Title

Essays in Macro-finance

Author

Sellemi, Victor

Publication Date

2024

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Essays in Macro-finance

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

in

Economics

by

Victor Sellemi

Committee in charge:

Professor Allan Timmermann, Chair
Professor Joseph Engelberg
Professor Valerie Ramey
Professor Alexis Akira Toda
Professor Rossen Valkanov

2024

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

University of California San Diego

2024

DEDICATION

To Mike, Salma, and Antonio.

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGEMENTS

First, I am grateful to my advisor Allan Timmermann for his professional and personal guidance and support throughout my graduate studies at UC San Diego. This dissertation would not have been possible without his support, and I am honored to become a scholar under his guidance.

For Chapter 1, I would like to acknowledge my other dissertation committee members, Joey Engelberg, Valerie Ramey, Alexis Toda, and Ross Valkanov for providing insightful feedback and guidance throughout the course of the project. I would also like to acknowledge helpful feedback and suggestions from Munseob Lee, Fabian Trottner, and Johannes Wieland.

For Chapter 2, I am particularly thankful to Valerie Ramey for her constant feedback and help thoughout the course of this work. I also acknowledge helpful comments and suggestions from Jim Hamilton, Juan Herreno, Nir Jaimovich, Marc Muendler, Fabian Trottner, and Johannes Wieland.

For Chapter 3, I acknowledge helpful feedback and suggestions from Alex Cloninger, Lily Xu, and Mo Zhou.

Finally, I also want to thank my former mentor Guido Kuersteiner, whose early support and encouragement at the University of Maryland made it possible for me to pursue a PhD in the first place.

Victor Sellemi is the sole author of Chapter 1.

Chapter 2 is coauthored with Edoardo Briganti. The dissertation author is one of the primary investigators of this material.

Chapter 3 is coauthored with Tjeerd de Vries. The dissertation author is one of the primary investigators of this material.

VITA

2018	B.S. in Physics, University of Maryland College Park
2018	B.S. in Mathematics, University of Maryland College Park
2019	M.A. in Economics, University of California San Diego
2024	Ph.D. in Economics, University of California San Diego

FIELD OF STUDY

Primary Field: Finance; Secondary Field: Macroeconomics

ABSTRACT OF THE DISSERTATION

Essays in Macro-finance

by

Victor Sellemi

Doctor of Philosophy in Economics

University of California San Diego, 2024

Professor Allan Timmermann, Chair

This dissertation is comprised of three free-standing chapters, each focused on topics in the intersection of financial economics and macroeconomics.

Chapter 1 examines the propagation of shocks in economic models with customer-supplier networks. I show that the common assumption of idiosyncratic shocks at the firm or industry levels implies empirically implausible sparsity restrictions on the input-output network structure. Moreover, I provide evidence that substitutability between trade partners is related to technological and product dispersion that is not captured by standard firm and industry definitions, and thus generates non-negligible correlation in shocks. Finally, I show that assets positively exposed to upstream and downstream shocks are useful hedges and earn lower average risk premia than less exposed peers. This is confirmed by statistically significant return spreads and a negative association between correlated shock

propagation and aggregate growth.

Chapter 2 studies the role of time-to-build in federal defense when estimating aggregate federal government spending multipliers. We find that the early impact of defense news shocks on GDP is due to a rise in business inventories, as contractors ramp up production for new defense contracts. These contracts do not affect government spending (G) until payment-on-delivery, which occurs 2-3 quarters later. Novel data on defense procurement obligations reveals that contract awards Granger-cause shocks to G identified via Cholesky decomposition, but not defense news shocks. We show that Cholesky shocks to G miss early changes in inventories, and thus result in lower multiplier estimates relative to the narrative method.

Chapter 3 explores the permanent-transitory decomposition of stochastic discount factor (SDF) processes in dynamic asset pricing models, in which the permanent component captures pricing at long payoff horizons. Analytic solutions for the permanent component are limited, and standard numerical methods are not well-suited to solve for them due to the curse of dimensionality and lack of boundary conditions and/or parametric assumptions. We propose a novel algorithm for computing the permanent-transitory decomposition for a general class of asset pricing models without such restrictions. We validate the algorithm's accuracy in several workhorse structural asset-pricing models, and argue that our approach applies to models whose state dynamics follow general and potentially high dimensional Lévy processes.

Chapter 1

Risk in Network Economies

Abstract

Economic models with input-output networks assume that firm or sector growth is driven by a combination of trade partners' growth and idiosyncratic shocks. This assumption generates unrealistic restrictions on network weights. Allowing for correlated shocks exposes units to additional risk that captures their ability to substitute away from supply and demand shocks. I provide evidence that substitutability between trade partners is related to technological and product dispersion that is not captured by standard firm and industry definitions. I propose a production-based asset pricing model in which supply chain substitutability depends on product/technology dispersion and shock correlation driven by shared suppliers and customers. The model predicts that assets positively exposed to upstream and downstream shocks are useful hedges and earn lower average risk premia than less exposed peers. This is confirmed by estimated return spreads of -11.4% and -4.2% and a negative association with aggregate growth.

1.1 Introduction

The total value of intermediate inputs flowing through production networks in the United States was more than \$81 billion in 2023.[1] Recent research finds that production networks play an important role in shock propagation, business cycles, and systematic risk in asset markets. However, the relationship between network linkages and comovement in economic risk is not yet entirely clear, especially at a granular level. Features of the input-output network are crucial to understanding the relationship between firm or industry-level risk and economy-wide aggregate risk.

The benchmark network model of the economy assumes that idiosyncratic shocks are drawn independently across units (i.e., firms or industries) before propagating to connected units as a function of network weights. Network weights capture the relative importance of each connection (edge) between units. As a result, shocks to any individual unit can have systematic effects. Acemoglu et al. (2012a) argues that heavy tails in the distribution of network weights can inhibit diversification and amplify the systematic effects of idiosyncratic shocks. Similarly, Gabaix (2011) shows that skewness in the firm size distribution also inhibits diversification away from shocks.

In this work, I show that realistic production network models of the economy must account for correlation in unit shocks which is not driven by network connections. In other words, the assumption of idiosyncratic shocks generates unrealistic restrictions on the sparsity of customer-supplier relationships (e.g., no bilateral trade between firms). I start from the reduced-form equation of static propagation of shocks, which links each unit's growth to the growth of its network connections plus a unit-specific shock. This reduced-form is consistent with a broad class of structural economic models that assume Cobb-Douglas aggregation of intermediate inputs in a production function.[2] In this equation, I show that imposing a diagonal structure on the variance-covariance matrix of

[1] This is equivalent to 74% of GDP according to data from the U.S. Bureau of Economic Analysis (See e.g., link to FRED)

[2] Such models are common in the production network literature (e.g., Acemoglu et al. (2012b), Acemoglu et al. (2016b), Ramìrez (2017), Herskovic (2018a), Herskovic et al. (2020b)).

shocks implicitly leads to sparsity restrictions on the network weights permitted in the model. More specifically, the set of permissible networks require either monopolistic trade (pairs of firms which trade with each other can only trade with each other) or unilateral trade (a firm is never both the customer and supplier of another firm). Neither restrictions holds in the data on industry and firm trade relationships in the U.S.

As a result, I argue that researchers should account for correlation in shocks when making use of network models. Practically speaking, there are several reasons why shocks to units in the input-output network might be correlated, especially as the unit definitions become less granular. For instance, two firms which produce the same goods should experience correlation in demand shocks at the product level. If the two firms also produce using similar inputs and/or technologies, then supply-side shocks are likely to be correlated as well. Hoberg and Phillips (2016) show from text data that firms that produce similar products often belong to different industries, which suggests that industries should also experience some degree of comovement in demand shocks. Along these lines, Hottman et al. (2016) show using scanner data that 69% of firms, which account for 99% of their industries' output, supply multiple and intersecting product varieties. In other words, industry and firm definitions are too broad for the idiosyncratic risk assumption to be reasonable.

Like product similarity, technological and geographic proximity might also generate comovement in supply and demand shocks. Bloom and Shankerman (2013) shows that regional shocks to research and development (R&D) incentives have correlated effects on the growth of firms who operate in closely related technology spaces. Similarly, firms operating in nearby locations are likely to be exposed to the same underlying geographic shocks. For instance, Autor et al. (2013) and Mian and Sufi (2014) provide evidence that local employment shocks have correlated effects within a region, and Tuzel and Zhang (2017) studies correlated exposure to regional risk associated with changes in local prices for factors of production. Even local climate risk could expose multiple firms to the same regional risks (see e.g., Barrot and Sauvagnat (2016) and Kruttli et al. (2019)).

When shocks are idiosyncratic, each unit's growth rate variance is the sum of unit-specific shock variance and a network-weighted sum of shock variances of its trade partners. Of course, the former term is unrelated to the presence of network connections. In the homoskedastic case, the second term simplifies to a constant times a concentration measure across each unit's trade partners. Acemoglu et al. (2012a) show that aggregate volatility shocks to this component decay at a rate slower than $\sqrt{n}$ when network weights follow a power law distribution. Herskovic et al. (2020b) focus on concentration of firm reliance on customers, and argue that increases in concentration are related to increases in firm size dispersion. Unlike these papers, this is the first work to investigate, both empirically and theoretically, the relationship between exposure to correlated shocks in the production network and realized variance.

In particular, when allowing for non-negligible correlation in unit-level shocks shocks, the expression for growth rate variance gains an additional covariance component, denoted concentration "between" trade partners. This new term captures the ability of each unit in the network to substitute away from correlated shocks to its trade partners. In particular, units are more substitutable (less concentrated) when they diversify trade between partners that are exposed to negatively correlated shocks. When units trade with partners that experience correlated shocks, they are more concentrated when the relative importance of those trade partners is similar.

Building on this intuition, I estimate concentration between trade partners using panel data at the industry level. Consistent with theory, I find that this new component explains a significant amount of variation in the panel of realized industry variance. More concentrated (less substitutable) industries are more volatile both in terms of market returns and output growth. This relationship is robust to controls for relevant industry characteristics such as size, centrality, concentration across trade partners, vertical position in the supply chain, and durability of output.

This finding alone does not provide any insight on the underlying source of correlated risks between trade partners. Diving deeper, I consider the results of Acemoglu et al.

(2016a), who argue that total factor productivity (TFP) shocks primarily propagate downstream while government spending shocks primarily propagate upstream from customers to suppliers. Consistent with this finding, I show that the elasticity of realized variance to concentration between trade partners is more precisely estimated on the supply-side when constructed using pairwise industry correlations in TFP growth. On the other hand, the elasticity of realized variance to concentration between customers is more precisely estimated when using correlations in federal procurement demand shocks.

Additionally, I suppose that correlation between upstream and downstream propagating shocks is driven by proximity of industries on a latent surface capturing final good varieties. As a tractable simplification, I assume that correlation between demand shocks is a function of product similarity, while correlation between supply shocks is a function of technological similarity.[3] I proxy product similarity using the text-based scores from Hoberg and Phillips (2016) and technological proximity following Bloom and Shankerman (2013). Similarly, I find that the elasticity of realized variance to between concentration is more precisely estimated on the demand-side using product similarity and on the supply-side using technological proximity.

These findings suggest one structural foundation for incorporating correlation in supply and demand shocks in network models of the economy. To fully investigate the risk implications of this correlation, I propose a production-based asset pricing model with input-output networks in which firm-level technology shocks propagate downstream from suppliers to customers and demand shocks propagate upstream from customers to suppliers. Firms are both customers and suppliers. Unlike most existing models, I account for both directions of propagation.[4] Additionally, I introduce a novel mechanism for correlation in shocks in which the propensity of shock transmission is a function of customer and supplier substitutability at the firm level.

[3]This is no the only valid source of correlation. In principal, other factors such as input/raw-material similarity or geographic proximity generate correlation in firm-level shocks. A more general model for this is left for future work.

[4]For example, Shea (2002) and Kramarz et al. (2020), andHerskovic et al. (2020b) focus on upstream propagation of demand shocks, while Acemoglu et al. (2012a) focus on downstream propagation.

In particular, I define substitutability as network weighted sum of latent distances between a firm's trade partners. Product distance characterizes customer substitutability, while technological distance characterizes supplier substitutability. Shared customers and suppliers between firms induce comovement in substitutability and thus correlation in propagated shocks. The proportion of firms that are affected by network propagated shocks is related to the changes in average propensity and average supply chain substitutability. Importantly, the model predicts that average propagation in the upstream and downstream directions represent distinct and negatively priced sources of systematic risk in the economy.

I test this prediction by calibrating the model and constructing empirical analogues of upstream and downstream network propagation risk factors. Consistent with theory, I confirm empirically that upstream and downstream propagation risk factors are negatively related to aggregate consumption growth, output growth, and dividend growth.[5] Moreover, a trading strategy which buys the highest and sells the lowest quintile upstream (downstream) propagation beta-sorted portfolio generates excess returns of -11.42% (-4.18%). These factors survive the standard set of robustness checks.

The paper is structured as follows. Section 1.2 provides theoretical evidence that the idiosyncratic risk assumption is overly restrictive in network models. Section 1.3 provides empirical evidence that correlated shocks can explain heterogeneity in the industry panel of realized variance. Section 1.4 proposes an asset pricing model which incorporates supply-chain substitutability. Section 1.5 verifies the predictions of the model in the data. Section 3.6 concludes.

1.1.1 Related Literature

This work most closely relates to the literature on static shock propagation through input-output networks, and more specifically recent work on asset pricing in the presence input-output networks. In both areas of research, one foundational assumption is that

[5]This links the model to a large class of consumption-based asset pricing models.

shocks originate idiosyncratically at the firm or industry level. As far as I know, this is the first work to directly challenge this assumption both theoretically and empirically.

Acemoglu et al. (2012a) focus on static production network economies and show that the rate of decay of idiosyncratic supply-side shocks depends on the structure of the input-output connections. Several studies also study upstream propagation of idiosyncratic shocks (Shea (2002), Ozdagli and Weber (2017), Kramarz et al. (2020), Herskovic et al. (2020b)). Acemoglu et al. (2016a) allow for both directions of propagation in the static model and argue both theoretically and empirically that idiosyncratic technology shocks mostly propagate downstream, while demand shocks mostly propagate upstream. This work is most similar to Acemoglu et al. (2016a), as it studies both directions of shock propagation. However, unlike these papers with idiosyncratic shocks, I allow for correlation in the firm or industry level shocks which propagate through the network. I am also the first to prove that the idiosyncratic shock assumption in the static network model implicitly lead to overly restrictive network weights which are inconsistent with the data.

Many papers study the role of the network structure in driving systematic fluctuations. Acemoglu et al. (2012a) focuses on volatility decay and the network weight distribution. Herskovic et al. (2020a) study the relationship between customer concentration and aggregate volatility. In financial markets, Herskovic (2018a) proposes two network-based systematic risk factors, termed concentration and sparsity. One general theme in this vein is that more concentrated networks leads to less diversification of risk.[6] However, these works do not account for correlation in shocks when proposing measures of network concentration. This work proposes a novel notion of "concentration and substitutability between" trade partners, which depends on both network weights and correlations in shocks between shared trade partners.

This work also relates to the production-based asset pricing literature with input-output networks. Herskovic (2018a) proposes asset pricing factors which depend directly

[6]Note that these papers assume networks are determined exogenously. This work does not really discuss the relationship between macroeconomic forces and endogenous network formation (e.g., Baqaee (2018), Taschereau-Dumouchel (2020))

on the network structure. Ramìrez (2017) proposes a model for shock propagation between firms and provides a theoretical foundation for Ahern (2013)'s result that more central networks earn higher risk premia. This work introduces upstream and downstream network propagation as systematic risk factors which capture the average ability of firms to diversify away from shocks propagating through the network. Like Ramìrez (2017), I explicitly introduce correlation in shocks, which is driven by product and technological dispersion.

This reasoning is closely tied to recent literature on new measures of firm-level similarity and competition. Hoberg and Phillips (2016) propose a text-based measure of product similarity, and find that product similarity implies firm-level clusters which do not always align with industry clusters. Similarly, Hottman et al. (2016) show from product-level data that most products or services are produced by multiple firms. On the technology side, Bloom and Shankerman (2013) show that shocks to research and development (R&D) have correlated effects on the productivity and growth of firms with similar production technology. Finally, several papers propose geographic location as a relevant source of firm-level correlations. These correlations might be generated by local labor markets (see e.g., Autor et al. (2013), Mian and Sufi (2014)), local factor prices (Tuzel and Zhang (2017), Grigoris (2019)), local technological progress (Oberfield (2018)), or local weather events (Barrot and Sauvagnat (2016), Kruttli et al. (2019)).

1.2 Idiosyncratic Risk in Input-Output Networks

In an economy where sectors or firms are connected through a network of input-output linkages, shocks to any individual unit might generate larger systematic effects. Intuitively, firms or industries with close trade relationships should also experience some degree of comovement in risk. Recent research proposes several approaches for modeling the spread of small shocks from firms or disaggregated sectors.[7] Since input-output networks are observed in the data, these approaches are often the foundation for empirical

[7] See e.g., Gabaix (2011), Acemoglu et al. (2012a), Taschereau-Dumouchel (2020), and Baqaee and Farhi (2019) for discussion on how microeconomic shocks can generate macroeconomic effects.

studies on the importance of various channels of shock propagation.[8]

The benchmark model studied in much of the literature involves static propagation of shocks through a deterministic network.[9] The main idea is that a sector or firm's growth rates depend on a network-weighted sum of growth rates of trade partners and an idiosyncratic shock which is drawn independently from the other units. Network weights capture the importance of direct trade relationships between sectors and are generally non-negative. Moreover, if the entries of the matrix are sales or purchase shares of inputs, these weights are also bounded above by 1 and in most cases assumed to sum to 1 or less than 1 for every unit. Typically no additional restrictions are imposed on the input-output network structure (e.g., symmetry or sparsity).

However, in this section I show that in this benchmark model, the assumption of idiosyncratic shocks across units is not consistent with such a general class of networks. In particular, new results show that stronger restrictions on the input-output network weights are required when the variance-covariance matrix of idiosyncratic shocks is diagonal. These additional restrictions are inconsistent with almost all empirically observed input-output networks at the industry or firm levels, and cannot be relaxed by adding omitted macroeconomic factors or by accounting for multiple networks. Additionally, even if these restrictions are satisfied, there is no definitive empirical evidence that supply and demand shocks have zero pairwise correlation across all pairs of units.

After formally establishing this result, I explore the implications of accounting for correlated shocks in this static framework. In this modified setting, sectors and firms are still exposed to risk from direct and indirect trade partners, but now can also substitute away from risk by having trade partners that are differentially exposed to supply and demand shocks. More specifically, the variance of a sector's growth rate inherits the standard network component which is related to the concentration of risk across trade partners, but also two additional components which capture a trade-off of concentration

[8] Generally the studies which rely on static models focus on propagation at business-cycle frequencies.

[9] Some examples include Acemoglu et al. (2012a), Acemoglu et al. (2016a), Ozdagli and Weber (2017), Herskovic (2018b), Herskovic et al. (2020b).

and substitutability *between* trade partners. Intuitively, higher concentration implies less diversification in supply-chains and should imply higher volatility. On the other hand, high substitutability implies that units can better diversify away the effects of shocks across customers or suppliers. In the following sections, I provide both theoretical and empirical motivation that researchers should account for correlated shocks when studying risk in network economies.

