for the models considered in Section IC.

1. The Precautionary Stockout-Avoidance Motive

Production consists of two stages. In the first stage, the same technology is used to produce a continuum of varieties $\{x_{kt}\}_{k\in[0,1]}$. These varieties enter input inventories after being produced. In the second stage, when idiosyncratic demand shocks are realized, varieties from the input inventories are aggregated into a sectoral good

$$Y_t = A_t \left(\int_0^1 \omega_{kt} s_{kt}^{\xi} \mathrm{d}k \right)^{\frac{1}{\xi}} \qquad \text{with } \Pr(\omega_{kt} > \omega) = \omega^{-\kappa} \text{ for } \omega > 1$$

where the idiosyncratic shocks ω_{kt} to varieties are independent and follow the same Pareto distribution. The law of motion of input inventories for each variety follows

$$h_{kt+1} = (1 - \delta^h) h_{kt} + x_{kt} - s_{kt}$$

where δ^h is the depreciation rate for inventories. Letting $H_t = \int_0^1 h_{kt} \mathrm{d}k$, $X_t = \int_0^1 x_{kt} \mathrm{d}k$ and $S_t = \int_0^1 s_{kt} \mathrm{d}k$, the law of motion for the aggregate quantities can be written as

$$H_{t+1} = (1 - \delta^h) H_t + X_t - S_t.$$

Let $F_j(\omega) \equiv 1 - \omega^{-\kappa_j}$ be the distribution function for the Pareto distribution. Leveraging the results from Wen (2011), the optimality conditions can be expressed in terms of the following functions of a cutoff level of idiosyncratic shock ω_t^*

$$\begin{aligned}
\mathcal{R}(\omega_t^*) &\equiv F_j(\omega_t^*) + \int_{\omega > \omega_t^*} \frac{\omega}{\omega_t^*} \mathrm{d}F_j(\omega) \\
\mathcal{S}(\omega_t^*) &\equiv \int_{\omega \leqslant \omega_t^*} \omega^{\frac{1}{1-\xi}} \mathrm{d}F_j(\omega) + \int_{\omega > \omega_t^*} (\omega_t^*)^{\frac{1}{1-\xi}} \mathrm{d}F_j(\omega) \\
\mathcal{H}(\omega_t^*) &\equiv \int_{\omega \leqslant \omega_t^*} \left[(\omega_t^*)^{\frac{1}{1-\xi}} - \omega^{\frac{1}{1-\xi}} \right] \mathrm{d}F_j(\omega) \\
\mathcal{G}(\omega_t^*) &\equiv \int_{\omega \leqslant \omega_t^*} \omega^{\frac{1}{1-\xi}} \mathrm{d}F_j(\omega) + \int_{\omega > \omega_t^*} \omega (\omega_t^*)^{\frac{\xi}{1-\xi}} \mathrm{d}F_j(\omega).
\end{aligned}$$

The optimality conditions relevant to the partial equilibrium outcomes can be written as

$$\begin{aligned}
\mu_t &= \beta(1+r_t)\mathbb{E}_t \mu_{t+1} && \text{(Consumption Euler Equation)} \\
\iota_t &= \mu_t A_t \mathcal{R}(\omega_t^*) \mathcal{G}(\omega_t^*)^{\frac{1-\xi}{\xi}} && \text{(Shadow Value of Inventories)} \\
A_t S_t &= Y_t \mathcal{S}(\omega_t^*) \mathcal{G}(\omega_t^*)^{-\frac{1}{\xi}} && \text{(Aggregate Demand of Varieties)} \\
X_t &= S_t \left(1 + \frac{\mathcal{H}(\omega_t^*)}{\mathcal{S}(\omega_t^*)} \right) - (1-\delta^h) H_{t-1} && \text{(Inventory Law of Motion)} \\
H_t &= S_t \frac{\mathcal{H}(\omega_t^*)}{\mathcal{S}(\omega_t^*)} && \text{(Optimal Inventory Level)} \\
\frac{\iota_t}{\mu_t} &= \frac{1-\delta^h}{1+r_t} \mathbb{E}_t \frac{\iota_{t+1}}{\mu_{t+1}} \mathcal{R}(\omega_t^*). && \text{(Inventory Euler Equation)}
\end{aligned}$$

2. Multistage Production

This model is based on Sarte, Schwartzman and Lubik (2015), which features production technology that combines output from multiple stages that are produced in different periods. Specifically, suppose there are $S+1$ production stages with output from each stage used as a component for the sectoral good in the sth period after being produced with s being $0, \ldots, S$. Production of components to be used in all of the S future periods and the current period $s = 0$ satisfies

$$X_t \equiv \sum_{s=0}^{S} X_{t|s} = A_t \left(K_t^{\alpha_j} L_t^{1-\alpha_j} \right)^{\theta_j} M_t^{1-\theta_j}$$

where $X_{t|s}$ is the amount of components produced with factors in period t to be used in period $t+s$.[33] The sectoral good is aggregated from the components produced across multiple periods

$$Y_t = \left(\sum_{s=0}^{S} \omega_s X_{t-s|s}^{\xi} \right)^{\frac{1}{\xi}}.$$

The law of motion of aggregate input inventories for a sector can be written as

$$H_{t+1} = H_t + \sum_{s=1}^{S} X_{t|s} - \sum_{s=1}^{S} X_{t-s|s}.$$

The first order conditions relevant to the partial equilibrium outcomes can be written as

$$\mu_t = \beta(1+r_t)\mathbb{E}_t \mu_{t+1} \qquad \text{(Consumption Euler Equation)}$$

$$\iota_t = \mu_t \omega_0 X_{t|0}^{\xi-1} Y_t^{1-\xi} \qquad \text{(Shadow Value of Inventories)}$$

$$\frac{\iota_t}{\mu_t} = \frac{1}{\prod_{s'=0}^{s-1}(1+r_{t+s'})} \mathbb{E}_t \left(\omega_s X_{t|s}^{\xi-1} Y_{t+s}^{1-\xi} \right), \qquad \forall s \in \{1, \ldots, S\}. \qquad \text{(Inventory Euler Equation)}$$

Appendix 1.C. Additional Model Details

This section provides additional model details.

A. The Inventory Problem

Here I consider the inventory problem in Section II.

1. Optimality Conditions

For producer i in stage s with $s > 1$, the first order condition for choosing the optimal order size $X_{it|s-1}^*$ at the beginning of period t can be written as

$$\mathbb{E}_{it} \sum_{l=0}^{\infty} \frac{(1-\lambda)^l}{\prod_{l'=0}^{l}(1+i_{t+l'})} \left[(1+i_{t+l}) P_{jt+l|s-1} - P_{jt+l+1}^{Z} \left(\frac{\upsilon Z_{it+l+1|s}}{(1-\upsilon) X_{it|s-1}^*} \right)^{\frac{1}{\kappa}} \right] = 0 \tag{1.11}$$

[33] Production of different stages uses the same technology and factors may be reallocated across stages without friction. However, once produced, the time when the component is used is fully determined by the production stage s and cannot be altered. Also notice that the notation used in subscripts differs from that used by Sarte, Schwartzman and Lubik (2015) in that the last index after the vertical bar is expressed as relative time instead of calendar time.

where

$$Z_{it+l+1|s} = \left(\frac{(Y_{it+l+1|s}/A_s)^{\frac{\kappa-1}{\kappa}} - \upsilon^{\frac{1}{\kappa}}(X^*_{it|s-1})^{\frac{\kappa-1}{\kappa}}}{(1-\upsilon)^{\frac{1}{\kappa}}} \right)^{\frac{\kappa}{\kappa-1}}.$$

Log-linearizing Equation (1.11) around a steady state with zero inflation yields

$$\frac{1}{\overline{m}} \sum_{l=0}^{\infty} \beta^l (1-\lambda)^l \left[\hat{P}_{jt+l|s-1} - \hat{P}^Z_{jt+l+1} + \hat{i}_{t+l} - \frac{1}{\kappa} \left(\hat{Z}_{it+l+1|s} - \hat{X}^*_{it|s-1} \right) \right] = 0$$

where $\overline{m} \equiv \sum_{l=0}^{\infty} \beta^l (1-\lambda)^l$ is a constant only depending on parameters and

$$\hat{Z}_{it+l+1|s} = \frac{1}{\chi_s} \hat{Y}_{it+l+1|s} - \frac{1-\chi_s}{\chi_s} \hat{X}^*_{it|s-1}$$

with

$$\chi_s \equiv \frac{(1-\upsilon)^{\frac{1}{\kappa}} Z_{*|s}^{\frac{\kappa-1}{\kappa}}}{\upsilon^{\frac{1}{\kappa}} X_{*|s-1}^{\frac{\kappa-1}{\kappa}} + (1-\upsilon)^{\frac{1}{\kappa}} Z_{*|s}^{\frac{\kappa-1}{\kappa}}}.$$

The above equation can be rewritten as

$$\hat{X}^*_{it|s-1} = \frac{1}{\overline{m}} \sum_{l=0}^{\infty} \beta^l (1-\lambda)^l \left[\kappa \chi_s (\hat{P}^Z_{jt+l+1} - \hat{i}_{t+l} - \hat{P}_{jt+l|s-1}) + \mathbb{E}_{it} \hat{Y}_{it+l+1|s} \right].$$

It follows that the aggregate order size for stage-$(s-1)$ output can be written as

$$\hat{X}_{t|s-1} = \lambda \sum_{l \geqslant 0} (1-\lambda)^l \overline{\mathbb{E}}_{it-l} \hat{X}^*_{it-l|s-1}. \tag{1.12}$$

Assuming that all producers have the same inventory-to-sales ratio h_* in the steady state, the aggregate input inventories across all stages satisfy

$$\hat{X}_t = \frac{1}{\overline{h}} \sum_{s=2}^{S} h_*^{S-s} \hat{X}_{t|s}$$

where $\overline{h} \equiv \sum_{s=2}^{S} h_*^{S-s}$.

2. Pricing

The log deviations in aggregate price index for stage-s output is characterized based on two observations. First, since the changes in aggregate sales effectively scale up or down the marginal cost of each individual producer by the same proportion, there is no need to keep track of the changes in sales across individual producers. Second, because of the Calvo-style adjustment, changes in marginal cost can be characterized recursively as follows. Let $\mathcal{M}^*_{tl|s}$ be the average marginal cost at time t among producers who have last adjusted order sizes l periods ago. Because the adjustment shock is independent from other changes, the average marginal cost across all producers can be written as

$$\begin{aligned}\mathcal{M}_{t|s} &= (1-\lambda)\mathcal{M}^{1-\eta}_{t-1|s} + \lambda\mathcal{M}^*_{t1|s} \\ &= \lambda\sum_{l=1}^{\infty}(1-\lambda)^{l-1}\mathcal{M}^*_{tl|s}.\end{aligned} \tag{1.13}$$

It follows from Equation (1.4) that

$$\hat{\mathcal{M}}^*_{tl|s} = \hat{P}^Z_t + \frac{1-\chi_s}{\kappa\chi_s}(\hat{Y}_{t|s} - \overline{\mathbb{E}}_{it-l}\hat{X}^*_{it-l|s-1}). \tag{1.14}$$

With the constant markup, combining Equations (1.12)–(1.14) yields

$$\hat{P}_{t|s} = \hat{P}^Z_t + \frac{1-\chi_s}{\kappa\chi_s}(\hat{Y}_{t|s} - \hat{X}_{t-1|s-1}).$$

B. Additional Model Outcomes under the Partial Equilibrium

As mentioned in Section III, the model responses produced there are based on three paths of changes that are fed into the model simultaneously. Figure 1.26 shows the model outcomes when these three paths are fed into the model separately. The three paths of input inventory responses can be interpreted as a decomposition of the total responses when only a single source of changes are allowed. As expected, the rise in interest rate discourages the use of input inventories because of the higher price of purchasing in advance. For the model to generate outcomes that are close to the empirical target, the effects from the changes in interest rate and prices need to be relatively small.

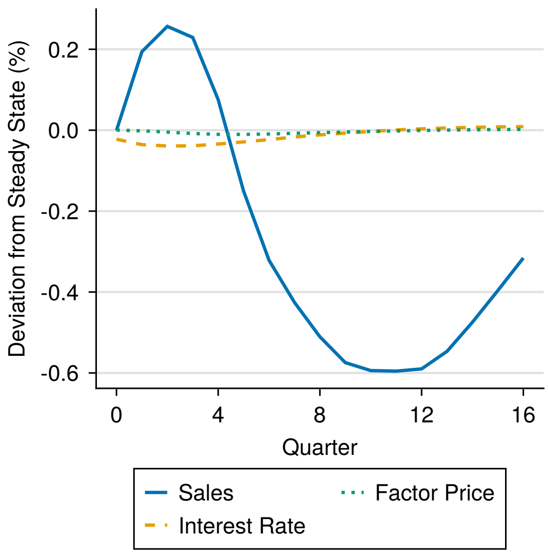

FIGURE 1.26: Input Inventory Responses by Source of Changes

Chapter 2. Aggregate Implications of the Lagged Inventory-Sales Comovement

What does inventory behavior tell us on the propagation of aggregate shocks? In the previous chapter, I have documented the lagged inventory-sales comovement. Namely, for durable input inventories, the movement of inventories lag sales movement by about three quarters. In this chapter, I argue that capturing this lead-lag relation is important because it provides crucial information on the timing of how the impact of an aggregate shock propagates across the economy and supports a new interpretation on how inventories participate in this process.

To that end, I embed the production model developed in the previous chapter, featuring frictional production activities along a serial supply chain, in a durable manufacturing sector of a New Keynesian framework. In this model, the decline in durable sales induced by a monetary policy shock is transmitted from the most downstream producers to their suppliers with delay due to frictions, generating lagged inventory-sales comovement as observed in data. General equilibrium predictions suggest that the delayed transmission of aggregate shock along the supply chain results in a timing difference in how the changes in interest rates are translated into aggregate GDP changes. Relative to a counterfactual scenario in which the movements of inventories and sales are fully synchronized, the baseline outcomes with lagged inventory-sales comovement demonstrate smaller aggregate GDP decline in the first year, and more gradual decline toward a trough, but slower recovery after that. Such differences are largely due to the information frictions faced by producers placing orders from suppliers that effectively lower the sensitivity of investment and utilization of capital and durable consumption goods with respect to changes in real interest rates. This is a new insight into how inventory behavior can affect aggregate outcomes that would not be recognized with only contemporaneous correlations.

The production problem is embedded into a fully-fledged multisector New Keynesian framework. The durable manufacturing sector is a special sector in which production happens along a supply chain consisting of three stages. Production in the other sectors is simply characterized by a Cobb-Douglas production function that converts inputs to outputs within the same period. The GE framework features linkages across six production sectors via the use of intermediate inputs (the input-output network) and the production of sector-specific investment goods (the investment network) as in Horvath (2000). The need for at least two production sectors is obvious as the durable input inventories are only relevant to durable goods industries but not the entire economy. The inclusion of the linkages across six sectors is motivated by the special role of the durable manufacturing sector in a production network as a hub for investment goods (vom Lehn and Winberry 2021). On top of this detailed multisector framework, I introduce real and nominal frictions following conventional wisdom as in Smets and Wouters (2007). Specifically, I allow external habit formation for the final consumption bundle, investment adjustment costs and variable capital utilization along with the sticky sectoral prices and sticky wages. This results in a GE framework that is flexible enough to generate hump-shaped impulse responses to aggregate shocks and allows directly matching model impulse responses with empirical counterparts.

Following Christiano, Eichenbaum and Evans (2005), the GE model is parameterized in two steps. First, I externally calibrate the majority of parameters for preferences and sector-specific production technology based on the input-output tables, investment flow data and convention. Second, I estimate the remaining parameters, which are all relevant to the dynamic behavior of the model outcomes, by solving a nonlinear least squares problem that minimizes the distance between the model impulse responses to monetary shocks and their empirical counterparts obtained from structural vector autoregressions.

With the calibrated GE model that incorporates the supply chain production problem for the durable manufacturing sector, I examine how the model outcomes would be different when the parameters affecting the inventory behavior are altered. Starting from the baseline specification with parameters that intend to capture the lagged inventory-sales comovement, I raise the relative precision of the signals on aggregate shocks and weaken the Calvo-style adjustment friction to synchronize the inventory and sales movements in model. Holding the path of real interest rates unchanged, I find that the aggregate GDP responses become substantially stronger in the first year and reach a trough earlier under this counterfactual scenario. I interpret this different result as suggesting that supply chain frictions can have a substantial impact on the transmission of monetary shocks.

What explains the different model outcomes? Following Auclert, Rognlie and Straub (2020), I conduct a general equilibrium decomposition of the GDP responses following the monetary shock.

In the GE model, the total effect of a monetary shock on an aggregate variable can be decomposed into two components: the intertemporal substitution effect and the user cost effect.[1] The former is attained via the impact of real interest rate changes on the consumption Euler equation. The latter is attained via the no-arbitrage condition for holding sectoral capital and durable consumption goods. The decomposition result shows that the different impact on aggregate GDP, especially over the short horizons, is mostly due to the smaller user cost effect under the baseline scenario. Capturing the durable input inventory behavior results in lower sensitivity of investment and utilization of capital and durable consumption goods with respect to changes in real interest rates. Intuitively, the aggregate shock's impact on sales affects the sales of suppliers with delay under the baseline scenario with lagged inventory-sales comovement. The gradual transmission of sales fluctuations upward along the supply chain in turn results in gradual adjustment of input factors among suppliers. In contrast, under the counterfactual scenario, changes in downstream sales immediately affect the sales of upstream suppliers and hence the factor demand among all producers simultaneously. This distinction on how the impact of the aggregate shock is transmitted upward along the supply chain alters the strength of the user cost effect. This is a novel model implication that has not been considered in prior works.

Related to the above finding, I additionally show that it is the information friction on aggregate sales that contributes to an intertemporal shift of the GDP impact, which results in the different timing for GDP to reach a trough. The Calvo-style adjustment affects the strength of the GDP responses over the later horizons, but does not mechanically alter the timing of the aggregate outcomes. Furthermore, counterfactual experiments suggest that increasing the average upstreamness of the supply chain members further strengthens the intertemporal shift of the GDP impact, due to the lowered rate at which changes in downstream sales spread across the producers.

This paper is related to the studies aiming at understanding the aggregate implications of inventory behavior. There are mainly two types of papers falling in this category. The first type of papers are mainly concerned with how inventory behavior has affected output volatility. This includes Khan and Thomas (2007), Wen (2011), Iacoviello, Schiantarelli and Schuh (2011) and Wang, Wen and Xu (2014). More recently, Alessandria et al. (2023) explore the aggregate impact

[1]Strictly speaking, there is also an effect due to changes in inflation rates. However, since this effect on GDP is small in the calibrated model, it is ignored.

of shipping delays in a setting where retailers running out of stock raise prices. The second type of papers are not concerned with inventory behavior per se but use inventory behavior to infer some other aspects of aggregate fluctuations. This includes Bils and Kahn (2000) and Kryvtsov and Midrigan (2012) who assess the relative importance of countercyclical markup and nominal cost rigidities based on the behavior of output inventories. My paper adds to this literature by demonstrating a novel implication that durable input inventory behavior dampens the real effect of monetary policy in the very short run by reducing the user cost sensitivity of investment decisions with respect to real interest rate changes. Additionally, all the papers mentioned above only consider the simplest form of supply chain consisting of only a single pair of supplier and customer (the dyadic structure). An exception is Ferrari (2023), who considers an industry-level production network with producers in each industry holding inventories. However, Ferrari (2023) abstracts away from the frictions involved in ordering intermediate inputs as I do and solely uses the procyclical inventory behavior to justify the amplification of final demand shocks along the supply chain.

Additionally, this paper is broadly related to the literature that investigates the relative importance of various frictions based on estimated dynamic stochastic general equilibrium (DSGE) models. Two well-known papers are Christiano, Eichenbaum and Evans (2005) and Smets and Wouters (2007). Relative to these papers, I focus on the relative importance of the durable input inventories and estimate the parameters for other model components only for obtaining realistic quantitative predictions. On the methodological front, I leverage the advances made by Auclert et al. (2021) and Auclert, Rognlie and Straub (2020) on the use of sequence-space characterization of model outcomes.

The paper proceeds as follows. Section I describes the general equilibrium framework for quantitative analysis. Section II considers the aggregate implications of the inventory behavior by comparing baseline model outcomes with counterfactuals. Section III conducts decompositions that help understand the counterfactual outcomes. Section IV discusses how the results depend on the distribution of producers along the supply chains. Section V concludes.

I. A Multisector New Keynesian Framework

In this section, I embed the supply chain production problem considered in Chapter 1 into a multisector New Keynesian framework. This general equilibrium environment can be viewed as an augmented version of the well-known one-sector framework considered in Smets and Wouters (2007) that combines a rich set of conventional ingredients for real and nominal frictions. Since inventories are not directly relevant to each sector, I consider six production sectors in the model: mining, construction, durable manufacturing, nondurable manufacturing, real estates and other services. Production in these sectors are interconnected via the use of intermediate inputs and an investment network as in Horvath (2000), with modifications made for introducing nominal rigidities. Only the durable manufacturing sector holds input inventories along the supply chain for critical inputs.

A. Households

A representative household consumes a consumption bundle and supplies labor subject to constraints. Preferences of the household can be expressed as[2]

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\log(C_t - \lambda_c C_{t-1}) - \frac{N_t^{1+\sigma_l}}{1+\sigma_l} \right]$$

where the utility flow from the individual consumption C_t is affected by the external habit formed from the lagged aggregate consumption with a strength parameter λ_c; σ_l is the reciprocal of the Frisch elasticity of labor supply. Following Erceg, Henderson and Levin (2000), the household supplies labor of amount N_t to a continuum of labor unions that differentiate the labor supply to a continuum of varieties. The labor unions have price-setting power under monopolistic competition. All varieties of labor are aggregated into a bundle using a Kimball (1995) aggregator that results in aggregate labor supply L_t for producers in the production sectors. The aggregate labor supply is

[2]The utility function can be viewed as a limiting case of the one considered in Smets and Wouters (2007)

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{(C_t - \lambda_c C_{t-1})^{1-\sigma_c}}{1-\sigma_c} \exp\left(\frac{\sigma_c - 1}{1+\sigma_l} N_t^{1+\sigma_l} \right)$$

with the intertemporal elasticity of substitution σ_c goes to one. Indices for individual households are omitted as the choices made by all households are identical.

allocated across N production sectors without friction such that

$$L_t = \sum_{i=1}^{N} L_{it}.$$

In addition to consumption and labor supply choices, the household holds a risk-free bond, and determines the investment and utilization of a durable consumption good and capital for each production sector. The budget constraint in period t can be written as

$$\begin{aligned} & P_t C_t + B_{t+1} + P_t^D I_t^D + \sum_{i=1}^{N} P_{it}^I I_{it} = \\ & \qquad W_t^h N_t + (1+i_t) B_t + [R_t^D Z_t^D - a_D(Z_t^D)] D_{t-1} + \sum_{i=1}^{N} [R_{it}^K Z_{it} - a(Z_{it})] K_{it-1} + \Pi_t \end{aligned}$$

where P_t is the price index of final consumption bundle; B_t is the nominal face value of the risk-free bond; P_t^D is the price index of the durable consumption good; I_t^D is the purchase of new durable consumption good; P_{it}^I is the price index of investment good for capital used in sector i; I_{it} is the investment in sector-i capital; W_t^h is the wage rate faced by the households; i_t is the risk-free nominal interest rate. Profits earned by producers in any activity are distributed to the households as dividends denoted as Π_t.

The investment and utilization of durable consumption good and capital goods are modeled in the same fashion.[3] With variable utilization level Z_t^D for durable consumption good and Z_{it} for sector-i capital, the rental income for these stocks depends on both the rental rates R_t^D and $\{R_{it}^K\}_i$ and the cost for setting the utilization levels that is captured by a function $a_D(\cdot)$ for durable consumption good and $a(\cdot)$ for all capital goods. The law of motion for the sector-specific capital can be written as

$$K_{it+1} = (1-\delta_i) K_{it} + \left[1 - S\left(\frac{I_{it}}{I_{it-1}}\right)\right] I_{it}$$

where δ_i is the depreciation rate for sector-i capital and $S(\cdot)$ is a function that captures investment

[3] The durable consumption good is effectively a capital used for producing the final consumption bundle. One can think of the final consumption bundle as the output from a hypothetical production sector that combines goods from other sectors without requiring labor inputs.

adjustment cost. The law of motion for the stock of durable consumption good is similarly written as

$$D_{t+1} = (1-\delta^D)D_t + \left[1 - S_D\left(\frac{I_t^D}{I_{t-1}^D}\right)\right] I_t^D.$$

The final consumption bundle is produced by combining sector-specific consumption goods from different sectors and the stock of durable consumption good

$$C_t = \left(Z_t^D D_t\right)^{\vartheta} \left(\prod_{i=1}^{n} C_{it}^{\eta_i}\right)^{1-\vartheta} \quad \text{with} \quad \sum_{i=1}^{N} \eta_i = 1. \tag{2.1}$$

Final use of output from the durable manufacturing sector directly contributes to the production of final consumption bundle only via the investment in the stock of durable consumption good D_t and is never directly consumed as C_{it}.[4]

B. Production

There are N production sectors in the economy. The durable manufacturing sector is populated by producers in S stages along a serial supply chain as described in Section II of Chapter 1. All the other sectors are modeled in the same fashion as described below without the supply chain structure.

Each production sector except the durable manufacturing sector is populated by a continuum of producers with the same sector-specific technology for differentiated varieties. For sector j (except the durable sector), the production technology is Cobb-Douglas and sectoral gross output satisfies

$$Y_{jt} = A_{jt}[(K_{jt}^s)^{\alpha_j}(L_{jt})^{1-\alpha_j} - \phi_p]^{\theta_j} M_{jt}^{1-\theta_j}$$

where Y_{jt} is gross output; A_{jt} is total factor productivity; K_{jt}^s is effective capital input; L_{jt} is labor input; M_{jt} is a bundle of intermediate inputs; ϕ_p is a fixed cost that lowers the effective value added component in production.[5] The parameters α_j and θ_j are between 0 and 1. In each period, the

[4] This is saying that η_i is zero for the durable manufacturing sector. Notice that the output from the durable manufacturing sector may also serve as intermediate inputs and ingredients of investment goods for the production of other sectors that affects the final consumption bundle indirectly. The investment in the stock of durable consumption good is one of the uses of the durable output that directly enters the production of final consumption bundle.

[5] The presence of ϕ_p effectively introduces a "wedge" into the first-order conditions for factor demand. One may think of ϕ_p as a flow of operational overhead that does not contribute to the output. Such fixed costs are also present in Christiano, Eichenbaum and Evans (2005) and Smets and Wouters (2007).

(variable) nominal cost incurred by the producers consists of three components and can be written as

$$R_{jt}^{K}K_{jt}^{s}+W_{t}L_{jt}+P_{jt}^{M}M_{jt}$$

where R_{jt}^{K} and W_t are nominal factor prices; P_{jt}^{M} is the price index for the intermediate inputs purchased in period t.