1.2.1 Networks and Risk Comovement

In this section, I argue that realistic input-output models of the economy should account for correlation in supply and demand shocks across units. In the benchmark static model of sectoral shock propagation, I find that the set of stable input-output networks that are consistent with the idiosyncratic risk assumption exhibit properties that are empirically implausible. Mathematically, in this broad class of reduced-form linear models, additional restrictions are required on the input-output network weights to be consistent with an arbitrary covariance matrix of sector or firm-level growth rates and an arbitrary diagonal covariance matrix of shocks. Although this result is not immediately intuitive, the assumption of idiosyncratic shocks implicitly generates a strict relationship on the interaction between network weights and elements of the variance-covariance matrix of growth rates.

To illustrate this point, I start from the general reduced-form model of shock propagation in which a firm's output growth is driven by a network component and firm-specific shocks.[10] In Appendix 1.7, I show that this model is consistent with the equilibrium outcome of a constant returns to scale economy in which Cobb-Douglas producers experience productivity shocks that propagate downstream from suppliers to customers and demand shocks that propagate upstream from customers to suppliers. In

[10] Similar models are used in Acemoglu et al. (2012a), Acemoglu et al. (2016a), Herskovic (2018a) and Herskovic et al. (2020b).

particular, for an n-firm economy, consider the commonly used static relationship:

$$\mathbf{y} = \mathbf{W}\mathbf{y} + \mathbf{u}, \tag{1.1}$$

where $\mathbf{y}$ is the $n \times 1$ vector of firm-level output growth, $\mathbf{W}$ is the $n \times n$ network matrix capturing interactions between industries, and $\mathbf{u}$ is the $n \times 1$ vector of firm-specific supply or demand shocks. This framework is compatible with either direction of propagation, upstream from customers to suppliers or downstream from suppliers to customers. The following two assumptions require that the propagation matrix $\mathbf{W}$ implies is stable, and that firm-specific shocks are idiosyncratic, respectively.

Assumption 1.2.1 (Stable Weighting Matrix)**.** *The weighting matrix $\mathbf{W} \in M_n$ is non-negative, and has bounded spectral radius $\rho(\mathbf{W}) \leq 1$.*

Assumption 1.2.2 (Idiosyncratic Shocks)**.** *Firm-specific shocks $u_i \sim \mathcal{P}_i(0, \sigma_i^2)$ are drawn independently across firms where $\mathcal{P}_i \in L^2$ has finite second moments. In other words, there exists a positive diagonal matrix $\mathbf{D} \in M_n$ such that $\mathbb{E}[\mathbf{u}\mathbf{u}^\top] = \mathbf{D}$.*

In the following proposition, I characterize a set of additional necessary restrictions on the matrix $\mathbf{W} \in M_n$ to satisfy Assumption 1.2.2 and (1.1). This is a novel result.

Proposition 1.2.3 (Necessary Restrictions on $\mathbf{W}$)**.** *Any weighting matrix $\mathbf{W}$ which satisfies 1.2.1 and 1.2.2 must have that for any pair of off-diagonal nodes w_{ij} with $i \neq j$, either $w_{ij}w_{ji} \geq 1$ or $w_{ij} > 0$ and $w_{ji} = 0$.*

Proof. See Appendix 1.9.1. ∎

This proposition highlights a key limitation of equation (1.1). To apply network models in a way that is consistent with industry and firm-level data, researchers must either restrict their focus to a very specific set of networks or allow for correlation in sectoral or firm-level shocks. Under Assumptions 1.2.1 and 1.2.2, consistent networks have only one-way connections or entries which are empirically implausible. In the structural model developed in Appendix 1.7, the entries of $\mathbf{W}$ are primitives of the production function and depend on each unit's sales and cost shares. For example, To capture the effect of demand shocks propagating from j to i, the implied weight is $w_{ij} = \frac{sales_{i\to j}}{sales_i}$. The first restriction in Proposition 1.2.3 is that $w_{ij}w_{ji} = \frac{sales_{i\to j}}{sales_i} \cdot \frac{sales_{j\to i}}{sales_j} \geq 1$ for all $i \neq j$, which implies that sectors which use each other's inputs can only use each other's inputs. If we believe that bilateral trade is possible and firms have multiple customers and suppliers, then this restriction is already violated. When weights are less than 1 (e.g., trade is shared across multiple partners), then the second restriction ($w_{ij} > 0$ and $w_{ji} = 0$) implies that no firm or industry can serve as both a customer and supplier. Input-output data in the US at both the industry and firm-levels do not satisfy these restrictions.

Intuitively, one might argue that the static network model in (1.1) is too parsimonious to capture all the sources of risk comovement in the economy. Although this is likely true, the restrictions on $\mathbf{W}$ cannot be relaxed by adding omitted macroeconomic factors driving common variation in risk nor by adding an omitted network component. Moreover, Proposition 1.2.3 implies that there is no sufficient statistic that can be obtained from $\mathbf{W}$ which fully characterizes cross-sectional variation in granular risk, even in a world where sectoral shocks are identically distributed. Another interpretation of this result is that firm and industry definitions are too broad for the idiosyncratic shock assumption to be viable. See Appendix 1.12 for supporting numerical evidence. In the remainder of this section, I explore the implications of allowing for correlation in demand and supply shocks across units.

1.2.2 Granular Variance with Correlated Shocks

I investigate the variance predictions of (1.1) when shocks u_i are allowed to be correlated across units i (i.e., $\mathrm{var}(\mathbf{u})$ is not diagonal). Practically speaking, there are several reasons why supply and demand shocks to units might be correlated, especially at the granular level. For example, two sectors or firms that produce related goods or services are likely to experience correlated demand shocks. If the two sectors produce using the same inputs, then supply-side shocks might be correlated as well. Hoberg and Phillips (2016) show that firms with similar products might belong to different industries (according to SIC or NAICS classifications).[11] In the even more simple setting where multiple firms produce the exact same goods and services, supply and demand shocks at the product level mechanically generate correlation in supply and demand shocks at the firm level. Hottman et al. (2016) provide empirical evidence that this is generally the case, with 69% of firms, which account for 99% of industrial output, supplying multiple (and intersecting) products.

Like product proximity, both technological and geographic proximity might also generate correlation in firm and sectoral shocks. For instance, Bloom and Shankerman (2013) show that shocks to research and development (R&D) have correlated effects on the productivity and growth of firms with similar technologies. Similarly, industries or firms operating in nearby geographies are exposed to the same underlying shocks associated with local labor markets (see e.g., Autor et al. (2013), Mian and Sufi (2014)), local factor prices (Tuzel and Zhang (2017), Grigoris (2019)), local technological progress (Oberfield (2018)), or local weather events (Barrot and Sauvagnat (2016), Kruttli et al. (2019)).

In the benchmark network model with uncorrelated shocks, the variance of growth rates depends solely on the concentration of risk across independent suppliers and/or customers. Herskovic et al. (2020b) provide theoretical and empirical evidence linking firm volatility and customer concentration in terms of size dispersion in this setting. However, allowing for correlated shocks implies two additional variance components. These components capture the concentration and substitutability of risk *between* trade partners, respectively. The distinction between concentration "across" and "between" trade partners is important. Concentration across refers to the composition of a unit's reliance on any particular customer or supplier, while concentration between refers to the distribution of reliance on a set of customers or suppliers that are exposed to the same shocks. On the

[11] More specifically, Hoberg and Phillips (2016) find that firms in the newspaper, printing, and publishing industry (SIC3 271) are similar to firms in the radio broadcasting industry (SIC3 483) and argue that this is driven by common customers who demand advertising services. They also find that Disney and Pixar have similar products (movies) although they are in different industries (business services (SIC3 737) and motion pictures (SIC3 781) industries, respectively). In this case, the differences in industry stem from the production method and not the product offering.

other hand, substitutability between customers and suppliers captures the distribution of reliance on a diversified set of customers or suppliers that are exposed to shocks of the opposite sign.

In other words, concentration between customers and suppliers captures compounding effects of positively related shocks to similar trade partners, while substitutability captures mitigating effects of spreading reliance on trade partners that are exposed to negatively related shocks. Intuitively, a supplier with major customers that tend to reduce demand at the same time is more risky than a supplier with some customers that increase demand when the others reduce it. To see this mathematically, define $[h_{ij}]_{ij}$ to be the set of entries in the Leontief inverse matrix $\mathbf{H} := (\mathbf{I}-\mathbf{W})^{-1}$ and recall that equation (1.1) can equivalently by written as $\mathbf{y} = \mathbf{Hu}$. Note that in this setup the element h_{ij} captures the percent change in unit i's growth after a 1% shock to unit j. Then the variance of unit i's growth rates can be written:

$$\text{var}(y_i) = \text{var}\left(\sum_{j=1}^{n} h_{ij}u_j\right) = \sum_{j=1}^{n} h_{ij}^2 \cdot \text{var}(u_j) + \sum_{j\neq k} h_{ij}h_{ik} \cdot \text{cov}(u_j, u_k).$$

The first term is the standard expression for variance in this network model (see e.g., Acemoglu et al. (2012a)), while the second term is only non-zero when inter-industry shocks are correlated. Next, I define the scalar s_{jk} to be the sign of the pairwise correlation between shocks to j and k (i.e., $s_{jk} := \text{sgn}(\sigma_{jk}) \equiv \text{sgn}(\text{cov}(u_j, u_k))$ where sgn(.) is the sign function). To build some more intuition on the additional terms, I can further decompose the covariance term as follows:

$$\text{var}(y_i) = \underbrace{h_{ii}^2 \cdot \sigma_i^2}_{\text{self}} + \underbrace{\sum_{j\neq i} h_{ij}^2 \cdot \sigma_j^2}_{\text{concentration “across”}} + \underbrace{\sum_{j\neq k, s_{jk}=1} h_{ij}h_{ik} \cdot \sigma_{jk}}_{\text{concentration “between”}} + \underbrace{\sum_{j\neq k, s_{jk}=-1} h_{ij}h_{ik} \cdot \sigma_{jk}}_{\text{substitutability}} \quad (1.2)$$

Consider a first order approximation of the Leontief inverse matrix such that $\mathbf{H} \approx \mathbf{I}+\mathbf{W}$, where the weights in the propagation matrix $\mathbf{W}$ are related to sales shares when modeling downstream propagation supply-side shocks, and purchase shares when modeling upstream propagation of demand-side shocks.[12] Suppose also that units are homoskedastic such that $\text{var}(u_j) = \sigma^2$ for all j and $\text{cov}(u_j, u_k) = \nu \cdot s_{jk}$ for all $j \neq k$, where σ and ν are positive scalars. In this case, the first term ($h_{ii}\sigma_i^2 = \sigma_i^2$) is unrelated to the network and captures the variance of supply or demand shocks to sector i. On the other hand, the second component (concentration across network linkages) is non-negative and large when reliance is highly concentrated across trade partners. Similarly, the third term (concentration between network linkages) is non-negative and large when reliance is concentrated between trade partners who experience positively correlated shocks.

Finally, the last term (substitutability of network linkages) is always non-positive and is large in magnitude when reliance is spread equally between trade partners who are likely to experience shocks of opposite sign. Additionally, the sum of the final two terms captures explicitly the trade-off between concentration and substitutability of correlated supply or demand shocks. Although this simplification is useful for building intuition, the more realistic version of the variance decomposition should also take into account unit heteroskedasticity. That is, two sectors with an equal set of input-output weights have different network-implied variance only if their trade partners are exposed to differential volatility in supply or demand shocks.

Consider for example the Printed Circuit Boards industry (SIC 3672), whose top 3 major manufacturing industry customers include Electronic Components (SIC 3679) and Electronic Computers (SIC 3571), and Communications Equipment (SIC 3669). At first glance, these customers appear very similar, and one might suspect that a negative demand shock to one customer is likely to be correlated with a negative shock to the other, which amplifies upstream propagation to their shared supplier. In other words, the Printed Circuit Boards industry has high concentration between customers and a harder time substituting

[12] More specifically, the weight of downstream propagation of supply-side shocks from supplier j to customer i is captured by $w_{ij}^d = sales_{j \to i}/purchases_i$ and the weight of upstream propagation of demand shocks from customer j to supplier i is captured by $w_{ij}^u = sales_{i \to j}/sales_i$. In general, these weights are both asymmetric (i.e., $w_{ij} \neq w_{ji}$) and different depending on the direction of propagation (i.e., $w_{ij}^u \neq w_{ij}^d$).

away from upstream effects demand shocks to its major customers.[13] On the other hand, the three most important customers of the Jewelry and Precious Metal industry (SIC 3911) include Watches, Clocks, and Clockwork Operated Devices (SIC 3873), Perfumes and Cosmetics (SIC 5048), and Drawing and Insulating of Nonferrous Wire (SIC 3357). In this case, customers produce seemingly unrelated goods (both durable and non-durable) and there is evidence the demand shocks have zero or negative pairwise correlation.[14] In other words, the Jewelry and Precious Metal industry is able able to substitute away from demand shocks propagating upstream from any individual customer.

There are similar examples of high concentration and substitutability on the supply-side. For instance, the Computer Storage Devices industry (SIC 3572) has a highly concentrated customer base composed of Electronic Components (SIC 3679), Electronic Coils, Transformers, and other Inductors (SIC 3677), Semiconductors and Related Devices (SIC 3674), and Electronic Connectors (SIC 3678). This industry is thus more exposed to correlated supply-side risk. On the contrary, the Meat Packing Plants industry (SIC 2011) can more easily substitute away from supply-side risk, with a more diversified set of major suppliers like Poultry Slaughtering and Processing (SIC 2015), Plastics Film and Sheet (3081), and Paper Mills (2621).

Although these network variance components are intuitive and theoretically justified if supply and demand shocks are correlated, an important practical concern is that granular supply and demand shocks are not easily identified from available data, especially at a high frequency. In the following section, I address this challenge and propose an empirical methodology for estimating customer and supplier concentration and substitutability at the industry and firm levels. I show that both supply and demand channels explain cross-sectional heterogeneity in risk exposure, beyond what can be explained by other determinants of variance identified by the literature.

[13] I find that the average product similarity score between these customer industries is in the top 10% (based on the similarity score developed in Hoberg and Phillips (2016)). Additionally, I find significant positive correlation in demand shocks to these industries such as changes of newly awarded federal defense procurement contracts.

[14] The average pairwise correlation in federal procurement shocks and Chinese import penetration shocks is -37% and -22%, respectively for the full set of Jewelry and Precious Metal customers.

1.3 Empirical Evidence

In this section, I provide empirical estimates of the network-implied variance components motivated in equation (1.2). This requires granular data on input-output relationships and estimates of the variance-covariance matrix of supply and demand shocks. Consistent with theory, I find that these additional components explain important variation in realized volatility, controlling for characteristics such as size, centrality, concentration across trade partners, vertical position in the supply chain, and durability of output. These results hold at both industry and firm levels. The main takeaway here is simple. When accounting for input-output linkages and non-negligible correlation in supply and demand shocks, heterogeneity in risk exposure is at least in part driven by differences in the ability of network units to substitute away from correlated supply and demand shocks.

1.3.1 Setup

Consider the n-sector network model from equation (1.1) and add a time subscript t. Suppose that in each period I obtain estimates for the $n \times n$ Leontief Inverse matrices $\hat{\mathbf{H}}_{q,t}$ and the variance-covariance matrices of supply and demand shocks $\hat{\boldsymbol{\Sigma}}_{q,t}$ where $q = u$ for upstream propagation and $q = d$ for downstream propagation. Then for both supply and demand-side shocks, I can compute three empirical network-implied variance components, denoted by “self-originating”, “across”, and “between” risk. The final component sums the final two covariance terms from (1.2) and captures the concentration/substitutability trade-off between trade partners. Low values of concentration between customers and suppliers implies high substitutability. Then for each industry, direction, and time triple (i, q, t), I compute self-originating risk as:

$$\hat{\sigma}^2_{iqt,self} = [\hat{h}_{qt}]^2_{ii} \cdot \hat{\sigma}^2_i, \tag{1.3}$$

and across and between risk as:

$$\hat{\sigma}^2_{iqt,acr} = \sum_{j \neq i} [\hat{h}_{qt}]^2_{ij} \cdot \hat{\sigma}^2_j, \tag{1.4}$$

$$\hat{\sigma}^2_{iqt,bet} = \sum_{j \neq k} [\hat{h}_{qt}]_{ij} \cdot [\hat{h}_{qt}]_{ik} \cdot \hat{\sigma}_{jk}. \tag{1.5}$$

In the next section, I provide details on the data sources, assumptions, and methodologies used to estimate $\hat{\mathbf{H}}_{qt}$ and $\hat{\mathbf{\Sigma}}_{qt}$. While the former can be observed directly, I need to make some assumptions to identify the latter from available data sources. Then I compute all three components at the industry-level and study their empirical relationship with realized industry variance. I find that the elasticity of realized variance to all three components is significant and positive for both directions of propagation, controlling for a variety of industry characteristics.

1.3.2 Upstream and Downstream Propagation Networks

I begin by constructing the network of input-output linkages at the disaggregated industry level from the Make and Use tables published by the Bureau of Economic Analysis (BEA). The goal is to build a directed weighted network which captures the importance of trade relationships over time and for the population of industries.[15] Network weights represent the strength of each unit's reliance on customers and suppliers, and the network is directed to capture differences in shock propagation in the upstream (customer to supplier) and downstream (supplier to customer) directions. In particular, the BEA publishes these tables annually between 1997-2020 for 66 industry groups.

More specifically, I construct downstream and upstream propagation matrices $\mathbf{W}_d = [w_d]_{ij}$ and $\mathbf{W}_u = [w_u]_{ij}$ with entries:

$$[w_d]_{ij} = \frac{sales_{j \to i}}{costs_i}, \qquad [w_u]_{ij} = \frac{sales_{i \to j}}{sales_i}, \tag{1.6}$$

[15] As far as I know, this is the most disaggregated database on the entire population of input-output relationships.

where $sales_{i \to j}$ represents gross trade flows from i to j, and $sales_i$ and $costs_i$ represent the total sales and costs of industry i, respectively. The downstream (upstream) weights are non-negative and capture the direct reliance of industry i on supplier (customer) j. When weights are large, direct effects of propagated shocks should also be large. To account for higher order (indirect) network effects as well, I calculate the strength of network propagation based on the Leontief inverse $\mathbf{H}_q := (\mathbf{I} - \mathbf{W}_q)^{-1}$ of the propagation matrices $\{\mathbf{W}_q : q \in u, d\}$.[16] The entries of $\mathbf{H}_q = [h_q]_{ij}$ capture the total percent effect on i of a 1% shock to j traveling in the q-stream direction when accounting for all weighted direct and indirect connections.

Appendix 1.10.2 reports summary statistics for observed input-output connections. At the 66-industry granularity, tables are updated annually. I find that both propagation and Leontief inverse weights are highly persistent with an average annual autocorrelation of more than 95% for each entry. The cross-sectional correlation is about 8.55% between upstream and downstream weights and about 11.08% between upstream and downstream Leontief matrix entries, suggesting that propagation occurs differently in either direction.