To introduce conventional nominal frictions into the model, the varieties in each sector, including those from the last stage in the durable manufacturing sector, are aggregated into a sectoral good by a representative sectoral good producer using a Kimball (1995) aggregator. The sectoral good from sector i is then traded on a competitive market for three types of usage: consumption by households denoted as C_{it}, intermediate use by variety producers in sector j denoted as M_{ijt} and production of investment goods for sector j denoted as I_{ijt}. For consumption use C_{it}, sectoral goods are aggregated into the final consumption bundle as described in Equation (2.1). For the other two types of usage, sectoral goods are aggregated using a Cobb-Douglas technology such that

$$M_{jt}=\prod_{i=1}^{N}M_{ijt}^{\gamma_{ij}},\quad I_{jt}=\prod_{i=1}^{N}I_{ijt}^{\lambda_{ij}}\quad \text{with}\quad \sum_{i=1}^{N}\gamma_{ij}=1,\quad \sum_{i=1}^{N}\lambda_{ij}=1.$$

C. The Rest of the Economy

The remaining components of the model economy are all standard in the literature, and hence only concise summaries are provided for the sake of completeness.

Nominal Rigidities Facing demand curves from the sectoral producers, variety producers in each sector set the optimal nominal prices for their output only periodically as in Calvo (1983*a*).[6] Labor unions set the wage rates for their differentiated labor in a similar staggered fashion. The resulting Phillips curves follow Smets and Wouters (2007), except that the relevant output prices are at the sectoral level because of the multisector environment.[7]

[6] For the durable manufacturing sector, the sluggish price setting can happen either among the producers in the last stage or producers in all stages. This does not seem to be crucial for the later quantitative analysis.

[7] It is possible to allow partial inflation indexation. However, whether there is partial inflation indexation or not does not seem to make a substantial difference on the outcomes and hence it is avoided.

Monetary Policy The central bank conducts monetary policy following an inertial Taylor rule for the risk-free nominal interest rate in response to changes in inflation, value added and value added growth that can be written as

$$1 + i_t = (1 + r_*)^{1-\rho_m}(1 + i_{t-1})^{\rho_m} \left[\left(\frac{P_t}{P_{t-1}} \right)^{\psi_1} \left(\frac{VA_t}{VA_*} \right)^{\psi_2} \right]^{1-\rho_m} \left(\frac{VA_t}{VA_{t-1}} \right)^{\psi_3} (1 + \varepsilon_t^i)$$

where r_* is the steady-state real interest rate; ρ_m is the persistence of the policy rate; VA_t is the aggregate value added; VA_* is the steady-state aggregate value added; ε_t^i is the monetary policy shock.

Equilibrium Given a path of monetary shocks $\{\varepsilon_t^i\}_t$, initial stocks B_{-1}, D_{-1}, $\{K_{i,-1}\}_i$, $\{X_{-1|s}\}_s$ initial prices $\{P_{i,-1}\}_i$, W_{-1}, and the Taylor rule for monetary policy, a competitive equilibrium consists of paths of prices $\{P_t,\ \{P_{it}\}_i,\ \{P_{it}^M\}_i,\ \{P_{it}^I\}_i,\ \{P_{t|s}\}_s,\ P_t^D,\ W_t,\ R_t^D,\ \{R_{it}^K\}_i,\ i_t,\ r_t,\ \{\pi_{it}\}_i,\ \pi_t^w\}_{t\geqslant 0}$ and aggregate quantities $\{Y_t,\ \{Y_{it}\}_i,\ VA_t,\ \{VA_{it}\}_i,\ C_t,\ \{C_{it}\}_i,\ B_t,\ D_t,\ Z_t^D,\ \{K_{it}\}_i,\ \{Z_{it}\}_i,\ \{M_{it}\}_i,\ \{N_t\},\ \{L_{it}\}_i,\ L_t,\ \{Y_{t|s}\}_s,\ \{\overline{\mathbb{E}}_{it}Y_{t|s}\}_s,\ \{X_{t|s}\}_s,\ \{Z_{t|s}\}_s,\ \Pi_t\}_{t\geqslant 0}$ such that households maximize their present value of utility given constraints; producers optimize their factor usage; a Phillips curve holds in each production sector and the labor market; a Fisher equation for the risk-free interest rate holds; labor market clears; and each sectoral goods market clears.

II. Aggregate Implications

In this section, I explore the aggregate implications of the lagged inventory-sales comovement based on the full general equilibrium model outlined in Section I.

A. Calibration and Estimation

The model parameters are divided into two groups. For the first group of parameters, their values are externally calibrated either based on sample moments or convention. For the second group, the values are chosen by minimizing the Euclidean distance between the model impulse responses and their empirical counterparts. The methodology follows Christiano, Eichenbaum and Evans (2005), except that for parameters directly involved in the production in durable manufacturing sector, I use the same parameter values obtained in Section IIIA of Chapter 1.

1. Calibrated Parameters

For some parameters, their values are directly chosen based on typical choices in prior work. They are collected in Table 2.1.

TABLE 2.1: Parameters Externally Calibrated By Convention

Discount rate (β)	0.99
Frisch elasticity of labor supply ($1/\sigma_l$)	0.40
Steady-state wage markup	1.50
Curvature parameter for Kimball aggregator	10.00

The technological parameters are determined based on BEA data. The parameters for the Cobb-Douglas production technology in each sector are calibrated based on the BEA input-output tables. From the annual input-output tables over 1967–2006, I aggregate the private industries to a level involving only the six sectors and compute the average shares of goods used as intermediate inputs from each source sector in each user sector. To determine the share of labor income in the value added components, I aggregate industry-level data from vom Lehn and Winberry (2021), which are also constructed from the BEA data. These data allow me to determine the values of the Cobb-Douglas technology parameters $\{\alpha_j, \theta_j\}_j$ and $\{\gamma_{ij}\}_{ij}$. From the investment network data and the industry-level capital depreciation rates constructed by vom Lehn and Winberry (2021) based on the BEA data, I compute the average expenditure shares of goods from each sector that are used as investment goods in each user sector and sector-specific capital depreciation rates at the level of the six sectors. This gives me the values of $\{\lambda_{ij}\}_{ij}$ and $\{\delta_j\}_j$, which govern the steady-state capital investment. From the information on final use for consumption in the input-output tables, I obtain average expenditure shares across consumption goods produced in different sectors. Combined with some other parameter values and a steady-state characterization of Euler equation, it is possible to back out parameter values for the final consumption bundle ϑ and $\{\eta_i\}_i$. Lastly, I assume that for all variety producers, the steady-state share of the fixed cost in their effective value added component is 0.6.[8]

[8]The share of 0.6 is the posterior estimate obtained by Smets and Wouters (2007).

2. Estimated Parameters

The remaining parameters are estimated by minimizing the Euclidean distance between the model impulse responses and their empirical counterparts. For the empirical counterparts, I interpret the Romer-Romer shocks as innovations to the monetary policy shock and obtain impulse response estimates via SVARs with the Romer-Romer shocks ordered the first in the same way as described in Section I of Chapter 1. I consider 10 outcome variables from aggregate time series data: nominal interest rate, GDP, nondurable consumption, durable consumption, investment, durable input inventory, durable sales, hours, wage and consumption price indices.[9] These 10 variables correspond to the model objects i_t, VA_t, C_{4t}, I_t^D, I_t, X_t, $Y_{t|S}$, L_t, W_t and P_t respectively.[10]

The nonlinear least squares problem for selecting the parameter values can be written as

$$\min_{\Theta}[\ J(\Theta) - \hat{J}\ ]^{\mathsf{T}} W^{-1} [\ J(\Theta) - \hat{J}\ ]$$

with $J(\Theta)$ being the model outcomes associated with a parameter vector Θ and $\hat{J}$ being the empirical counterpart used as targets. W is a diagonal weight matrix with weights proportional to the length of the bootstrap confidence intervals obtained for the impulse response estimates.[11] The estimates are additionally subject to a set of constraints that ensure them to fall in a range of reasonable values. Standard errors of the estimates are obtained from the asymptotic variance-covariance matrix

$$\hat{\mathbf{V}} = \left[\left(\frac{\partial J}{\partial \Theta}(\hat{\Theta}) \right)' W^{-1} \frac{\partial J}{\partial \Theta}(\hat{\Theta}) \right]^{-1}$$

where $\frac{\partial J}{\partial \Theta}(\hat{\Theta})$ is a Jacobian matrix of J with respect to Θ evaluated at the vector of point estimates.

Table 2.2 collects the estimated parameter values. Figure 2.1 compares the model impulse responses under these parameter values with the empirical targets. The model impulse responses match their empirical counterparts reasonably well, although there is some difficulty in capturing the somewhat delayed trough of the investment responses. Regarding the parameter values, some

[9] Data sources and additional details for these variables are in Appendix 2.A.

[10] The subscript 4 in C_{4t} refers to the nondurable manufacturing sector.

[11] These weights are proportional to the estimates for standard error rather than the variance, which is different from the popular practice. This change is made in order to assign more weights on impulse response estimates from the median horizons where the the confidence intervals are wider but are more useful to distinguish the dynamic behavior of different variables.

TABLE 2.2: Estimated Parameter Values

Parameter	Estimate	Std. Error
Household habit (λ_c)	0.97	
Calvo price stickiness	0.70	
Calvo wage stickiness	0.86	0.002
Capital utilization adjustment cost elasticity ($a''(1)/a'(1)$)	0.25	
Durable goods utilization adjustment cost elasticity ($a''_D(1)/a'_D(1)$)	0.33	0.032
Capital investment adjustment cost sensitivity (S''_*)	20.00	
Durable goods investment adjustment cost sensitivity (S''_{D*})	13.11	0.257
Taylor rule inertia (ρ_m)	0.67	0.025
Taylor rule coefficient on inflation (ψ_1)	1.81	0.274
Taylor rule coefficient on value added (ψ_2)	0.02	
Taylor rule coefficient on value added growth (ψ_3)	0.40	0.230

Notes: A parameter value estimate without an accompanying standard error estimate is constrained by an imposed boundary condition instead of an interior solution from the nonlinear least squares problem.

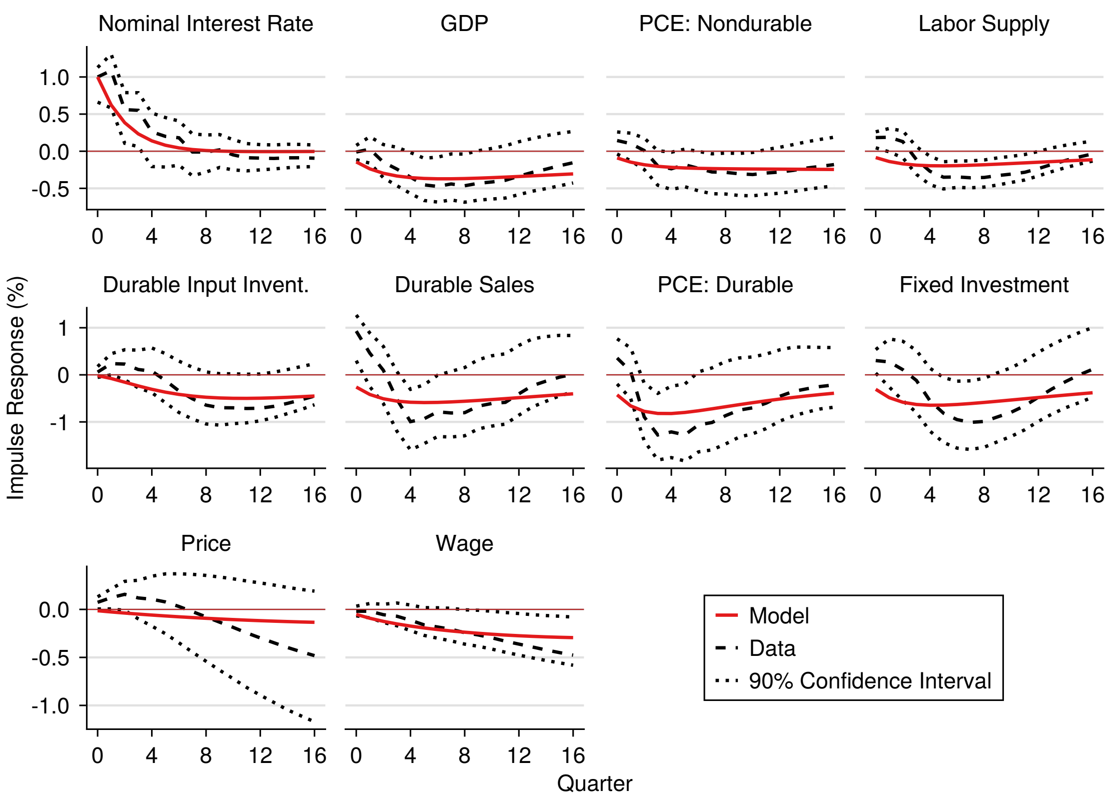

FIGURE 2.1: Model Impulse Responses vs Empirical Counterparts

Note: Red solid lines represent model impulse responses under estimated parameter values. Black dashed lines represent impulse response estimates obtained from SVARs. The dotted lines show the 90% bootstrap percentile intervals. The magnitude of shocks are normalized so that the response of nominal interest rate is 1% on impact.

of them have to be restricted by externally imposed constraints for the solver to reach a more reasonable solution.[12] The model requires rather strong habit for consumption for fitting the empirical counterparts. The solver also tends to select relatively small value of $\sigma_a \equiv a''(1)/a'(1)$, implying a low cost for adjusting capital utilization. This is accompanied by the tendency to select rather large S''_*, which implies substantial investment adjustment cost. Combined together, the solver attempts to fit the empirical counterpart by restricting the investment in capital in favor of adjusting the capital utilization. Such choices are conceivably due to the negligible investment responses on impact.[13] With the imposed restrictions on these two parameters, the model generates responses in investment of more reasonable magnitude, which would have been too small in the absence of the restrictions.[14] Regarding the strength of nominal rigidity, an upper bound of 0.7 is imposed for the probability of keeping the same sectoral good price.[15]

B. Aggregate Impact of Monetary Shocks

Does the lag in comovement between durable input inventories and durable sales matter for aggregate outcomes? With the estimated model, it is possible to answer this question by conducting counterfactual exercises in which certain model parameters are altered. To this end, it is useful to first recognize that whenever a parameter value is changed, there are two reasons for the model outcomes following the same innovation to the monetary policy shock ε_t^i to differ. First, the path of real interest rates can be different in the general equilibrium due to the different endogenous responses following the Taylor rule.[16] Second, holding the path of real interest rates unchanged (by letting the path of monetary shocks accommodate), the general equilibrium outcomes differ because of the changed properties of the model under alternative parameters. For the sake of understanding

[12] In some sense, such parameters are calibrated because there is no interior solution and the constraints are externally specified.

[13] A possibly related issue has also been encountered by Christiano, Eichenbaum and Evans (2005). They restrict the value of σ_a to be 0.01 to prevent the solver from searching an even smaller value.

[14] Nonetheless, the model responses in investment still reach a trough of a smaller absolute magnitude than what the data suggest. Model implications related to the magnitude of investment responses may have been understated because of this.

[15] This value is slightly larger than what Smets and Wouters (2007) and Christiano, Eichenbaum and Evans (2005) use in their benchmark models. However, it is smaller to what Auclert, Rognlie and Straub (2020) obtain (0.926) for their model.

[16] Although monetary shocks directly affect nominal interest rates, it is more useful to focus on changes in real interest rates because their impact on aggregate outcomes mainly takes effect via changes in real interest rates under the estimated model. Their impact via changes in inflation expectation and their direct impact on the ordering of the critical inputs are small.

the role of inventory behavior, it seems useful to hold the path of real interest rates unchanged and hence focus on the second way for aggregate outcomes to differ. I therefore conduct counterfactual experiments as follows. From the model with baseline parameter values, I obtain the real interest rate path following a one-time monetary shock that raises the nominal interest rate by 1%. I then alter parameter values and obtain model impulse responses when the same path of real interest rates is fed into the model.

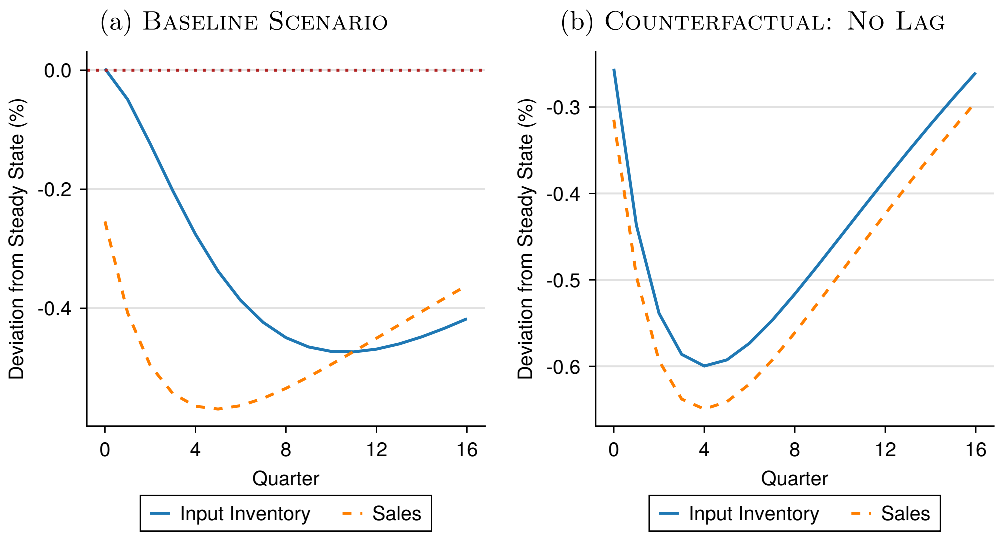

FIGURE 2.2: Inventory-Sales Comovement with and without the Lag

A specific counterfactual experiment of interest is to explore how aggregate outcomes vary when there is no lag in the inventory-sales comovement. That is, the cross-correlation between input inventory and sales attains the highest value when considering their contemporaneous deviations from the steady-state levels. For this purpose, I alter two parameters considered in Section III of Chapter 1, the precision of the signals received by the producers relative to the precision of prior τ_v/τ_ε and the probability of adjusting order size in a quarter λ. As shown in Figure 2.2b, with τ_v/τ_ε being 5 and λ being 0.6, the input inventories and sales move in the same direction throughout the horizons. In the general equilibrium under the same real interest rate path, the responses in input inventory and sales additionally demonstrate larger magnitude when they reach a trough, relative to the baseline scenario shown in Figure 2.2a. Such differences are expected as the responses encapsulate all GE forces with sales determined endogenously.

With the above changes to the two parameters governing inventory behavior in the durable

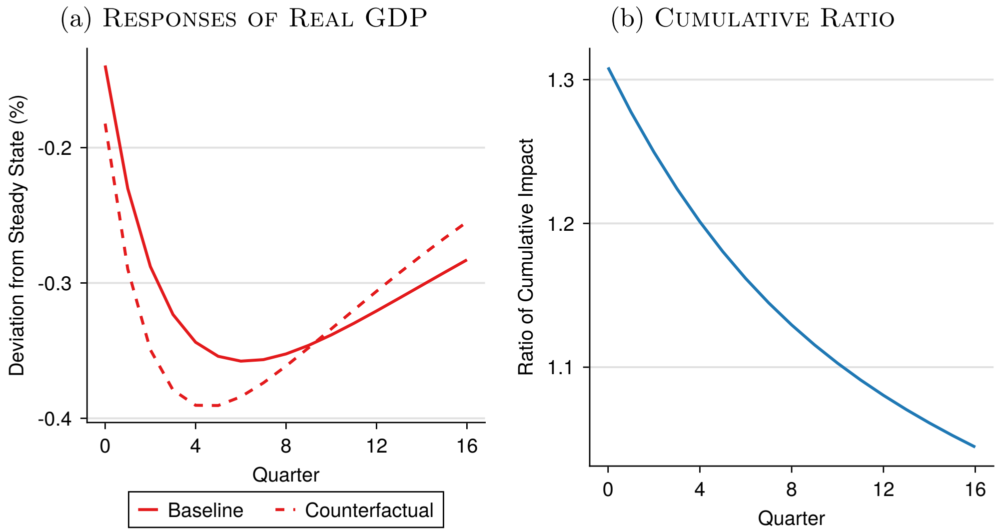

FIGURE 2.3: Impulse Responses of Real GDP: Baseline vs Counterfactual

manufacturing sector, Figure 2.3a shows the impulse responses of real GDP under the baseline and counterfactual scenarios while holding the path of real interest rates unchanged. Although the two parameters affected are only directly relevant to a single sector of the economy, their changes result in substantial impact on the aggregate real GDP of the entire economy. In particular, the responses are stronger in the first year and reach a trough earlier under the counterfactual scenario. After reaching the trough, the responses also revert back at a higher rate, resulting in GDP changes that are less persistent than the baseline scenario. Figure 2.3b highlights the differences by showing the ratios of the cumulative impact on real GDP over different horizons. It is clear that under the counterfactual scenario, the effects on real GDP are more front loaded, with cumulative effects being around 25% larger in the first year and the gap gradually declining towards zero by the end of the fourth year.

III. Decompositions

In this section, I conduct decompositions that help understand the results from the counterfactual experiment in Section II.

A. GE Decompositions by the Effects of Real Interest Rates

To gain further insights on how the model predictions are altered when the lag in inventory-sales comovement is taken out, I conduct GE decompositions following Auclert, Rognlie and Straub (2020). Specifically, changes in real interest rates demonstrate their impact on aggregate outcomes by perturbing two types of equilibrium conditions in the model: the Euler equation that governs the intertemporal substitution of consumption and the no-arbitrage conditions that govern the investment and utilization of durable consumption goods and sectoral capital. In the linearized model, the former can be written as

$$(1+\lambda_c)\hat{C}_t = \lambda_c\hat{C}_{t-1} + \mathbb{E}_t\hat{C}_{t+1} - (1-\lambda_c)\hat{r}_t$$

while for sector-j capital, the latter can be written as

$$\hat{r}_t = \beta(1-\delta_j)\mathbb{E}_t\hat{Q}_{jt+1} + [1-\beta(1-\delta_j)]\mathbb{E}_t\hat{R}^K_{jt+1} - \hat{Q}_{jt}$$

with $\hat{Q}_{jt}$ being the log deviations in Tobin's Q.[17] For these two sets of equilibrium conditions, I separately feed the real interest rate path into one of them while leaving the other unaffected. The resulting two sets of GE outcomes represent the intertemporal substitution effect of real interest rate changes and the user cost effect of real interest rate changes under GE. I repeat this process under the baseline and counterfactual scenarios with the same path of real interest rates obtained under the baseline scenario to examine the changes in aggregate outcomes solely driven by the changes in model properties.[18]

Figure 2.4 shows the GE decompositions of GDP responses for the aggregate economy and the durable manufacturing sector. The user cost effect of real interest rate changes takes the largest share of the total effect in both cases. More importantly, it explains most of the differential outcome

[17] In the model, Tobin's Q is the ratio between the shadow value of sector-j capital and the marginal value of relaxing the household budget constraint. The condition for durable consumption goods is similar.

[18] For the baseline scenario, summing up the outcomes obtained from both effects yields the total effect that is virtually equal to what would be obtained by directly solving the impulse responses following the monetary shock. For the counterfactual scenario, the total effect is equal to what would be obtained when holding the real interest rate path unchanged as shown in Figure 2.3a. The total effect is almost the same as the effect from the one-time monetary shock because the remaining effects from nominal interest rates that are not already captured by changes in real interest rates are very small.

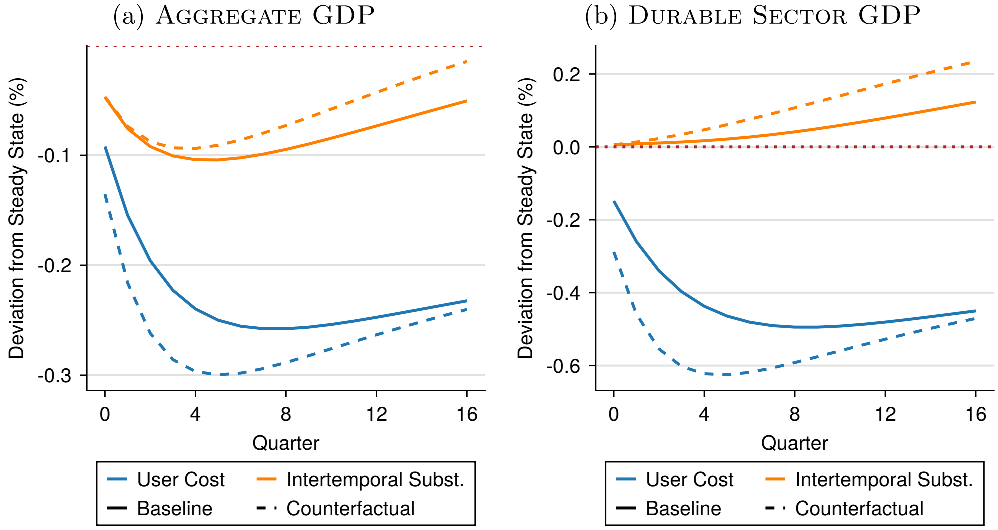

FIGURE 2.4: GE Decompositions by the Effects of Real Interest Rate Changes

paths between the baseline and counterfactual scenarios over the short horizons. In particular, under the counterfactual scenario in which inventory movements do not lag sales movements, there is a deeper decline in GDP due to the user cost effect. This decline from the durable manufacturing sector is so strong that it largely drives the differences in the aggregate GDP responses between the two scenarios.