1.3.3 Variance-Covariance Matrix of Supply and Demand Shocks

Unlike the sectoral input-output network, there is no definitive data source on supply and demand shocks and their variance-covariance matrix. As a baseline, I implement an empirical analog of the reduced-form equation in (1.1). In particular, consider the spatial panel regression:

$$\tilde{y}_{it} = \delta_t + \phi \cdot \tilde{y}_{i,t-1} + \beta_u \cdot \sum_{j=1}^{n} w_{u,ij} \tilde{y}_{jt} + \beta_d \cdot \sum_{j=1}^{n} w_{d,ij} \tilde{y}_{jt} + \varepsilon_{it}, \tag{1.7}$$

where $\tilde{y}_{it} := \log(y_{it}/y_{i,t-1})$ is output growth in industry i at quarter t and $w_{q,ij}$ is the (i,j) entry of the q-stream propagation matrix $\mathbf{W}_q$. Assuming the variance-covariance matrix of residuals is static, then $\hat{\mathbf{\Sigma}}$ is the empirical variance-covariance matrix of estimated residuals

[16] See e.g., Baqaee and Farhi (2019) and Herskovic et al. (2020a) for discussion on the importance of higher order network effects.

$\hat{\varepsilon}_{it}$. To ensure that my estimates for network components (1.4) and (1.5) are robust to estimation error in $\hat{\boldsymbol{\Sigma}}$, I calculate the average value over samples in which I randomly drop 10% of pairwise non-zero correlations.[17] See Appendix 1.11 for more details and alternative specifications.

Pairwise correlation in residuals is centered with a mean value of 0.5% (0.4%) and a standard deviation of 25% (26%). The largest positive pairwise correlation is 82% between the Primary Metals (BEA Code 331) and Wholesale Trade (BEA Code 42) and 81% between Housing (HS) and Educational Services (61). On the other hand, the largest negative pairwise correlation is -80% between Primary Metals (331) and Federal Reserve Banks, Credit Intermediation, and Related Activities (521CI) and -72% between Food and Beverage and Tobacco Products (311FT) and Wholesale Trade (42).

1.3.4 Network Determinants of Realized Variance

After relaxing the idiosyncratic shock assumption, the benchmark input-output propagation model predicts that realized variance should depend positively on three network risk components: risk that is self-originating, risk across trade partners, and risk between trade partners. In the baseline setup, this might hold mechanically for self-originating risk since it is estimated from the variance of residual output growth in equation (1.7). However, both risk across and between trade partners contain only variance-covariance information associated with other industries. I verify these predictions empirically using panel regressions of the log of realized industry variance on the log of network components, controlling for a variety of characteristics such as size, centrality, durability of output, and industry cluster and time fixed effects.[18] I measure realized industry variance using both stock market and output growth data. I define market variance as the annual return variance of an equal-weighted industry portfolio and fundamental variance as the variance of quar-

[17]Note that estimation error from $\hat{\boldsymbol{\Sigma}}_q$ is magnified in estimated network components (1.4) and (1.5) at a rate proportional to the number of nonzero row entries in the Leontief inverse matrix $\mathbf{H}_q$.

[18]I adjust network components by a constant to ensure that the minimum value is positive so the log is well defined. Industry clusters are defined by major industry groups (2-digit NAICS code).

terly year-on-year output growth. I obtain similar results when using idiosyncratic variance as the dependent variable.[19] Although the annual variance across quarterly year-on-year output growth and monthly returns are fairly noisy proxies for true realized cash-flow variance, the results are robust for several specifications.

I summarize the main results in Table 1.1. Consistent with theoretical predictions in equation (1.2), the elasticity of realized industry variance to concentration across and between customers are both positive and significant in all specifications. This holds for both market and output growth measures of variance. Conditional on both directions of propagation and all controls, increasing concentration between customers from the median to the 90^{th} percentile increases industry sales growth variance by over 45% (about 0.37 standard deviations) and market variance by 15% (about 0.09 standard deviations). Similarly, increasing concentration between suppliers from the median to the 90^{th} percentile increases industry sales growth variance by over 20% (about 0.17 standard deviations) and market variance by 19% (about 0.11 standard deviations). Without controls, downstream network risk explains 23% of time series variation in market variance and 22% of time series variation in output growth variance. Similarly, upstream network risk explains 22% and 31% of market and output growth variance, respectively. Both directions of propagation are important for explaining the panel dynamics of industry variance.

Consistent with the firm-level findings of Herskovic et al. (2020b), I find that industry variance has a positive elasticity to concentration across customers and a negative elasticity to average size. A new but related result is the positive elasticity of variance to concentration across suppliers. Additionally, Ahern (2013) argues that more central industries have greater market risk since they are more exposed to aggregate shocks, and thus earn higher returns on average. On the other hand, my results suggest that more central industries have less volatile stock returns, but also have less exposure to aggregate

[19] I define idiosyncratic market variance as the variance of equal weighted residual returns from a Fama and French three-factor model. Similarly, I define idiosyncratic output growth as the residual of industry output growth after a regression on aggregate output growth. Results also replicate for value-weighted industry portfolios, or industry sales growth, which is constructed as the year-on-year change in the sum of quarterly sales (reported on Compustat) for all public firms in the industry.

volatility risk and lower idiosyncratic volatility.[20] My results are thus consistent with Ahern (2013), since stocks with lower exposure to aggregate volatility risk and lower idiosyncratic volatility earn higher returns on average (see e.g., Ang et al. (2006)). Table 1.7 shows that there is no significant relationship between centrality and concentration between or across trade partners.

1.3.5 Sources of Correlation in Network Propagating Shocks

So far, I have established both theoretically and empirically the importance of accounting for correlation in shocks that propagate through the input-output network. In particular, I show that concentration between trade partners explains a large amount of variation in the industry panel of realized variance. However, statistical estimates for the variance-covariance matrix of shocks do not provide much insight on the underlying sources of correlation between industries. In this section, I argue that correlation in supply-side shocks that propagate downstream can be explained by technological proximity between sectors, while correlation in demand-side shocks that propagate upstream can be explained by product similarity.

Observed Supply and Demand Shocks

Acemoglu et al. (2016a) argue that productivity shocks primarily propagate downstream while government spending and trade shocks primarily propagate upstream. In this case, these shocks might help to capture differences in inter-industry correlations which are specific to the direction of propagation. Along these lines, I construct an annual industry panel of 5-factor total factor productivity (TFP) growth between 1959-2018 from the NBER-CES Database (Becker et al., 2016). Since this measure of TFP controls for materials, it does not mechanically encode any information related to downstream effects

[20] I find that industries in the highest average upstream (downstream) centrality decile have 31% (21%) less exposure to systematic volatility risk than the lowest decile. Average upstream and downstream centrality are positively correlated (56% cross-sectionally), and industries who are in the top decile for both average centrality measures have a 52% lower exposure to aggregate volatility risk than industries in the bottom decile for both. Moreover, top centrality decile stocks have 25% lower idiosyncratic volatility, on average.

such as changes in price and/or quantity. Similarly, I construct a monthly panel of newly awarded federal procurement contracts between Jan 2000-Jan 2021 from the universe of contracts published in the Federal Procurement Data System (FPDS).[21].

To focus on inter-industry correlations which are unrelated to common aggregate factors (e.g., the secular decline in several manufacturing industries), I estimate the variance-covariance matrix of residuals after an OLS regression on the cross-sectional average of shocks.[22] I then estimate the corresponding network components using equations (1.4) and (1.5) and study their relationship with realized variance. Table 1.8 shows that the elasticity of realized variance to supplier concentration is positive and more precisely estimated when calculating supply-side shock covariance as a function of productivity growth. On the other hand, Table 1.9 reports more precise estimates for the elasticity of variance to customer concentration when calculating demand shock covariance as a function of federal procurement shocks. This suggests that productivity growth is more informative about upstream network risk, while changes in government demand are more informative about downstream network risk. This is consistent with the results of Acemoglu et al. (2016a).

Technological and Product Proximity

Suppose now that correlation between supply and demand-side shocks is a function of underlying firm and industry characteristics. Intuitively, I might assume that correlation in demand shocks propagating upstream is driven by product similarity and/or geographic proximity of customers, while correlation in supply shocks propagating downstream is driven by technological similarity and/or geographic proximity of suppliers.

More generally, I assume that each industry is associated with a vector of positions in some latent surface $\mathbf{z}_{it} \in \Omega_z \subset \mathbb{R}^d$ and that the correlation between industry shocks can be

[21] I also consider other observed shocks in Appendix 1.11

[22] The cross-sectional mean approximates the first principal component of shocks when there are missing values. For shocks dz_t, I calculate the covariance between sectors k and j as $\text{cov}(u_{kt}, u_{jt})$ where u_{kt} is the residual in the regression $dz_{kt} = \alpha + \beta \cdot \bar{dz}_{.,t} + u_{kt}$. Endogeneity of shocks is not a major concern assuming any confounding shocks largely propagate in the same direction in the network. Given such a confounder, my estimate for the variance-covariance matrix of shocks can be written as the true estimate plus some measurement error.

written as a function of the distance between these latent vectors.[23] Under the assumption of a factor structure of risk (see e.g., Herskovic et al. (2016)), aggregate movements in this latent space be interpreted as factor innovations. Following McCormick and Zheng (2015), I suppose that industry positions $\mathbf{z}_{it}$ lie on the surface of a p-dimensional latent surface on the $p+1$-dimensional unit hypersphere $\mathcal{S}^{p+1}$. This implicitly implies that latent positions follow a uniform distribution across the sphere's surface. Moreover, since the hypersphere has bounded surface area, the distance between any two points is bounded. I further assume that points in the same position have correlation 1 and points on opposite sides of the sphere have correlation -1.

In practice, I experiment with constructing latent positions of industries using several combinations of industry variables. For simplicity, my main results rely on univariate distances in product and technology space.[24] I measure product distance using use the text-based scores developed in Hoberg and Phillips (2016) and technology distance using patent-based technological proximity scores along the lines of Bloom and Shankerman (2013). Since both of these scores are available at the firm-level, I first construct a firm-by-firm product distance network where distances are inversely related to proximity. To get the distance between sectors, I use the median length of the shortest weighted path between firms in the two sectors, rescaled such that the furthest pairwise distance is 1 and the shortest pairwise distance is zero. I calculate the shortest pairwise path between any two nodes using Dijkstra's Algorithm. I calculate these measures annually.

Transforming distances to correlations, I rescale by the variance estimates from residuals in equation (1.7) and recompute network components. When using product distance to calculate network risk, the elasticity of realized variance to concentration between customers is significant and positive. On the other hand, the analogous elasticity to concentration between suppliers is significantly negative, which suggests that product similarity across suppliers actually indicates better substitutability away from supply-

[23]Latent surface models are often used to impute network relationships in microeconomic applications (see e.g., McCormick and Zheng (2015), Breza et al. (2020)).

[24]Moreover, contours of the sphere present some calibration difficulties in higher dimension.

side shocks. When approximating correlations based on technological distance, realized variance has a positive elasticity to concentration between suppliers and customers, but the elasticity is more precisely estimated on the supply side. Taken together, these results suggest that technological proximity is a good proxy for correlated exposure to supply shocks propagating downstream, while product proximity is a good proxy for correlated exposure to demand shocks propagating upstream. Along these lines, Table 1.2 shows that average technological proximity between sectors is closely related to correlation in TFP growth shocks, while product similarity is closely related to correlation in federal defense procurement shocks.

Accounting for Dynamics

To account for potential time-variation in industry correlations, I also compute pairwise inter-industry correlations using the dynamic conditional correlation (DCC) estimator from Engle (2002). In particular, I estimate bivariate DCC models for all pairs of industries using Bayesian MCMC following Fioruci et al. (2013). Since product and technological proximity are computed at an annual frequency, I am thus able to obtain a one-to-one comparison between correlation in spatial panel residuals, observed shocks, and distance based measures. Table 1.2 shows that accounting for dynamic changes in industry correlations actually strengthens the consistency between network components estimated in different ways.

More specifically, I find that networks components constructed using the dynamic TFP variance-covariance matrix are more highly correlated with technological-proximity-based components than their static variance-covariance counterparts. Similarly, network components constructed using the dynamic procurement variance-covariance matrix are more highly correlated with product proximity than their static counterparts. This suggests that there are meaningful dynamics underlying these components which show up both in observed shocks and fluctuations in latent product and technology space.

Table 1.1: Network Determinants of Industry Variance

Panel A: Market Return Variance

	(1)	(2)	(3)	(4)	(5)	(6)
Self-origin (demand)	0.055** (0.022)	0.056** (0.024)			0.003 (0.028)	0.017 (0.028)
Across (demand)	0.094** (0.036)	0.083** (0.030)			0.081** (0.024)	0.072** (0.021)
Between (demand)	0.147*** (0.051)	0.122** (0.049)			0.197*** (0.053)	0.089** (0.038)
Self-origin (supply)			0.072*** (0.021)	0.073*** (0.023)	0.067*** (0.022)	0.083** (0.023)
Across (supply)			0.216*** (0.056)	0.156*** (0.054)	0.160*** (0.060)	0.154* (0.067)
Between (supply)			0.416*** (0.093)	0.317*** (0.095)	0.210** (0.088)	0.196** (0.073)
Size		-0.378*** (0.094)		-0.361*** (0.091)		-0.289*** (0.104)
Upstream centrality		-0.182* (0.103)		-0.289*** (0.095)		-0.232** (0.101)
Downstream centrality		-0.051 (0.054)		-0.086** (0.043)		-0.023 (0.054)
Durability		-0.167 (0.573)		-0.404 (0.662)		-0.637 (0.615)
Vertical position		1.550** (0.701)		-0.756** (0.332)		1.970*** (0.710)
Constant	-6.279	-3.26	-4.692	-1.082	-5.005	-3.608
Obs	1484	1484	1484	1484	1484	1484
Adj R2	0.231	0.292	0.159	0.223	0.245	0.330

Panel B: Output Growth Variance

	(1)	(2)	(3)	(4)	(5)	(6)
Self-origin (demand)	0.034 (0.026)	0.006 (0.026)			0.026 (0.030)	0.006 (0.028)
Across (demand)	0.159** (0.069)	0.074 (0.024)			0.157** (0.070)	0.095 (0.057)
Between (demand)	0.210*** (0.064)	0.281*** (0.078)			0.196*** (0.065)	0.280*** (0.082)
Self-origin (supply)			0.081** (0.021)	0.072 (0.022)	0.017 (0.024)	0.042 (0.025)
Across (supply)			0.128** (0.053)	0.114** (0.033)	0.098** (0.021)	0.091** (0.031)
Between (supply)			0.332** (0.098)	0.241*** (-0.079)	0.221** (0.103)	0.198** (0.092)
Size		-0.046 (0.099)		-0.185* (0.098)		-0.110 (0.106)
Upstream centrality		-0.127 (0.108)		-0.096 (0.095)		-0.131 (0.111)
Downstream centrality		-0.218*** (0.056)		-0.067 (0.042)		-0.222*** (0.056)
Durability		0.115 (0.647)		-0.038 (0.643)		0.024 (0.703)
Vertical position		1.545** (0.654)		0.181 (0.355)		1.375** (0.682)
Constant	-5.088	-6.954	-6.818	-5.327	-6.17	-6.718
Obs	1484	1484	1484	1484	1484	1484
Adj R2	0.221	0.382	0.198	0.319	0.277	0.412

Notes: This table reports panel regressions of realized industry variance on a variety of characteristics, including the log variance of supply and demand shocks (self-origin), log concentration across trade partners, log concentration between trade partners, log total output (size), log centrality of the upstream and downstream propagation networks, durability of output, vertical position in the supply chain, and industry cluster and year fixed effects. In Panel A, the dependent variable is the log variance of annualized monthly returns on an equal-weighted industry portfolio. In Panel B, the dependent variable is the log variance of total quarterly year-on-year industry sales growth. I obtain return data from CRSP and GDP data from the BEA. Concentration between and across trade partners is calculated as the average value of estimates obtained using equations. 1.4 and 1.5. I calculate the variance-covariance matrix of shocks using residuals from equation 1.7 and calculate the average value of components over 1000 random samples each randomly dropping 10% of pairwise correlations. Following Ahern (2013), I compute industry centrality as the eigenvector centrality of upstream and downstream propagation adjacency matrices. I calculate durability as the proportion of sub-industries classified as durable by Gomes et al. (2009), and I calculate vertical position of each industry following Antràs et al. (2012) and Gofman et al. (2020). $***$, $**$, and $*$ indicate significance at the 1%, 5%, and 10% levels, respectively. Standard errors are clustered at the BEA major 15 major industry group level. Sample is at an annual frequency from 1997 to 2019 for 66 BEA non-government industries.

Table 1.2: Sources of Correlation in Supply and Demand Substitutability

Panel A: Upstream Supplier Substitutability (static)

Covariance Method	Spatial	TFP	Procurement	Prod Similarity	Tech Proximity
Spatial	1	0.59	0.37	0.30	0.48
TFP		1	0.31	0.35	0.46
Procurement			1	0.62	0.39
Prod Similarity				1	0.27
Tech Proximity					1

Panel B: Downstream Customer Substitutability (static)

Covariance Method	Spatial	TFP	Procurement	Prod Similarity	Tech Proximity
Spatial	1	0.55	0.61	0.43	0.35
TFP		1	0.27	0.20	0.48
Procurement			1	0.69	0.40
Prod Similarity				1	0.24
Tech Proximity					1

Panel C: Upstream Supplier Substitutability (dynamic)

Covariance Method	Spatial	TFP	Procurement	Prod Similarity	Tech Proximity
Spatial	1	0.61	0.40	0.34	0.49
TFP		1	0.35	0.39	0.48
Procurement			1	0.65	0.42
Prod Similarity				1	0.25
Tech Proximity					1

Panel D: Downstream Customer Substitutability (dynamic)

Covariance Method	Spatial	TFP	Procurement	Prod Similarity	Tech Proximity
Spatial	1	0.57	0.68	0.45	0.30
TFP		1	0.22	0.27	0.52
Procurement			1	0.70	0.38
Prod Similarity				1	0.30
Tech Proximity					1

Notes: This table reports the correlation across different measures of upstream and downstream substitutability measures (negative of concentration "between"). Substitutability is calculated as the negative value of the log of (1.5) (plus a large enough constant) and correlation between two measures x_{it} and y_{it} is defined by $\hat{\rho}$ from the regression $y_{it} = \hat{\rho} x_{it} + u_{it}$, where x_{it} and y_{it} are transformed to have mean zero and standard deviation one. Panels A and B rely no a static assumption for the variance-covariance matrix across shocks, while Panels C and D estimate a dynamic conditional variance-covariance matrix à la Fioruci et al. (2013). Panels A and C report results in the upstream (supply-side) direction and Panels B and D report results in the downstream (customer-side) direction. Spatial covariance is based on the panel model in equation (1.7). TFP covariance is based on TFP growth measured in Becker et al. (2016), procurement based on federal government shares interacted with the procurement proxy in Briganti and Sellemi (2022), product similarity using the latent distance method and scores from Hoberg and Phillips (2016), and tech proximity constructed following Bloom and Shankerman (2013).