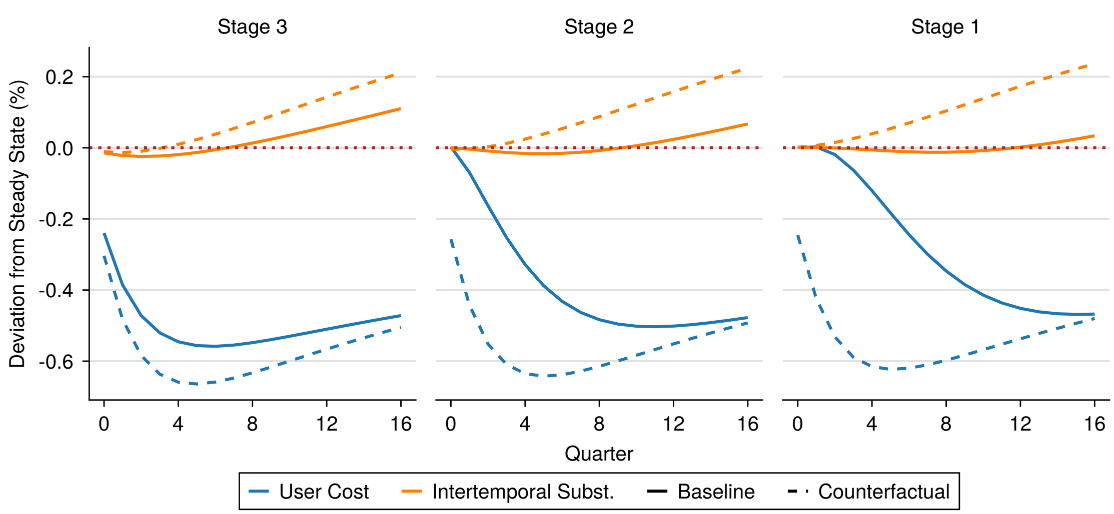

FIGURE 2.5: GE Decompositions of Sales across Production Stages of Durable Sector

Because of the importance of how production in the durable manufacturing sector affects the aggregate GDP, it is useful to dive deeper into the changes occurred within this sector. Figure 2.5

shows the decompositions of stage-specific sales across the three production stages in the durable manufacturing sector. Moving from the downstream stage 3 towards the upstream stage 1, we see that the discrepancies between the user cost effects get larger and larger between the baseline and counterfactual scenarios. For the baseline scenario, the decline in sales in the downstream stage 3 happens immediately following the increase in real interest rates; while the decline in sales starts with a delay for the other two stages. When the lag in inventory-sales comovement is removed under the counterfactual scenario, we see similar immediate decline of sales across all three stages.

What exactly are the different user cost effects that drive the bulk of the differences in the effects on GDP? By definition, the user cost effects are the changes induced by the real interest rate changes when households choose their optimal allocations of resources on the durable consumption goods and sectoral capital stocks. GDP changes as households adjust the investment and utilization of these assets because of the changes in how they value these assets under the new path of real interest rates. The supply chain frictions introduced in Section II of Chapter 1 reduce the sensitivity of such adjustment with respect to real interest rate changes. More specifically, as illustrated in Figure 2.5, this reduction in sensitivity can be attributed to the delayed responses of sales (and hence production) as the impact on downstream sales gradually spreads toward the more upstream stages along the supply chain. Putting together, we may interpret the different aggregate impact as follows. In the baseline scenario where inventory movements lag sales movements, the decline of durable sales as real interest rate goes up passes upward along the supply chain with delays. Because of these delays, the changes in input factor demand for production induced by the decline in sales happen more gradually over time, resulting in smaller changes in the effective capital usage and hence weaker sensitivity of the user cost to real interest rate changes. The presence of supply chain frictions effectively act as a mechanism for the shock to propagate across the production sectors at a lower rate.

B. Contributions of Alternative Frictions

Recall that, as discussed in Section III of Chapter 1, the model involves three types of frictions. The one that suppliers have to place their orders for the critical inputs one quarter in advance is held unchanged. Changes in the other two about the information friction on aggregate sales and the Calvo-style adjustment of order size jointly affect the lag in inventory-sales comovement in the counterfactual experiment conducted above. In particular, to generate the comovement shown in Figure 2.2b without the lag, the information friction is removed and the Calvo friction is weakened. Since reductions in both types of frictions are involved, it is instructive to isolate their contributions to the changes in aggregate outcomes and compare their relative importance.

To that end, I conduct a different type of decompositions over the parameter space. The idea is to leverage the model feature that the strength of information friction is determined by the relative signal precision τ_v/τ_ε and the strength of Calvo friction is determined by the probability of adjustment λ. Specifically, let $\mathcal{Y}_h(\tau_v/\tau_\varepsilon, \lambda)$ denote the cumulative impulse response of GDP in the hth quarter after the monetary shock as a function of τ_v/τ_ε and λ. The parameter values under the baseline and counterfactual scenarios are denoted as $(\tau_{v0}/\tau_{\varepsilon 0}, \lambda_0)$ and $(\tau_{v1}/\tau_{\varepsilon 1}, \lambda_1)$ respectively. With these notations, the changes in the cumulative effects on GDP between the two scenarios can be expressed as

$$\begin{aligned}\mathcal{Y}_h(\tau_{v1}/\tau_{\varepsilon 1}, \lambda_1) - \mathcal{Y}_h(\tau_{v0}/\tau_{\varepsilon 0}, \lambda_0) = \; & \underbrace{\mathcal{Y}_h(\tau_{v1}/\tau_{\varepsilon 1}, \lambda_0) - \mathcal{Y}_h(\tau_{v0}/\tau_{\varepsilon 0}, \lambda_0)}_{\text{Information Friction}} + \\ & \underbrace{\mathcal{Y}_h(\tau_{v0}/\tau_{\varepsilon 0}, \lambda_1) - \mathcal{Y}_h(\tau_{v0}/\tau_{\varepsilon 0}, \lambda_0)}_{\text{Calvo Adjustment}} + \underbrace{\text{residual}}_{\text{Interaction}} \end{aligned} \tag{2.2}$$

where the total difference is decomposed into three components representing the contributions from each type of frictions along with a residual term that can be interpreted as an interaction between the two types of frictions. The first two components are obtained by recomputing the GE impulse responses with only a single parameter value changed.

Figure 2.6 shows the decomposition results based on Equation (2.2) for both the aggregate GDP and the sectoral GDP of the durable manufacturing sector. Notice how the contributions of the frictions vary across horizons. Removing the information friction alone results in more negative cumulative effects that are diverting away from the baseline levels over the first two years, but are reverting back since then, resulting in hump-shaped curves in Figure 2.6. This implies that the presence of information friction affects the GDP in such a way as if there is an intertemporal shift that moves some of the negative effects of monetary shock towards later periods. The interaction between the two types of frictions further strengthens this intertemporal shift. In contrast, weakening the Calvo friction alone results in constantly strengthening gaps of the cumulative GDP responses

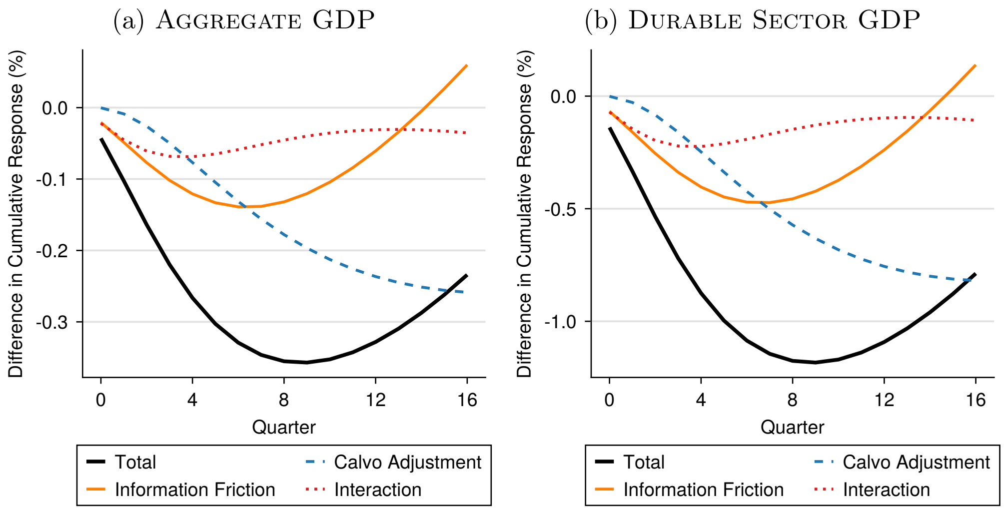

FIGURE 2.6: Contributions of Alternative Frictions to Changes in GDP Responses

between the two scenarios, which are depicted by the monotonically declining curves in Figure 2.6. Despite the substantial contribution from the Calvo friction in explaining the different aggregate outcomes under the counterfactual scenario, only the removal of the information friction and its interaction affect the timing of the aggregate outcomes. This is an important observation because it shows that the model prediction that the GDP responses would reach a trough earlier under the counterfactual scenario is not mechanically driven by the higher rate at which producers adjust their order size. It is the different optimal order size they choose when they make the adjustment that leads to the different aggregate outcomes.

IV. Upstreamness of Supply Chain

The supply chain structure has been fixed unchanged so far. How would the model predictions on aggregate outcomes vary if production involves a different number of stages on average? To answer this question, I examine the cumulative GDP responses across different horizons when there is a change in the average upstreamness of the supply chain, with an upstreamness measure proposed by Antràs and Chor (2013). Since the supply chain in the model is serial, the upstreamness measure for each production stage s is simply $S - s + 1$, the distance towards the final demand. I therefore define the average upstreamness of the supply chain as

$$\overline{\mathcal{U}} = \frac{\sum_{s=1}^{S}(S-s+1)P_{*|s}Y_{*|s}}{\sum_{s=1}^{S}P_{*|s}Y_{*|s}}$$

where $P_{*|s}Y_{*|s}$ is the steady-state sales from stage-s producers. Intuitively, $\overline{\mathcal{U}}$ measures the sales-weighted average distance to the final demand. An increase in $\overline{\mathcal{U}}$ by one can be interpreted as a change in supply chain structure such that one more stage is required on average for each product to arrive at the final users. Notice that this does not mean that the total number of stages existing in the supply chain increases by one (a change in S), as the upstreamness depends on how the production activities are distributed along the supply chain. The same increase in $\overline{\mathcal{U}}$ can be achieved by shifting production towards the upstream stages in different ways.

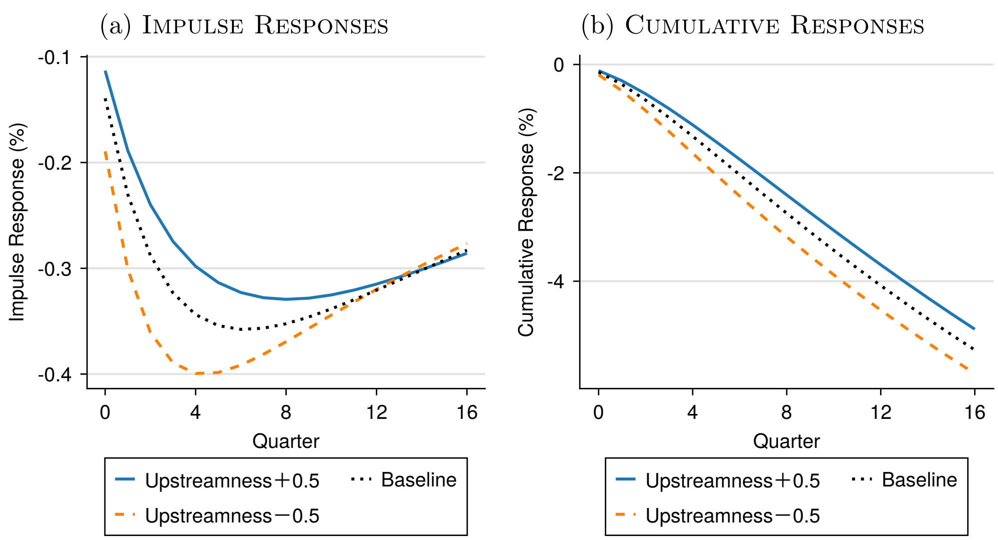

FIGURE 2.7: Aggregate GDP Responses under Different Supply Chain Upstreamness

For the specific quantitative exercise conducted here, I alter $\overline{\mathcal{U}}$ by only changing the steady-state sales ratio between the adjacent production stages $\frac{P_{*|s-1}Y_{*|s-1}}{P_{*|s}Y_{*|s}}$. In the baseline scenario, this ratio takes the same value for all $s > 1$ denoted as ι. It is clear that $\overline{\mathcal{U}}$ is strictly increasing in ι. A change in ι therefore defines a specific way for $\overline{\mathcal{U}}$ to increase. To quantify the aggregate impact of a change in the number of production stages, I compute the cumulative impulse responses of GDP under alternative levels of $\overline{\mathcal{U}}$ attained by changing ι while holding all other parameters unchanged. As before, I feed the same path of real interest rates into the model when solving the GE impulse responses.

Figure 2.7 shows the results when ι is changed in such a way that $\overline{\mathcal{U}}$ either increases by 0.5 or decreases by 0.5. Notice that the aggregate GDP responses across the different scenarios diverge over the first year but start to converge since then. This is reflected in the cumulative responses as a widening gap over the initial horizons that remains stable in later horizons. Intuitively, this observation reflects a feature of model outcomes that is captured by Figure 2.5. Namely, since producers in each stage only make their choices based on what they observe from historical order sizes, producers in different stages start to make noticeable contributions to changes in aggregate outcomes in different periods. Under alternative supply chain structures that vary by average upstreamness, a higher concentration of producers in the upstream further delays the spread of sales impact. In contrast, a shift of producers towards the downstream accelerates this process. A change in the average upstreamness of the supply chain therefore affects the strength of the intertemporal shift of GDP impact, which is the most noticeable over the short horizons.

V. Conclusions

To assess the quantitative significance of the lagged inventory-sales comovement, I embed the supply chain production problem from Chapter 1 into the durable manufacturing sector of a multisector New Keynesian framework with input-output relations. I calibrate and estimate model parameters to allow the model to generate hump-shaped impulse responses following a monetary shock that look similar to their empirical counterparts obtained from SVARs. With the estimated model, I compare baseline aggregate outcomes with those obtained under a counterfactual scenario in which input inventory movements do not lag sales movements. The counterfactual experiment suggests that, holding the path of real interest rates unchanged, altering the supply chain frictions in such a way that removes the lag in comovement results in stronger responses of aggregate GDP in the first year that reach a trough earlier but also revert back faster. General equilibrium decompositions suggest that the different outcomes are mostly accounted by the user cost effect of real interest rates, which in turn can be largely attributed to the delayed transmission of sales changes across production stages in the durable manufacturing sector. Furthermore, decomposing the contributions of the two types of model frictions governing the counterfactual outcomes suggest that the information friction on aggregate sales underlies the intertemporal shift of the GDP impact. Increasing the average upstreamness of the supply chain further strengthens such intertemporal shift.

Despite the stylized nature of the model economy, the findings in this paper suggest the potential of fruitful future research on the interaction between the supply chain structure of an economy and its aggregate outcomes. This paper has only focused on a positive aspect of the economy. Policy implications are left for future work.

Acknowledgements

Chapter 2, in part, is currently being prepared for submission for publication of the material. The dissertation author was the primary author of this paper.

Appendix 2.A. Additional Details on Data

Time series data used for estimation are all obtained from standard sources listed below.

TABLE 2.3: Time Series Data

Variable	Source	Series Code	Frequency
Nominal value added			
Gross domestic product	NIPA	A191RC	Quarterly
Nondurable consumption	NIPA	DNDGRC	Monthly
Durable consumption	NIPA	DDURRC	Monthly
Fixed investment	NIPA	A007RC	Quarterly
Real inventories and gross sales			
Numerous series from NIPA underlying detail tables			Monthly
Employment			
Weekly hours	CES	AWHNONAG (FRED)	Monthly
Employment level	CPS	CE16OV (FRED)	Monthly
Unemployment rate	CPS	UNRATE (FRED)	Monthly
Population	NIPA	B230RC	Quarterly
Prices			
Consumption price deflator	NIPA	DPCERG	Monthly
Producer price index: All commodities	PPI	PPIACO (FRED)	Monthly
Producer price index: Intermediate demand, Materials and components for manufacturing	PPI	WPSID611 (FRED)	Monthly
Hourly earnings	CES	AHETPI (FRED)	Monthly
Federal funds rate	H15	FEDFUNDS (FRED)	Monthly
Other			
Industrial production: Total index	G17	INDPRO (FRED)	Monthly

Notes: NIPA—National Income and Product Accounts from US Bureau of Economic Analysis; CES—Current Employment Statistics from US Bureau of Labor Statistics; CPS—Current Population Survey from US Census Bureau; PPI—Producer Price Indexes from US Bureau of Labor Statistics; H15—H.15 Selected Interest Rates from Board of Governors of the Federal Reserve System; G17—G.17 Industrial Production and Capacity Utilization from Board of Governors of the Federal Reserve System. Series code is either assigned by the original data source or Federal Reserve Economic Data (FRED).

For variables that are involved in the impulse response estimation, they are constructed as follows. Since the SVARs are estimated at the monthly frequency, I transform the data that are only available at the quarterly frequency to the monthly frequency by linear interpolation. All the nominal variables are converted to real quantities using the same consumption price deflator (NIPA DPCERG) and divided by population. For labor supply, I adjust the weekly hours data by multiplying the employment rate that is computed as the ratio between employment and population levels. I use the federal funds rate as the nominal interest rate. Inflation rate is computed as the annual growth rate of the consumption price deflator. Real interest rate is computed as the difference between the nominal interest rate and the one-month-ahead inflation rate. I use the ratio between hourly earnings and the consumption price deflator as real wage. For estimation, I take the log transformation for all variables except the interest rates and inflation rates that are already in percentage terms.

Chapter 3. Dynamic Adjustment to Trade Shocks

Innovations and disruptions to global supply chains lead to gradual adjustments in international trade flows. It has long been recognized that the trade elasticity, a key parameter that captures the substitution between imported goods from different countries in response to trade costs, varies by time horizon (e.g. Dekle, Eaton and Kortum 2008). Boehm, Levchenko and Pandalai-Nayar (2023) use plausibly exogenous tariff changes to measure the trade elasticity by time horizon and find that the short-run trade elasticity is about half the size of the long-run elasticity. This differential implies substantial frictions in trade adjustment that a static trade model cannot account for. A dynamic framework is needed to provide a rigorous and plausible quantification of the transitory and lasting impacts of shocks to global supply chains.

This paper proposes a dynamic general-equilibrium model of trade with many countries and many industries, where staggered sourcing decisions give rise to horizon-specific trade elasticities. Under the Ricardian trade tenet, products are sourced from the least expensive global supplier. However, the opportunity to switch to a new supplier only arrives randomly following a Poisson process. As a consequence, only some buyers respond to a trade disruption by adjusting to optimal sourcing relations. Other buyers endure a suboptimal sourcing choice until they can adjust. In this framework, disruptions put the world economy through a sustained period of adjustment.

The model preserves the analytical tractability of a class of quantitative Ricardian models based on Eaton and Kortum (2002, henceforth EK). We characterize impulse responses in the model using the dynamic hat algebra method. We establish a closed-form expression for the horizon-specific trade elasticity, showing that our model rationalizes empirical estimates of the trade elasticity at different time horizons as a convex combination of short- and long-run elasticity parameters, linked by transitory weights that shift at a constant rate of decay. Furthermore, we derive a novel characterization of the horizon-specific gains from trade that sheds light on the importance of sourcing frictions. Our model shows how the original static welfare formula based on Arkolakis, Costinot and Rodríguez-Clare (2012) can be augmented to account for dynamic adjustment so it delivers welfare predictions at any time horizon under a time-varying trade elasticity.

Specifically, we assume that intermediate goods are produced using constant returns-to-scale technologies and producers differ by productivity drawn from a country-sector specific Fréchet distribution. Trade is subject to iceberg trade costs. An assembler of an industry's final good at a destination d seeks to buy from the least expensive global supplier, but may not be able to instantaneously switch from one supplier to another. The assembler's sourcing decision is governed by a binary random process: an assembler is either in a position to choose the least expensive global supplier of an intermediate good from any source-industry, or the assembler has to continue purchasing from the same producer as in the preceding period. We can therefore characterize equilibrium as a set of measurable partitions of the space of intermediate goods for each supplier, and then derive the equilibrium distributions. An intermediate good's price at a moment in time equals the initial destination price adjusted for the cumulative changes in marginal costs since the supplier was last selected. We show that a destination country's expenditure shares by source country across intermediate goods take an analytic form as in EK and similar Ricardian frameworks that are consistent with the gravity equation of trade.

The expenditure shares in the augmented gravity equation encode the price that a buyer paid at the time of the last supplier change. Through this unmoved component, while a buyer-supplier relationship lasts, cross price effects of substitution are governed by the short-run trade elasticity, similar to an Armington (1969) model. When all supplier-buyer relationships are reset optimally, the gravity expression simplifies to the common gravity equation in an EK framework, so that the long-run trade elasticity prevails. With the equilibrium relationships at hand, we compute impulse responses recursively, and we analytically derive the trade elasticity ε_i^h for each time horizon h after a shock to the global supply network at time $t = 0$:

$$\varepsilon_i^h \equiv \frac{\partial \log \lambda_{sdi,h}}{\partial \log \tau_{sdi,0}} = -\theta_i \left[1 - (1-\zeta_i)^{h+1}\right] - (\sigma_i - 1)(1-\zeta_i)^{h+1},$$

where $\lambda_{sdi,h}$ is destination country d's expenditure share falling on intermediate goods from source country s in industry i in the hth period after the shock, $\tau_{sdi,0}$ is the trade cost component that is shocked at time $t = 0$, θ_i is the long-term trade elasticity as in EK, $\sigma_i - 1$ is the short-term trade elasticity as in Armington, and $\zeta_i \in (0, 1)$ is a parameter that describes the frequency at which buyers of intermediate goods from industry i can switch suppliers. The prevailing trade elasticity ε_i^h increases over time in absolute value from the short-run to the long-run level (for the common parametrization $\theta_i > \sigma_i - 1$).

In the long-run, the trade elasticity converges to the familiar Fréchet parameter θ_i as in EK. The rate of convergence depends on the frequency at which buyers can establish a new sourcing relationship ζ_i. The key parameters of our model are therefore identifiable from reduced-form estimates of the trade elasticity at varying time horizons as in Boehm, Levchenko and Pandalai-Nayar (2023). This characterization of the horizon-specific trade elasticity also implies a horizon-specific welfare formula, which we derive in closed form. The horizon-specific welfare formula features a dynamic adjustment component, which fades over time, and nests the well-known formula from Arkolakis, Costinot and Rodríguez-Clare (2012) as the limiting case in the long run.

We show how the above results can be used to derive a set of estimation equations for the relevant parameters governing short and long-run trade elasticities, document how existing results from Boehm, Levchenko and Pandalai-Nayar (2023) can be employed, and quantify our trade model. With the tractability of our model and data on input-output relations, we consider a model world economy consisting of 32 industries across 77 regions. We apply the model to the episode of the US-China trade war started in 2018 and show that rich industry-level dynamics can result, with consequential changes in welfare implications. First, despite the low trade elasticity in the short-run, the United States main suffer a smaller welfare loss over the short run relative to the long-run outcome, when the sourcing frictions are no longer relevant. China, on the other hand, may suffer a short-run welfare loss that exceeds the long-run loss. A lower short-run trade elasticity therefore does not necessarily imply a larger short-run welfare impact in this world economy. Second, a direct application of the static welfare formula from Arkolakis, Costinot and Rodríguez-Clare (2012), using realized domestic trade shares, can result in qualitatively misleading predictions over finite time horizons. The reason is that sourcing frictions, and the resulting time-varying trade elasticities, can induce substantive and shifting deviations from the long-term welfare outcome. Third, gains from trade can differ between the short and the long run in both sign and magnitude. In the short-run, price disruptions caused by the US-China trade war propagate through the network of existing supply relationships, leading to a global reduction in economic welfare. Those short-run losses, in part, reflect the limited scope for third-party countries to gain from the trade dispute by forming new supply relationships with the United States or China. Gains for third countries may materialize in the medium to long term, however. As a consequence, countries whose previous trade linkages leave them most exposed to the US-China trade war, such as Mexico and Vietnam, experience large initial welfare losses in the short-run, but marked and potentially sizeable increases in welfare in the long-run.

The wide discrepancy between a low (short-run) trade elasticity in international macroeconomics and a high (long-run) trade elasticity in international trade has been documented in, for example, Ruhl (2008, who calls the discrepancy an "international elasticity puzzle") and Fontagné, Martin and Orefice (2018). Fontagné, Guimbard and Orefice (2022), Boehm, Levchenko and Pandalai-Nayar (2023) and Anderson and Yotov (2022) offer estimation procedures to separately identify short- and long-run trade elasticities. de Souza et al. (2022) obtain horizon-specific trade elasticity estimates in a difference-in-differences design for anti-dumping tariff changes. Anderson and Yotov (2022) rationalize their estimation procedure with firm heterogeneity in lag times from recognition to action in the spirit of Lucas and Prescott (1971). In an alternative approach from a macroeconomic perspective, Yilmazkuday (2019) proposes a framework with nested CES models and derives the trade elasticity as the weighted average of macro elasticities. Our general equilibrium model offers a rationalization for the existing estimation methods with a mixture of the Armington and EK elasticities.

The importance of staggered contracts for trade and exchange rate dynamics has been recognized since at least Kollintzas and Zhou (1992) and shares features with staggered pricing (Calvo 1983*b*). We generalize deterministic contract ages to supplier relationships that end stochastically. In a related approach, Arkolakis, Eaton and Kortum (2011) embed a consumer with no knowledge of the identity of source countries into an EK model. The consumer can switch to the lowest-cost supplier at random intervals but cannot act strategically because the supplier is unknown. We rationalize consumer behavior by introducing assemblers that operate similar to a wholesale or retail firm in that they source bundles of goods at lowest cost while the consumer cannot unbundle the assembled final good. An assembler, in turn, cannot incur losses in imperfect capital markets and thus sources from the current lowest-cost supplier. Our model allows us to derive a stationary equilibrium distribution of supplier prices by age of contract beyond a binary characterization in Arkolakis, Eaton and Kortum (2011).[1] Based on the mixture of the stationary equilibrium distributions of prices by

[1] The underlying stochastic process shares features with the so-called Sisyphos Process (Montero and Villarroel 2016).

contract age, we can fully characterize steady states as well as transition dynamics. As a result, we obtain the original EK model as the limit of the equilibria along the transition path. Our welfare formula therefore endogenously inherits the long-run elasticity as a special case when all supplier contracts are optimally set.

The remainder of the paper is organized as follows. We present the model in Section I, with details on mathematical derivations relegated to the Appendix. In Section II we turn to the dynamic analysis of the model. Estimation of the key parameters follows in Section III. To illuminate the novel dynamic features of the model for the allocation of economic activities during the adjustment path and the welfare consequences, we present a case study of the US-China trade war in Section IV. Section V concludes.

I. Model

A. Fundamentals

Consider a world economy with N destination countries $d \in \mathcal{D} := \{1, 2, \cdots, N\}$, $s \in \mathcal{D}$ source countries of trade flows, and I industries $i, j \in \mathcal{I} := \{0, 1, 2, \cdots, I\}$. Time t is discrete. Subscripts sdi, t denote a trade flow from source region s to destination d in industry i at time t. Households inelastically supply a single production factor (labor) to domestic firms, and markets are perfectly competitive.

Households. In each period t, a mass of L_d infinitely-lived households in country d inelastically supplies one unit of the production factor to domestic firms at a competitive wage $w_{d,t}$. Household utility in country d at time t is given by $u(C_{d,t})$, where $C_{d,t}$ is the final good: a Cobb-Douglas aggregate over the composite goods $C_{di,t}$ from each industry with

$$C_{d,t} = \prod_{i \in \mathcal{I}} (C_{di,t})^{\eta_{di}} .$$

The coefficient η_{di} is the consumption expenditure share of industry $i's$ composite good, with $\sum_{i \in \mathcal{I}} \eta_{di} = 1$. Let $P_{di,t}$ denote the price index of the industry i good in d at time t. Country $d's$ consumer price index is then given by $P_{d,t} = \prod_{i \in \mathcal{I}} (P_{di,t}/\eta_{d,i})^{\eta_{di}}$. We assume that households consume their income in every period and discount future utility flows at rate $\beta \in (0, 1)$.

Intermediate Goods. Every industry i consists of a continuum of producers of intermediate goods $\omega \in [0,1]$. For each intermediate good, there is a large set of potential producers in each country with different technologies to produce the good. In each industry, producers of an intermediate good ω have an individual productivity z and operate a constant-returns-to-scale technology to produce the good using domestic labor ℓ and composite goods M_{ji} sourced from other industries:

$$y_i(\omega) = z\,(\ell)^{\alpha_{di}} \prod_{j \in \mathcal{I}} (M_{ji})^{\alpha_{dji}}.$$

where $y_i(\omega)$ is the output of good ω. The coefficient α_{di} is the value-added share of industry i and the parameters $\alpha_{dij} \geqslant 0$ are such that $\alpha_{di} = 1 - \sum_{j \in \mathcal{J}} \alpha_{dji}$.

We assume that intermediate goods can be traded across countries subject to an iceberg transportation cost, which implies that shipping one unit of a good in industry i from country s to country d at time t requires producing $d_{sdi,t} \geqslant 1$ units in s, where $d_{ddi,t} = 1$ for all d. Moreover, goods imported by d from s at t may be subject to an ad-valorem tariff $\bar{\tau}_{sdi,t}$. We combine both trade costs into one parameter $\tau_{sdi,t} \equiv d_{sdi,t}\bar{\tau}_{sdi,t}$.

Given this formulation of trade costs and technologies, there is a *common unit cost component* at destination d for all intermediate goods produced in country s, which we denote with

$$c_{sdi,t} \equiv \Theta_{sj}\tau_{sdi,t}\,(w_{s,t})^{\alpha_{si}} \prod_{j \in \mathcal{J}} (P_{sj,t})^{\alpha_{sji}}, \tag{3.1}$$

where Θ_{sj} is a collection of Cobb-Douglas coefficients. The resulting unit cost of good ω at destination d produced in country s with a productivity $z(\omega)$ is given by $c_{sdi,t}/z(\omega)$.

Production technologies for intermediate goods arrive stochastically and independently at a rate that varies by country and industry. In particular, we follow Eaton and Kortum (2012) in assuming that the mass of intermediate goods ω in country s's industry i that can be produced with a productivity higher than z to be distributed Poisson with mean $A_{si}z^{-\theta_i}$.

Assembly of Composite Goods. In each industry, assemblers bundle intermediate goods into a composite good for consumption or production. An assembler procures intermediate goods at the lowest possible price and costlessly aggregates the sourced intermediates into $Y_{di,t}$ units of industry

i's composite good using the technology

$$Y_{di,t} = \left(\int_{[0,1]} y_{di,t}(\omega)^{(\sigma_i - 1)/\sigma_i} \mathrm{d}\omega \right)^{\frac{\sigma_i}{\sigma_i - 1}}, \tag{3.2}$$

where $y_{di,t}(\omega)$ is the quantity purchased of an intermediate good ω by an assembler in country d, and σ_i is the elasticity of substitution between intermediate goods in industry i. We let $p_{di,t}(\omega)$ denote the lowest possible price at which an intermediate good ω can be purchased at destination d. We will explain the exact price at which this intermediate good is available in greater detail below. As we elaborate in Appendix A, cost minimization given (3.2) implies that the price of industry i's composite good at destination d satisfies

$$P_{di,t} = \left(\int_{[0,1]} p_{di,t}(d\omega)^{-(\sigma_i - 1)} \mathrm{d}\omega \right)^{-\frac{1}{\sigma_i - 1}}.$$

B. Sourcing Decisions and Trade Flows

Under the Ricardian trade tenet, assemblers seek to source an intermediate good from the least expensive global supplier. However, an assembler may not have the opportunity to adjust its choice of suppliers at any given time due to a sourcing friction, which we describe now. For every intermediate good ω, there is a continuum of producers in every country. Under perfect competition, an assembler optimally sources any given intermediate good ω from only one source country when given the choice.

The assemblers' choice of source country for any given intermediate good ω is governed by an i.i.d. random variable $x_{i,t}(\omega) \in \{0, 1\}$ for each industry. If $x_{i,t}(\omega) = 1$, that is if the global draw for an intermediate good ω from industry i gives all assemblers worldwide the green light to switch to their preferred source country, then all assemblers optimally choose to purchase from the least costly source country for variety ω in industry i at time t. Between assemblers in different countries the optimal source country can vary because of different trade costs. Else, if $x_{i,t}(\omega) = 0$, that is if the global draw for intermediate ω turns to red for all assemblers worldwide, then all assemblers must purchase their intermediate goods ω in industry i from the same producer as in the preceding period $t - 1$. While the identity of the source country does not change, the quantity procured and the price that the assembler pays can differ from the preceding period if the factory gate price moves (because of changing factor costs) or the currently prevailing trade cost moves.

This formulation of sourcing frictions captures search costs and other types of impediments that prevent the optimal rematch of supply relationships at a moment in time. An implication of the sourcing friction is that price elasticities of demand will differ across intermediate goods according to when their suppliers were last chosen. Let $\Omega^k_{j,t}$ denote the set of industry j goods whose supplier at time t was last chosen k periods ago:

$$\Omega^k_{i,t} = \left\{\omega : x_{di,t-k}(\omega) = 1, \, \prod_{\varsigma=t-k+1}^{t} x_{di,\varsigma}(\omega) = 0\right\},$$

where $\cup_k \Omega^k_{j,t} = [0,1]$. The sets $\Omega^k_{i,t}$ mutually exclusively and exhaustively partition the unit interval of intermediate goods for each industry i.

1. Demand for Intermediate Goods with Newly Formed Supply Relationships

We now describe the global demand for intermediate goods in each of these sets, beginning with those that are concurrently formed, $\omega \in \Omega^0_{dj,t}$.

If country s is chosen by an assembler in destination d to supply industry i's intermediate good ω at time t, the combination of the producer's productivity ω, factor cost in source country s and the trade cost between s and d in industry i must make the intermediate good the least expensive.

Let $z_{si}(\omega)$ denote the highest realized productivity by any producer in country-industry si. Similar to Eaton and Kortum (2002), our distributional assumptions imply that z_{si} has a country-industry specific Fréchet distribution given by[2]

$$\Pr\left[z_{si}(\omega) \leqslant z | A_{si}, \theta_i\right] = \exp\left\{-A_{si} z^{-\theta_i}\right\}.$$

For an assembler in destination d the price of an intermediate good ω from the cheapest available

[2]Our model could also accommodate productivity change over time with a country-industry-time specific Fréchet distribution and resulting $z_{si,t}(\omega)$ realizations that vary over time. To focus most sharply on adjustment to trade shocks, we do not specify productivity shocks.

source country at time t is

$$p_{di,t}(\omega) = \min_{s \in \mathcal{D}} \left\{ \frac{c_{sdi,t}}{z_{si}(\omega)} \right\}$$

for the common unit cost component $c_{sdi,t}$ given by (3.1) and the producer with the highest realized productivity $z_{si}(\omega)$ in country-industry si.

As in Eaton and Kortum (2002), the distribution of paid prices across intermediate goods in the set $\Omega^0_{i,t}$ in destination d at time t satisfies

$$G^0_{di,t}\left[p_{di,t}(\omega) \leqslant p\right] \equiv \Pr\left[p_{di,t}(\omega) \leqslant p \middle| x_{i,t}(\omega) = 1\right] = 1 - \exp\left\{-\Phi^0_{di,t} p^{-\theta_i}\right\},$$

where

$$\Phi^0_{di,t} \equiv \sum_{n \in \mathcal{N}} A_{ni}[c_{ndi,t}]^{-\theta_i} \tag{3.3}$$

is a measure of destination d's market access for intermediate goods $\omega \in \Omega^0_{i,t}$, given trade cost and factor prices behind the common unit cost component $c_{ndi,t}$ by (3.1). We relegate the derivation of these results to Appendix B. To guarantee that the distribution of paid prices has a finite mean later, we impose the standard parametric restriction that $\theta_i > \sigma_i - 1$ for all $i \in \mathcal{I}$.

The properties of the Fréchet distribution imply that $G^0_{di,t}$ also equals the distribution of prices for intermediate goods $\omega \in \Omega^0_{i,t}$ sourced from any source country s. As a result, country d's expenditure share for each potential source country s across intermediate goods $\omega \in \Omega^0_{i,t}$ must equal the probability that this source country offers the lowest global price:

$$\lambda^0_{sdi,t} = \frac{A_{sj}[c_{sdi,t}]^{-\theta^i}}{\Phi^0_{di,t}}. \tag{3.4}$$

with the common unit cost component $c_{sdi,t}$ given by (3.1).

Within the set of intermediate goods that are sourced through concurrently and optimally formed supply relationships, the partial equilibrium elasticity of trade flows with respect to trade cost is governed by the familiar Fréchet parameter:

$$\left.\frac{\partial \log \lambda^0_{sdi,t}}{\partial \log \tau_{sdi,t}}\right|_{\Phi^0_{di,t}} = -\theta_j.$$

2. Demand for Intermediate Goods with Continuing Supply Relationships

Intermediate goods $\omega \in \Omega^k_{j,t}$ are purchased from a supplier that was chosen at time $t-k$. To characterize prices and expenditure allocations across these intermediate goods at time t, we denote changes over time for a variable x_t succinctly by $\hat{x}_t \equiv x_t/x_{t-1}$.

Suppose an assembler in d first sourced an intermediate good ω from s at time $t-k$ under the unit input cost $c_{sdi,t-k}/z_{si}(\omega)$, which depends on equilibrium factor prices and parameters by the common unit cost component (3.1). If the intermediate good is still sourced from the same producer at time t, its price will then equal:[3]

$$p^k_{sdj,t}(\omega) = \frac{c_{sdi,t}}{z_{si}(\omega)} = \frac{c_{sdi,t-k}\prod_{\varsigma=t-k+1}^{t}\hat{c}_{sid,\varsigma}}{z_{si}(\omega)},$$

which is the initial destination price adjusted for the cumulative changes in iceberg trade costs and factor cost since $t-k$.

We show in Appendix C that country d's expenditure share by source country across intermediate goods $\omega \in \Omega^k_{i,t}$ equals

$$\lambda^k_{sdi,t} = \frac{\lambda^0_{sdi,t-k}\left(\prod_{\varsigma=t-k+1}^{t}\hat{c}_{sid,\varsigma}\right)^{1-\sigma_i}}{\Phi^k_{di,t}}, \tag{3.5}$$

where

$$\Phi^k_{di,t} \equiv \sum_{n\in\mathcal{N}}\lambda^0_{ndi,t-k}\left(\prod_{\varsigma=t-k+1}^{t}\hat{c}_{nid,\varsigma}\right)^{1-\sigma_i}$$

reflects the mean price that a buyer pays for the set of intermediate goods $\Omega^k_{i,t}$ at time $t-k$ through the trade shares $\left\{\lambda^0_{nid,t-k}\right\}_{n\in\mathcal{N}}$.

Comparing Equations (3.4) and (3.5) shows how cross-price effects differ across intermediate goods depending on when a supply relationship is formed. If assemblers can source from the least expensive global supplier of an intermediate good at time t, cross-price demand effects are governed the Fréchet parameter θ_i, and trade is governed by comparative advantage.

Conversely, if an assembler is unable to switch suppliers, then the extensive margin is shut down. The only margin of adjustment is the intensive margin, which is captured by the terms that

[3] Note that $x_t = x_{t-k}\frac{x_{t-k+1}}{x_{t-k}}\cdots\frac{x_t}{x_{t-1}} \equiv x_{t-k}\hat{x}_{t-k+1}\cdots\hat{x}_t$. For a composite variable such as $c_{sdi,t} = \tau_{sdi,t}\,w_{s,t}$, the change over time is $\hat{c}_{sdi,t} = \hat{\tau}_{sdi,t}\,\hat{w}_{s,t}$.

collect the product of changes in unit input costs. Effectively, over those partitions, trade happens as if varieties were differentiated across countries with the measure of varieties of each source defined at the last period of adjustment —i.e. at period $t-k$ for partition $\Omega_{i,t}^{k}$.

In order words, for each partition $\Omega_{i,t}^{k}$, trade happens under Armington forces. Intuitively, the price elasticity of demand is governed by the elasticity of substitution $\sigma_i - 1$, which captures Armington trade:

$$\left.\frac{\partial \log \lambda_{sdi,t}^{k}}{\partial \log \tau_{sdi,\varsigma}}\right|_{\Phi_{di,t}^{k}} = -\left(\sigma_i - 1\right) \quad \text{for } t-k < \varsigma < t.$$

To close the model, we now show how aggregate global demand for industry i's composite good follows from aggregating the trade shares in Equations (3.4) and (3.5).

C. Aggregation

To find aggregate demand, we leverage the homotheticity of assembly. The partial price index for the composite of intermediate goods purchased at time t from suppliers chosen $t-k$ periods ago satisfies $(P_{di,t}^{k})^{1-\sigma_j} = \int_{\omega \in \Omega_{i,t}^{k}} p(\omega)_{di,t}^{1-\sigma_j} \mathrm{d}\omega$. The sets $\left\{\Omega_{i,t}^{k}\right\}_{k=0}^{\infty}$ form a partition of industry i's product space, so we can obtain country d's price index for industry i goods at time t by aggregating these partial price indices over all partitions and find $P_{di,t}^{1-\sigma_j} = \sum_{k=0}^{\infty} \left(P_{di,t}^{k}\right)^{1-\sigma_j}$.

We establish in Appendix B that the partial price index for the set of intermediate goods whose suppliers are being chosen at time t takes the familiar form

$$P_{di,t}^{0} = \gamma_i \, \mu_{i,t}(0)^{1/(1-\sigma_j)} \left(\Phi_{di,t}^{0}\right)^{-\frac{1}{\theta_i}}, \tag{3.6}$$

where $\gamma_i \equiv \Gamma\left([\theta_i - \sigma_i + 1]/\theta_i\right)^{1-\sigma_i}$ is a constant, $\Phi_{di,t}^{0}$ is given by (3.3), and $\mu_{i,t}(0)$ denotes the measure of the set $\Omega_{i,t}^{0}$. Following the previous discussion, the endogenous market access term $\Phi_{di,t}^{0}$ represents the mean price of intermediate goods whose suppliers are chosen at time t. The measure $\mu_{i,t}(0)$ accounts for gains from variety. This measure recursively evolves over time according to the stochastic process that governs sourcing decisions, given by

$$\mu_{i,t}(k) = \begin{cases} \zeta_i, & k = 0 \\ (1-\zeta_i)\mu_{i,t-1}(k-1), & k > 0. \end{cases}$$

As we show in Appendix C, the partial price index across intermediate goods whose suppliers were last chosen at time $t-k$ is given by

$$P^k_{di,t} = P^0_{di,t-k}\left(\frac{\mu_{i,t}(k)}{\mu_{i,t-k}(0)}\Phi^k_{di,t}\right)^{1/(1-\sigma_i)}, k > 1 \tag{3.7}$$

which is the period $t-k$ price index of the basket of intermediate goods Ω^0_{t-k}, adjusted for the subsequent change in variety composition, captured by $\mu_{i,t}(k)/\mu_{i,t-k}(0)$, and prices, captured by $\Phi^k_{di,t}$.

Given Equations (3.6) and (3.7), we can solve for the composite price index of industry i goods in country d at time t:

$$P_{di,t} = \gamma_i \left(\Phi^0_{di,t}\right)^{-\frac{1}{\theta_i}} \left[\mu_{i,t}(0) + \sum_{k=1}^{\infty} \mu_{i,t}(k) \left(\frac{\Phi^0_{di,t}}{\Phi^0_{di,t-k}}\right)^{\frac{1-\sigma_i}{\theta_i}} \Phi^k_{di,t}\right]^{\frac{1}{1-\sigma_j}} \tag{3.8}$$

The term $\gamma_i \left(\Phi^0_{di,t}\right)^{-1/\theta_i}$ on the right-hand-side of Equation (3.8) captures the prices paid under flexible supplier choice. The term in brackets quantifies the extent to which current aggregate demand is affected by the stickiness of supply relationships. The term $\Phi^k_{di,t}$ captures differences in demand across intermediate goods driven by differences in the age of their supply relationships and reflect their impact on aggregate demand at time t. The terms $(\Phi^0_{di,t}/\Phi^0_{di,t-k})^{(1-\sigma_i)/\theta_i}$ measure the current demand of a buyer whose supplier relationship from k periods ago differs from that of a buyer who just updated its supplier.

Using the above price indices, we can readily derive country d's expenditure share on industry i goods sourced from country s

$$\lambda_{sdi,t} = \sum_{k=0}^{\infty} \lambda^k_{sdi,t} \left(\frac{P^k_{di,t}}{P_{di,t}}\right)^{1-\sigma_i}. \tag{3.9}$$

where $\lambda^k_{sdi,t}$ is given by Equation (3.4) if $k = 0$ and (3.5) if $k > 0$.

The set of trade shares $\{\lambda_{sdi,t}\}_{s,d\in\mathcal{N},i\in\mathcal{I}}$ fully characterize demand in the world economy at time t. To close the model, we now describe the conditions for market clearing and define a general equilibrium.

D. Equilibrium

Denote the total revenue of an industry i in a source country s at time t by $X_{si,t}$. To define equilibrium, we express each industry's revenue in terms of trade shares, given by Equation (3.9), and total expenditures on consumption, $E_{d,t}$, and intermediate inputs in the rest of the world:

$$X_{si,t} = \sum_{d \in \mathcal{N}} \lambda_{sdi,t} \left[\eta_{di} E_{d,t} + \sum_{j \in \mathcal{I}} \alpha_{dij} X_{dj,t} \right]. \tag{3.10}$$

A country's national consumption spending is the sum of its factor income and trade deficit, $E_{d,t} = w_{d,t} L_{d,t} + D_{d,t}$, with $\sum_{d \in \mathcal{N}} D_{d,t} = 0$. We follow the conventional approach in the international trade literature and treat aggregate trade deficits as exogenous. To clear the factor market, wages then adjust to ensure that expenditures equal disposable income,

$$w_{d,t} L_{d,t} = \sum_{i \in \mathcal{I}} (1 - \alpha_{di}) X_{di,t}, \tag{3.11}$$

and goods market clearing is guaranteed by Walras' law.

We are now ready to define a dynamic general equilibrium and a steady state.

Definition 1. An economy is described by a set of time-invariant parameters summarizing technologies, preferences and factor endowments, $\boldsymbol{\Theta} = \{\theta_i, \sigma_i, \{\alpha_{dji}\}_{j \in \mathcal{I}}, \varphi_{di}, A_{di}, \eta_{di}, L_d\}_{d \in \mathcal{N}}\}_{i \in \mathcal{I}}$, sourcing frictions $\boldsymbol{\zeta} = \{\zeta_i\}_{i \in \mathcal{I}}$, as well as a measure $\boldsymbol{\mu}_{t_0} = \{\mu_{t_0}(k)\}_{k \in \{0,1,\cdots\}}$ for some t_0. Given histories of trade costs $\boldsymbol{\tau}_{t-1} \equiv \{\tau_t\}_{\varsigma < t} = \{\tau_{sid,\varsigma}\}_{s,d \in \mathcal{N}, i \in \mathcal{I}, \varsigma < t}$ and their changes $\hat{\tau}_t \equiv \{\hat{\tau}_{sdi,t}\}_{s,d \in \mathcal{N}, i \in \mathcal{I}}$ as well as nominal wages $\boldsymbol{w}_{t-1} = \{w_\varsigma\}_{\varsigma < t} = \{w_{d,\varsigma}\}_{d \in \mathcal{N}, \varsigma < t}$:

1. A static equilibrium at time t is a vector of wages $w(\hat{\tau}_t \times \tau_{t-1} \cup \boldsymbol{\tau}_{t-1}, \boldsymbol{w}_{t-1}, \zeta, \Theta) = w_t$ that jointly solves Equations (3.9)–(3.11) for all $s, d \in \mathcal{N}$ and $i \in \mathcal{I}$.

2. A dynamic equilibrium at time t is a history of wages $\boldsymbol{w}_t$ so that, for all $w_\varsigma \in \boldsymbol{w}_t$, $w_\varsigma = w(\hat{\tau}_{\varsigma-1} \times \tau_{\varsigma-1} \cup \boldsymbol{\tau}_{\varsigma-1}, w_{\varsigma-1} \cup \boldsymbol{w}_{\varsigma-2}, \zeta, \Theta)$.

3. A dynamic equilibrium at time t is a steady state if $w(\mathbf{1}_{N \times N \times I} \times \tau_t \cup \boldsymbol{\tau}_{t-1}, w_t \cup \boldsymbol{w}_{t-1}, \zeta, \Theta) = w_t$.

E. Steady-State Properties

In the following, we show that our model preserves the class of quantitative trade models based on Eaton and Kortum (2002) in the limit when the economy is in steady state, irrespective of the magnitude of the frictions underlying imperfect supplier adjustment, $\zeta_i \in (0,1)$. Intuitively, the transitory effects of trade disruptions that arise in our model reflect how opportunities for finding new suppliers are limited in the short-run but increasing over time. As assemblers get to adjust all supply relationships in the long-run, we then obtain the EK-model as the limit of the equilibria along the transition path.

More formally, let $w^{EK}(\hat{\tau}_t \times \tau_{t-1} \cup \boldsymbol{\tau}_{t-1}, \boldsymbol{w}_{t-1}, 1, \Theta)$ represent the equilibrium allocation in an economy in which suppliers can be flexibly adjusted for all goods, $\zeta_i = 1$ for all i. We can then establish

Proposition 1. *If w_{t*} is a steady state equilibrium, then*

1. *For any ζ, $w_{t*} = w(\mathbf{1}_{N\times N\times I} \times \tau_{t*} \cup \boldsymbol{\tau}_{t*-1}, w_{t*} \cup \boldsymbol{w}_{t*-1}, \zeta, \Theta) = w^{EK}(\mathbf{1}_{N\times N\times I} \times \tau_{t*} \cup \boldsymbol{\tau}_{t*-1}, w_{t*} \cup \boldsymbol{w}_{t*-1}, 1, \Theta)$.*

2. *For all $k \in \{0,1,...\}$, the measure of goods $\omega \in \Omega^k_{i,t}$ equals $\mu_{i,t*}(k) = (1-\zeta_i)^k \zeta_i$, and trade flows are given by $\lambda^k_{sdi,t*} = \lambda_{sdi,t} = \lambda^{EK}_{sid}$ where λ^{EK}_{sid} denotes the trade shares in the frictionless economy.*

Proposition 1 provides numerous useful insights. The first part makes clear that the tools developed by the literature studying the equilibrium properties of static quantitative trade models can be deployed to establish the existence and uniqueness of steady states in our model.

The second part of Proposition 1 highlights properties of the steady states that we later leverage to quantify the model. In particular, it shows that the process governing the evolution of the age distribution of supply relationships over time has a simple geometric stationary distribution. Further, it shows that steady state expenditure allocations are equalized across goods within an industry, irrespective of when their supplier was chosen.

II. Dynamic Adjustment to Trade Shocks

In this section, we theoretically characterize the economy's dynamic response to trade disruptions. In particular, we derive a new structural estimating equation for the trade elasticity at different time horizons, and show that transitional dynamics can be characterized using the dynamic hat-algebra. Finally, we provide a new formula for characterizing the horizon-specific gains from trade.

A. Trade Elasticity by Time Horizon

We begin by showing how the trade elasticity, that is the elasticity of trade flows with respect to transport cost, varies over time. To do so, we let $\varepsilon^h_{sdi,t}$ denote the trade elasticity at horizon h, which we define by:

$$\varepsilon^h_{sdi,t-1} \equiv \left.\frac{\partial \log X_{sdi,t+h}}{\partial \log \tau_{sdi,t}}\right|_{\{\Phi^k_{di,t+\varsigma}\}_{t\leqslant\varsigma\leqslant h,k}}, \tag{3.12}$$

which is the elasticity of trade flows in industry i from country s to d at time $t+h$, $X_{sdi,t+h}/X_{sdi,t-1}$ with respect to change in trade costs at t, $\mathrm{d}\log\tau_{sdi,t} = \log\hat{\tau}_{sdi,t}$, holding fixed the general equilibrium terms that summarize changes in market access for industry i goods in destination d. The following derives a closed-form expression for this elasticity.

Proposition 2. *Suppose that the economy is in steady state at $t = -1$. Then, up to a first order, the horizon-h response of trade flows to a shock to trade cost at time $t = 0$ is given by:*

$$\varepsilon^h_i = -\theta_i\left[1-(1-\zeta_i)^{h+1}\right] - (\sigma_i - 1)(1-\zeta_i)^{h+1}. \tag{3.13}$$

If $\zeta_i \in (0,1)$, $\lim_{h\to\infty}\varepsilon^h_i = -\theta_i$, where the rate of convergence equals

$$\lim_{h\to\infty}\frac{\varepsilon^{h+1}_j + \theta_j}{\varepsilon^h_i + \theta_i} = \log(1-\zeta_i).$$

Following Proposition 2, the trade elasticity increases over time if $\theta_i > \sigma_i - 1$. In the long-run, it is equal to the Fréchet parameter θ_i, where the rate of convergence, intuitively, depends on the frequency at which buyers can establish a new sourcing relationship ζ_i.

It is worth noting that Equation (3.12) is consistent with reduced-form estimates of the trade elasticity at varying time horizons as in Boehm, Levchenko and Pandalai-Nayar (2023). Later, we leverage this equivalence to identify the key structural parameters in our model. The horizon-specific formulation of the trade elasticity implied by our model also induces a horizon-specific welfare formula, which we provide next.

B. The Horizon-Specific Welfare Gains from Trade

When supply relationships are slow to adjust to shocks, trade disruptions can put the economy through a sustained period of readjustment. The following proposition shows that our framework yields a simple formula for welfare analysis, giving changes in real wages associated with an initial set of foreign shocks over varying time horizons.