1.4 Dynamic Network Model of Supply-Chain Substitutability

In this section, I incorporate the supply chain substitutability-concentration trade-off and my empirical results from the previous section in a structural dynamic asset pricing model detailed in Appendix 1.7. This model builds on existing production-based models with input output networks (see e.g., Ramìrez (2017), Herskovic (2018b), or Gofman et al. (2020)). Unlike existing models, I introduce a correlation structure in shocks to firm growth rates which propagate both upstream and downstream in an input-output network.[25] Firms are subject to both productivity shocks which propagate downstream and demand shocks which propagate upstream. Shocks are drawn from a joint distribution with finite second moments and in which correlation across firm-level shocks is induced by shared variation in firms' input-output substitutability. More specifically, firms' ability to substitute away from productivity (demand) shocks is inversely related to concentration of trade partners in latent technology (product) space and correlated between firms who share trade partners.

1.4.1 Setting

Consider a discrete-time economy with n distinct goods and n firms. Output goods are characterized by vector in product-technology space, which is fixed exogenously for each good. Firm output (cash flow) depends on aggregate economic conditions and the cash flows of its customers and suppliers. There are two kinds of random shocks in this economy, productivity shocks which propagate downstream from suppliers to customers, and demand shocks which propagate upstream from customers to suppliers. The input-output network is captured by two sequences of graphs with n nodes for each firm and weighted directed edges capturing the importance of firm-to-firm trade relationships. In the customer (supplier) network, the edge from i to j represents the relative reliance of j

[25]To my knowledge, this is the first network model to feature both correlated shocks and two directions of propagation.

on customer (supplier) i.

For tractability, I assume that trade relationships are exogenously determined at the start of each period. Additionally, the model features a representative investor with constant relative risk aversion (CRRA) preferences who owns all firms and lives off labor wages and dividends. Next, I describe the process for firm cash flows, the network structure, and the mechanism of shock propagation through the input-output network. Then I derive equilibrium consumption growth and asset prices. For ease of exposition, I provide details on the production side of the economic since that is the primary source of risk. Further details are left to Appendix 1.7.

1.4.2 Substitutability and Firms' Cash Flows

Firms are exposed to undiversifiable aggregate risk factors and risk from trade partners which can be mitigated with diversification of customers and suppliers. Every firm is both a customer who purchases inputs from other firms, and a supplier who produces a single final good. Final goods are characterized by a latent position in technology-product space, which fluctuates according to a persistent stationary process discussed in Appendix 1.8.[26] Latent position dynamics are exogenous to firm and household decisions and can be interpreted as random changes in product differentiation. For example, Syverson (2004) argues that the same products might be perceived differently as a result of intangible factors like delivery speed, documentation, product support, or branding and advertising.

In reduced-form, firm cash flows are determined by random shocks each period which propagate stochastically both downstream to customers and upstream to suppliers. The probability that a shock propagates through the supply chain is a function of firms' customer and supplier substitutability. Consistent with the empirical results from Section 1.3, I assume that a firm's customer substitutability depends on the product diversity of the goods sold by its customers. Likewise on the supply-side, a firm's suppliers substitutability depends on the technological diversity of its suppliers. When supply chains are highly

[26] This assumption is justified empirically by the results of Section 1.3, which suggest both that distances in product space and technology space change over time.

substitutable, shocks are less likely to propagate.

In particular, firm cash-flow growth has the following reduced-form equation:

$$\Delta y_{i,t+1} = \Delta z_{i,t+1} + \Delta g_{i,t+1}, \tag{1.8}$$

where $\Delta z_{i,t+1} = \log(z_{i,t+1}/z_{it})$ is a shock to productivity and $\Delta g_{i,t+1} = \log(g_{i,t+1}/g_{it})$ is a shock to government demand. I assume that dependence across shocks is determined by both the firm's input-output network and the relative location of its final good in product-technology space. Productivity growth follows the process:

$$\Delta z_{i,t+1} = \gamma_u \cdot a_{t+1} - \beta_u \cdot \varepsilon_{iu,t+1}, \tag{1.9}$$

where $a_t \sim_{iid} \mathcal{N}(0, \sigma_a^2)$ is aggregate productivity growth at time t, γ_u and β_u are positive scalars, and ε_{iut} is a Bernoulli shock that negatively affects productivity and originates upstream. Similarly, government demand growth follows the process:

$$\Delta g_{i,t+1} = \gamma_d \cdot g_{t+1} - \beta_d \cdot \varepsilon_{id,t+1}, \tag{1.10}$$

where $g_t \sim_{iid} \mathcal{N}(0, \sigma_g^2)$ is aggregate growth in government spending at time t, γ_d and β_d are positive scalars, and ε_{idt} is a Bernoulli shock which negatively affects demand and originates downstream. In other words, ε_{idt} (ε_{iut}) is equal to one when firm i experiences a demand (supply) shock which originates at i and/or propagates from its downstream customers (upstream suppliers). Shocks propagating in different directions are independent (i.e., $\varepsilon_{idt} \perp \varepsilon_{iut}$ for all i and t).

1.4.3 Network Structure

The propagation of shocks depends on the sequence of input-output network connections between firms, defined as follows. The sequence of upstream and downstream graphs $(\mathcal{G}_{n,u,t})_n$ and $(\mathcal{G}_{n,d,t})_n$ are n-node graphs with weighted edges given by w^u_{ijt} and

w_{ijt}^{d}, respectively. Weights capture the importance of the directed relationship $i \to j$ from the perspective of i and are fixed exogenously at the start of period t.

To ensure that the input-output network is realistic, I assume that all weights are between 0 and 1 and introduce some additional restrictions on the growth rates of input-output connections relative to the number of firms. In particular, I assume that the number of shared customers and suppliers between two firms cannot grow at a rate faster than the total number of firms in the economy n, and that the maximum number of firm suppliers or customers must grow slower than the total number of possible edges. First consider the following definitions.

Definition 1.4.1 (Paths)**.** A k-path between nodes i and j in graph $\mathcal{G}$ is a length k-sequence $\{a_\ell\}_{\ell=1}^{k}$ where $a_1 = i$, $a_k = j$, and $w_{a_\ell a_{\ell+1}} > 0$ for all $\ell = 1, ..., k-1$. Denote by $A_{ij}(\mathcal{G})$ the set of paths between nodes i and j and by $A_i := \{k : A_{ki}(\mathcal{G}) \neq \emptyset\}$ the set of nodes for which a path to i exists.

Definition 1.4.2 (Maximal Dependency)**.** The maximal dependency of an n-vertex graph $\mathcal{G}_n$ is given by:

$$\bar{M}_n(\mathcal{G}_n) := \sup_{i,j} \left[\text{card}\big(A_i(\mathcal{G}_n) \cap A_j(\mathcal{G}_n)\big) \right] \tag{1.11}$$

Definition 1.4.3 (Maximal Degree)**.** The maximal (unweighted) degree in an n-vertex graph $\mathcal{G}_n$ is given by:

$$\bar{D}_n(\mathcal{G}_n) = \sup_{i} \left[\sum_{j=1}^{n} \text{card}\big(A_{ji}(\mathcal{G}_n)\big) \right] \tag{1.12}$$

If only direct connections exist, then $\bar{M}_n(\mathcal{G}_n) = \sup_{i,j} \sum_{k=1}^{n} \mathbb{1}_{\{w_{ki}>0\}} \mathbb{1}_{\{w_{kj}>0\}}$ and $\bar{D}_n(\mathcal{G}_n) = \sup_i \sum_{j=1}^{n} \mathbb{1}_{\{w_{ji}>0\}}$. Given these definitions, the following assumptions formally restricts the growth rate of input-output connections as the number of firms n grows. These assumptions are fairly general and relevant for deriving tractable theoretical properties of the model.

Assumption 1.4.4 (Bounded Growth Rate of Maximal Degree Sequence)**.** *For all q and t, the maximal degree sequence grows at a rate strictly less than n^2:*

$$\bar{D}_{nq} = o(n^2)$$

These assumptions are intuitive and weaker than the restriction that no firms can serve as a customer or supplier to all other firms. In this case, both the maximal dependency and the maximal degrees must grow at a rate slower than n.

Assumption 1.4.5 (Bounded Growth Rate of Maximal Dependency)**.** *For all q and t, the maximal dependency sequence grows at a rate strictly less than n:*

$$\bar{M}_{nq} = o(n)$$

1.4.4 Shock Propagation Mechanism

For tractability, productivity and demand shocks propagate in a single direction within period t and die out in the following period. Network connections induce correlation across firm-level shocks. At the start of period t, shocks are drawn from distributions $\varepsilon_{idt} \sim \text{Bernoulli}(p_{idt})$ and $\varepsilon_{iut} \sim \text{Bernoulli}(p_{iut})$ where p_{idt} and p_{iut} represent time-varying propensities for firms to experience downstream (demand-side) or upstream (supply-side) shocks, respectively. Propensities are a function of the network structure and firm substitutability, both of which are fixed exogenously at the start of each period.

Intuitively, firms with more substitutability across customers (suppliers) should have a lower average propensity p_{idt} (p_{iut}) to experience shocks. Mathematically, I assume propensities follow a logistic (sigmoid) curve:

$$p_{iqt} = g(s_{iqt}; k_{iq}, x_{iq}) = \frac{1}{1 + \exp\left\{k_{iqt} \cdot (s_{iqt} - x_{iqt})\right\}}, \quad q \in \{u, d\}, \tag{1.13}$$

where s_{iut} (s_{idt}) is the supply-side (demand-side) substitutability of firm i, k_{iqt} is the sensitivity (steepness) of firm propagation to substitutability, and x_{iqt} is a scalar midpoint. The cross-sectional normalization ensures that firms with substitutability x_{iqt} have a 50% chance of being shocked. Substitutability captures network-weighted dispersion in i's supplier-technology (customer-product) space, while x_{iqt} and k_{iqt} jointly characterize the firm-specific risk of firm i. Inverting terms in the "between" concentration measure from (1.5), I assume substitutability can be written:

$$s_{iqt} = \log \sum_{j \neq k} w^q_{ijt} \cdot w^q_{ikt} \cdot \delta^q_{jkt}, \tag{1.14}$$

where w^q_{ijt} represents the importance of trade between j and i in the q-stream direction and δ^q_{jkt} is normalized distance between industries j and k in latent product ($q = d$) or technological ($q = u$) space. See Appendix 1.8 for details. Shared customer and supplier connections induce correlation in substitutability s_{iqt} across firms. This also implies the shock transmission propensities p_{iqt} are also correlated. Time-variation in firm product differentiation generates correlated changes in substitutability across firms who share customers and suppliers. When there are no network connections, firms are hit by shocks with probability $p_{iqt} = 1/(1+\exp(-k_{iqt}x_{iqt}))$. For remaining sections, I assume that $k_{iqt} = k_{iq}$ and $x_{iqt} = x_{iq}$ are time invariant.

1.4.5 Consumption Growth and the Stochastic Discount Factor

I assume that representative households in this economy own shares in each firm and have the following preferences:

$$u(c_{1t}, \ldots, c_{nt}, \ell_t) = \frac{1}{1-\gamma} \cdot \left(\prod_{i=1}^{n} c_{it}^{\beta_i} \right)^{1-\gamma} \cdot \varphi(\ell_t), \tag{1.15}$$

where c_{it} is the consumption of good i with preference weights β_i such that $\sum_i \beta_i = 1$, γ is risk aversion, and $\varphi(.)$ is a decreasing and differentiable function capturing disutility of labor ℓ_t. In Appendix 1.7, I show that equilibrium consumption growth and output growth are equal such that $\Delta\tilde{c}_{i,t+1} := \log(c_{i,t+1}/c_{it}) = \Delta\tilde{y}_{i,t+1}$ for all i and t. I also derive an appropriate price normalization such that equilibrium consumption expenditure is given by $C_t = \prod_i c_{it}^{\beta_i} = \sum_i p_{it} c_{it}$ for a given set of positive prices p_{it}. Finally, the following proposition derives the expression for growth in aggregate consumption expenditure under the same assumptions.

Proposition 1.4.6 (Aggregate Consumption and Output Growth)**.** *Assuming* $\beta_i = 1/n$ *for all i and under the price normalization in Appendix 1.7, aggregate consumption growth can be written:*

$$\Delta\tilde{c}_{t+1} = \gamma_u \cdot a_{t+1} + \gamma_d \cdot g_{t+1} - \beta_u \cdot W_{u,t+1} - \beta_d \cdot W_{d,t+1}, \tag{1.16}$$

where $W_{ut} = \frac{1}{n}\sum_{i=1}^{n} \varepsilon_{iut}$ *and* $W_{dt} = \frac{1}{n}\sum_{i=1}^{n} \varepsilon_{idt}$ *and* $\gamma_u, \gamma_d, \beta_u, \beta_d$ *are positive scalars.*

Proof. See Appendix 1.9.2. ■

This proposition decomposes aggregate consumption growth into four components. The first two components capture innovations to aggregate productivity and demand growth (a_t and g_t, respectively), both of which are positively related to output and consumption growth. On the other hand, the next two components are negatively related to output and consumption growth and capture the average impact of bad shocks to productivity originating upstream (W_{ut}), and the average impact of bad shocks to demand originating downstream (W_{dt}). Combining this result with (1.35), the log stochastic discount factor (SDF) can be written:

$$m_{t+1} = \log\beta - \gamma\big(\gamma_u \cdot a_{t+1} + \gamma_d \cdot g_{t+1} - \beta_u \cdot W_{u,t+1} - \beta_d \cdot W_{d,t+1}\big), \tag{1.17}$$

where β is the intertemporal discount factor and γ is risk aversion. This implies that aggregate productivity and demand growth have a positive price of risk while average upstream and downstream propagation have a negative price of risk.

1.4.6 Additional Theoretical Results

This section summarizes some additional relevant theoretical results from the model. The following proposition states that the conditional distribution of consumption growth in this model is asymptotically normal as the number of firms grows.

Proposition 1.4.7 (Distribution of Consumption Growth)**.** *Under Assumption 1.4.4, the sequence of consumption growth is asymptotically normal as $n \to \infty$, conditional on time t for all t:*

$$\Delta \tilde{c}_{t+1} \xrightarrow{d} \mathcal{N}(\mu_{c,t+1}, \sigma^2_{c,t+1}), \tag{1.18}$$

where:

$$\mu_{ct} := \mathbb{E}_t[\Delta \tilde{c}_t] = \frac{1}{n}\sum_{i=1}^{n}(p_{iut} + p_{idt}),$$
$$\sigma^2_{ct} := var_t[\Delta \tilde{c}_t] = \sigma^2_g + \sigma^2_a + var_t(W_{ut} + W_{dt}).$$

Proof. See Appendix 1.9.4. ∎

Although the conditional mean of consumption growth is known in this model, there is no closed form expression for the conditional variance term. This follows from the fact that shock transmission propensities follow a logistic normal distribution (see Appendix 1.8). After deriving the asymptotic distribution of consumption growth, the next corollary characterizes the probability the W_{nqt} deviates from its cross-sectional mean when propensities are known.

Corollary 1.4.8 (Concentration of Network Factors)**.** *Under Assumption 1.4.5 and if $\bar{M}_{nqt} > 1$, the propagation factor W_{qt} can be written:*

$$W_{nqt} = \mu_{nqt|t} + \varepsilon_{nqt|t} \tag{1.19}$$

where $\mu_{nqt|t} = \frac{1}{n}\sum_i p_{iqt|t}$ *and* $\varepsilon_{nqt|t} \sim \mathcal{N}(0, \sigma^2_{nqt|t})$ *where:*

$$\sigma^2_{nqt|t} \leq \frac{\bar{M}_{nqt}}{n} = o(1) \tag{1.20}$$

Moreover, for any $k > 0$*, the magnitude of* $\varepsilon_{nqt|t}$ *can be upper bounded as follows:*

$$Pr\big(|\varepsilon_{nqt|t}| \geq 2k(\bar{M}_{qt}/n)\big) \leq \frac{1}{k^2}$$

Proof. See Appendix 1.9.3. ■

1.5 Testable Implications

In this section, I verify the main quantitative predictions of the model using financial and macroeconomic data. According to equation (1.17), innovations in average supply and demand shock propagation have a negative price of risk. In addition, level changes in these components should be negatively correlated with aggregate consumption growth.

1.5.1 Data and Calibration

I construct a panel of firms between 1997-2019 whose North American Industry Classification System (NAICS) are in the set of industries for which BEA Input-Output accounts are available. I obtain annual and quarterly firm variables from Compustat and stock return data from CRSP for share codes 10, 11, and 12.[27] I obtain aggregate time series of Total Factor Productivity growth from Fernald (2012a), government demand growth from the procurement proxy in Briganti and Sellemi (2022), and annual market and risk-free returns from Kenneth French's Website.

I begin by computing input-output propagation factors, denoted by $\hat{W}_{ut}$ and $\hat{W}_{dt}$. In Section 1.3, I introduce a latent distance approach to compute the panel of industry concentration and substitutability between customers and suppliers from equation (1.5). Assuming

[27] Firm and return variables are winsorized at the 1% level unless otherwise specified.

that substitutability is the same for firms in a given industry, I can then directly compute $\hat{s}_{iqt}$ for any firm with industry data available. The expression for $\hat{p}_{iqt} = g(\hat{s}_{iqt}; k_{iq}, x_{iq})$ follows directly from equation (1.13) conditional on scalar parameters k_{iq} and x_{iq}. To calibrate these parameters, I first estimate the following panel regression:

$$\Delta \tilde{y}_{i,t+1} = \gamma_u a_{t+1} + \gamma_d g_{t+1} + \text{controls} + \varepsilon_{i,t+1}, \tag{1.21}$$

where $\Delta \tilde{y}_{i,t+1}$ is year-on-year sales growth, a_{t+1} is TFP growth, and g_{t+1} is growth in the federal defense. Controls include year and industry fixed effects, lagged firm size, age, and return on assets to ensure that changes in $\varepsilon_{i,t+1}$ is unrelated to aggregate economy-wide or industry-level forces or trends in large, young, or profitable firms.[28] Then let $\hat{\varepsilon}_{i,t+1}$ denote residual sales growth, and let ω_{iu} (ω_{id}) denote the average cost share (sales share) of intermediate inputs in i's industry, and choose values of $k_{iq} \geq 0$ and $s_{iq} \in \mathbb{R}$ such that:

$$\widehat{\text{var}}(\hat{\varepsilon}_{i,t+1}) = \frac{\exp(k_{iq}(\bar{s}_{iq} - x_{iq}))}{\left(1 + \exp(k_{iq}(\bar{s}_{iq} - x_{iq}))\right)^2}, \quad \text{and} \quad \omega_{iq} = \frac{\left(1 + \exp(k_{iq}(\bar{s}_{iq} - x_{iq}))\right)^2}{\left(1 + \exp(-k_{iq} x_{iq})\right)^2},$$

for $q \in \{u, d\}$ where $\bar{s}_{iq} = \frac{1}{T} \sum_t s_{iqt}$ is firm i's average substitutability over time. The first restriction is based on equation (1.8) and ensures that the variance of a typical Bernoulli(p_{iqt}) shock is equal to residual sales growth variance, while the second restriction requires ω_{iq} proportion of this variance to be attributed to network propagation. Together, the system of equations uniquely identify k_{iq} and x_{iq}. Table **??** summarizes the calibrated parameter values. I then approximate each realized network propagation factor with its cross-sectional empirical mean as follows:

$$\hat{W}_{qt} \approx \frac{1}{n} \sum_{i=1}^{n} \hat{p}_{iqt}, \quad q \in \{u, d\}, \tag{1.22}$$

where $\hat{p}_{iqt}$ is the empirical propensity. I use the cross-sectional mean since realized shocks cannot be identified even when firm propensities p_{iqt} are known. In practice, this is not

[28] Industry fixed effects are at the two-digit NAICS granularity.

a large concern, as Proposition 1.4.8 shows that the measurement error can be bounded arbitrarily by increasing the sample size.[29] I plot the estimated series in Figure 1.1 and report summary statistics in Table 1.3. See Appendix 1.14 for more details.