PROPOSITION 3. *Suppose the economy is in steady state at $t = -1$. Then, the change in real wages in country d at time $h = \{0, 1, ...\}$, $\hat{W}_d^h = C_{d,h}/C_{d,-1}$, that follows a set of arbitrary shocks to trade cost at time at $t = 0$, is given by*

$$\hat{W}_d^h = \prod_{j \in \mathcal{I}} \left[\left(\frac{\lambda_{ddj,h}}{\lambda_{ddj,-1}} \right)^{-\frac{1}{\theta_j}} (\Xi_{dj,h})^{\frac{1}{\sigma_j - 1}} \right]^{\sum_{i \in \mathcal{I}} \bar{a}_{dji} \eta_i}, \tag{3.14}$$

where

$$\Xi_{dj,h} \equiv \zeta_j \left(\frac{\lambda_{ddj,h}}{\lambda_{ddj,h}^{k=0}} \right)^{\frac{\sigma_j - 1 - \theta_j}{\theta_j}} + (1 - \zeta_j)^{h+1} \left(\frac{\lambda_{ddj,h}}{\lambda_{ddj,-1}} \right)^{\frac{\sigma_j - 1 - \theta_j}{\theta_j}} + \sum_{\varsigma=1}^{h} \zeta_j (1 - \zeta_j)^k \left(\frac{\lambda_{ddj,h}}{\lambda_{ddj,h-\varsigma}^{k=0}} \right)^{\frac{\sigma_j - 1 - \theta_j}{\theta_j}}$$

and $\bar{a}_{dji}$ is the (j, i)-th element of the Leontief inverse $(\boldsymbol{I}d - \boldsymbol{A}_d)^{-1}$, with the elements of $\boldsymbol{A}_d$ given by α_{dji}. If $\zeta_i \in (0, 1)$, then $\lim_{h \to \infty} \hat{W}_d^h = \lim_{h \to \infty} \prod_{j \in \mathcal{I}} (\lambda_{ddj,t+h}/\lambda_{ddj,-1})^{-\sum_{i \in \mathcal{I}} \bar{a}_{dji} \eta_i / \theta_j}$.

Although our model features transition dynamics on the supply side, Equation (3.14) shows that welfare analysis can still be conducted using only a few sufficient statistics. These statistics delineate how the impact of trade shocks on real wages varies over time due to staggered sourcing decisions, decomposing the change in real wages associated with foreign shocks into two effects.

The first effect is captured by the terms $(\lambda_{ddj,h}/\lambda_{ddj,-1})^{-1/\theta_j}$ on the right-hand-side of Equation (3.14). Because the Fréchet parameter θ_j gives the price elasticity of trade flows sourced from the currently cheapest global supplier and the share of domestic expenditures the response of trade to prices, each of these terms would give the change in a particular industry j's domestic price index if all goods were optimally sourced. Because all supply relationships are flexible in the long-run, i.e., when $h \to \infty$, changes in aggregate home expenditure shares and the long-run trade elasticity, thus, remain sufficient for long-run welfare analysis in our model, as in Eaton and Kortum (2002). However, staggered sourcing decisions spell additional welfare effects in the short-run, i.e., when not all goods can be sourced optimally.

Staggered adjustment of suppliers spells time-varying distortions in prices and terms-of-trade, captured by the terms $(\Xi_{dj,h})^{1/(\sigma_j-1)}$ in Equation (3.14). Intuitively, these distortions manifest via expenditure allocations, and will vary across goods depending on when their current supplier was chosen. If a good was last optimally sourced k periods ago, the resulting distortion in its price at horizon h can be informed by the difference between the share of domestic expenditures on all goods time h and on optimally sourced goods at time $h-k$, $(\lambda_{ddj,h}/\lambda^{k=0}_{ddj,h-k})^{(\sigma_j-1-\theta_j)/\theta_j}$. Intuitively, a decrease in $\lambda_{ddj,h}/\lambda^{k=0}_{ddj,h-k}$ indicates that suppliers that were chosen k periods ago are now, at horizon h, less competitive; the implied deterioration in a country's aggregate terms-of-trade is decreasing in the elasticity of substitution, σ_j, and increasing in the share of goods sourced from these suppliers, $\zeta_j \cdot (1-\zeta_j)^k$, is higher.

As an implication of Proposition 3, the trade elasticity relevant for welfare analysis varies over time. To further illustrate this point, it is useful to approximate changes in industry-level prices up to a first-order, which yields

$$\log(\frac{\lambda_{ddj,h}}{\lambda_{ddj,-1}})^{-\frac{1}{\theta_j}}(\Xi_{dj,h})^{\frac{1}{\sigma_j-1}} \approx -\frac{1}{\theta_j}[1-(1-\zeta_j)^{h+1}]\log\frac{\lambda^{k=0}_{ddj,h}}{\lambda_{ddj,-1}} - \frac{1}{\sigma_j-1}(1-\zeta_j)^{h+1}\log\frac{\lambda^{k=h+1}_{ddj,h}}{\lambda_{ddj,-1}} - \mathcal{E}^h_{dj},$$

where $\mathcal{E}^h_{dj} = \sum_{\varsigma=1}^{h+1}(1-\zeta)^{\varsigma}\zeta\left[\frac{1}{\sigma_j-1}\log\frac{\lambda^{k=\varsigma}_{ddj,h}}{\lambda^{k=0}_{ddj,h-\varsigma_j+1}} - \frac{1}{\theta_j}\log\frac{\lambda^{k=0}_{ddj,h+1-\varsigma}}{\lambda^{0}_{ddj,-1}}\right]$.

The first term on the right captures how changes in the prices of goods that were procured optimally at least once contribute to the overall change in prices at horizon h, assuming that past changes in factor prices were equal to those observed h periods after the shock. The second term, in contrast, captures changes in aggregate prices due to changes in the prices of goods whose suppliers have never been adjusted. The relative importance of these two effects varies over time, in tandem with the structural trade elasticity.

The last term, $\mathcal{E}^h_{dj}$, captures how suboptimal sourcing decisions from the past continue to distort prices at horizon h by distorting the equilibrium adjustment of factor prices relative to the long-run. Such distortions are reflected in price differences between goods whose suppliers were adjusted before and those that are procured optimally at horizon h.

Staggered sourcing decisions, hence, imply that the trade elasticity relevant for welfare analysis differs from the structural elasticity in Equation (3.13) due the dynamic interaction of sourcing decisions and factor prices. Due to these interactions, the welfare effects of trade shocks may, then, vary both quantitatively and qualitatively over time, even conditional on the structural parameters underlying the time variation in the trade elasticity. Viewed through this lens, Proposition 3 is fortunate in that it allows us to summarize these dynamic effects in terms of a few statistics, which, as we will now describe, also enables us deploy familiar tools from the international trade literature to solve exactly for the equilibrium response of prices and wages to trade shocks implied by the model.

C. Characterization of Impulse Responses

We now show that solving for the responses of trade and production to shocks does not require knowledge of the economy's structural fundamentals (productivities, and trade costs). As an implication, the so-called "hat algebra" of Dekle, Eaton and Kortum (2007) can be deployed to characterize impulse responses in our model.

Absent inter-sectoral linkages, trade flows at time t can be expressed in terms of succinct changes in trade costs and wages, as well as past changes in trade flows for optimally sourced goods, trade costs and wages:

$$\lambda_{sdi,t} = \frac{\left[1 + \left(\hat{\tau}_{sdi,t}\hat{w}_{s,t}/\hat{w}_{d,t}\right)^{1-\sigma_i+\theta_i} \omega_{sdi,t-1}\right] \lambda^{k=0}_{sdi,t-1} \left(\hat{\tau}_{sdi,t}\hat{w}_{s,t}\right)^{-\theta_i}}{\sum_{s'\in\mathcal{N}} \left[1 + \left(\hat{\tau}_{s'id,t}\hat{w}_{s',t}/\hat{w}_{d,t}\right)^{1-\sigma_i+\theta_i} \omega_{s'id,t-1}\right] \lambda^{k=0}_{s'id,t-1} \left(\hat{\tau}_{s'id,t}\hat{w}_{s',t}\right)^{-\theta_i}},$$

where the wedges

$$\omega_{sdi,t-1} \equiv \frac{\mu_{i,t}(1)}{\mu_{i,t}(0)} + \sum_{k'=2}^{\infty} \frac{\mu_{i,t}(k')}{\mu_{i,t}(0)} \left(\frac{\lambda^{k=0}_{ddi,t-1}}{\lambda^{k=0}_{ddi,t-k'}}\right)^{\frac{\sigma_i-1}{\theta_i}} \frac{\lambda^{k=k'}_{sdi,t-1}}{\lambda^{k=0}_{sdi,t-1}} \prod_{\varsigma=t-k''+1}^{t-1} \left(\hat{\tau}_{sid,\varsigma}\frac{\hat{w}_{s,t}}{\hat{w}_{d,\varsigma}}\right)^{1-\sigma_i},$$

summarize how prior distortions in factor prices continue to impact trade flows at time t by distorting the terms of trade.

Now suppose that the economy was in steady state at some time prior to t. Then, given bilateral country-sector trade flows, industry-level consumption and intermediate good expenditure shares as well as per-capita GDP, the only additional industry-level parameters that are required to recursively compute changes in trade flows at increasing time horizons are given by $\{\zeta_i, \theta_i, \sigma_i\}$. Given this recursive formulation for trade flows, we can express the market clearing conditions (3.11) in terms of changes in trade costs and factor prices, as in Dekle, Eaton and Kortum (2007), and, hence, solve for the period-by-period change in wages associated with (a sequence of) trade shocks.

III. Estimation

We now turn to exploring the quantitative implications of our theory for the response of production and welfare to trade shocks. In this section, we outline and implement our approach to estimating the structural parameters that govern the time variation of the trade elasticity. In the next section, we will use these estimates for a quantitative re-evaluation of how the 2018 US-China trade war impacted trade, production and welfare.

A. Approach

Proposition 2 implies that we can express the trade elasticity at varying time horizons h as a function of the set of structural parameters $\Theta_i \equiv \{\theta_i, \sigma_i, \zeta_i\}$:

$$f_i^h(\Theta_i) \equiv \varepsilon_i^h = \frac{\partial \log X_{sdi,t+h}}{\partial \log \tau_{sdi,t}} = -\theta_i \left[1 - (1-\zeta_i)^{h+1}\right] + (1-\sigma_i)(1-\zeta_i)^{h+1}.$$

Our approach to recovering these structural involves, as a first step, obtaining reduced-form estimates of the trade elasticity over varying horizons. Such estimates can be obtained from the following specification using local projection methods:

$$\log\left(\frac{X_{sdi,t+h}}{X_{sdi,t-1}}\right) = \beta_i^h \log\left(\frac{\bar{\tau}_{sdi,t}}{\bar{\tau}_{sdi,t-1}}\right) + \delta_{si,t+h} + \delta_{di,t+h} + u_{sdi,t+h},$$

where $X_{sdi,t}$ denotes the exports of industry i goods from s to d at time t,and $t_{sdi,t}$ is the associated gross ad valorem tariff. The remaining terms denote source- or destination-industry-year-specific country fixed effects, and $u_{sdi,t}$ is an idiosyncratic error term. The coefficient β_X^h captures the change in trade flows h periods ahead that follows an initial one-period change in tariffs. Suppose that tariff changes were always one-time permanent shocks. Then a consistent estimate of β_i^h would yield an estimate of the structural trade elasticity at horizon h, ε_i^h. We now show how to recover the structural parameters governing the trade elasticity in our model, given a set of reduced-form estimates its behavior at varying time horizons h. With a slight abuse of notation, let $\{\hat{\beta}_i^h\}_{h=0}^H$ denote a set of such estimates ranging up to horizon $H > 0$.

Intuitively, the parameter σ_i governs the behavior of the trade elasticity in the short-run, while θ_i pins down its long-run value. The rate at which the trade elasticity converges to its long-run value, in turn, depends on how fast buyers form new supply relationships, ζ_i. More formally, we can use the structural expression for the trade elasticity to show that ζ_i, at any time $h > 0$, satisfies

$$\log(1 - \zeta_i) = \frac{1}{h} \log \left(\frac{f_i^H(\Theta) - \theta_i}{f_i^0(\Theta) - \theta_i} \right),$$

which captures the rate at which the process governing the trade elasticity converges to its long-run limit. Given a set of reduced-form estimates $\hat{\boldsymbol{\beta}}_i \equiv \{\hat{\beta}_i^h\}_{h=0}^H$, we recover our structural parameters by minimum distance:

$$\hat{\Theta}_i(\hat{\boldsymbol{\beta}}_i) = \arg\min_{\Theta} (f_i^h(\Theta) - \hat{\beta}_i^h)_{i \in \mathcal{I}}^T W \left(f_i^h(\Theta) - \hat{\beta}_i^h \right)_{i \in \mathcal{I}},$$

where W is a H-dimensional weighting matrix. Provided that the estimates of the trade elasticity are consistent, the continuous mapping theorem implies that $\hat{\Theta}_i(\hat{\boldsymbol{\beta}}_i)$ will provide a consistent estimate of Θ.

B. Implementation and Results

To implement our estimation approach, we leverage a set of comprehensive reduced-form estimates of the trade elasticity at different time horizons by Boehm, Levchenko and Pandalai-Nayar (2023). Following the reduced-form empirical approach outlined above, they find that arguably exogenous tariff changes in third countries predict a short-run trade elasticity that is substantially lower over shorter compared to longer horizons h, where $h = 0, 1, ..., 10$. To recover our set of

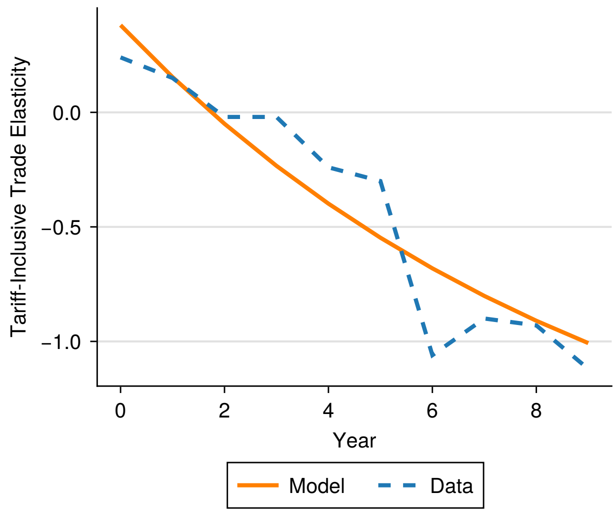

FIGURE 3.1: Horizon-Specific Trade Elasticity

TABLE 3.1: Trade Elasticity Parameter Estimates for the Manufacturing Industry

Parameter		Estimate
Supplier adjustment probability	ζ	0.10
Long-run Trade Elasticity	θ	1.89
Short-run Trade Elasticity	$\sigma - 1$	-0.63

structural parameters, we focus on matching the implied empirical behavior of the trade elasticity within the first two years, as well as at horizons $h = \{8, 9, 10\}$. Specifically, we set the weighting matrix W so that our estimator targets the response of trade flows to an initial change in tariffs Table 3.1 presents the results.

We find that supply relationships reset at an annual rate of about 9 percent, indicating substantial stickiness in supply relationships. The long-run trade elasticity across manufacturing industries, on average, equals 3.2, consistent with estimates in the literature on gravity. Our estimate of the elasticity of substitution equals 1.145, suggesting that trade elasticity, in the short-run, will be substantially lower, given the stickiness of supply relationships.

Figure 1 graphs the structural trade elasticity implied by these parameter estimates, along with the reduced-form elasticity estimates by Boehm, Levchenko and Pandalai-Nayar (2023). On impact ($h = 0$), the structural trade elasticity is close to zero. Over time, it smoothly increases in absolute value, reflecting the gradual resetting of supply relationship and reaching a level of -2.2

after 10 years. Reassuringly, the structural trade elasticity matches the behavior of its empirical counterpart also at horizons that were not explicitly targeted by our estimator.

IV. Quantitative Application: The 2018 US-China Trade War

We now apply our model to examine the general equilibrium responses of trade and production to the 2018 US-China trade war.

A. Calibration of the Initial Steady State

We assume that the world economy is in a steady state prior to the announcement of tariff changes due to the trade war. The remaining model parameters and initial levels of certain quantities are therefore calibrated so that trade activities implied by the model equilibrium in the absence of any shock match the data in 2017. For this purpose, we utilize the 2017 table from the 2023 edition of OECD Inter-Country Input-Output (ICIO) tables (OECD 2023). Table 3.2 summarizes the parameters and initial levels obtained for the calibration.

TABLE 3.2: Model Parameters and Variable Levels for Initial Steady State

Parameters or Initial Levels	Notation	Level of Variation
Matching Input-Output Data Exactly		
Producer expenditure shares across inputs	α_{sij}, α_{sj}	Producer region-sector
Initial import shares by source region	λ_{sdi}	User region
Initial level of bilateral trade flows	X_{sdi}	Sector-specific bilateral pair
Derived from Model Equilibrium		
Household expenditure shares across sectors	η_{dj}	User region
Initial aggregate labor income	$w_s L_s$	Producer region
Deficit (difference between expenditure and income)	D_d	User region

The ICIO table covers 45 industries in 76 economies along with the constructed rest of the world (ROW). In the model, we allow 77 economies corresponding to each of these in the data. For China (mainland) and Mexico, the data additionally record the input-output relations for a subset of manufacturing activities only intended for goods to be exported separately.[4] To take advantage of these additional details for China and Mexico in the model, we view these two economies as consisting

[4]These are available in the extended version of the ICIO tables for addressing heterogeneity of producers that do not directly deliver output in the domestic markets.

of two types of producers for each industry respectively. Namely, for each industry-specific good in these two economies, there is a set of regular producers delivering output for both domestic and foreign use; an additional set of producers produce special varieties that are only delivered abroad.[5] The technological parameters including those governing trade shares are allowed to be different across these two types of producers. However, the value added generated from all these producers are pooled together for computing the aggregate income in these two countries. Furthermore, labor inputs are assumed to be perfectly mobile across the two types of producers. These producers therefore face possibly different prices for intermediate inputs but identical wages.

Among the 45 industries, we exclude three of them that are primarily for public expenditure or services that are hard to classify. We further aggregate the remaining 42 industries into 32 sectors by combining certain non-manufacturing industries. A list of the sectors can be found in Appendix. Since we do not cover all industries in the ICIO table, the remaining data values no longer satisfy all restrictions imposed by accounting identities exactly.[6] For this reason, we need to take a stance on how we recover the identities. In other words, it is impossible to match the original ICIO table in every aspect. We choose certain dimensions of the data that we target exactly but use the model equilibrium conditions to derive those that cannot be targeted simultaneously.

We set the technological parameters $\{\{\alpha_{sij}\}_i, \alpha_{sj}\}_{sj}$ so that the expenditure shares across each production inputs match those in the data exactly.[7] We also match the initial import shares $\{\lambda_{sdi}\}_{sdi}$ exactly. Notice that these parameters already determine a complete input requirement matrix for the world economy. However, we still need to determine the relative levels of output across all region-sector pairs and there are alternative approaches. Since the exposure of each economy to the trade war depends on the initial levels of bilateral trade flows, we choose to target the levels of bilateral trade flows exactly, which include self trade.[8] From these bilateral trade flows, we immediately obtain the levels of output from each region-sector pair and the total expenditure on each sectoral good in each region.[9] From the levels of output and the technological parameters, we

[5]Depending on the calibration procedure, any (residual) domestic final use generated for output from the second type of producers are treated as arising only from exogenous deficit but not labor income.

[6]The output from a region-sector pair must be identical to the sum of intermediate or final use of the region-sector good around the world.

[7]The ICIO tables contain the margins for taxes or subsidies. These margins are treated as special expenditures that are not contributing to any part of the disposable income. For this reason, the sum of the expenditure shares across inputs are smaller than one.

[8]Alternatives include targeting the levels of sectoral final use, or sectoral value added, etc.

[9]Again, because the input-output relations no longer hold after excluding some industries, these values implied by

obtain the level of total expenditure on each intermediate input in each region. Subtracting these levels of intermediate use from total expenditures yield the levels of final use for each sector in each region that ensure all accounting identities hold. We set the household expenditure shares across goods from different sectors based on these derived final use. Lastly, with the region-sector specific value added shares, we compute the initial levels of aggregate labor income.[10] The discrepancy between total expenditure and total labor income in each country is treated as exogenous deficit that our model does not address.

B. Measuring the Tariff Changes

The tariff changes associated with the trade war are obtained from Fajgelbaum et al. (2020). Since the tariffs are determined at a detailed Harmonized System (HS) code level, we compute weighted averages of these tariff changes within each of the model sectors across different years. The tariff changes are aggregated both across different HS code and across months when they take into effect. For the aggregation across product categories, we determine the most relevant model sector based on their associated industry classifications and use the annual bilateral trade volume of each product in 2017 as weights. For the temporal aggregation across months, it has already been implemented by Fajgelbaum et al. (2020) using the shares of months within a year for which the tariff changes are in effect as weights. Table 3.3 collects the aggregated tariff changes on US imports from China along with the sectoral composition of US imports from China in 2017. Table 3.4 collects the aggregated retaliatory tariff changes on US exports to China. Notice that for model calibration, we have relied on the OECD ICIO table, which reconciles trade data with national accounts. However, for aggregating tariff changes, we require the 10-digit HS-code level data from US Census. It is therefore inevitable to see some discrepancies of the relative importance of sectoral imports or exports between the two types of data. Fortunately, for most of the sectors, the discrepancies seem to be small. For the tariff changes, we see that since many products were affected only after the second half of 2018, the aggregate changes at the annual level in 2018 are much smaller than those in 2019. By the end of 2019, all tariff changes associated with the trade war had been in place.

the subset of bilateral trade flows can be different from the original values in the ICIO table which cover all industries.

[10] Without dealing with other factor income, we abstract away from heterogeneity in the labor income shares within the value added components.

TABLE 3.3: US Tariff Increases on Imports from China

	2017 Imports in Total (%)		Cumulative Increases in Tariffs (%)		
Affected Sector in Model	OECD ICIO	US Census	2018	2019	2020–
Agriculture, forestry and fishing	0.5	0.6	2.5	14.7	20.6
Mining and quarrying	0.0	0.1	1.0	5.6	7.4
Food products, beverages and tobacco	1.9	0.8	2.6	15.5	22.3
Textiles, textile products, leather and footwear	17.7	12.8	0.6	6.6	13.8
Wood and products of wood and cork	1.3	0.8	2.9	16.4	22.1
Paper products and printing	1.2	1.3	2.1	11.0	15.8
Coke and refined petroleum products	0.2	0.1	2.4	14.2	20.5
Chemical and chemical products	3.2	3.1	2.7	12.7	17.7
Pharmaceuticals, medicinal and botanical products	1.3	0.5	0.0	0.1	0.1
Rubber and plastics products	3.7	3.6	2.2	10.9	15.1
Other non-metallic mineral products	2.4	1.7	2.1	12.3	17.4
Basic metals	1.0	0.9	8.8	22.4	24.5
Fabricated metal products	3.8	4.1	3.4	15.0	20.0
Computer, electronic and optical equipment	29.0	36.3	2.0	8.1	11.2
Electrical equipment	9.1	8.9	3.9	14.9	18.8
Machinery and equipment, nec	6.9	7.3	6.1	18.3	22.3
Motor vehicles, trailers and semi-trailers	4.2	3.2	4.7	19.3	24.7
Other transport equipment	0.8	0.7	7.2	20.5	24.0
Furniture and other manufacturing	11.7	13.3	1.1	7.1	11.0

Notes: "Imports in Total" are the shares of sectoral US imports in total imports from China. "OECD ICIO" refers to the input-output table used for calibrating the model. "US Census" refers to the HS-level bilateral trade data accessed via USA Trade Online. The tariff changes are aggregated based on weights derived from the US Census data. Tariff changes are obtained from Fajgelbaum et al. (2020).

TABLE 3.4: Retaliatory Tariff Increases on US Exports to China

	2017 Exports in Total (%)		Cumulative Increases in Tariffs (%)		
Affected Sector in Model	OECD ICIO	US Census	2018	2019	2020–
Agriculture, forestry and fishing	14.7	15.1	11.9	31.1	31.3
Mining and quarrying	7.7	7.0	3.5	11.2	14.0
Food products, beverages and tobacco	4.2	2.8	10.4	19.9	21.0
Textiles, textile products, leather and footwear	0.4	0.9	2.5	12.1	15.3
Wood and products of wood and cork	1.5	1.5	2.7	12.9	16.3
Paper products and printing	2.0	2.5	2.1	7.7	8.8
Coke and refined petroleum products	2.6	1.0	10.2	26.0	26.0
Chemical and chemical products	10.6	9.6	3.7	12.5	14.3
Pharmaceuticals, medicinal and botanical products	2.5	2.8	0.3	1.6	2.7
Rubber and plastics products	1.4	1.3	2.3	10.0	12.4
Other non-metallic mineral products	0.6	0.8	4.0	13.7	15.8
Basic metals	10.7	1.9	4.1	15.4	18.9
Fabricated metal products	1.2	1.4	2.5	10.9	13.3
Computer, electronic and optical equipment	10.4	13.8	2.4	9.3	11.2
Electrical equipment	1.4	2.5	3.8	15.8	19.4
Machinery and equipment, nec	6.0	8.0	2.1	9.1	11.1
Motor vehicles, trailers and semi-trailers	8.6	11.0	10.5	21.5	21.7
Other transport equipment	12.4	13.3	0.0	0.1	0.1
Furniture and other manufacturing	1.1	2.8	4.1	12.8	14.2

Notes: The same notes for Table 3.3 apply.

C. The General Equilibrium Impact of the Trade War

We are interested in how the 2018 US-China trade war had affected the aggregate economic activities around the world and how its welfare impact had evolved over time. To that end, we conduct a general equilibrium counterfactual experiment in which we compute how the model outcomes evolve following the tariff changes we measure in Section B relative to a hypothetical scenario in which the trade war had not happened.

Trade Flows Figure 3.2 plots the changes in the tariff-inclusive US (China) imports from China (US) among the 19 sectors listed in Tables 3.3 and 3.4 that are directly affected by the trade war. With higher tariff payments and low short-run trade elasticity, the bilateral trade flows increase, as predicted by both the full GE model outcomes and the PE trade elasticity. The bilateral trade flows start to drop below the initial levels only since year 4 for US and year 2 for China. As time goes, the trade flows keep declining as the relevant trade elasticity shifts towards the long-run level. Notice that for US, the trade flows fall by less than what the structural trade elasticity predict due to the changes in factor prices. The discrepancies between what the full GE model predicts and what the PE trade elasticity predicts for US demonstrate the need of taking into account the GE effects.

Prices The sluggish short-run response of US demand to the rise in trade costs induces a substantial rise in its domestic price level. As shown in Figure 3.3, aggregate price indices faced by US consumers and producers rise across all industries. Some industries, notably textiles, basic metals, and electrical equipment, see prices rise by over 4% as all retaliatory tariffs are in place. As sourcing decisions gradually adjust to the initial rise in trade cost, prices decline although they remain high. In contrast to the substantial and uneven price hikes in US, domestic prices in China decline across all industries. Intuitively, a rise in trade barriers can temporarily improve a region's terms-of-trade when trade adjustment is not primarily driven by comparative advantage. Figure 3.4 shows the price impact on the remaining sectors that are not directly exposed to the tariff changes.

Real Wages and Welfare Figure 3.5 traces the counterfactual response of real wages, as well as nominal wages and consumer prices in the US and China. In the long run, the trade war reduces real income in both countries by a similar magnitude. However, its short-run impact differs substantially

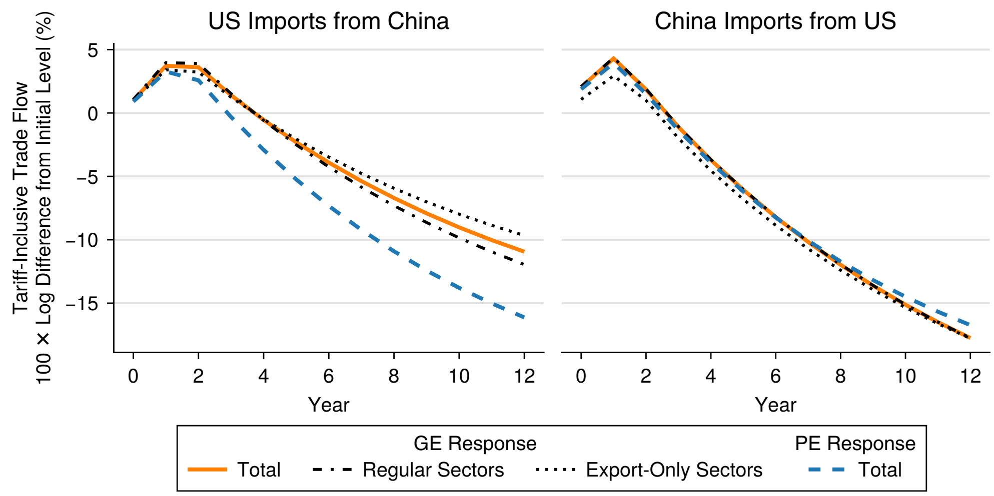

FIGURE 3.2: Changes in Tariff-Inclusive Trade Flows

Note: Tariff changes are gradually implemented over the first two years. The model determines changes in trade shares at the sector level. The country-level outcomes are based on aggregate trade flows summed across sectors. "GE Response" and "PE Response" refer to results generated from the full model involving factor price changes and results only based on the PE trade elasticity respectively. "Regular Sectors" and "Export-Only Sectors" are only relevant to China due to the feature of ICIO tables explained in Section A.

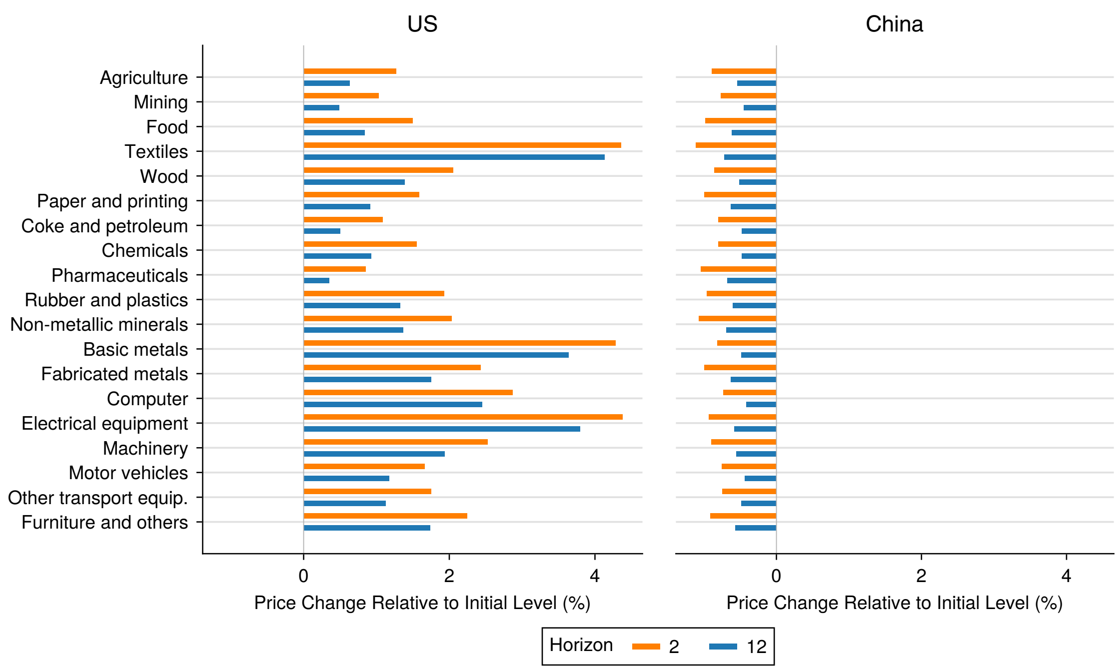

FIGURE 3.3: Changes in Price Indices Among Directly Affected Sectors

between the US and China. In the US, the real wage responds gradually, with a moderate decline within the first two years of the trade war (-0.1%) that corresponds to about 50% of the overall

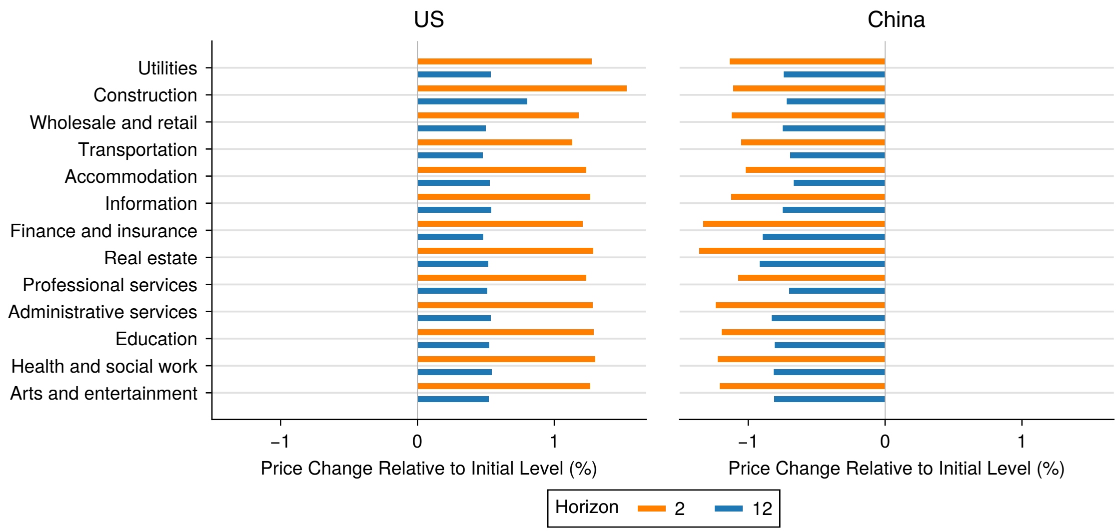

FIGURE 3.4: Changes in Price Indices Among Sectors Not Directly Affected

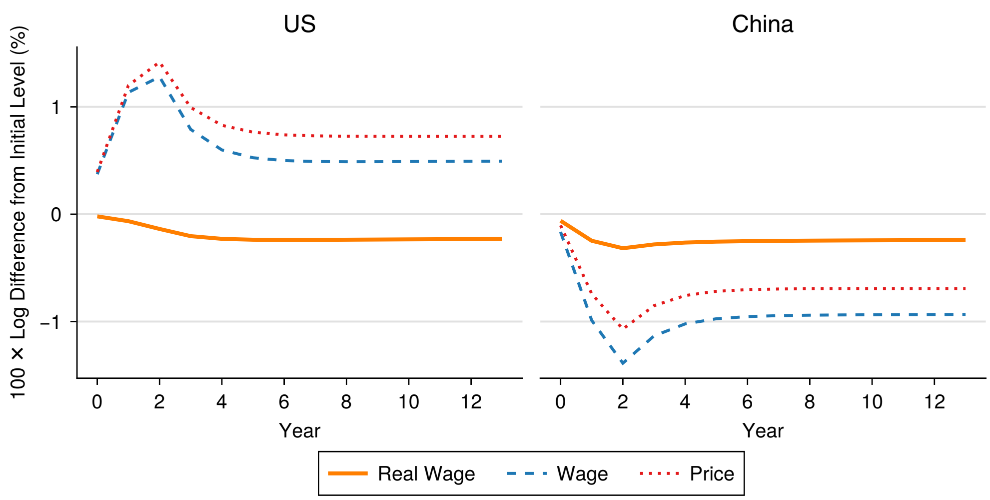

FIGURE 3.5: Changes in Real Wages, Wages and Consumer Prices

Note: Tariff changes are gradually implemented over the first two years. "Real Wage" in year t refers to real wage changes between t and the initial steady state generated by the full model. "Wage" refers to the corresponding change in nominal wage. "Price" is the change in aggregate consumer price index.

effect (-0.22%). In contrast, while the long-run costs of the trade war in China are similar to those in the US, China also experiences a substantially larger decline in real income by the end of the second year (-0.3%).

In Figure 3.6, we leverage the ACR-style welfare formula shown in Proposition 3 to elucidate how the presence of adjustment frictions alter the transitory dynamics of real wages. Interestingly, if one applies a multisector version of the original welfare formula from Arkolakis, Costinot and

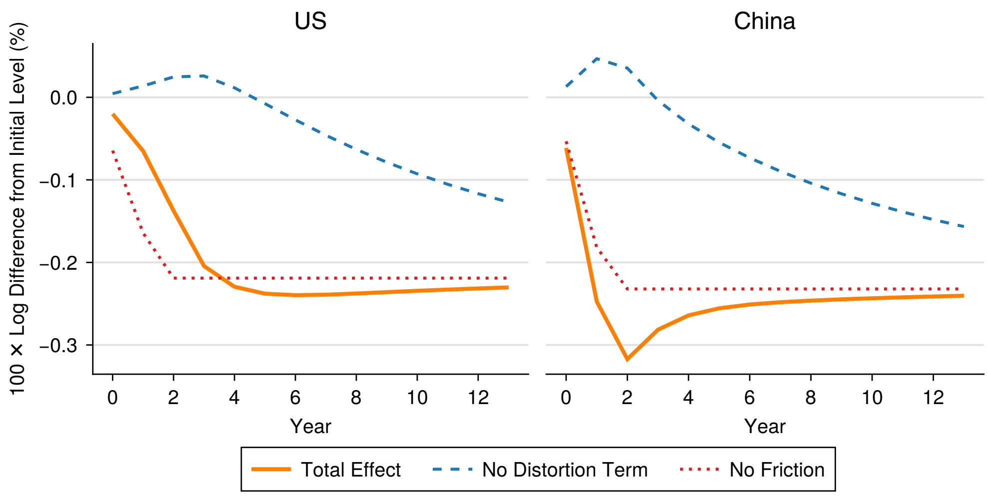

FIGURE 3.6: Horizon-Specific Welfare Impact

Note: Tariff changes are gradually implemented over the first two years. "Total Effect" in year t refers to the real wage changes between t and the initial steady state generated by the full model. "No Distortion Term" refers to the (partial) welfare impact when applying the ACR formula to the model-implied domestic trade shares while ignoring the distortion term Ξ. "No Friction" refers to the change in real wages when assuming that the economy reaches the long-run outcomes instantly ($\zeta = 1$).

Rodríguez-Clare (2012) to the economy with adjustment frictions, the results will be very misleading, as illustrated by the curve labeled as "No Distortion Term". The reason is that in addition to using the less appropriate long-run trade elasticity, the observed changes in aggregate domestic trade shares go in the opposite direction relative to the actual changes in real wages over the short run. With substantial adjustment frictions among trade partners, the changes in aggregate domestic trade shares over the initial years are not in line with the long-run Ricardian forces that govern the original welfare formula. In fact, the distortion term Ξ highlighted in Proposition 3 is quantitatively substantial and drives most of the short-run welfare impact.

As another illustration, we compare how the paths of welfare impact differ from the hypothetical scenario in which there is no adjustment friction ($\zeta = 1$). In this case, the economy jumps to the long-run outcomes instantly.[11] For the US, we see the model predicts short-run welfare impact that is smaller than the long-run outcome over the initial years. However, for China, the short-run welfare impact is noticeably larger than the long-run level. In particular, the much lower trade elasticity in the short run does not mechanically imply larger welfare impact over the short

[11]The smaller impact over the first two years are merely from the fact that the tariff changes are not fully implemented until the end of 2019.

run.[12] The asymmetry of the responses of real income over time illustrates that sourcing frictions may mitigate or amplify the costs of trade disruptions in the short run. Note that over the initial years, prices and wages rise by a similar magnitude in the US; while in China, domestic wages fall substantially more than consumer prices. The intuition is that sluggishness in the response of trade flows helps smooth the transition for the US: It benefits not only from additional tariff revenues generated by the fact that producers continue to import goods from China but also from a limited response of its export demand to the rise in its export prices. In contrast, adjustment frictions pose additional short-run costs for China as they impede its ability to leverage the decline in domestic wages to increase exports, while only generating limited additional tariff revenue (due to the fact that it imports relatively little from the US to begin with). The gradual realignment of trade flows with comparative advantage over time therefore ameliorates the welfare loss for the US but exacerbates the real impact in China. Moving toward the future horizons, the welfare impact gets closer to the long-run levels while being slightly lower for both countries, due to the persistent effects of the distortions on the prices and allocations.

Effects on Third-Party Countries We conclude by highlighting how accounting for short-run adjustment frictions affects the welfare implications of the US-China trade war for third-party countries. In Figures 3.7 and 3.8, we present the counterfactual responses of real income in Mexico and Vietnam. Notably, both countries experience benefits from the tariff increases in the long run, while also facing losses in the short run. For Mexico, this short-run income loss ranks among the largest for all third-party countries; however, it distinguishes itself as one of the few countries that benefit from the trade war in the long run.

Therefore, the welfare impact of trade disruptions can qualitatively differ over time. Intuitively, when trade adjustments are subject to frictions, disruptions negatively affect all countries in the short run. In the long run, however, realignments of supply relationships may benefit some countries. In the context of the US-China trade war, both Mexico and Vietnam experience a sustained increase in their domestic wages, reflecting both the reallocation of US and Chinese demand as well as their favorable positions in the international production network.

[12] Again, with the adjustment friction, changes in aggregate domestic shares alone are not sufficient for accounting the welfare impact.

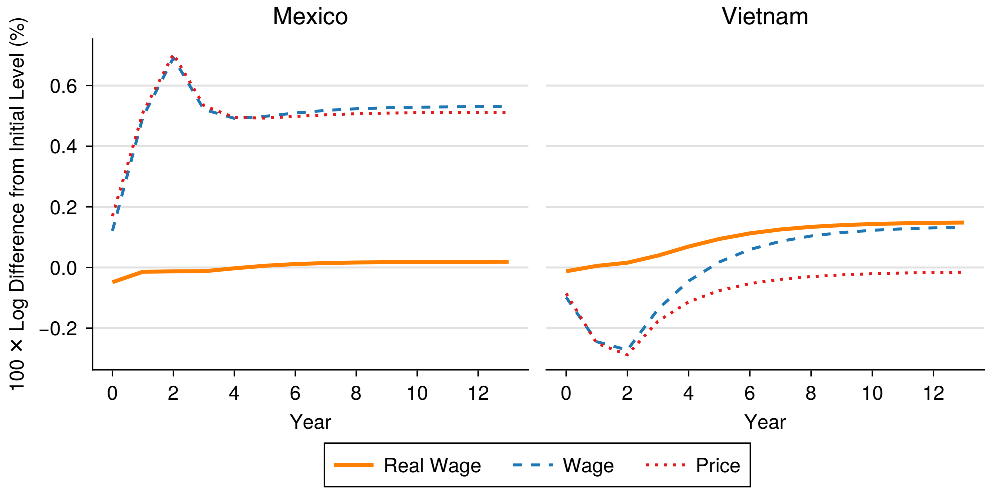

FIGURE 3.7: Changes in Prices and Wages in Mexico and Vietnam

Note: The same notes for Figure 3.5 applies.

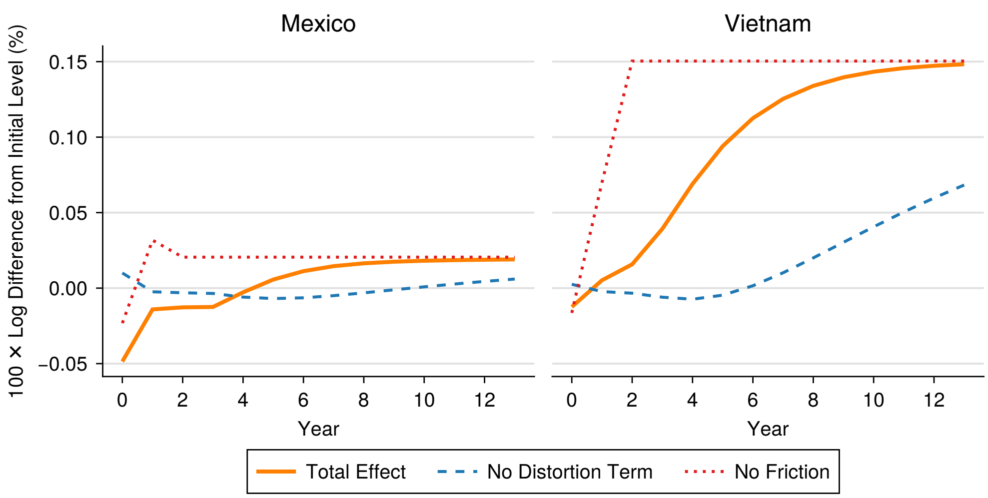

FIGURE 3.8: Welfare Impact in Mexico and Vietnam

Note: The same notes for Figure 3.6 applies.

V. Concluding Remarks

To account for imperfect adjustment to global supply chain shocks, we develop a Ricardian trade framework with frictions that result from infrequent decisions of producers to change global suppliers. We obtain novel formulas for accounting welfare changes to trade openness and trade shocks, derive novel estimation equations for trade elasticity estimation at varying time horizons, and quantify the model. Counterfactual experiments of the US-China trade war suggests that rich sectoral dynamics ensue, resulting in considerable short-term reallocations and substantive welfare fluctuations that are not captured by a standard welfare formula as in Arkolakis, Costinot and Rodríguez-Clare (2012).

Acknowledgements

Chapter 3, in full, is currently being prepared for submission for publication of the material. Chapter 3 is coauthored with Carlos Góes, Marc-Andreas Muendler, and Fabian Trottner. The dissertation author was one of the primary authors of this paper.

Appendix 3.A. Additional Details on Model and Proofs

A. Ideal Price Indexes and Generic Trade Shares

The composite good in industry j is

$$Y_{dj,t} \equiv \left(\int_{[0,1]} y_{dj,t}(\bar{\omega})^{\frac{\sigma_j - 1}{\sigma_j}} \mathrm{d}\bar{\omega} \right)^{\frac{\sigma_j}{\sigma_j - 1}}.$$

Product space $\Omega_j = [0,1]$ can be partitioned into disjoint sets with $\Omega_j = \bigcup_{k=0}^{\infty} \Omega_{j,t}^k$, so we can rewrite the composite good as

$$Y_{dj,t} \equiv \left(\sum_{k=0}^{\infty} \int_{\Omega_{j,t}^k} y_{dj,t}(\bar{\omega})^{\frac{\sigma_j - 1}{\sigma_j}} \mathrm{d}\bar{\omega} \right)^{\frac{\sigma_j}{\sigma_j - 1}}.$$

The assembler's associated cost minimization problem is

$$\begin{aligned}
\min_{\{y_{dj,t}(\bar{\omega})\}_{\bar{\omega} \in \Omega_{j,t}}, \{Y_{dj,t}^k\}} P_{dj,t} Y_{dj,t} \quad &= \quad \sum_{k=0}^{\infty} P_{dj,t}^k Y_{dj,t}^k \\
s.t. \quad & \quad Y_{dj,t} = \left(\sum_{k=0}^{\infty} \left(Y_{dj,t}^k \right)^{\frac{\sigma_j - 1}{\sigma_j}} \right)^{\frac{\sigma_j}{\sigma_j - 1}}, Y_{dj,t}^k \equiv \left(\int_{\Omega_{j,t}^k} y_{dj,t}(\bar{\omega})^{\frac{\sigma_j - 1}{\sigma_j}} \mathrm{d}\bar{\omega} \right)^{\frac{\sigma_j}{\sigma_j - 1}}, \\
& \quad P_{dj,t}^k Y_{dj,t}^k = \int_{\Omega_{j,t}^k} p_{dj,t}(\bar{\omega}) y_{dj,t}(\bar{\omega}) \mathrm{d}\bar{\omega},
\end{aligned}$$

where we define the partial composite good $Y_{dj,t}^k \equiv \left(\int_{\Omega_{j,t}^k} y_{dj,t}(\bar{\omega})^{\frac{\sigma_j - 1}{\sigma_j}} \mathrm{d}\bar{\omega} \right)^{\frac{\sigma_j}{\sigma_j - 1}}$ for each partition k as a helpful construct for derivations and implicity define the associated partial ideal price index $P_{dj,t}^k$ that satisfies $P_{dj,t}^k Y_{dj,t}^k = \int_{\Omega_{j,t}^k} p_{dj,t}(\bar{\omega}) y_{dj,t}(\bar{\omega}) \mathrm{d}\bar{\omega}$.

Under homotheticity of the assembler's production, this problem can be solved in two steps. First, the assembler decides which share of cost it allocates to each partial composite good $Y_{dj,t}^k$. Given those choices, the assembler then decides the optimal cost for each intermediate good $y_{dj,t}(\bar{\omega})$. Optimal demand satisfies

$$Y_{dj,t}^k = \left(\frac{P_{dj,t}^k}{P_{dj,t}} \right)^{-\sigma_j} Y_{dj,t} \quad \text{and} \tag{3.15}$$

$$y_{dj,t}^k(\bar{\omega}) = \left(\frac{p_{dj,t}(\bar{\omega})}{P_{dj,t}^k} \right)^{-\sigma_j} Y_{dj,t}^k = \left(\frac{p_{dj,t}(\bar{\omega})}{P_{dj,t}} \right)^{-\sigma_j} Y_{dj,t} \quad \text{for each } \bar{\omega} \in \Omega_{j,t}^k, \tag{3.16}$$

where the last equality also shows that the partitioned solution equals the standard solution under a constant elasticity of substitution. Replacing the demand functions above in the definition of the budget constraint results in the expressions for the ideal price indices:

$$P_{dj,t} = \left(\int_{[0,1]} p_{dj,t}(\bar{\omega})^{1-\sigma_j} \mathrm{d}\bar{\omega} \right)^{\frac{1}{1-\sigma_j}}, \qquad P^k_{dj,t} = \left(\int_{\Omega^k_{j,t}} p_{dj,t}(\bar{\omega})^{1-\sigma_j} \mathrm{d}\bar{\omega} \right)^{\frac{1}{1-\sigma_j}}.$$

We have now established that partitioning the product space into disjoint sets results in well-behaved demand functions such that, given optimal choices within each set, we can analyze demand for each intermediate good independently and then aggregate. In subsequent derivations, expenditure shares within each partition k will play a crucial role, so we state a general definition here:

$$\begin{aligned} \lambda^k_{sdj,t} \equiv \frac{X^k_{sdj,t}}{X^k_{dj,t}} &\equiv \frac{\int_{\Omega^k_{j,t}} \mathbf{1}\left\{s \text{ is } \omega\text{'s source country}\right\} p_{dj,t}(\omega) y_{dj,t}(\omega)\, \mathrm{d}\omega}{\int_{\Omega^k_{j,t}} p_{dj,t}(\omega) y_{dj,t}(\omega)\, \mathrm{d}\omega} \\ &= \frac{\int_{\Omega^k_{j,t}} \mathbf{1}\left\{s \text{ is } \omega\text{'s source country}\right\} p_{dj,t}(\omega) y_{dj,t}(\omega)\, \mathrm{d}\omega}{\sum_n \int_{\Omega^k_{j,t}} \mathbf{1}\left\{n \text{ is } \omega\text{'s source country}\right\} p_{dj,t}(\omega) y_{dj,t}(\omega)\, \mathrm{d}\omega}. \end{aligned} \tag{3.17}$$

B. Trade Shares When Firms Are Sourcing Optimally ($k = 0$)

Under perfect competition, the destination price for intermediate good $\omega \in \Omega^0_{j,t}$ offered by country s to country d is $p_{sdj,t}(\omega) = c_{sdj,t}/z_{sj}(\omega)$ for the common unit cost component $c_{sdj,t}$ by (3.1) and supplier ω's productivity $z_{si}(\omega)$. Under the EK assumptions, the cumulative distribution function of prices is therefore

$$\tilde{F}_{sdj,t}(p) = \mathbb{P}\left[p_{sdj,t}(\omega) < p\right] = 1 - F_{sj}\left(\frac{c_{sdj,t}}{p}\right) = 1 - \exp\left\{-A_{sj}(c_{sdj,t})^{-\theta_j} p^{\theta_j}\right\}.$$

The resulting probability that country d sources an intermediate good $\omega \in \Omega^0_{j,t}$ from country s is

$$\mathbb{P}\left[s = \arg\min_n \left\{p_{ndj,t}(\omega)\right\}\right] = \int_0^\infty \prod_{n \neq s} \left[1 - \tilde{F}_{ndj,t}(p)\right] \mathrm{d}\tilde{F}_{sdj,t}(p) = \frac{A_{sj}(c_{sdj,t})^{-\theta_j}}{\Phi_{dj,t}},$$

where $\Phi_{dj,t} \equiv \sum_n A_{sj}(c_{sdj,t})^{-\theta_j}$.

For products in $\Omega^0_{j,t}$, the distribution of prices $G^0_{sdj,t}(p)$ paid in country d on products sourced from country s equals the overall distribution of prices paid in country d: $G^0_{dj,t}(p)$. For any given source country s:

$$G^0_{sdj,t}(p) = \mathbb{P}\left[p_{dj,t}(\omega) \leqslant p | s = \arg\min_n \left\{p_{ndj,t}(\omega)\right\}\right] = 1 - \exp\left\{-\Phi_{dj,t} p^{\theta_j}\right\}.$$

The unconditional distribution is the same as the distribution conditional on each source country, so

$$\begin{aligned} G^0_{dj,t}(p) &= \sum_s \mathbb{P}\left[p_{dj,t}(\omega) \leqslant p | s = \arg\min_n \left\{p_{ndj,t}(\omega)\right\}\right] \mathbb{P}\left[s = \arg\min_n \left\{p_{ndj,t}(\omega)\right\}\right] \\ &= \sum_s \left(1 - \exp\left\{-\Phi_{dj,t} p^{\theta_j}\right\}\right) \lambda^0_{sdj,t} = 1 - \exp\left\{-\Phi_{dj,t} p^{\theta_j}\right\}, \end{aligned} \tag{3.18}$$

where the last equality follows from the fact that $\sum_s \lambda^0_{sdj,t} = 1$.