Table 1.3: Descriptive Statistics Network Propagation Factors

	W_{ut}	S_{ut}	W_{dt}	S_{dt}	g_t	a_t	σ_t^{civ}	σ_t^{mkt}	AC(1)
W_{ut}	1	-0.81	0.13	0.17	0.25	-0.20	-0.38	-0.31	0.26
S_{ut}		1	-0.36	-0.03	-0.21	-0.16	0.62	0.52	0.20
W_{dt}			1	-0.44	-0.05	-0.13	-0.22	-0.06	0.01
S_{dt}				1	-0.15	-0.17	0.13	-0.14	0.04
g_t					1	0.05	0.00	0.26	0.62
a_t						1	-0.19	-0.43	0.30
σ_t^{civ}							1	0.53	-0.30
σ_t^{mkt}								1	-0.21

Notes: This table reports the time-series correlation and first-order autocorrelation of network propagation factors W_{ut} and W_{dt}, average industry substitutability S_{ut} and S_{dt}, procurement demand growth g_t, productivity growth a_t, innovations to common idiosyncratic volatility σ_t^{civ}, and innovations to market volatility σ_t^{mkt}. I calculate g_t as the first log difference in the federal procurement proxy from Briganti and Sellemi (2022), a_t as the first difference in the TFP series from Fernald (2012a), innovations in common idiosyncratic volatility as the first log difference in the first principal component of firm volatility following Herskovic et al. (2016), and innovations in market volatility as the first log difference in market return volatility.

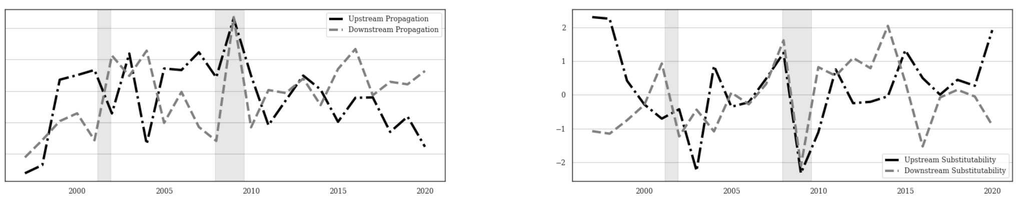

(a) Panel A: Propagation Factors (b) Panel B: Innovations in Avg Substitutability

Figure 1.1: Network Propagation Risk Factors

Notes: This figure plots the time series of network propagation risk factors (Panel A), the cross-sectional average industry substitutability (Panel B). Shaded regions indicate NBER-dated recession periods.

[29] As a heuristic evaluation of this bound, suppose I restrict our sample to only firms that show up in the Customer Segments database ($\bar{M}_{nqt} = 149$ and $n = 12489$), then the probability that the measurement error more than 10% is less than 1%.

1.5.2 Asset Pricing Tests

To verify the prices of risk predicted in (1.17), I sort stocks based on their exposure to factors and form quintile-sorted portfolios. In particular, for every stock i I regress annual excess returns $r_{it} - r_{ft}$ on a constant, aggregate demand and productivity growth, and additional controls.[30] The main regression is given by:

$$r_{it} - r_{ft} = \alpha_i + \beta_{ia} a_t + \beta_{ig} g_t + \beta_{iu} W_{ut} + \beta_{id} W_{dt} + \varepsilon_{it}, \tag{1.23}$$

where equation (1.17) implies that stocks with high β_{ia} and β_{ig} should have higher expected excess returns and stocks with high β_{iu} and β_{id} should have lower expected excess returns. For each year t, I compute stock exposure to factors on a 15-year rolling window from $t-14$ to t using (1.23), and then sort stocks into five portfolios on each beta both separately (one-way sort) and pairwise (two-way sort). Then I construct value and equal-weighted portfolios over the subsequent year $t+1$ and compute average out-of-sample excess returns for each portfolio.

Table 1.4 provides evidence of a significant return spread in one-way beta sorted portfolios. In particular, the highest quintile upstream propagation beta portfolio earns -11.42% lower annual returns than the lowest quintile portfolio, while the highest quintile downstream propagation beta portfolio earns -4.18% lower annual returns than the lowest quintile portfolio. Both return spreads are statistically significant, although more pronounced for upstream propagation beta sorted portfolios.[31] This is consistent with Herskovic et al. (2020b), who argue that upstream propagation is the more important channel.

I also observe a return spread in post-sample alphas from the CAPM and Fama and French (FF3) three factor models, which implies that network propagation risk is not

[30] I test several specifications including controlling for lag factor levels. Results are robust to several specifications on the set of controls, including the baseline without controls.

[31] I also test for monotonicity of returns in upstream and downstream propagation betas, following Patton and Timmermann (2010). I reject this null hypothesis at the 10% level for upstream beta sorted portfolios, but fail to reject for downstream beta sorted portfolios.

captured by market returns or FF3 factors. In light of the variance results of Section 1.3, I also verify that return spreads are not explained by market volatility or idiosyncratic volatility factors in Table 1.16.[32] Additionally, return spreads cannot be explained by differences in return volatility, average size, or average book-to-market ratios. Finally, the average correlation between upstream and downstream propagation betas is 8.6%, suggesting that the two network factors are distinct sources of risk.

Return spreads are robust to the choice of trailing window length, equal or value weighting in portfolios, control variables, and show up in double-sorted portfolios as well. See Appendix 1.14 for more details.

[32] I measure market volatility as the annual volatility of market returns and idiosyncratic volatility following Herskovic et al. (2016).

Table 1.4: One-Way Sorted Portfolios on Network Propagation Factors

Panel A: One-way sorts on upstream propagation beta (controlling for a_t and g_t)

	1 (Low)	2	3	4	5 (High)	H-L	t(H-L)	MR p-val
$\mathbb{E}[r]-r_f$	18.10	12.81	10.59	9.42	6.69	-11.42	-13.22	0.07
α_{capm}	0.29	-0.1	-0.23	-0.31	-1.02	-1.32	-15.71	0.05
α_{ff3}	0.08	-0.09	-0.22	-0.29	-0.54	-0.63	-8.61	0.09
Volatility (%)	15.54	13.89	13.59	13.03	19.66	-	-	-
Book-to-market	0.52	0.56	0.53	0.55	0.50	-	-	-
Market value ($bn)	6.46	16.99	10.62	15.15	9.11	-	-	-

Panel B: One-way sorts on downstream propagation beta (controlling for a_t and g_t)

	1 (Low)	2	3	4	5 (High)	H-L	t(H-L)	MR p-val
$\mathbb{E}[r]-r_f$	13.54	13.23	11.02	9.77	9.36	-4.18	-7.56	0.25
α_{capm}	-0.04	-0.18	-0.28	-0.38	-0.60	-0.56	-4.78	0.00
α_{ff3}	-0.11	-0.14	-0.23	-0.28	-0.36	-0.25	-3.62	0.03
Volatility (%)	15.44	13.95	18.58	12.99	13.88	-	-	-
Book-to-market	0.52	0.56	0.55	0.53	0.51	-	-	-
Market value ($bn)	15.84	7.45	4.54	17.6	12.72	-	-	-

Panel C: One-way sorts on upstream propagation beta (no controls)

	1 (Low)	2	3	4	5 (High)	H-L	t(H-L)	MR p-val
$\mathbb{E}[r]-r_f$	15.15	12.61	11.46	9.44	7.23	-7.91	-11.57	0.07
α_{capm}	0.09	-0.15	-0.17	-0.28	-1.09	-1.18	-17.53	0.26
α_{ff3}	-0.12	-0.14	-0.18	-0.22	-0.58	-0.46	-9.96	0.31
Volatility (%)	15.26	14.23	13.57	12.61	20.97	-	-	-
Book-to-market	0.54	0.58	0.52	0.52	0.50	-	-	-
Market value ($bn)	6.87	17.48	10.92	16.38	6.56	-	-	-

Panel D: One-way sorts on downstream propagation beta (no controls)

	1 (Low)	2	3	4	5 (High)	H-L	t(H-L)	MR p-val
$\mathbb{E}[r]-r_f$	12.66	11.94	11.8	8.34	5.13	-7.53	-8.65	0.42
α_{capm}	-0.15	-0.18	-0.19	-0.4	-0.64	-0.49	-11.37	0.41
α_{ff3}	-0.10	-0.21	-0.21	-0.29	-0.32	-0.22	-4.76	0.44
Volatility (%)	14.09	13.9	14.44	13.22	31.95	-	-	-
Book-to-market	0.54	0.49	0.58	0.51	0.54	-	-	-
Market value ($bn)	15.97	12.79	6.33	16.88	6.34	-	-	-

Notes: This table reports average excess returns and post-sample alphas in annual percentages for value-weighted portfolios sorted into quintiles on annual upstream and downstream propagation factors. Sample is between 1997-2021 for more than 10,000 stocks belonging to the BEA 66 non-government industry classifications. Panels A and B control for productivity growth and federal procurement demand growth, while Panels C and D have no controls. I also report average return volatility, book-to-market ratio and market value for each portfolio. To test for significant return spreads, I report t-statistics for the null hypothesis $H_0 : xr_5 = xr_1$, where xr_q is the average return of the q^{th} quintile single sorted portfolio. Moreover, I report p-values for the test $H_0 : xr_{q+1} < xr_q \forall q \leq 4$, calculated via bootstrap following Patton and Timmermann (2010).

1.5.3 Verifying Macroeconomic Predictions

Equation (1.16) predicts that upstream and downstream propagation factors should be negatively correlated with consumption, output growth, and aggregate dividend growth. To test this, I construct aggregate series between 1997-2021 for consumption and output growth from the National Income and Product Accounts (NIPA) and corporate dividend growth from BEA data. Then I regress each outcome on network propagation factors, controlling for aggregate productivity and federal procurement demand growth. I standardize each variable to have zero mean and unit standard deviation. Consistent with the predictions of the model, Table 1.5 reports negative and statistically significant coefficients on both upstream and downstream propagation risk factors. The factors explain a large portion of time variation in consumption, output, and dividend growth with R^2 values of 56%, 68%, and 26%, respectively.

The coefficients on downstream propagation are -0.17 ($t = -1.89$), -0.60 ($t = -2.57$), and -0.01 ($t = -0.52$) for aggregate consumption, output, and dividend growth regressions, respectively. On the other hand, the coefficients on upstream propagation are not significant for aggregate consumption and output growth regressions, although the coefficient in the dividend growth regression is -0.03 ($t = -1.80$). Additionally, the coefficients on upstream propagation are significant when the dependent variable is limited to only durable consumption or output growth, -0.111 ($t = -2.39$) and -1.38 ($t = -2.39$), respectively. This suggests that durable consumption is more sensitive to upstream (supply-side) risk.

Table 1.5: Network Propagation and Macroeconomic Factors

Variable	Δc_t	Δc_t^{dur}	Δc_t^{nondur}	Δy_t	Δy_t^{dur}	Δy_t^{nondur}	ΔD_t
W_{ut}	0.003	-0.111**	-0.158	-0.112	-1.383**	-0.416	-0.033*
	(0.103)	(0.048)	(0.087)	(0.275)	(0.607)	(0.229)	(0.021)
W_{dt}	-0.174*	-0.031	-0.022	-0.598**	-0.387	-0.058	-0.010
	(0.074)	(0.054)	(0.101)	(0.222)	(0.684)	(0.267)	(0.021)
a_t	0.321**	0.128**	0.197**	1.018**	1.607**	0.520**	0.018
	(0.114)	(0.055)	(0.108)	(0.271)	(0.689)	(0.285)	(0.010)
g_t	0.271	-0.056	0.604	-0.011	-0.704	1.592	0.029
	(0.332)	(0.276)	(0.638)	(0.956)	(3.450)	(1.682)	(0.109)
Intercept	-0.485	-0.111	-0.205	1.218	4.09	2.023	0.042
Obs	24	24	24	24	24	24	24
R^2	0.56	0.537	0.376	0.679	0.537	0.376	0.257

Notes: This table reports results of OLS regressions of aggregate consumption growth, output growth, and dividend growth on input-output network propagation risk factors, controlling for productivity and federal procurement demand growth. The columns represent different dependent variables corresponding to aggregate PCE growth, durable consumption growth, non-durable consumption growth, output growth, durable output growth, non-durable output growth, and dividend growth. All series are standardized to have zero mean and unit variance. Sample is at an annual frequency between 1997-2021.

1.6 Conclusion

In this work, I prove that the idiosyncratic shock assumption in network models with static shock propagation generates empirically implausible restrictions on the network structure. Empirical evidence confirms the importance of correlation in shocks for explaining variance dynamics. As an alternative to the standard model, I propose a production-based asset pricing model with input-output networks, in which correlation in firm level supply and demand shocks is driven by technological and product proximity. This model implies upstream and downstream propagation systematic risk risk factors which I find are negatively priced in the cross-section of returns, and are associated with lower aggregate consumption, output, and dividend growth.

Future work might investigate in more depth the sources of correlation in supply and demand shocks and generate more granular estimates of the agents' ability to substitute away from them. Moreover, an open question is whether there exist firm-level clustering in which the idiosyncratic shock assumption might be consistent with empirical data.

Appendices

1.7 General Equilibrium Model of Input-Output Linkages

In this section, I show that (1.1) can be cast as an outcome of a production-based asset pricing model. This model provides a structural foundation for the theoretical contributions of this work and is closely related to Acemoglu et al. (2012a), Acemoglu et al. (2016a), Ramìrez (2017), and Herskovic (2018b). Consider a competitive economy with n production units (firms or industries) with Cobb-Douglas technology, representative households with constant relative risk aversion (CRRA) preferences over a basket of goods and who work and own shares in all firms and live off wages and dividends, and a government that finances purchases with a lump-sum tax. In this economy, Hicks-neutral productivity shocks propagate downstream from suppliers to customers, while government demand shocks propagate upstream from customers to suppliers.

1.7.1 Production

Production unit i's output is a constant returns to scale function of labor and intermediate inputs:

$$y_{it} = \exp(z_{it})\ell_{it}^{\alpha_{i\ell}} \prod_{j=1}^{n} x_{ijt}^{w_{ijt}}, \tag{1.24}$$

where x_{ijt} is the amount of product j used as input by industry i at time t, ℓ_{it} is labor input, and z_{it} is a Hicks-neural productivity shock, respectively. I assume that for all i and t, the labor share of production is positive (i.e., $\alpha_{i\ell} > 0$) and intermediate input shares are non-negative ($w_{ijt} \geq 0$) and sum to the capital share of production (i.e., ($\sum_{j=1}^{n} w_{ijt} = 1 - \alpha_{i\ell}$).

Taking both spot market prices and input shares as given, production units optimize dividends (denoted D_{it}) as a function of input and labor purchases:

$$D_{it} = \max_{\{x_{ijt}\}_{j=1}^{n}, \ell_{it}} p_{it} y_{it} - \sum_{j=1}^{n} p_{jt} x_{ijt} - p_{wt} \ell_{it} \tag{1.25}$$

subject to (1.24) and $\ell_{it} \in (0,1)$. Suppose further that M_{t+1} is the stochastic discount factor (SDF) that prices all assets in the economy. Then the cum dividend value of firm i (denoted V_{it}) satisfies the following Bellman equation:

$$V_{it} = D_{it} + \mathbb{E}_t[M_{t+1} V_{i,t+1}]. \tag{1.26}$$

1.7.2 Government

The government purchases goods G_{it} from each unit i at time t and finances them via a lump-sum tax T_t. Taking prices as given, the government's financing constraint implies that $T_t = \sum_{i=1}^{n} p_{it} G_{it}$.

1.7.3 Households

Assume that the representative household owns shares in each unit and has the following preferences:

$$u(c_{1t}, ..., c_{nt}, \ell) = \frac{1}{1-\gamma} \cdot \left(\prod_{i=1}^{n} c_{it}^{\beta_i} \right)^{1-\gamma} \cdot g(\ell_t), \tag{1.27}$$

where c_{it} is the consumption of good i with preference weights β_i such that $\sum_i \beta_i = 1$ and $g(.) = (1 - \ell_t)^{\nu}$ is a decreasing and differentiable function capturing disutility of labor.

Households also have a time-discount factor of β and cannot store goods from one period to another. In equilibrium, households hold a zero net position in a risk-free asset and choose to own ϑ_{it} in each unit according to the following budget constraint:

$$T_t + p_{wt}\ell_t + \sum_{i=1}^{n} p_{it}c_{it} + \sum_{i=1}^{n} \vartheta_{i,t+1}(V_{it} - D_{it}) = \sum_{i=1}^{n} \vartheta_{i,t}V_{it}, \tag{1.28}$$

where the right hand side is total value of investments and the left hand side is the sum of taxes paid, wages earned, cost of consumption, and unrealized capital gains, respectively. The household's optimization problem satisfies the Bellman equation:

$$U_t = \max_{\{c_{it},\vartheta_{i,t+1},\ell\}_{i=1}^{n}} u(.) + \beta\mathbb{E}_t[U_{t+1}], \tag{1.29}$$

subject to (1.28).

1.7.4 Equilibrium

The competitive equilibrium of the economy consists of spot market prices $\{p_{it}\}_{i=1}^{n}$, consumption bundles $\{c_{it}\}_{i=1}^{n}$, share holdings $\{\vartheta_{it}\}_{i=1}^{n}$, labor supply ℓ_t, wages p_{wt}, and input bundles $\{x_{ijt}\}_{i,j=1}^{n}$ such that both production units and households exhibit optimal behavior and good/asset markets clear.

Market Clearing

In equilibrium, all good markets clear such that:

$$y_{it} = \underbrace{c_{it}}_{\text{final consumption demand}} + \underbrace{\sum_{j=1}^{n} x_{jit}}_{\text{intermediate demand}} + \underbrace{G_{it}}_{\text{government consumption}},$$

and all asset markets clear $\vartheta_{it} = 1$ for all i and t.

Producer Optimality

Taking prices as given, unit i's first order dividend maximizing conditions satisfy

$$w_{ijt} = \frac{p_{jt}x_{ijt}}{p_{it}y_{it}} \equiv \frac{sales_{j\to i}}{sales_i} \tag{1.30}$$

and

$$\alpha_{i\ell} = \frac{p_{wt}\ell_{it}}{p_{it}y_{it}} \tag{1.31}$$

Consumer Optimality

Given the Cobb-Douglas aggregation in preferences over goods (i.e., $C_t := \prod_i c_{it}^{\beta_i}$), utility maximizing households consume β_i of income on good i and hold shares fixed at $\vartheta_{it} = 1$. More specifically, letting λ_t be the Lagrange multiplier for the period t household budget constraint, the first-order condition for consumption is written:

$$\lambda_t = \frac{C_t^{-\gamma}}{p_{it}} \cdot \frac{\partial C_t}{\partial c_{it}}. \tag{1.32}$$

This implies that equilibrium consumption satisfies:

$$p_{it}c_{it} = \beta_i \left(p_{wt}\ell_t^* + \sum_{j=1}^{n} D_{jt} - T_t \right). \tag{1.33}$$

where ℓ_t^* solves:

$$\frac{p_{wt}\ell_t^*}{p_{wt}\ell_t^* + \sum_{j=1}^{n} D_{jt} - T_t} = -\frac{\ell_t^* g'(\ell_t^*)}{g(\ell_t^*)} \tag{1.34}$$

Asset Prices

From (1.32), the stochastic discount factor can be written:

$$M_{t+1} = \beta \frac{\lambda_{t+1}}{\lambda_t} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \frac{p_{it} \cdot \partial C_{t+1} / \partial c_{i,t+1}}{p_{i,t+1} \cdot \partial C_t / \partial c_{it}}. \tag{1.35}$$

Following Herskovic (2018a), I assume that prices are normalized such that $p_{it} = \partial C_t / \partial c_{it}$, or equivalently that $\prod_j p_{jt}^{\beta_j} = \prod_j \beta_j^{\beta_j}$ for all i and t. This implies the the utility aggregator is equal to the household's consumption expenditure $C_t = \sum_{i=1}^n p_{it} c_{it}$.[33]. Then (1.35) simplifies to:

$$M_{t+1} = \beta \left(\frac{\sum_{i=1}^n p_{i,t+1} c_{i,t+1}}{\sum_{i=1}^n p_{it} c_{it}} \right)^{-\gamma}. \tag{1.36}$$

Shock Propagation

I now derive closed form expressions for the effects of productivity and government demand shocks on output growth in this model. The main takeaway is that output growth is captured by the following reduced form expression:

$$\mathbf{d} \log \mathbf{y}_t = \mathbf{H}_{down,t} \mathbf{dz}_t + \mathbf{H}_{up,t} \mathbf{dG}_t,$$

where $\mathbf{H}_{\cdot,t}$ are $n \times n$ Leontief inverse propagation matrices. I provide a derivation for each component separately.

Productivity Shocks Totally differentiate the expression in (1.24) to obtain:

$$d \log y_{it} = dz_{it} + \alpha_{i\ell} d \log \ell_{it} + \sum_{j=1}^n w_{ijt} d \log x_{ijt} \tag{1.37}$$

[33] I further assume that $C_t = p_{wt} \ell_t$.

Totally differentiating (1.30), (1.31), and (1.33) and plugging in to this expression yields:

$$\begin{aligned} d\log y_{it} &= dz_{it} + \alpha_{i\ell} d\log \ell_{it} + \sum_{j=1}^{n} w_{ijt}(d\log y_{it} + d\log p_{it} - d\log p_{jt}) \\ &= dz_{it} + \alpha_{i\ell}(d\log y_{it} - d\log c_{it}) + \sum_{j=1}^{n} w_{ijt}(d\log y_{it} - d\log c_{it} + d\log c_{jt}). \end{aligned}$$

Given constant returns to scale ($\alpha_{i\ell} + \sum_j w_{ij} = 1$), this expression can be further simplified as follows:

$$d\log c_i = dz_{it} + \sum_{j=1}^{n} w_{ijt} d\log c_j$$

or in vector notation:

$$\mathbf{d}\log \mathbf{c}_t = \mathbf{dz}_t + \mathbf{W}_t \mathbf{d}\log \mathbf{c}_t,$$

where $\mathbf{W}_t$ has entries w_{ijt}. Note that market clearing and profit maximization conditions together imply that:

$$\frac{y_{jt}}{c_{jt}} = 1 + \sum_{i=1}^{n} w_{ijt} \frac{\beta_i y_{it}}{\beta_j c_{it}}$$

which implies that equilibrium consumption growth is equal to equilibrium output growth:

$$\mathbf{d}\log \mathbf{c}_t = \mathbf{d}\log \mathbf{y}_t \tag{1.38}$$

and thus that:

$$\mathbf{d}\log \mathbf{y}_t = (\mathbf{I} - \mathbf{W}_t)^{-1}\mathbf{dz}_t, \tag{1.39}$$

where $\mathbf{H}_{down,t} := (\mathbf{I} - \mathbf{W}_t)^{-1}$ is the Leontief inverse of $\mathbf{W}_t$. Here, the $\mathbf{W}_t$ matrix determines the strength of downstream propagation of productivity shocks.

Demand Shocks To study the effects of government spending shocks in the model, normalize $\mathbf{z}_t = 0$ and consider the unit cost function for i:

$$C_{it}(\mathbf{p}_t, p_{wt}) = A_{it} p_{wt}^{\alpha_{i\ell}} \prod_{j=1}^{n} p_{jt}^{w_{ijt}}, \quad \text{where } A_{it} = \alpha_{i\ell}^{-\alpha_{i\ell}} \prod_{j=1}^{n} w_{ijt}^{-w_{ijt}}.$$

The zero productivity normalization implies zero dividends for production units, and combined with the price normalization for wages, this implies that:

$$\log p_{it} = \log A_{it} + \sum_{j=1}^{n} w_{ijt} \log p_{jt}$$

Conditional on productivity vector $\mathbf{z}_t$ and defining the vector $\mathbf{a}_t$ with entries $\log A_{it}$, prices are a function of the network and cost but not government purchases:

$$\log \mathbf{p}_t = (\mathbf{I} - \mathbf{W}_t)^{-1} \mathbf{a}_t.$$

Setting $\nu = 1$ Then (1.34) and the fact that $T_t = \sum_i p_{it} G_{it}$ implies that:

$$\ell_t = \frac{1}{2} + \frac{1}{2} \sum_{i=1}^{n} p_{it} G_{it} \tag{1.40}$$

and thus:

$$p_{it} c_{it} = \beta_i [p_{wt} \ell_t - T_t] = \frac{\beta_i}{2} \left(1 - \sum_{j=1}^{n} p_{jt} G_{jt} \right).$$

Differentiating and combining with the resource constraint and profit maximization conditions yields:

$$\frac{d(p_{it} y_{it})}{p_{it} y_{it}} = \sum_{j=1}^{n} w_{jit} \frac{d(p_{jt} y_{jt})}{p_{it} y_{it}} + \frac{dG_{it}}{y_{it}} - \frac{\beta_i}{2} \sum_{j=1}^{n} \frac{d(p_{jt} G_{jt})}{p_{it} y_{it}}.$$

Since prices are constant (i.e., $d(p_{it}y_{it})/p_{it}y_{it} = d\log y_{it}$), I can write in vector notation:

$$\mathbf{d}\log\mathbf{y}_t = \mathbf{H}_{up,t}\mathbf{dG_t}, \tag{1.41}$$

where $\mathbf{H}_{up,t} = (\mathbf{I} - \mathbf{W}_t^\top)^{-1}\boldsymbol{\Lambda}_t$ is the upstream propagation Leontief inverse and $\boldsymbol{\Lambda}_t$ is a scaling matrix with diagonal entries $(1-\beta_i/2)/p_{it}y_{it}$ and off-diagonal entries $-(\beta_i/2)/p_{it}y_{it}$ for row indices i.

1.8 Product Varieties in Latent Space

Suppose that each industry (or firm) is associated with a random position $z_{it} := (\cos\theta_{it}, \sin\theta_{it})$ on a circular surface on the 3-dimensional hypersphere $\mathcal{S}^{p+1}$. Suppose that the surface represents the space of varieties in production technology space. I also assume each unit corresponds to a position in latent product variety space, and that product positions are independent of technological positions. The stochastic process for positions is the same in both spaces and depends on changes in the angle θ_{it} as follows:

$$\theta_{it} = \rho \cdot \theta_{i,t-1} + \varepsilon_{it}, \qquad \varepsilon_{it} \sim_{iid} \mathcal{N}(0, \sigma_\theta^2), \tag{1.42}$$

where θ_{it} is measured in radians. The distance between two points i and j can then be written:

$$\delta_{ijt} = \frac{1}{2\pi}|\theta_{it} - \theta_{jt}|. \tag{1.43}$$

This implies a correlation structure in the distances between units as follows:

$$\operatorname{corr}(\delta_{ijt}, \delta_{kmt}) = \begin{cases} 0 & \text{if } i \notin \{k,m\} \cap j \notin \{k,m\}, \\ 1 & \text{if } i \in \{k,m\} \cup j \in \{k,m\}, \end{cases}, \tag{1.44}$$

or equivalently that:

$$\operatorname{cov}(\delta_{ijt}, \delta_{kmt}) = \begin{cases} 0 & \text{if } i \notin \{k,m\} \cap j \notin \{k,m\}, \\ \sqrt{\operatorname{var}(\delta_{ijt})\operatorname{var}(\delta_{kmt})} & \text{if } i \in \{k,m\} \cup j \in \{k,m\}, \end{cases} \tag{1.45}$$

where:

$$\sigma_d^2 := \operatorname{var}(\delta_{ijt}) = \operatorname{var}(\theta_{it}) + \operatorname{var}(\theta_{jt}) - \mathbb{E}[|\theta_{it} - \theta_{jt}|]^2 = \frac{2\sigma_\theta^2}{1-\varphi^2} - \frac{(4/\pi)\sigma_\theta^2}{1-\varphi^2},$$

Define the set of i's q-stream located trade partners by $A_{iqt} := \{k : A_{ki}(\mathcal{G}_{qt}) \neq \emptyset\}$. Note that in technology space, q refers to upstream propagation. Combining with equation (1.14) yields:

$$\begin{aligned} \operatorname{cov}(s_{iqt}, s_{jqt}) &= \sum_{k \neq \ell} \sum_{m \neq p} w_{ikt} w_{i\ell t} w_{jmt} w_{jpt} \cdot \operatorname{cov}(\delta_{k\ell t}, \delta_{mpt}) \\ &= \sum_{k,\ell \in A_{iqt}; m,p \in A_{jqt}} w_{ikt} w_{i\ell t} w_{jmt} w_{jpt} \cdot \sigma_d^2 \cdot \mathbb{1}\{(m \in A_{iqt}) \cup (p \in A_{iqt})\} \end{aligned}$$

Additionally, s_{iqt} are jointly distributed as a folded truncated normal with variance-covariance matrix $\mathbf{\Sigma}_{qt} = [\operatorname{cov}(s_{iqt}, s_{jqt})]_{ij}$ and mean vector $\boldsymbol{\mu}_{qt} = [\mu_{iqt}]_i$ with entries:

$$\mu_{iqt} := \mathbb{E}[s_{iqt}] = -\sigma_d \sqrt{8/\pi} \cdot \sum_{j<k; j,k \in A_{iqt}} w_{ijt} w_{ikt}.$$

When $i = j$, I can further simplify as follows:

$$\operatorname{var}(s_{iqt}) = 4\sigma_d^2 \cdot \sum_{j<k; j,k \in A_{iqt}} (w_{ijt} w_{ikt})^2.$$

Next, using equation (1.13), I can write the log-odds function as follows:

$$l_{iqt} := \log\left(\frac{p_{iqt}}{1-p_{iqt}}\right) = k\left(s_{iqt} - \frac{1}{n}\sum_{i=1}^{n} s_{iqt}\right) = k\left(\hat{e}_{in}^\top - \frac{1}{n}\iota_n^\top\right)\mathbf{s}_{qt},$$

where $\hat{e}_{in}$ is the ith column of an $n \times n$ identity matrix and ι_n is an $n \times 1$ vector of ones. Equivalently, the vector $\mathbf{l}_{qt} := (l_{1qt}, ..., l_{nqt})^\top$ can be written:

$$\mathbf{l}_{qt} = k\left(\mathbf{I}_n - \frac{\iota_n \iota_n^\top}{n}\right)\mathbf{s}_{qt} \sim \mathcal{N}\big(\mathbf{B}_{k,n}\boldsymbol{\mu}_{qt}, \mathbf{B}_{k,n}^\top \boldsymbol{\Sigma}_{qt} \mathbf{B}_{k,n}\big),$$

where $\mathbf{B}_{k,n} := k(\mathbf{I}_n - \iota_n \iota_n^\top / n)$. Notice that $p_{iqt} = F(l_{iqt})$ where $F(x) = (1 + \exp(-x))^{-1}$ and thus the vector $\mathbf{p}_{qt}$ has a logistic normal distribution and thus no closed form representation for its mean vector and variance-covariance matrix. However, I can write $\text{median}(\mathbf{p}_{qt}) = F(\mathbf{B}_{k,n}\boldsymbol{\mu}_{qt})$.

1.9 Proofs

1.9.1 Proof of Proposition 1.2.3

I begin by characterizing the family of matrices $\mathbf{W} \in M_n$ that are consistent with Assumptions 1.2.1 and 1.2.2 with a sequence of if and only if relationships. First define $\mathbf{\Sigma}_y := \mathbb{V}[\mathbf{y}]$ to be the variance-covariance matrix of sectoral production, and notice that Assumption 1.2.2 implies that $\mathbf{\Sigma}_y = (\mathbf{I}-\mathbf{W})^{-1}\mathbf{D}(\mathbf{I}-\mathbf{W}^\top)^{-1}$ is non-negative. This is because $\mathbf{\Sigma}_y$ is the product of non-negative matrices by definition since $(\mathbf{I}-\mathbf{W})^{-1} = \mathbf{I}+\mathbf{W}+\mathbf{W}^2+\mathbf{W}^3+... \geq 0$ and $\mathbf{D} \geq 0$. The former decomposition is only possible when $\rho(\mathbf{W}) \leq 1$. Rearranging terms, I obtain:

$$\mathbf{D} = \mathbb{V}[(\mathbf{I}-\mathbf{W})\mathbf{y}] = (\mathbf{I}-\mathbf{W})\mathbf{\Sigma}_y(\mathbf{I}-\mathbf{W}^\top).$$

For ease of notation, let $h_{ij} := \mathbb{1}_{\{i=j\}} - w_{ij}$ and $\rho_{ij}\sigma_i\sigma_j$ to be the (i,j) entries of $\mathbf{H} := \mathbf{I}-\mathbf{W}$ and $\mathbf{\Sigma}_y$, respectively. Then the matrix $\mathbf{Q}_n := \mathbf{H}\mathbf{\Sigma}_y\mathbf{H}^\top = [q_{ij}]$ is diagonal with entries:

$$q_{ij} = \sum_{k=1}^{n}\sum_{m=1}^{n} h_{ik}h_{jm}\rho_{km}\sigma_k\sigma_m$$

Notice that $\mathbf{Q}_n$ is symmetric and since it is also diagonal, then the following must hold:

$$\det(\mathbf{Q}_n) = \prod_{i=1}^{n}\left(\sum_{k=1}^{n}\sum_{m=1}^{n} h_{ik}h_{im}\rho_{km}\sigma_k\sigma_m\right).$$

Additionally, $\mathbf{P}\mathbf{Q}_n\mathbf{P}^\top$ is also diagonal for any permutation matrix $\mathbf{P} \in M_n$ so I can exchange the ordering of units without loss of generality. Consider the Laplace expansion of $\mathbf{Q}_n$ by minors along row i. Set $i = 1$ arbitrarily and notice that:

$$\det(\mathbf{Q}_n) = q_{11}\cdot\det(\mathbf{Q}_n[1,1]) + \sum_{k=2}^{n}(-1)^{1+k}q_{1k}\cdot\det(\mathbf{Q}_n[1,k]),$$

where $\mathbf{Q}_n[i,j]$ is a size $n-1$ sub-matrix of $\mathbf{Q}_n$ obtained by deleting row i and column j. Since the principal sub-matrix $\mathbf{Q}_n[1,1]$ is also diagonal, its determinant is $\prod_{i=2}^{n} q_{ii}$ and thus the second term in the above expression must be zero. This implies that the determinant for any size 2 principal sub-matrix of $\mathbf{Q}_n$ with size $n-2$ index set α and entries a_{ij} must satisfy:

$$\det\big(\mathbf{Q}_n[\alpha]\big) = a_{11}a_{22} - a_{12}a_{21} = a_{11}a_{22}$$

Note that the indices (1,2) refer without loss of generality to any arbitrary pair of sectors in the set $\{1,...,n\}$, and all of the following results must hold for all nC2 pairwise combinations of units. Substitute terms and notice that all principal sub-matrices of a symmetric matrix must also be symmetric (i.e., $a_{12} = a_{21}$) to obtain the equivalent restriction:

$$0 = a_{12} = w_{12}\sigma_2^2 - \rho_{12}\sigma_1\sigma_2(1 + w_{21}w_{12}) + w_{21}\sigma_1^2$$

Or identically:

$$w_{21}\mathbb{V}[y_1] + w_{12}\mathbb{V}[y_2] - \text{cov}(y_1, y_2)(1 + w_{12}w_{21}) = 0, \tag{1.46}$$

where $\mathbf{V}[y_i]$ and $\text{cov}(y_i, y_j)$ are the (i,i)th and (i,j)th entries of $\mathbf{\Sigma}_y$, respectively. I consider two different cases.

Case 1. At least one of w_{21} or w_{12} is zero. If both are zero, then (1.46) holds trivially. However, if $w_{ij} = 0$ for all $i \neq j$, then $\mathbf{W}$ no longer has full rank. If only one is zero (e.g., $w_{12} = 0$), this implies that $\text{cov}(y_1, y_2) = w_{21}\mathbb{V}[y_1]$ which is consistent with (1.1).

Case 2. Both w_{21} and w_{12} are positive. Redefine $\tilde{y}_1 := y_1/\sqrt{w_{21}}$ and $\tilde{y}_2 := y_2/\sqrt{w_{12}}$ and apply the Cauchy-Schwarz Inequality:

$$\frac{(\mathbb{V}[\tilde{y}_1] + \mathbb{V}[\tilde{y}_2])^2}{(w_{21}w_{12}) \cdot (1 + w_{21}w_{12})^2} = \text{cov}(\tilde{y}_1, \tilde{y}_2)^2 \leq \mathbb{V}(\tilde{y}_1)\mathbb{V}(\tilde{y}_2) \tag{1.47}$$

For now write $c := w_{21}w_{12}(1+w_{21}w_{12})^2$:

$$\begin{aligned}\mathbb{V}[\tilde{y}_1]^2+\mathbb{V}[\tilde{y}_2]^2+2\mathbb{V}[\tilde{y}_1]\mathbb{V}[\tilde{y}_2] &\leq c\mathbb{V}(\tilde{y}_1)\mathbb{V}(\tilde{y}_2)\\ \mathbb{V}[\tilde{y}_1]^2+\mathbb{V}[\tilde{y}_2]^2+(2-c)\mathbb{V}[\tilde{y}_1]\mathbb{V}[\tilde{y}_2] &\leq 0\end{aligned}$$

The left hand side is a quadratic equation in $\mathbf{R}^2_+$ of the variables $\mathbf{V}[\tilde{y}_1]$ and $\mathbf{V}[\tilde{y}_2]$. Suppose without loss of generality that $\mathbb{V}[\tilde{y}_1] = \mathbb{V}[\tilde{y}_2] > 0$, then c must satisfy:

$$(4-c)\mathbb{V}[\tilde{y}]^2 \leq 0,$$

which implies that $c \geq 4$, or equivalently:

$$w_{21}w_{12}(1+w_{12}w_{21})^2 \geq 4$$

Recall that since $\mathbf{W}$ is a non-negative matrix, any pair of weights which satisfy this inequality must have $w_{12}w_{21} \geq 1$. Note that due to (1.47), this restriction is necessary but not sufficient. Moreover, since the spectral radius $\rho(\mathbf{W}) \leq 1$, there must be at least one row or column bounded above by one (see e.g., Theorem 8.1.22 in Horn and Johnson (2013)).