Putting these results together, we can now solve for the expenditure share within partition 0. Starting from the definition of expenditure shares,

$$\begin{aligned} \lambda^0_{sdj,t} &\equiv \frac{\int_{\Omega^0_{j,t}} \mathbf{1}\left\{s = \arg\min_m \left\{p_{mdj,t}(\omega)\right\}\right\} \left(p_{sdj,t}(\omega)\right)^{1-\sigma_j} \mathrm{d}\omega}{\sum_n \int_{\Omega^0_{j,t}} \mathbf{1}\left\{n = \arg\min_m \left\{p_{mdj,t}(\omega)\right\}\right\} \left(p_{ndj,t}(\omega)\right)^{1-\sigma_j} \mathrm{d}\omega} \\ &= \frac{\int_{\Omega^0_{j,t}} \mathbf{1}\left\{s = \arg\min_m \left\{p_{mdj,t}(\omega)\right\}\right\} \int_0^\infty (p)^{1-\sigma_j} \mathrm{d}G_{sdj,t} \mathrm{d}\omega}{\sum_n \int_{\Omega^0_{j,t}} \mathbf{1}\left\{n = \arg\min_m \left\{p_{mdj,t}(\omega)\right\}\right\} \int_0^\infty (p)^{1-\sigma_j} \mathrm{d}G_{ndj,t} \mathrm{d}\omega} \\ &= \frac{\int_{\Omega^0_{j,t}} \mathbf{1}\left\{s = \arg\min_m \left\{p_{mdj,t}(\omega)\right\}\right\} \mathrm{d}\omega \int_0^\infty (p)^{1-\sigma_j} \mathrm{d}G_{dj,t}}{\sum_n \int_{\Omega^0_{j,t}} \mathbf{1}\left\{n = \arg\min_m \left\{p_{mdj,t}(\omega)\right\}\right\} \mathrm{d}\omega \int_0^\infty (p)^{1-\sigma_j} \mathrm{d}G_{dj,t}} \\ &= \frac{\int_{\Omega^0_{j,t}} \mathbf{1}\left\{s = \arg\min_m \left\{p_{mdj,t}(\omega)\right\}\right\} \mathrm{d}\omega}{\int_{[0,1]} \mathbf{1}\left\{\omega \in \Omega^0_{j,t}\right\} \mathrm{d}\omega} \\ &= \frac{\mu_{j,t}(0) \mathbb{P}\left[s = \arg\min_m \left\{p_{mdj,t}(\omega)\right\}\right]}{\mu_{j,t}(0)} \\ &= \frac{A_{sj}(c_{sdj,t})^{-\theta_j}}{\Phi_{dj,t}}, \end{aligned} \tag{3.19}$$

where $\mu_{i,t}(0)$ is the measure of the set $\Omega^0_{i,t}$. The third line uses the fact again that the distribution of prices conditional on the source country is the same as the unconditional distribution of prices, and the last equality uses the probability that a given source country hosts the lowest-cost supplier.

We can derive the corresponding ideal price indices using

$$\begin{aligned}\left(P^0_{dj,t}\right)^{1-\sigma_j} &= \int_{\Omega^0_{j,t}} p_{dj,t}(\bar{\omega})^{1-\sigma_j}\,\mathrm{d}\bar{\omega} = \int_{\Omega^*_{j,t}0}\int_0^\infty (p)^{1-\sigma_j}\,\mathrm{d}G_{dj,t}\,\mathrm{d}\bar{\omega} \\ &= \int_{\Omega^0_{j,t}}\int_0^\infty (p)^{1-\sigma_j}\,\theta_j\Phi_{dj,t}p^{\theta_j-1}\exp\left\{-\Phi_{dj,t}p^{\theta_j}\right\}\,\mathrm{d}p\,\mathrm{d}\bar{\omega}.\end{aligned}$$

For a change of variables, define $x \equiv p_j^\theta\Phi_{dj,t}$, which implies that $\mathrm{d}x = \theta_j\Phi_{dj,t}p^{\theta_j-1}\mathrm{d}p$ and $p = (x/\Phi_{dj,t})^{1/\theta_j}$. Denoting $\gamma_j \equiv \Gamma\left([\theta_j + 1 - \sigma_j]/\theta_j\right)$, we can then rewrite the integral above as

$$\left(P^0_{dj,t}\right)^{1-\sigma_j} = \int_{\Omega^0_{j,t}}\int_0^\infty \left(\frac{x}{\Phi_{dj,t}}\right)^{\frac{1-\sigma_j}{\theta_j}}\exp\{-x\}\,\mathrm{d}x\,\mathrm{d}\bar{\omega} = \gamma_j\,\mu_{j,t}(0)\left(\Phi_{dj,t}\right)^{-\frac{1-\sigma_j}{\theta_j}}, \tag{3.20}$$

$\mu_{j,t}(0)$ denotes the measure of the set $\Omega^0_{j,t}$. The results show that, when firms are adjusting, trade shares operate as in the frictionless economy of EK.

Using standard hat algebra for changes in the common unit cost component $\hat{c}_{sdj,t} \equiv c_{sdj,t}/c_{sdj,t-1}$, we can express trade shares and price levels within partition $k = 0$ as:

$$\lambda^0_{sdj,t} = \frac{\lambda^0_{sdj,t-1}\hat{c}^{-\theta_j}_{sdj,t}}{\sum_n \lambda^0_{ndj,t-1}(\hat{c}_{ndj,t})^{-\theta_j}} \tag{3.21}$$

$$P^0_{dj,t} = P^0_{dj,t-1}\left[\sum_s \lambda^0_{sdj,t-1}(\hat{c}_{sdj,t})^{-\theta_j}\right]^{-\frac{1}{\theta_j}}. \tag{3.22}$$

We next derive an analogous result for partitions $k > 0$ when firms are not adjusting their extensive margin of suppliers.

C. Trade Shares When Firms Are Not Adjusting ($k > 0$)

For intermediate goods $\omega \in \Omega^k_{j,t}$, assemblers last adjusted the least-cost supplier $t - k$ periods ago. In order to account for changes in trade shares and price levels, we therefore need to recall optimal sourcing choices at period $t - k$ and trace changes in parameters and prices since $t - k$.

Suppose that in period $t - k$ intermediate good ω was optimally sourced from country s to country d in industry j. Then the destination price in period t for this intermediate good will be:

$$p_{sdj,t}(\omega) = \frac{c_{sdj,t}}{z_{sj}(\omega)} = \frac{\prod_{\varsigma=t-k+1}^t c_{sdj,t-k}\,\hat{c}_{sdj,\varsigma}}{z_{sj}(\omega)} = p_{sdj,t-k}(\omega)\prod_{\varsigma=t-k+1}^{t}\left(\hat{c}_{sdj,\varsigma}\right),$$

which is the initial destination price adjusted for the cumulative changes in trade costs and factor costs. Using this result, we can derive country d's expenditure share by source country across intermediate goods $\omega \in \Omega^k_{j,t}$

$$
\begin{aligned}
\lambda^k_{sdj,t} &\equiv \frac{\int_{\Omega^k_{j,t}} \mathbf{1}\left\{s = \arg\min_m \left\{p_{mdj,t-k}(\omega)\right\}\right\} \left(p_{sdj,t-k}(\omega) \prod_{\varsigma=t-k+1}^{t} \hat{c}_{sdj,\varsigma}\right)^{1-\sigma_j} \mathrm{d}\omega}{\sum_n \int_{\Omega^k_{j,t}} \mathbf{1}\left\{n = \arg\min_m \left\{p_{mdj,t-k}(\omega)\right\}\right\} \left(p_{ndj,t-k}(\omega) \prod_{\varsigma=t-k+1}^{t} \hat{c}_{ndj,\varsigma}\right)^{1-\sigma_j} \mathrm{d}\omega} \\
&= \frac{\int_{\Omega^k_{j,t}} \mathbf{1}\left\{s = \arg\min_m \left\{p_{mdj,t-k}(\omega)\right\}\right\} \int_0^\infty (p)^{1-\sigma_j}\, \mathrm{d}G_{sdj,t-k} \mathrm{d}\omega \left(\prod_{\varsigma=t-k+1}^{t} \hat{c}_{sdj,\varsigma}\right)^{1-\sigma_j}}{\sum_n \int_{\Omega^k_{j,t}} \mathbf{1}\left\{n = \arg\min_m \left\{p_{mdj,t-k}(\omega)\right\}\right\} \int_0^\infty (p)^{1-\sigma_j}\, \mathrm{d}G_{ndj,t-k} \mathrm{d}\omega \left(\prod_{\varsigma=t-k+1}^{t} \hat{c}_{ndj,\varsigma}\right)^{1-\sigma_j}} \\
&= \frac{\int_{\Omega^k_{j,t}} \mathbf{1}\left\{s = \arg\min_m \left\{p_{mdj,t-k}(\omega)\right\}\right\} \mathrm{d}\omega \int_0^\infty (p)^{1-\sigma_j}\, \mathrm{d}G_{dj,t-k} \left(\prod_{\varsigma=t-k+1}^{t} \hat{c}_{sdj,\varsigma}\right)^{1-\sigma_j}}{\sum_n \int_{\Omega^k_{j,t}} \mathbf{1}\left\{n = \arg\min_m \left\{p_{mdj,t-k}(\omega)\right\}\right\} \mathrm{d}\omega \int_0^\infty (p)^{1-\sigma_j}\, \mathrm{d}G_{dj,t-k} \left(\prod_{\varsigma=t-k+1}^{t} \hat{c}_{ndj,\varsigma}\right)^{1-\sigma_j}} \\
&= \frac{\int_{\Omega^k_{j,t}} \mathbf{1}\left\{s = \arg\min_m \left\{p_{mdj,t-k}(\omega)\right\}\right\} \mathrm{d}\omega \left(\prod_{\varsigma=t-k+1}^{t} \hat{c}_{sdj,\varsigma}\right)^{1-\sigma_j}}{\sum_n \int_{\Omega^k_{j,t}} \mathbf{1}\left\{n = \arg\min_m \left\{p_{mdj,t-k}(\omega)\right\}\right\} \mathrm{d}\omega \left(\prod_{\varsigma=t-k+1}^{t} \hat{c}_{ndj,\varsigma}\right)^{1-\sigma_j}} \\
&= \frac{\mu_{j,t}(k)\lambda_{sdj,t-k} \left(\prod_{\varsigma=t-k+1}^{t} \hat{c}_{sdj,\varsigma}\right)^{1-\sigma_j}}{\sum_n \mu_{j,t}(k)\lambda_{ndj,t-k} \left(\prod_{\varsigma=t-k+1}^{t} \hat{c}_{ndj,\varsigma}\right)^{1-\sigma_j}} \\
&= \frac{\lambda^0_{sdj,t-k} \left(\prod_{\varsigma=t-k+1}^{t} \hat{c}_{sdj,\varsigma}\right)^{1-\sigma_j}}{\sum_n \lambda^0_{ndj,t-k} \left(\prod_{\varsigma=t-k+1}^{t} \hat{c}_{ndj,\varsigma}\right)^{1-\sigma_j}}, \qquad (3.23)
\end{aligned}
$$

where $\mu_{i,t}(k)$ is the measure of the set $\Omega^k_{i,t}$. The third line again uses the fact that, at $t-k$, the distribution of prices conditional on the source is the same as the unconditional distribution; and the last line uses the result from the previous section that $\lambda^0_{sdj,t-k} = \mathbb{P}\left[s = \arg\min_s \left\{p_{sdj,t-k}(\omega)\right\}\right]$.

We can derive the corresponding ideal price indices using

$$
\begin{aligned}
\left(P^k_{dj,t}\right)^{1-\sigma_j} &= \int_{\Omega^k_{j,t}} p_{dj,t}(\bar{\omega})^{1-\sigma_j} \mathrm{d}\bar{\omega} \\
&= \sum_s \int_{\Omega^k_{j,t}} \mathbf{1}\left\{s = \arg\min_m \left\{p_{mdj,t-k}(\omega)\right\}\right\} \left(p_{sdj,t-k}(\omega) \prod_{\varsigma=t-k+1}^{t} \hat{c}_{sdj,\varsigma}\right)^{1-\sigma_j} \mathrm{d}\omega \\
&= \sum_s \int_{\Omega^k_{j,t}} \mathbf{1}\left\{s = \arg\min_m \left\{p_{mdj,t-k}(\omega)\right\}\right\} \int_0^\infty (p)^{1-\sigma_j}\, \mathrm{d}G_{sdj,t-k} \mathrm{d}\omega \left(\prod_{\varsigma=t-k+1}^{t} \hat{c}_{sdj,\varsigma}\right)^{1-\sigma_j} \\
&= \int_0^\infty (p)^{1-\sigma_j}\, \mathrm{d}G_{dj,t-k} \sum_s \int_{\Omega^k_{j,t}} \mathbf{1}\left\{s = \arg\min_m \left\{p_{mdj,t-k}(\omega)\right\}\right\} \mathrm{d}\omega \left(\prod_{\varsigma=t-k+1}^{t} \hat{c}_{sdj,\varsigma}\right)^{1-\sigma_j} \\
&= \frac{\mu_{j,t}(k)}{\mu_{j,t-k}(0)} \left(P^0_{dj,t-k}\right)^{1-\sigma_j} \sum_s \lambda^0_{sdj,t-k} \left(\prod_{\varsigma=t-k+1}^{t} \hat{c}_{sdj,\varsigma}\right)^{1-\sigma_j} \qquad (3.24)
\end{aligned}
$$

The price level change in partition 0 satisfies $P^0_{dj,t} = P^0_{dj,t-1}\left[\sum_s \lambda^0_{sdj,t-1}(\hat{c}_{sdj,t})^{-\theta_j}\right]^{-\frac{1}{\theta_j}}$ by (3.20), so we can rewrite the ideal price for composite goods with the last supplier selection k periods ago

$$\left(P^k_{dj,t}\right)^{1-\sigma_j} = \frac{\mu_{j,t}(k)}{\mu_{j,t-k}(0)}\left(P^0_{dj,t-k-1}\right)^{1-\sigma_j}\left[\sum_n \lambda^0_{ndj,t-k-1}\hat{c}^{-\theta_j}_{ndj,t-k}\right]^{-\frac{1-\sigma_j}{\theta_j}}\sum_s \lambda^0_{sdj,t-k}\left(\prod_{\varsigma=t-k+1}^{t}\hat{c}_{sdj,\varsigma}\right)^{1-\sigma_j}.$$

Denoting $\gamma_j \equiv \Gamma\left([\theta_j + 1 - \sigma_j]/\theta_j\right)$ and using the fact that $\left(P^0_{dj,t}\right)^{1-\sigma_j} = \mu_{j,t}(0)\left(\Phi_{dj,t}\right)^{-\frac{1-\sigma_j}{\theta_j}}\gamma_j$, we can rewrite the expression above as:

$$\left(P^k_{dj,t}\right)^{1-\sigma_j} = \gamma_j\mu_{j,t}(k)\left(\Phi_{dj,t-k}\right)^{-\frac{1-\sigma_j}{\theta_j}}\sum_s\left[\lambda^0_{sdj,t-k-1}\hat{c}^{-\theta_j}_{sdj,t-k}\right]^{-\frac{1-\sigma_j}{\theta_j}}\left(\prod_{\varsigma=t-k+1}^{t}\hat{c}_{sdj,\varsigma}\right)^{1-\sigma_j} \tag{3.25}$$

after expressing $\lambda^0_{sdj,t-k}$ recursively.

D. Aggregation Over Partitions

The aggregate ideal price level of the final good can be rewritten as a combination of the price levels of the partial price indices for the composites of intermediate goods purchased at time t from suppliers chosen $t-k$ periods ago:

$$\left(P_{dj,t}\right)^{1-\sigma_j} = \int_{[0,1]} p_{dj,t}(\bar{\omega})^{1-\sigma_j}\mathrm{d}\bar{\omega} = \sum_{k=0}^{\infty}\int_{\Omega^k_{j,t}} p_{dj,t}(\bar{\omega})^{1-\sigma_j}\mathrm{d}\bar{\omega} = \sum_{k=0}^{\infty}\left(P^k_{dj,t}\right)^{1-\sigma_j}.$$

Using the price index expressions (3.20) and (3.25) from the preceding subsections yields

$$\begin{aligned}
\left(P_{dj,t}\right)^{1-\sigma_j} &= \gamma_j\sum_{k=0}^{\infty}\mu_{j,t}(k)\left(\Phi_{dj,t-k}\right)^{-\frac{1-\sigma_j}{\theta_j}}\sum_s\left[\lambda^0_{sdj,t-k-1}\,\hat{c}^{-\theta_j}_{sdj,t-k}\right]^{-\frac{1-\sigma_j}{\theta_j}} \\
&\quad\times\exp\left\{\mathbf{1}\{k>0\}\log\left(\prod_{\varsigma=t-k+1}^{t}\hat{c}_{sdj,\varsigma}\right)^{1-\sigma_j}\right\} \\
&= \sum_{k=0}^{\infty}\frac{\mu_{j,t}(k)}{\mu_{j,t-k}(0)}\left(P^0_{dj,t-k-1}\right)^{1-\sigma_j}\sum_n\left[\lambda^0_{ndj,t-k-1}\hat{c}^{-\theta_j}_{ndj,t-k}\right]^{-\frac{1-\sigma_j}{\theta_j}} \\
&\quad\times\exp\left\{\mathbf{1}\{k>0\}\log\left[\sum_s\lambda^0_{sdj,t-k}\left(\prod_{\varsigma=t-k+1}^{t}\hat{c}_{sdj,\varsigma}\right)^{1-\sigma_j}\right]\right\}.
\end{aligned} \tag{3.26}$$

Recall that, by optimal demand, expenditure shares of each partition relative to total expenditures are

$$\frac{P^k_{dj,t} Y^k_{dj,t}}{P_{dj,t} Y_{dj,t}} = \left(\frac{P^k_{dj,t}}{P_{dj,t}} \right)^{1-\sigma_j}$$

Total expenditure shares are therefore simply the weighted average of trade shares across partitions

$$\lambda_{sdj,t} \equiv \sum_{k=0}^{\infty} \frac{P^k_{dj,t} Y^k_{dj,t}}{P_{dj,t} Y_{dj,t}} \lambda^k_{sdj,t} = \sum_{k=0}^{\infty} \left(\frac{P^k_{dj,t}}{P_{dj,t}} \right)^{1-\sigma_j} \lambda^k_{sdj,t},$$

which can also be stated as

$$\lambda_{sdj,t} = \left(\frac{P^0_{dj,t}}{P_{dj,t}} \right)^{1-\sigma_j} \frac{\lambda^0_{sdj,t-1} \hat{c}^{-\theta_j}_{sdj,t}}{\sum_n \lambda^0_{ndj,t-1} \hat{c}^{-\theta_j}_{ndj,t}} + \sum_{k=1}^{\infty} \left(\frac{P^k_{dj,t}}{P_{dj,t}} \right)^{1-\sigma_j} \frac{\lambda^0_{sdj,t-k} \left(\prod_{\varsigma=t-k+1}^{t} \hat{c}_{sdj,\varsigma} \right)^{1-\sigma_j}}{\sum_n \lambda^0_{ndj,t-k} \left(\prod_{\varsigma=t-k+1}^{t} \hat{c}_{ndj,\varsigma} \right)^{1-\sigma_j}}.$$

Writing $\lambda^0_{sdj,t-k}$ and $\lambda^0_{ndj,t-k}$ recursively, we can express trade shares compactly as

$$\lambda_{sdj,t} = \sum_{k=0}^{\infty} \left(\frac{P^k_{dj,t}}{P_{dj,t}} \right)^{1-\sigma_j} \frac{\lambda^0_{sdj,t-k-1} \hat{c}^{-\theta_j}_{sdj,t-k} \exp\left\{ \mathbf{1}\{k>0\} \log \left(\prod_{\varsigma=t-k+1}^{t} \hat{c}_{sdj,\varsigma} \right)^{1-\sigma_j} \right\}}{\sum_n \lambda^0_{ndj,t-k-1} \hat{c}^{-\theta_j}_{ndj,t-k} \exp\left\{ \mathbf{1}\{k>0\} \log \left(\prod_{\varsigma=t-k+1}^{t} \hat{c}_{ndj,\varsigma} \right)^{1-\sigma_j} \right\}}.$$

E. Convergence

Results in the preceding subsection imply that trade shares can be expressed a sum over infinitely many partitions. We now establish regularity conditions for convergence.

LEMMA 1 (Convergence). *If cumulative changes in trade costs are finite-valued*

$$\lim_{k\to\infty} \left| \prod_{\varsigma=t-k+1}^{t} \hat{c}_{ndj,\varsigma} \right| < \infty,$$

then price levels $P^k_{dj,t} < \infty$ and trade shares $0 < \lambda_{dj,t} < 1$ are finite-valued.

Proof. Note that $(\Phi_{dj,t-k})^{(\sigma_j-1)/\theta_j} < \infty$ and $\sum_s \left[\lambda^0_{sdj,t-k-1} A_{sj} \hat{c}^{-\theta_j}_{sdj,t-k} \right]^{(\sigma_j-1)/\theta_j} < \infty$ are both finite-valued, because they are equilibrium objects of a static equilibrium of the model. Also note that, for any $k > m$, if $|\prod_{\varsigma=t-k+1}^{t} \hat{c}_{ndj,\varsigma}| < \infty$, then $|\prod_{\varsigma=t-m+1}^{t} \hat{c}_{ndj,\varsigma}| < \infty$, since the product up to k includes every term in the product up to m. Therefore, if $\lim_{k\to\infty} |\prod_{\varsigma=t-k+1}^{t} \hat{c}_{ndj,\varsigma}| < \infty$, then, for every $k < \infty$, the product will also be finite. It follows that $P^k_{dj,t} < \infty$ is finite valued for every k. Given that $\lim_{k\to\infty} \mu_{j,t}(k) = \lim_{k\to\infty} (1-\zeta_j)^k \zeta_j = 0$. These findings also guarantee that $P_{dj,t} < \infty$. $\square$

F. Proofs

1. Proof of Proposition 1.

When the economy is in steady state, then for any $t <$ changes must satisfy $\hat{\mathbf{F}}_t = \hat{\mathbf{F}}_{\not k}$ and $\hat{\mathbf{w}}_t = \hat{\mathbf{w}}_{\not k}$ so that $\hat{c}_{s,t} = 1$ for all $s \in \mathcal{D}$. For the firms that are adjusting at t ($k = 0$), evaluating Equation (3.9) at those values, $\lambda^0_{sdj,t} = \lambda^0_{sdj,t-1} = \cdots = \lambda^0_{sdj,0}$ for all t. For the firms that are not adjusting at t ($k > 0$), we have $t - k > 0$ in equilibrium as long as the partition exists and can evaluate Equation (3.9) using the same logic as above: $\lambda^k_{sdj,t} = \lambda^0_{sdj,t-k} = \lambda^0_{sdj,0}$ for all t. From Equation (3.9), it is easy to see that $\lambda_{sdj,t} = \lambda^0_{sdj,t}$, which shows that $\boldsymbol{\lambda}_t = \boldsymbol{\lambda}^{EK}$ in steady state.

To derive the stationary distribution of contract lengths, begin by noting that the case $k = 0$ is trivial, since $\mu(0) = \mathbb{P}[K_t = 0] = \zeta_j$ does not vary. Now consider the case $k > 0$. Note that:

$$\begin{aligned}
\mathbb{P}[K_t = k, k > 0] &= \sum_{l=0}^{\infty} \mathbb{P}[K_t = k, k > 0 | K_{t-1} = l]\, \mathbb{P}[K_{t-1} = l] \\
&= (1 - \zeta_j)\mathbb{P}[K_{t-1} = k - 1]
\end{aligned}$$

The remaining proof for $k > 0$ then follows by induction. For $K_t = 1$, $\mathbb{P}[K_t = 1] = (1 - \zeta_j)\zeta_j$, and for $K_t = 2$, $\mathbb{P}[K_t = 2] = (1 - \zeta_j)\mathbb{P}[K_{t-1} = 1] = (1 - \zeta_j)^2\zeta_j$, and so forth recursively, for an arbitrary $K_t = k$ we must have $\mathbb{P}[K_t = k] = (1 - \zeta_j)^k\zeta_j$. This is the probability density function of a geometric distribution with mean $(1 - \zeta_j)/\zeta_j$ and standard deviation $\sqrt{1 - \zeta_j}/\zeta_j$.

Finally, using the definition of the measure μ, $\mu_{j,t}(k) = \mathbb{P}[K_t = k]$ for $t \geqslant k$. Given the Markov property of K_t, the following distribution will be stationary for all $k \in \mathbb{N}_0$:

2. Proof of Proposition 2.

For ease of notation, we suppress sector subscripts throughout the derivations. Consider a one-time permanent change in trade costs such that $\hat{\tau}_{sd,t} \neq 1$ and $\hat{\tau}_{sd,t+h} = 1\ \forall h > 0$. To characterize the partial trade elasticity at horizon h, we first characterize the elasticity for trade shares of each partition, then aggregate them up using the consumption shares derived from the CES preferences over partitions. The change in expenditure shares on intermediate goods in the kth partition in period $t+h$, relative to period $t-1$ is given by

$$\log \frac{\lambda^k_{sd,t+h}}{\lambda^k_{sd,t-1}} = \begin{cases} -(\sigma-1)\log \hat{\tau}_{sd,t} + \log \frac{\lambda^0_{sd,t+h-k}}{\lambda^k_{sd,t-1}} \left(\frac{(c_{s,t+h}/P^k_{d,t+h})}{(c_{s,t+h-k}/P^k_{d,t+h-k})} \right)^{1-\sigma} & , k \geqslant h \\ \log \frac{\lambda^0_{sd,t+h-k}}{\lambda^k_{sd,t-1}} \left(\frac{(c_{s,t+h}/P^k_{d,t+h})}{(c_{s,t+h-k}/P^k_{d,t+h-k})} \right)^{1-\sigma} & , 1 \leqslant k < h \\ \log \frac{\lambda^0_{sd,t+h-1}}{\lambda^k_{sd,t-1}} \left(\frac{(c_{s,t+h}/P^0_{d,t+h})}{(c_{s,t-1}/P^0_{d,t-1})} \right)^{\theta} & , k = 0 \end{cases}$$

The first line denotes intermediate goods that have not updated suppliers since the shock arrived. For such intermediate goods, changes in expenditure shares still explicitly depend on the shock to trade costs. The remaining intermediate goods have updated at least once, and a "new" optimal sourcing share $\lambda^0_{sd,t+h-k}$ from a time period between t and $t+h$ encodes the "initial price index" relative to which changes in expenditure shares are updated as well as the effect of the shock in trade costs. Unit costs are the relevant GE variables.

Denote

$$\Delta \boldsymbol{G}^{EK}_{sd,t,t+h} = -\theta \log \prod_{k=1}^{h} \frac{\hat{c}_{sd,t+k}}{\hat{P}^0_{sd,t+k}}$$

and

$$\Delta \boldsymbol{G}^{k}_{sd,\varsigma,t+h} = (1-\sigma) \log \prod_{\varsigma'=\varsigma+1}^{t+h} \frac{\hat{c}_{sd,\varsigma'}}{\hat{P}^k_{sd,\varsigma'}}$$

Then we can solve backwards to express all changes in trade shares above in terms of $\lambda^0_{sd,t-1}$, if possible:

$$\log \frac{\lambda^k_{sd,t+h}}{\lambda^k_{sd,t-1}} = \begin{cases} -(\sigma-1)\log \hat{\tau}_{sd,t} + \log \frac{\lambda^0_{sd,t+h-k}}{\lambda^k_{sd,t-1}} + \Delta \boldsymbol{G}^k_{sd,t,t+h} & , k \geqslant h \\ -\theta \log \hat{\tau}_{sd,t} + \log \frac{\lambda^0_{sd,t-1}}{\lambda^k_{sd,t-1}} + \Delta \boldsymbol{G}^{EK}_{sd,t,t+h-k} + \Delta \boldsymbol{G}^k_{sd,t+h-k,t+h} & , 1 \leqslant k < h \\ -\theta \log \hat{\tau}_{sd,t} + \Delta \boldsymbol{G}^{EK}_{sd,t,t+h} & , k = 0 \end{cases}$$

Use the fact that outcomes determined at t and earlier do not respond to the change in trade costs.

Hence, the elasticity of $\lambda^k_{sd,t+h}$ with respect to a change in trade costs at t, is hence given by,

$$\frac{\mathrm{d}\log(\lambda^k_{sd,t+h}/\lambda^k_{sd,t})}{\mathrm{d}\log\tau_{sd,t}} = \begin{cases} -(\sigma-1) + \frac{\mathrm{d}\Delta\boldsymbol{G}^k_{sd,t,t+h}}{\mathrm{d}\log\tau_{sd,t}} & , k \geqslant h \\ -\theta + \frac{\mathrm{d}\Delta\boldsymbol{G}^{EK}_{sd,t,t+h-k}}{\mathrm{d}\log\tau_{sd,t}} + \frac{\mathrm{d}\Delta\boldsymbol{G}^k_{sd,t+h-k,t+h}}{\mathrm{d}\log\tau_{sd,t}} & , 1 \leqslant k < h \\ -\theta + \frac{\mathrm{d}\Delta\boldsymbol{G}^{EK}_{sd,t,t+h}}{\mathrm{d}\log\tau_{sd,t}} & , k = 0 \end{cases}$$

To a first order, the change in overall expenditures at time $t+h$ caused by a one-time permanent shock to trade costs at t is given by

$$\begin{aligned}
\frac{\mathrm{d}\log(\lambda_{sd,t+h}/\lambda_{sd,t})}{\mathrm{d}\log\tau_{sd,t}} &= \sum_{k=0}^{\infty}\omega_k\left\{\frac{\mathrm{d}\log\lambda^k_{sd,t+h}/\lambda^k_{sd,t}}{\mathrm{d}\log\tau_{sd,t}} + (1-\sigma)\frac{\mathrm{d}\log\frac{P^k_{sd,t+h}P_{sd,t}}{(P^k_{sd,t}P_{sd,t+h})}}{\mathrm{d}\log\tau_{sd,t}}\right\} \\
&= \sum_{k=0}^{h-1}\omega_k\left\{-\theta + \frac{\mathrm{d}\Delta\boldsymbol{G}^{EK}_{sd,t,t+h}}{\mathrm{d}\log\tau_{sd,t}} + \frac{\mathrm{d}\Delta\boldsymbol{G}^k_{sd,t+h-k,t+h}}{\mathrm{d}\log\tau_{sd,t}} + (1-\sigma)\frac{\mathrm{d}\log\frac{P^k_{sd,t+h}P_{sd,t}}{(P^k_{sd,t}P_{sd,t+h})}}{\mathrm{d}\log\tau_{sd,t}}\right\} \\
&\quad + \sum_{k=h}^{\infty}\omega_k\left\{(1-\sigma) + \frac{\mathrm{d}\Delta\boldsymbol{G}^k_{sd,t,t+h}}{\mathrm{d}\log\tau_{sd,t}} + (1-\sigma)\frac{\mathrm{d}\log\frac{P^k_{sd,t+h}P_{sd,t}}{(P^k_{sd,t}P_{sd,t+h})}}{\mathrm{d}\log\tau_{sd,t}}\right\} \\
&= -\theta\sum_{k=0}^{h-1}\omega_k + (1-\sigma)\sum_{k=h}^{\infty}\omega_k \\
&\quad + \sum_{k=0}^{h-1}\omega_k\frac{\mathrm{d}\Delta\boldsymbol{G}^{EK}_{sd,t,t+h}}{\mathrm{d}\log\tau_{sd,t}} \\
&\quad + \sum_{k=0}^{h-1}\omega_k(1-\sigma)\left\{\frac{\sum_{\varsigma=t+h-k+1}^{t+h}\mathrm{d}\log c_{sd,\varsigma}}{\mathrm{d}\log\tau_{sd,t}} + \frac{\sum_{\varsigma=t}^{t+h-k}\mathrm{d}\log P^k_{sd,\varsigma}}{\mathrm{d}\log\tau_{sd,t}}\right\} \\
&\quad + \sum_{k=h}^{\infty}\omega_k(1-\sigma)\left\{\frac{\sum_{i=0}^{t+h}\mathrm{d}\log c_{sd,t+i}}{\mathrm{d}\log\tau_{sd,t}}\right\} \\
&\quad - (1-\sigma)\frac{\sum_{i=0}^{h}\mathrm{d}\log P_{sd,t+i}}{\mathrm{d}\log\tau_{sd,t}}
\end{aligned}$$

where $\omega_k \equiv \frac{\left(\frac{P^k_{dj,t}}{P_{dj,t}}\right)^{1-\sigma}\lambda^k_{sdj,t}}{\sum_k\left(\frac{P^k_{dj,t}}{P_{dj,t}}\right)^{1-\sigma}\lambda^k_{sdj,t}} = \frac{\mu_t(k)\lambda^k_{sdj,t}}{\sum_k\mu_t(k)\lambda^k_{sdj,t}}$. If t was a steady state, then $\omega_k = \mu(k)$, and the partial horizon-h trade elasticity equals:

$$\varepsilon^{t+h}_{sd} \equiv \frac{\partial\log\lambda_{sdj,t+h}}{\partial\log\tau_{sd,t}} = -\theta\sum_{k=0}^{h-1}\mu(k) + (1-\sigma)\sum_{k=h}^{\infty}\mu(k).$$

Using the stationary distribution of $\mu_t(k)$ to substitute for $\mu(k)$, we obtain the expression stated in the main text.

G. Proof of Proposition 3.

We begin by rearranging Equation (3.9) to express the prices of composite goods in terms of home expenditure shares

$$\begin{aligned}
\lambda_{ddi,t} P_{di,t}^{1-\sigma_i} &= \gamma_i \mu_i(0) \left(\Phi_{di,t}^0\right)^{-\frac{1-\sigma_i}{\theta_i}} \lambda_{ddi,t}^0 + \sum_{k\geqslant 1} \gamma_i \mu_i(k) \left(\Phi_{di,t-k}^0\right)^{-\frac{1-\sigma_i}{\theta_i}} \Phi_{di,t}^k \lambda_{ddi,t}^k \\
&= \gamma_i \mu_i(0) \left(\Phi_{di,t}^0\right)^{-\frac{1-\sigma_i}{\theta_i}} \lambda_{ddi,t}^0 + \sum_{k\geqslant 1} \gamma_i \mu_i(k) \left(\Phi_{di,t-k}^0\right)^{-\frac{1-\sigma_i}{\theta_i}} \lambda_{ddi,t-k}^0 \left(\frac{c_{dd,t}}{c_{dd,t-k}}\right)^{1-\sigma_i} \\
&= \gamma_i \mu_i(0) \left(\frac{c_{dd,t}^{-\theta_i}}{\lambda_{ddi,t}^0}\right)^{-\frac{1-\sigma_i}{\theta_i}} \lambda_{ddi,t}^0 + \sum_{k\geqslant 1} \gamma_i \mu_i(k) \left(\frac{c_{dd,t-k}^{-\theta_i}}{\lambda_{ddi,t-k}^0}\right)^{-\frac{1-\sigma_i}{\theta_i}} \lambda_{ddi,t-k}^0 \left(\frac{c_{dd,t}}{c_{dd,t-k}}\right)^{1-\sigma_i}.
\end{aligned}$$

It follows that

$$P_{di,t}^{1-\sigma_i} = c_{dd,t}^{1-\sigma_i} \left(\lambda_{ddi,t}^0\right)^{\frac{1-\sigma_i}{\theta_i}} \frac{1}{\lambda_{ddi,t}} \gamma_i \left[\mu_i(0)\lambda_{ddi,t}^0 + \sum_{k\geqslant 1} \mu_i(k) \left(\frac{\lambda_{ddi,t}^0}{\lambda_{ddi,t-k}^0}\right)^{-\frac{1-\sigma_i}{\theta_i}} \lambda_{ddi,t-k}^0\right]$$

where the price index is expressed in terms of unit cost and domestic trade shares. With the unit cost under Cobb-Douglas technology, the above equation can be rewritten as

$$\frac{P_{di,t}}{w_{d,t}} = \left(\lambda_{ddi,t}^0\right)^{\frac{1}{\theta_i}} \left(\lambda_{ddi,t}\right)^{1/(\sigma_i-1)} \left(\gamma_i \xi_{di,t}\right)^{1/(1-\sigma_i)} \alpha_{di}^{-\alpha_{di}} \prod_j \left(\frac{P_{dj,t}}{\alpha_{dji} w_{d,t}}\right)^{\alpha_{dji}}$$

where

$$\xi_{di,t} \equiv \mu_i(0)\lambda_{ddi,t}^0 + \sum_{k\geqslant 1} \mu_i(k) \left(\frac{\lambda_{ddi,t}^0}{\lambda_{ddi,t-k}^0}\right)^{\frac{\sigma_i-1}{\theta_i}} \lambda_{ddi,t-k}^0.$$

Taking logs yields

$$\log \frac{P_{di,t}}{w_{d,t}} = \log B_{si,t} + \sum_j \alpha_{sji} \log \frac{P_{sj,t}}{w_{s,t}}$$

where $B_{di,t} \equiv \alpha_{di}^{-\alpha_{di}} \left(\prod_j \alpha_{dji}^{-\alpha_{dji}}\right) \left(\lambda_{ddi,t}^0\right)^{\frac{1}{\theta_i}} \left(\lambda_{ddi,t}\right)^{1/(\sigma_i-1)} \left(\gamma_i \xi_{di,t}\right)^{1/(1-\sigma_i)}$. In matrix notation, this leads to

$$(\mathbf{I} - A_d) \log \hat{\boldsymbol{P}}_{d,t} = \log \boldsymbol{B}_{d,t},$$

where $A_d = \{\alpha_{dji}\}$ and $\log \hat{\boldsymbol{P}}_{d,t}$ and $\log \boldsymbol{B}_{d,t}$ are $I \times 1$ vectors. Inverting this system of equations, we obtain

$$\frac{P_{di,t}}{w_{d,t}} = \prod_j B_{dj,t}^{\bar{a}_{dji}},$$

where $\bar{a}_{dji}$ is the (j,i) entry of the Leontief matrix $(\mathbf{I} - A_d)^{-1}$. The consumer price index in country d can be written as

$$P_{d,t} = \prod_i (P_{di,t})^{\eta_i} = w_{d,t} \prod_{i,j} B_{dj,t}^{\bar{a}_{dji}\eta_i} = w_{d,t} \prod_j B_{dj,t}^{\sum_i \bar{a}_{dji}\eta_i}$$

It follows that the real wage is given by

$$W_{d,t} \equiv \frac{w_{d,t}}{P_{d,t}} = \prod_j B_{dj,t}^{-\sum_i \bar{a}_{dji}\eta_i}.$$

Taking the ratio between real wages in $t-1$ and $t+h$ yields

$$\frac{W_{d,t+h}}{W_{d,t-1}} = \prod_j \left[\left(\frac{\lambda^0_{ddj,t+h}}{\lambda^0_{ddj,t-1}} \right)^{-\frac{1}{\theta_j}} \left(\frac{\lambda_{ddj,t+h}}{\lambda_{ddj,t-1}} \right)^{-\frac{1}{\sigma_j - 1}} \left(\frac{\xi_{dj,t+h}}{\xi_{dj,t-1}} \right)^{\frac{1}{\sigma_j - 1}} \right]^{\sum_i \bar{a}_{dji}\eta_i},$$

If $t-1$ is a steady state, then $\lambda^k_{ddj,t-1} = \lambda_{ddj,t-1}$ for all $k \in \{0, 1, 2, ...\}$ and the above expression simplifies to

$$\begin{aligned}
\frac{W_{d,t+h}}{W_{d,t-1}} &= \prod_j \left[\left(\frac{\lambda^0_{ddj,t+h}}{\lambda^0_{ddj,t-1}} \right)^{-\frac{1}{\theta_j}} \left(\frac{\lambda_{ddj,t+h}}{\xi_{dj,t+h}} \right)^{-\frac{1}{\sigma_j - 1}} \right]^{\sum_i \bar{a}_{dji}\eta_i} \\
&= \prod_j \left[\left(\frac{\lambda_{ddj,t+h}}{\lambda_{ddj,t-1}} \right)^{-\frac{1}{\theta_j}} \left(\frac{\lambda^0_{ddj,t+h}}{\lambda_{ddj,t+h}} \right)^{-\frac{1}{\theta_j}} \left(\frac{\lambda_{ddj,t+h}}{\xi_{dj,t+h}} \right)^{-\frac{1}{\sigma_j - 1}} \right]^{\sum_i \bar{a}_{dji}\eta_i} \qquad (3.27) \\
&= \prod_j \left[\left(\frac{\lambda_{ddj,t+h}}{\lambda_{ddj,t-1}} \right)^{-\frac{1}{\theta_j}} (\Xi_{dj,h})^{\frac{1}{\sigma_j - 1}} \right]^{\sum_i \bar{a}_{dji}\eta_i},
\end{aligned}$$

where

$$\Xi_{dj,t} \equiv \zeta_j \left(\frac{\lambda_{ddj,t+h}}{\lambda^0_{ddj,t+h}} \right)^{\frac{\sigma_j - 1 - \theta_j}{\theta_j}} + \sum_{k=1}^{h} \zeta_j (1 - \zeta_j)^k \left(\frac{\lambda_{ddj,t+h}}{\lambda^0_{ddj,t+h-k}} \right)^{\frac{\sigma_j - 1 - \theta_j}{\theta_j}} + (1 - \zeta_j)^{h+1} \left(\frac{\lambda_{ddi,t+h}}{\lambda_{ddj,t-1}} \right)^{\frac{\sigma_j - 1 - \theta_j}{\theta_j}}$$

is obtained after combining the last two factors in Equation (3.27).

REFERENCES

Alessandria, George, Joseph P. Kaboski, and Virgiliu Midrigan. 2010. "Inventories, Lumpy Trade, and Large Devaluations." *American Economic Review*, 100(5): 2304–39.

Alessandria, George, Shafaat Yar Khan, Armen Khederlarian, Carter Mix, and Kim J. Ruhl. 2023. "The Aggregate Effects of Global and Local Supply Chain Disruptions: 2020–2022." *Journal of International Economics*, forthcoming.

Anderson, James E., and Yoto V. Yotov. 2022. "Estimating Gravity from the Short to the Long Run: A Simple Solution to the 'International Elasticity Puzzle'." *NBER Working Paper*, 30809.

Angeletos, George-Marios, and Zhen Huo. 2021. "Myopia and Anchoring." *American Economic Review*, 111(4): 1166–1200.

Antràs, Pol, and Davin Chor. 2013. "Organizing the Global Value Chain." *Econometrica*, 81(6): 2127–2204.

Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare. 2012. "New Trade Models, Same Old Gains?" *American Economic Review*, 102(1): 94–130.

Arkolakis, Costas, Jonathan Eaton, and Samuel Kortum. 2011. "Staggered Adjustment and Trade Dynamics." Yale University, unpublished manuscript.

Armington, Paul S. 1969. "A Theory of Demand for Products Distinguished by Place of Production." *International Monetary Fund Staff Papers*, 16(1): 159–178.

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub. 2021. "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models." *Econometrica*, 89(5): 2375–2408.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2020. "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model."

Aviv, Yossi. 2001. "The Effect of Collaborative Forecasting on Supply Chain Performance." *Management Science*, 47(10): 1326–1343.

Baxter, Marianne, and Robert G. King. 1999. "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series." *The Review of Economics and Statistics*, 81(4): 575–593.

Bils, Mark, and James A. Kahn. 2000. "What Inventory Behavior Tells Us about Business Cycles." *American Economic Review*, 90(3): 458–481.

Blinder, Alan S. 1981. "Retail Inventory Behavior and Business Fluctuations." *Brookings Papers on Economic Activity*, 1981(2): 443–520.

Blinder, Alan S., and Douglas Holtz-Eakin. 1986. "Inventory Fluctuations in the United States since 1929." In *The American Business Cycle: Continuity and Change*. Vol. 25 of *Studies in Business Cycles*, , ed. Robert J. Gordon, Chapter 3, 183–236. University of Chicago Press.

Boehm, Christoph E., Andrei A. Levchenko, and Nitya Pandalai-Nayar. 2023. "The Long and Short (Run) of Trade Elasticities." *American Economic Review*, 113(4): 861–905.

Cachon, Gérard P., and Marshall Fisher. 2000. "Supply Chain Inventory Management and the Value of Shared Information." *Management Science*, 46(8): 1032–1048.

Calvo, Guillermo A. 1983*a*. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics*, 12(3): 383–398.

Calvo, Guillermo A. 1983*b*. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics*, 12(3): 383–398.

Caplin, Andrew S. 1985. "The Variability of Aggregate Demand with (s, S) Inventory Policies." *Econometrica*, 53(6): 1395–1409.

Chatfield, Dean C., Jeon G. Kim, Terry P. Harrison, and Jack C. Hayya. 2004. "The Bullwhip Effect—Impact of Stochastic Lead Time, Information Quality, and Information Sharing: A Simulation Study." *Production and Operations Management*, 13(4): 340–353.

Christiano, Lawrence J. 1988. "Why Does Inventory Investment Fluctuate so Much?" *Journal of Monetary Economics*, 21(2): 247–280.

Christiano, Lawrence j., Martin Eichenbaum, and Charles l. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy*, 113(1): 1–45.

Davis, Steven J., and James A. Kahn. 2008. "Interpreting the Great Moderation: Changes in the Volatility of Economic Activity at the Macro and Micro Levels." *Journal of Economic Perspectives*, 22(4): 155–80.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum. 2007. "Unbalanced Trade." *American Economic Review: Papers and Proceedings*, 97(2): 351–55.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum. 2008. "Global Rebalancing with Gravity: Measuring the Burden of Adjustment." *IMF Staff Papers*, 55(3): 511–540.

de Souza, Gustavo, Naiyuan Hu, Haishi Li, and Yuan Mei. 2022. "(Trade) War and Peace: How to Impose International Trade Sanctions." Federal Reserve Bank of Chicago.

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, Geography, and Trade." *Econometrica*, 70(5): 1741–79.

Eaton, Jonathan, and Samuel Kortum. 2012. "Putting Ricardo to Work." *Journal of Economic Perspectives*, 26(2): 65–90.

Eichenbaum, Martin. 1983. "A Rational Expectations Equilibrium Model of Inventories of Finished Goods and Employment." *Journal of Monetary Economics*, 12(2): 259–277.

Eichenbaum, Martin. 1984. "Rational Expectations and the Smoothing Properties of Inventories of Finished Goods." *Journal of Monetary Economics*, 14(1): 71–96.

Eichenbaum, Martin. 1989. "Some Empirical Evidence on the Production Level and Production Cost Smoothing Models of Inventory Investment." *The American Economic Review*, 79(4): 853–864.

Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin. 2000. “Optimal Monetary Policy with Staggered Wage and Price Contracts.” *Journal of Monetary Economics*, 46(2): 281–313.

Fajgelbaum, Pablo D., Pinelopi K. Goldberg, Patrick J. Kennedy, and Amit K. Khandelwal. 2020. “The Return to Protectionism.” *Quarterly Journal of Economics*, 135(1): 1–55.

Ferrari, Alessandro. 2023. “Inventories, Demand Shocks Propagation and Amplification in Supply Chains.” University of Zurich.

Fontagné, Lionel, Houssein Guimbard, and Gianluca Orefice. 2022. “Tariff-based Product-level Trade Elasticities.” *Journal of International Economics*, 137: 103593.

Fontagné, Lionel, Philippe Martin, and Gianluca Orefice. 2018. “The International Elasticity Puzzle Is Worse Than You Think.” *Journal of International Economics*, 115(0): 115–29.

Giard, Vincent, and Mustapha Sali. 2013. “The Bullwhip Effect in Supply Chains: A Study of Contingent and Incomplete Literature.” *International Journal of Production Research*, 51(13): 3880–3893.

Gonçalves, Sílvia, and Lutz Kilian. 2004. “Bootstrapping Autoregressions with Conditional Heteroskedasticity of Unknown Form.” *Journal of Econometrics*, 123(1): 89–120.

Hamilton, James D. 2018. “Why You Should Never Use the Hodrick-Prescott Filter.” *The Review of Economics and Statistics*, 100(5): 831–843.

Horvath, Michael. 2000. “Sectoral Shocks and Aggregate Fluctuations.” *Journal of Monetary Economics*, 45(1): 69–106.

Iacoviello, Matteo, Fabio Schiantarelli, and Scott Schuh. 2011. “Input and Output Inventories in General Equilibrium.” *International Economic Review*, 52(4): 1179–1213.

Kahn, James A. 1987. “Inventories and the Volatility of Production.” *The American Economic Review*, 77(4): 667–679.

Khan, Aubhik, and Julia K. Thomas. 2007. “Inventories and the Business Cycle: An Equilibrium Analysis of (S, s) Policies.” *American Economic Review*, 97(4): 1165–1188.

Kimball, Miles S. 1995. “The Quantitative Analytics of the Basic Neomonetarist Model.” *Journal of Money, Credit and Banking*, 27(4): 1241–1277.

Kollintzas, Tryphon, and Ruilin Zhou. 1992. “Import Price Adjustments with Staggered Import Contracts.” *Journal of Economic Dynamics and Control*, 16(2): 289–315.

Kryvtsov, Oleksiy, and Virgiliu Midrigan. 2012. “Inventories, Markups, and Real Rigidities in Menu Cost Models.” *The Review of Economic Studies*, 80(1): 249–276.

Kydland, Finn E., and Edward C. Prescott. 1982. “Time to Build and Aggregate Fluctuations.” *Econometrica*, 50(6): 1345–1370.

Lee, Hau L., V. Padmanabhan, and Seungjin Whang. 1997. “Information Distortion in a Supply Chain: The Bullwhip Effect.” *Management Science*, 43(4): 546–558.

Li, Dake, Mikkel Plagborg-Møller, and Christian K. Wolf. 2021. “Local Projections vs. VARs: Lessons from Thousands of DGPs.”

Lucas, Robert E. Jr.., and Edward C. Prescott. 1971. "Investment Under Uncertainty." *Econometrica*, 39(5): 659–81.

Mcconnell, Margaret M., and Gabriel Perez-Quiros. 2000. "Output Fluctuations in the United States: What Has Changed since the Early 1980's?" *American Economic Review*, 90(5): 1464–1476.

Metters, Richard. 1997. "Quantifying the Bullwhip Effect in Supply Chains." *Journal of Operations Management*, 15(2): 89–100.

Montero, Miquel, and Javier Villarroel. 2016. "Directed Random Walk with Random Restarts: The Sisyphus Random Walk." *Physical Review E*, 94(3): 1–10. 032132, arXiv:1603.09239v2.

Morris, Stephen, and Hyun Song Shin. 2002. "Social Value of Public Information." *American Economic Review*, 92(5): 1521–1534.

Newey, Whitney K., and Kenneth D. West. 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55(3): 703–708.

Nimark, Kristoffer. 2008. "Dynamic Pricing and Imperfect Common Knowledge." *Journal of Monetary Economics*, 55(2): 365–382.

OECD. 2023. "OECD Inter-Country Input-Output Database."

Ouyang, Yanfeng. 2007. "The Effect of Information Sharing on Supply Chain Stability and the Bullwhip Effect." *European Journal of Operational Research*, 182(3): 1107–1121.

Plagborg-Møller, Mikkel, and Christian K. Wolf. 2021. "Local Projections and VARs Estimate the Same Impulse Responses." *Econometrica*, 89(2): 955–980.

Ramey, Valerie A. 1989. "Inventories As Factors of Production and Economic Fluctuations." *The American Economic Review*, 79(3): 338–354.

Ramey, Valerie A. 2016. "Macroeconomic Shocks and Their Propagation." In *Handbook of Macroeconomics*. Vol. 2, , ed. John B. Taylor and Harald Uhlig, Chapter 2, 71–162. Elsevier.

Ramey, Valerie A., and Kenneth D. West. 1999. "Inventories." In *Handbook of Macroeconomics*. Vol. 1, Chapter 13, 863–923. Elsevier.

Romer, Christina D., and David H. Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review*, 94(4): 1055–1084.

Ruhl, Kim J. 2008. "The International Elasticity Puzzle." University of Texas at Austin, unpublished manuscript.

Sarte, Pierre-Daniel, Felipe Schwartzman, and Thomas A. Lubik. 2015. "What Inventory Behavior Tells Us about How Business Cycles Have Changed." *Journal of Monetary Economics*, 76: 264–283.

Simchi-Levi, David, and Yao Zhao. 2003. "The Value of Information Sharing in a Two-Stage Supply Chain with Production Capacity Constraints." *Naval Research Logistics (NRL)*, 50(8): 888–916.

Smets, Frank, and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review*, 97(3): 586–606.

Stock, James H, and Mark W Watson. 2002. "Has the Business Cycle Changed and Why?" National Bureau of Economic Research Working Paper 9127.

vom Lehn, Christian, and Thomas Winberry. 2021. "The Investment Network, Sectoral Comovement, and the Changing U.S. Business Cycle." *The Quarterly Journal of Economics*, 137(1): 387–433.

Wang, Pengfei, Yi Wen, and Zhiwei Xu. 2014. "What Inventories Tell Us about Aggregate Fluctuations—A Tractable Approach to (S,s) Policies." *Journal of Economic Dynamics and Control*, 44: 196–217.

Wen, Yi. 2005. "Understanding the Inventory Cycle." *Journal of Monetary Economics*, 52(8): 1533–1555.

Wen, Yi. 2011. "Input and Output Inventory Dynamics." *American Economic Journal: Macroeconomics*, 3(4): 181–212.

Wieland, Johannes F., and Mu-Jeung Yang. 2020. "Financial Dampening." *Journal of Money, Credit and Banking*, 52(1): 79–113.

Yilmazkuday, Hakan. 2019. "Understanding the International Elasticity Puzzle." *Journal of Macroeconomics*, 59(0): 140–53.