1.9.2 Proof of Proposition 1.4.6

Assuming $\beta_i = 1/n$ for all i, equilibrium consumption expenditure is given by $C_t = \sum_{i=1}^n p_{it} c_{it} = \prod_{i=1}^n c_{it}^{1/n}$. Given the price normalization in Section A.4.4. from Appendix 1.7, I can write:

$$\begin{aligned}
\Delta c_{t+1} := \log(C_{t+1}/C_t) &= \log\left(\prod_{i=1}^n \left(\frac{c_{i,t+1}}{c_{it}}\right)^{1/n}\right) \\
&= \frac{1}{n}\sum_{i=1}^n \log(c_{i,t+1}/c_{it}) \\
&= \frac{1}{n}\sum_{i=1}^n \log(y_{i,t+1}/y_{it})
\end{aligned}$$

Plugging in the reduced form from (1.8) yields:

$$\begin{aligned}
\Delta c_{t+1} &= \frac{1}{n}\sum_{i=1}^n \left(\Delta z_{i,t+1} + \Delta g_{i,t+1}\right) \\
&= \gamma_u \cdot \frac{1}{n}\sum_{i=1}^n a_{t+1} + \gamma_d \cdot \frac{1}{n}\sum_{i=1}^n g_{t+1} - \beta_u \cdot \frac{1}{n}\sum_{i=1}^n \varepsilon_{iu,t+1} - \beta_d \cdot \frac{1}{n}\sum_{i=1}^n \varepsilon_{id,t+1}
\end{aligned}$$

1.9.3 Proof of Proposition 1.4.8

The proof of Proposition 1.4.7 established the asymptotic normality of W_{nqt} for all q and t. I omit subscripts and write:

$$\begin{aligned}
\sigma_n^2 = \operatorname{var}(W_n) &= \frac{1}{n^2}\sum_{i=1}^n \operatorname{var}[\varepsilon_i] + \frac{1}{n^2}\sum_{i\neq j} \operatorname{cov}(\varepsilon_i, \varepsilon_j) \\
&= \frac{1}{n^2}\sum_{i=1}^n p_i(1-p_i) + \frac{1}{n^2}\sum_{i\neq j} \operatorname{cov}(\varepsilon_i, \varepsilon_j) \cdot \mathbb{1}\{\mathbb{1}\{i, j \in A_i(\mathcal{G}_{nq}) \cap A_j(\mathcal{G}_{nq})\}\}
\end{aligned}$$

When $n > 1$. This quantity can be bounded above by:

$$\sigma_n^2 \leq \frac{1}{n} + \frac{\bar{M}_n}{n} \leq 2\frac{\bar{M}_n}{n} = o(1),$$

where the final equality follows from Assumption 1.4.5. Moreover, for $k > 0$ Chebyshev's Inequality yields:

$$\Pr\big(|W_n - \mathbb{E}[W_n]| \geq 2k\bar{M}_n/n\big) \leq \Pr\big(|W_n - \mathbb{E}[W_n]| \geq k\sigma_n\big) \leq \frac{1}{k^2}$$

1.9.4 Proof of Proposition 1.4.7

This proof is an application of Theorem 2 from Janson (1988). Fix t and q and omit the time-subscript without loss of generality. Since $\varepsilon_{iq} \sim \text{Bernoulli}(p_{iq})$, the mean and variance of firm cash flow shocks conditional on p_{iq} can be written:

$$\mathbb{E}[\varepsilon_{iq}] = p_{iq}, \quad \text{var}(\varepsilon_{iq}) = p_{iq}(1 - p_{iq}),$$

and the covariance between firm shocks can be written:

$$\text{cov}(\varepsilon_{iq}, \varepsilon_{jq}) = \mathbb{E}[\varepsilon_{iq}\varepsilon_{jq}] - \mathbb{E}[\varepsilon_{iq}]\mathbb{E}[\varepsilon_{jq}] = \underbrace{\Pr(\varepsilon_{iq} = 1, \varepsilon_{jq} = 1)}_{=:p_{ijq}} - p_{iq}p_{jq} > 0.$$

Therefore, the mean and variance of $W_{nq} = \frac{1}{n}\sum_{i=1}^{n} \varepsilon_{iq}$ can be written:

$$\begin{aligned}
\mu_{nq} := \mathbb{E}[W_{nq}] &= \frac{1}{n}\sum_{i=1}^{n} \mathbb{E}[\varepsilon_{iq}] = \frac{1}{n}\sum_{i=1}^{n} p_{iq} \\
\sigma_{nq}^2 := \text{var}\big[W_{nq}\big] &= \frac{1}{n^2}\sum_{i=1}^{n} \text{var}[\varepsilon_{iq}] + \frac{1}{n^2}\sum_{i \neq j} \text{cov}(\varepsilon_{iq}, \varepsilon_{jq}) \\
&\geq \frac{1}{n^2}\sum_{i=1}^{n} \text{var}[\varepsilon_{iq}], \\
&= \frac{1}{n^2}\sum_{i=1}^{n} p_{iq}(1 - p_{iq}) \\
&\geq \frac{1}{n} \cdot \min_i p_{iq}
\end{aligned}$$

Let $\bar{D}_{nq}$ denote the maximal number of edges incident to a single vertex in graph $\mathcal{G}_{nq}$. Theorem 2 from Janson (1988) requires that $\mathcal{X}_{n,m} = o(1)$ for some integer m, where

$\mathbb{X}_{n,m}$ is given by:

$$\mathcal{X}_{n,m} := \left(\frac{n}{\bar{D}_{nq}}\right)^{1/m} \cdot \frac{\bar{D}_{nq}}{n\sigma_{nq}} = \frac{\bar{D}_{nq}^{1-1/m}}{n} \cdot \frac{1}{n\sigma_{nq}} \leq \frac{\bar{D}_{nq}^{1-1/m}}{n}.$$

Assumption 1.4.4 implies:

$$\mathcal{X}_{n,2} = o(1),$$

and therefore:

$$(W_{nq} - \mu_{nq})/\sigma_{nq} \xrightarrow{d} \mathcal{N}(0,1)$$

Combine this result with the assumptions that $a_t \sim \mathcal{N}(0, \sigma_a^2)$ and $g_t \sim \mathcal{N}(0, \sigma_g^2)$ to get the final distribution of consumption growth.

1.10 Upstream and Downstream Propagation Matrices

1.10.1 Construction from BEA Data

In this section, I discuss a simple method for constructing from the data the upstream and downstream propagation matrices developed in the previous section.[34] These matrices capture the strength of a connection between an industry and its customer or supplier industries. Shocks transmitted from customer i to supplier j (supplier i to customer j) should depend on the strength of the connection in the upstream (downstream) direction. I construct these matrices directly from the BEA make and use tables described in Horowitz, Planting, et al. (2006). Consider again an economy with n industries.

The Make Table I extract from the BEA make table an $n \times n$ industry-by-commodity matrix with entries:

$$(MAKE)_{ij} = OUT_{i \to j} \equiv \text{dollar value of commodity } j \text{ produced by industry } i$$

Note that the BEA makes a slight distinction between commodities and industries, since in principle an industry might produce another industry's commodity as a by-product of its own output. Next, I denote the total production of commodity j by $OUT_j := \sum_{i=1}^{n} OUT_{i \to j}$. Using the notation in Horowitz, Planting, et al. (2006), I define the *market share matrix* with the following entries:

$$(MKTSHARE)_{ij} = \frac{(MAKE)_{ij}}{OUT_j} = \frac{OUT_{i \to j}}{OUT_j}$$

Here, the (i, j) entry describes the share of industry i in the total production of commodity j. Equivalently, I can write $MKTSHARE = MAKE \odot (\boldsymbol{\iota}_n \cdot \mathbf{S})$ where $\odot$ denotes the Hadamard (elementwise) product and $S = (OUT_1^{-1}, \ldots, OUT_n^{-1})$ is the $1 \times n$ scaling vector.

[34] Similar procedures are discussed in Acemoglu et al. (2016a), Ozdagli and Weber (2017), and Gofman et al. (2020).

The Use Table Similarly, I extract from the BEA use table an $n \times n$ commodity-by-industry matrix with entries:

$$(USE)_{ij} = IN_{i \to j} \equiv \text{dollar value of commodity } i \text{ used as input by industry } j$$

Define the total output of industry i by y_i.[35] Then I construct the *input requirement matrix* by rescaling the value of an industry's inputs by the industry's total value as measured by output. The entries of this matrix are given by:

$$(INPUTREQ)_{ij} = \frac{(USE)_{ij}}{y_j} = \frac{IN_{i \to j}}{y_j}. \tag{1.48}$$

The (i, j) entry of the above matrix describes the importance of industry j's inputs from industry i relative to j's total size.

Scrap Adjustment The BEA input-output tables include scrap as a commodity which includes any by-products of production with zero market demand. I therefore redefine the total output of an industry as the non-scrap output. Mathematically, this adjustment is implemented using the non-scrap ratio, calculated as follows:

$$\theta_i = \frac{y_i - scrap_i}{y_i},$$

where $scrap_i$ denotes the total scrap produced by industry i. I then write the entries of the scrap-adjusted market share matrix as follows:

$$(MKT\tilde{S}HARE)_{ij} = \frac{OUT_{i \to j}}{OUT_j} \cdot \frac{1}{\theta_j}$$

[35] I calculate total industry output from the BEA use table as the sum of total intermediates, scrap, and value added. An industry's value added is defined by the BEA as the "market value it adds in production, or the difference between the price at which it sells its products and the cost of the inputs it purchases from other industries".

Direct Requirements From the market share and input requirements matrices, I then construct the industry-by-industry direct requirement table, denoted by $\mathbf{W}$:

$$\underbrace{\mathbf{W}}_{\text{industry}\times\text{industry}} = \underbrace{MKT\tilde{S}HARE}_{\text{industry}\times\text{commodity}} \cdot \underbrace{(INPUTREQ)}_{\text{commodity}\times\text{industry}} .$$

To understand this construction, consider the entries of $\mathbf{W}$:

$$(\mathbf{W})_{ij} = \frac{1}{y_j}\sum_{k=1}^{n} \underbrace{\frac{OUT_{i\to k}}{OUT_k \cdot \theta_k}}_{\text{share of } i \text{ in producing } k} \cdot \underbrace{IN_{k\to j}}_{\text{use of } k \text{ by } j} .$$

Here, the (i,j) entry captures industry j's total dependence on inputs from industry i relative to it's total output y_j. The key assumption here is as follows. If industry j purchases X of commodity K, then the proportion of K coming from industry i is equal to i's adjusted market share of production of K. This is a reasonable assumption on average. Given this assumption, the following identity holds:

$$(\mathbf{W})_{ij} \equiv \frac{SALES_{i\to j}}{SALES_j} \iff (\mathbf{W}^\top)_{ij} = \frac{SALES_{j\to i}}{SALES_i}$$

The downstream weighting matrix is thus defined:

$$\mathbf{A} \equiv \mathbf{W}_{down} := \mathbf{W}^\top,$$

whose (i,j) entry in the above matrix represents the dependence of industry i on input from industry j (i.e., shocks to supplier j propagate downstream to customer i according to the corresponding downstream weight). The sum of row i in this matrix is equal to x_i/y_i where x_i is industry i's total input purchases relative to its size (normalized by i's total output). For the upstream weighting matrix, I require the intermediate rescaling matrix $\mathbf{R}$

with entries $(\mathbf{R})_{ij} = y_j/y_i$. The upstream weighting matrix is thus defined:

$$\hat{\mathbf{A}}^\top \equiv \mathbf{W}_{up} := \mathbf{W} \odot \mathbf{R},$$

where $\odot$ is the Hadamard (elementwise) product. The (i, j) entry in the above matrix represents the dependence of industry i on sales to industry j (i.e., shocks to customer j propagate upstream to supplier i according to the corresponding upstream weight)

$$(\mathbf{W} \odot \mathbf{R})_{ij} = \frac{SALES_{i \to j}}{SALES_i}.$$

1.10.2 Descriptive Statistics

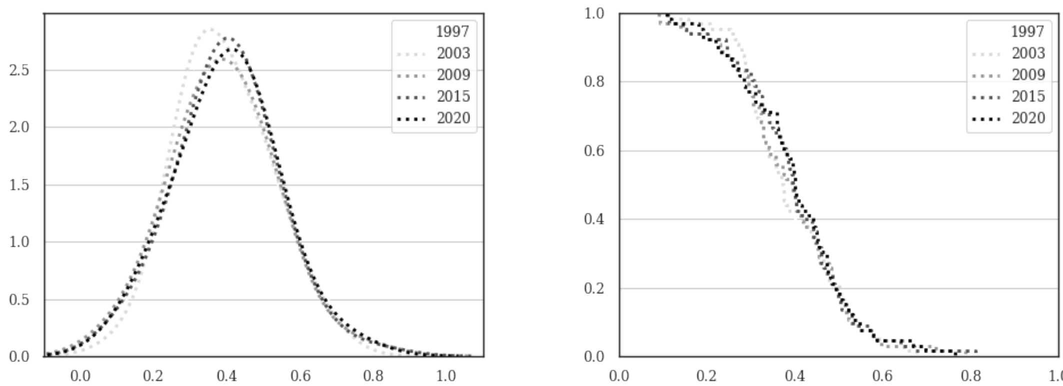

(a) Kernel Density (b) Empirical CCDF

Figure 1.2: Weighted In-Degree Distribution

Notes: This figure visualizes the empirical distribution of weighted in and out-degrees across 66 non-government industries as defined in the Bureau of Economic Analysis (BEA) Make and Use Tables. The left panel shows the Gaussian kernel density estimate of the distribution, and the right panel shows the empirical counter-cumulative distribution function.

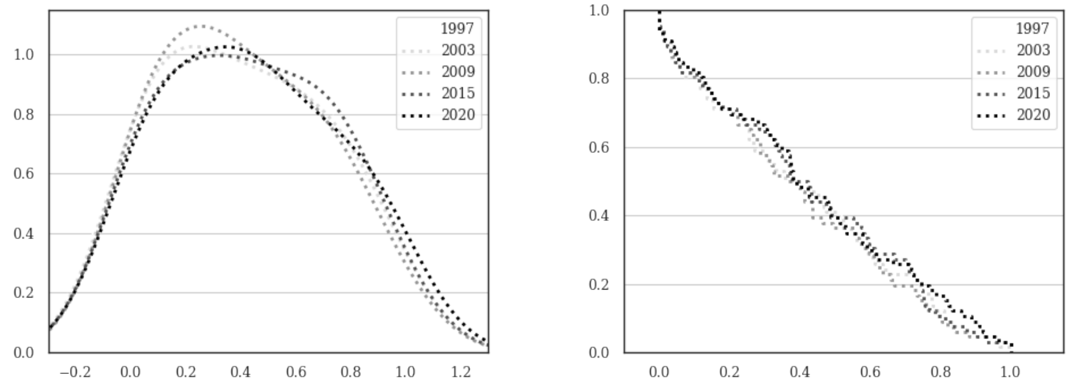

(a) Kernel Density (b) Empirical CCDF

Figure 1.3: Weighted Out-Degree Distribution

Notes: This figure visualizes the empirical distribution of weighted in and out-degrees across 66 non-government industries as defined in the Bureau of Economic Analysis (BEA) Make and Use Tables. The left panel shows the Gaussian kernel density estimate of the distribution, and the right panel shows the empirical counter-cumulative distribution function.

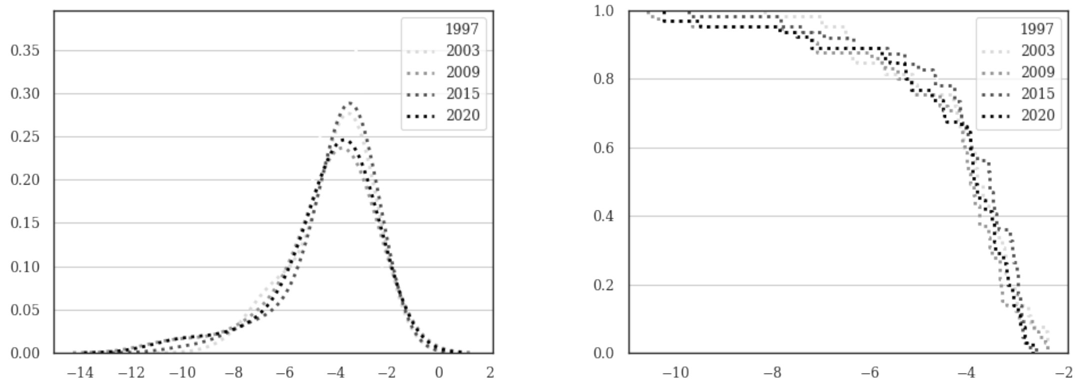

(a) Kernel Density (b) Empirical CCDF

Figure 1.4: Distribution of Downstream Centrality

Notes: This figure visualizes the empirical distribution of weighted log eigenvector centrality of the upstream propagation matrix across 66 non-government industries as defined in the Bureau of Economic Analysis (BEA) Make and Use Tables. The left panel shows the Gaussian kernel density estimate of the distribution, and the right panel shows the empirical counter-cumulative distribution function.

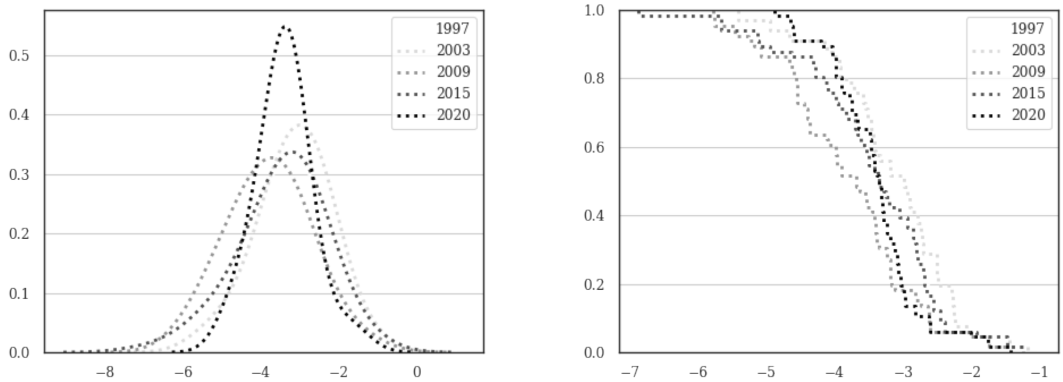

(a) Kernel Density (b) Empirical CCDF

Figure 1.5: Distribution of Upstream Centrality

Notes: This figure visualizes the empirical distribution of weighted log eigenvector centrality of the upstream propagation matrix across 66 non-government industries as defined in the Bureau of Economic Analysis (BEA) Make and Use Tables. The left panel shows the Gaussian kernel density estimate of the distribution, and the right panel shows the empirical counter-cumulative distribution function.

Table 1.6: Correlation of Input-Output Shares

Entry	$w_{d,ij,t-1}$	$w_{u,ij,t-1}$	$h_{d,ij,t-1}$	$h_{u,ij,t-1}$	$w_{d,ij,t}$	$w_{u,ij,t}$	$h_{d,ij,t}$	$h_{u,ij,t}$
$w_{d,ij,t-1}$	1	0.086	0.970	0.096	0.991	0.086	0.961	0.096
$w_{u,ij,t-1}$		1	0.094	0.967	0.084	0.994	0.093	0.961
$h_{d,ij,t-1}$			1	0.110	0.962	0.095	0.990	0.111
$h_{u,ij,t-1}$				1	0.095	0.962	0.109	0.994
$w_{d,ij,t}$					1	0.086	0.970	0.096
$w_{u,ij,t}$						1	0.095	0.967
$h_{d,ij,t}$							1	0.111
$h_{u,ij,t}$								1

Notes: This table report correlations between entries in the upstream and downstream propagation matrices W_{qt} and their Leontief inverses H_{qt}. I construct annual matrices from the BEA Input-Output Accounts for 66 non-government industries between 1997-2020.

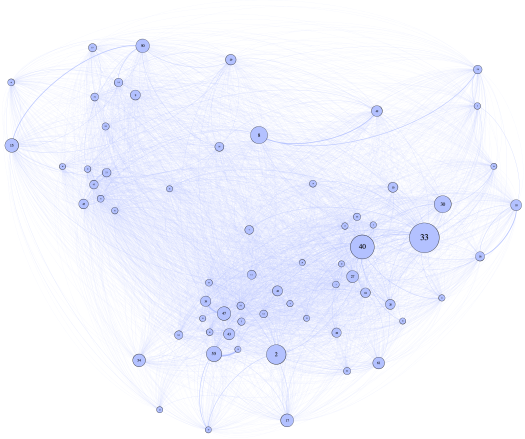

(a) In-Degree (downstream)

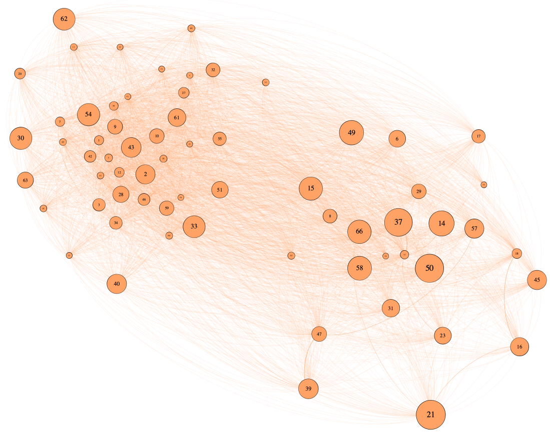

(b) Out-Degree (upstream)

Figure 1.6: Visualizing Weighted Degree by Industry

Notes: The left panel shows the average weighted in-degree by sector as constructed from the downstream propagation network $\mathbf{W}_{down}$. The right panel shows the average weighted in-degree by sector as constructed from the upstream propagation network $\mathbf{W}_{up}$. The sectors are numbered identically in both panels, according to the BEA 66 non-government industry classification.

1.11 Determinants of Network Variance

Table 1.7: Correlation in Predictors of Realized Variance

Predictor	β_{vol}	C_u	C_d	Between (d)	Between (s)	Size	Across (s)	Across (d)	VP	mkt ivol	mkt vol
β_{vol}	1	-0.329	-0.130	-0.308	0.072	0.285	0.061	-0.128	-0.225	0.287	0.406
C_u		1	0.284	0.356	-0.251	0.131	-0.441	0.319	0.423	-0.383	-0.396
C_d			1	0.588	-0.018	-0.128	-0.116	0.288	0.289	-0.003	-0.022
Between (d)				1	-0.205	-0.113	-0.156	0.704	0.708	-0.270	-0.311
Between (s)					1	-0.133	0.056	-0.163	-0.313	0.073	0.054
Size						1	-0.230	0.018	-0.111	-0.386	-0.275
Across (s)							1	-0.088	-0.159	0.205	0.255
Across (d)								1	0.627	-0.224	-0.249
VP									1	-0.279	-0.322
mkt ivol										1	0.965
mkt vol											1

Notes: This table reports the average correlation between predictors in the panel regression from Table **??**. C_u and C_d refer to upstream and downstream centrality, the (d) and (s) labels denote demand and supply-side concentration, size is average output, and VP is vertical position. Idiosyncratic volatility is calculated from FF3 residual returns.

Table 1.8: Network Determinants of Industry Variance (TFP growth shocks)

Panel A: Market Return Variance

	(1)	(2)	(3)	(4)	(5)	(6)
Self-origin (demand)	0.365*** (0.038)	0.246*** (0.039)			0.217*** (0.021)	0.119*** (0.022)
Across (demand)	0.100*** (0.020)	0.084*** (0.018)			0.085*** (0.020)	0.119*** (0.019)
Between (demand)	0.043** (0.018)	0.010 (0.022)			0.091*** (0.019)	-0.076*** (0.020)
Self-origin (supply)			0.039* (0.029)	0.149*** (0.038)	0.217*** (0.021)	0.119*** (0.022)
Across (supply)			0.119*** (0.021)	0.147*** (0.026)	0.270*** (0.027)	0.170*** (0.037)
Between (supply)			0.085*** (0.019)	0.042** (0.019)	0.086*** (0.025)	0.179*** (0.034)
Size		-0.202*** (0.030)		-0.325*** (0.023)		-0.183*** (0.034)
Upstream centrality		-0.145** (0.062)		0.129** (0.051)		-0.238*** (0.066)
Downstream centrality		1.956*** (0.145)		0.296** (0.125)		2.137*** (0.165)
Durability		1.088*** (0.153)		0.389*** (0.108)		1.164*** (0.156)
Vertical position		3.890*** (0.116)		0.243*** (0.083)		4.063*** (0.134)
Constant	-2.123	-0.634	-2.821	-0.096	-1.473	-0.693
Obs	2359	1626	2861	2138	2114	1471
Adj R^2	0.225	0.607	0.083	0.191	0.288	0.633

Panel B: Cash Flow Variance

	(1)	(2)	(3)	(4)	(5)	(6)
Self-origin (demand)	0.363*** (0.072)	0.009 (0.093)			0.190*** (0.037)	0.017 (0.049)
Across (demand)	0.038 (0.039)	-0.050* (0.039)			0.011 (0.040)	0.099** (0.042)
Between (demand)	0.136*** (0.037)	0.128*** (0.042)			0.156*** (0.039)	0.184*** (0.047)
Self-origin (supply)			0.026 (0.065)	0.286*** (0.076)	0.190*** (0.037)	0.017 (0.049)
Across (supply)			0.250*** (0.051)	0.384*** (0.062)	0.309*** (0.061)	0.107* (0.091)
Between (supply)			0.093** (0.044)	0.169*** (0.055)	0.202*** (0.057)	0.162** (0.086)
Size		-0.161** (0.064)		-0.333*** (0.052)		-0.155** (0.075)
Upstream centrality		-0.069 (0.150)		0.562*** (0.132)		-0.100 (0.162)
Downstream centrality		2.193*** (0.343)		-0.070 (0.305)		2.315*** (0.379)
Durability		2.204*** (0.272)		1.172*** (0.182)		2.307*** (0.285)
Vertical position		3.502*** (0.268)		1.142* (0.181)		3.590*** (0.320)
Constant	-4.201	-0.711	-5.246	0.023	-4.123	-0.75
Obs	2359	1626	2861	2138	2114	1471
Adj R^2	0.074	0.248	0.017	0.080	0.087	0.271

Notes: This table reports panel regressions of realized industry variance on a variety of characteristics including the average log variance of supply and demand shocks, log concentration across trade partners, log concentration between trade partners, log total employment (size), log centrality of the upstream and downstream propagation networks, durability of output, vertical position in the supply chain, and industry cluster and year fixed effects. I calculate network components as the average value of 1000 bootstrap samples that randomly drop 10% of estimated pairwise non-zero correlations. In Panel A, the dependent variable is the log variance of annualized monthly returns on an equal-weighted industry portfolio. In Panel B, the dependent variable is the log variance of total quarterly year-on-year industry sales growth. I obtain return data from CRSP and sales data from Compustat. Concentration between and across trade partners are calculated as in 1.4 and 1.5, where I calculate the variance-covariance matrix of supply and demand shocks directly from four-factor TFP growth in the NBER-CES Database (Becker et al., 2016). Following Ahern (2013), I compute industry centrality as the eigenvector centrality of upstream and downstream propagation adjacency matrices. I obtain durability classifications from Gomes et al. (2009) and calculate vertical position of each industry as in Antràs et al. (2012) and Gofman et al. (2020). ***, ** and * indicate significance at the 1%, 5%, and 10% levels, respectively. Standard error are clustered at the BEA 15 major industry group level. Sample is at an annual frequency from 1988 to 2017 for 479 BEA manufacturing industries.

Table 1.9: Network Determinants of Industry Variance (federal procurement shocks)

Panel A: Market Return Variance

	(1)	(2)	(3)	(4)	(5)	(6)
Self-origin (demand)	1.058*** (0.157)	0.903*** (0.200)			0.166* (0.111)	0.558*** (0.121)
Across (demand)	0.088** (0.041)	0.314*** (0.059)			0.062* (0.044)	0.356*** (0.070)
Between (demand)	0.149*** (0.040)	0.153*** (0.039)			0.437*** (0.055)	0.261*** (0.077)
Self-origin (supply)			0.219*** (0.044)	0.363*** (0.056)	0.166* (0.111)	0.558*** (0.121)
Across (supply)			0.142*** (0.036)	0.094* (0.052)	0.072 (0.072)	0.088 (0.079)
Between (supply)			0.150*** (0.034)	0.062 (0.039)	-0.063 (0.075)	-0.13 (0.077)
Size		0.059 (0.053)		-0.242*** (0.047)		-0.033 (0.053)
Upstream centrality		-2.434** (1.053)		-4.885*** (1.116)		-2.140* (1.269)
Downstream centrality		10.76*** (2.686)		13.06*** (2.501)		10.55*** (3.043)
Durability		4.128*** (0.511)		1.991*** (0.407)		3.884*** (0.460)
Vertical position		11.74*** (0.915)		4.334*** (0.595)		10.66*** (0.930)
Constant	-6.278	-3.489	-5.461	-4.234	-6.781	-3.419
Obs	839	666	1003	828	839	666
Adj R2	0.595	0.76	0.071	0.331	0.666	0.762

Panel B: Cash Flow Variance

	(1)	(2)	(3)	(4)	(5)	(6)
Self-origin (demand)	0.425 (0.421)	1.853* (0.951)			0.267 (0.293)	1.025** (0.488)
Across (demand)	0.057 (0.090)	0.790*** (0.215)			0.183* (0.121)	0.768*** (0.207)
Between (demand)	0.309*** (0.091)	0.461*** (0.121)			0.524*** (0.146)	0.471** (0.199)
Self-origin (supply)			0.310*** (0.109)	0.004 (0.185)	0.267 (0.293)	1.025** (0.488)
Across (supply)			-0.184* (0.101)	0.002 (0.139)	0.136 (0.193)	0.327* (0.258)
Between (supply)			0.235** (0.092)	0.028 (0.112)	-0.247* (0.192)	-0.268 (0.244)
Size		0.183 (0.209)		-0.406*** (0.122)		-0.141 (0.214)
Upstream centrality		1.430 (3.309)		3.297* (2.975)		4.077* (3.421)
Downstream centrality		5.582 (8.923)		-5.585 (6.799)		0.810 (8.718)
Durability		5.285*** (1.493)		1.175 (0.847)		5.536*** (1.677)
Vertical position		12.20*** (3.184)		0.948 (1.236)		12.73*** (4.096)
Constant	-1.223	-1.809	-7.991	1.81	-10.389	-0.263
Obs	839	666	1003	828	839	666
Adj R2	0.239	0.211	0.043	0.08	0.248	0.215

Notes: This table reports panel regressions of realized industry variance on a variety of characteristics including the average log variance of supply and demand shocks, log concentration across trade partners, log concentration between trade partners, log total employment (size), log centrality of the upstream and downstream propagation networks, durability of output, vertical position in the supply chain, and industry cluster and year fixed effects. I calculate network components as the average value of 1000 bootstrap samples that randomly drop 10% of estimated pairwise non-zero correlations In Panel A, the dependent variable is the log variance of annualized monthly returns on an equal-weighted industry portfolio. In Panel B, the dependent variable is the log variance of total quarterly year-on-year industry sales growth. I obtain return data from CRSP and sales data from Compustat. Concentration between and across trade partners are calculated as in 1.4 and 1.5, where I calculate the variance-covariance matrix of changes in total obligations from newly awarded federal procurement contracts measured from FPDS. Following Ahern (2013), I compute industry centrality as the eigenvector centrality of upstream and downstream propagation adjacency matrices. I obtain durability classifications from Gomes et al. (2009) and calculate vertical position of each industry as in Antràs et al. (2012) and Gofman et al. (2020). ***, ** and * indicate significance at the 1%, 5%, and 10% levels, respectively. Standard error are clustered at the BEA 15 major industry level. Sample is at an annual frequency from 1991 to 2011 for 479 BEA manufacturing industries.

Table 1.10: Network Determinants of Industry Variance (technological proximity)

Panel A: Market Return Variance

	(1)	(2)	(3)	(4)	(5)	(6)
Self-origin (demand)	0.257*** (0.046)	0.352*** (0.043)			0.286*** (0.050)	0.308*** (0.047)
Across (demand)	0.102*** (0.011)	0.120*** (0.013)			0.124*** (0.011)	0.145*** (0.013)
Between (demand)	0.246*** (0.013)	0.117*** (0.018)			0.227*** (0.013)	0.058*** (0.018)
Self-origin (supply)			0.290*** (0.035)	0.258*** (0.039)	0.229*** (0.040)	0.419*** (0.043)
Across (supply)			0.009 (0.027)	0.032 (0.036)	0.223*** (0.031)	-0.035 (0.042)
Between (supply)			0.190*** (0.031)	0.168*** (0.043)	0.284*** (0.034)	0.143** (0.051)
Size		-0.343*** (0.031)		-0.312*** (0.027)		-0.334*** (0.030)
Upstream centrality		-0.755*** (0.108)		-0.102* (0.087)		-0.543*** (0.104)
Downstream centrality		3.933*** (0.279)		1.052*** (0.190)		3.632*** (0.275)
Durability		-0.195* (0.112)		-0.177* (0.098)		-0.076 (0.108)
Vertical position		1.669*** (0.152)		0.130 (0.086)		2.160*** (0.172)
Constant	-7.926	-1.275	-3.85	-0.341	-8.468	-1.177
Obs	1994	1467	1994	1467	1994	1467
Adj R2	0.593	0.642	0.127	0.279	0.61	0.665

Panel B: Cash Flow Variance

	(1)	(2)	(3)	(4)	(5)	(6)
Self-origin (demand)	0.391*** (0.103)	0.521*** (0.103)			0.473*** (0.102)	0.548*** (0.103)
Across (demand)	0.065** (0.028)	0.061* (0.033)			0.050* (0.029)	0.049* (0.034)
Between (demand)	0.253*** (0.030)	0.139*** (0.045)			0.172** (0.077)	0.217* (0.122)
Self-origin (supply)			-0.019 (0.076)	-0.245*** (0.087)	-0.162** (0.079)	-0.185* (0.097)
Across (supply)			0.175*** (0.067)	0.138* (0.100)	0.072 (0.068)	0.160* (0.102)
Between (supply)			0.307*** (0.076)	0.228** (0.119)	0.271*** (0.031)	0.179*** (0.048)
Size		-0.298*** (0.070)		-0.332*** (0.066)		-0.298*** (0.071)
Upstream centrality		0.113 (0.267)		0.355* (0.276)		-0.068 (0.280)
Downstream centrality		3.217*** (0.665)		0.311 (0.599)		3.555*** (0.680)
Durability		0.574*** (0.210)		0.265* (0.211)		0.481** (0.218)
Vertical position		1.922*** (0.323)		-0.311 (0.208)		1.493*** (0.373)
Constant	-10.991	-1.043	-5.631	-0.101	-11.776	-1.152
Obs	1994	1467	1994	1467	1994	1467
Adj R2	0.21	0.23	0.011	0.051	0.216	0.233

Notes: This table reports panel regressions of realized industry variance on a variety of characteristics including the average log variance of supply and demand shocks, log concentration across trade partners, log concentration between trade partners, log total employment (size), log centrality of the upstream and downstream propagation networks, durability of output, vertical position in the supply chain, and industry cluster and year fixed effects. I calculate network components as the average value of 1000 bootstrap samples that randomly drop 10% of estimated pairwise non-zero correlations. In Panel A, the dependent variable is the log variance of annualized monthly returns on an equal-weighted industry portfolio. In Panel B, the dependent variable is the log variance of total quarterly year-on-year industry sales growth. I obtain return data from CRSP and sales data from Compustat. Concentration between and across trade partners are calculated as in 1.4 and 1.5, where I calculate the variance-covariance matrix of supply and demand shocks using the weighted technological similarity scores following Bloom and Shankerman (2013). Following Ahern (2013), I compute industry centrality as the eigenvector centrality of upstream and downstream propagation adjacency matrices. I obtain durability classifications from Gomes et al. (2009) and calculate vertical position of each industry as in Antràs et al. (2012) and Gofman et al. (2020). ***, ** and * indicate significance at the 1%, 5%, and 10% levels, respectively. Standard errors are clustered at the BEA 15 major industry group level. Sample is at an annual frequency from 1988 to 2018 for 479 BEA manufacturing industries.

Table 1.11: Network Determinants of Industry Variance (product similarity)

Panel A: Market Return Variance

	(1)	(2)	(3)	(4)	(5)	(6)
Self-origin (demand)	0.136** (0.056)	0.354*** (0.040)			0.303*** (0.046)	0.321*** (0.041)
Across (demand)	0.103*** (0.013)	0.071*** (0.011)			-0.065*** (0.013)	0.051*** (0.012)
Between (demand)	0.131*** (0.013)	0.081** (0.012)			0.384*** (0.014)	0.058*** (0.016)
Self-origin (supply)			0.096*** (0.027)	0.056* (0.033)	-0.047 (0.033)	0.207*** (0.034)
Across (supply)			0.147*** (0.015)	0.137*** (0.020)	0.186*** (0.020)	-0.000 (0.021)
Between (supply)			-0.086*** (0.012)	-0.106*** (0.016)	-0.362*** (0.017)	-0.071*** (0.020)
Size		-0.361*** (0.025)		-0.257*** (0.021)		-0.310*** (0.024)
Upstream centrality		-0.149*** (0.053)		0.031 (0.049)		-0.176*** (0.055)
Downstream centrality		2.660*** (0.174)		0.811*** (0.117)		2.891*** (0.180)
Durability		0.216** (0.102)		0.265*** (0.096)		0.286*** (0.098)
Vertical position		3.414*** (0.086)		0.368*** (0.076)		3.336*** (0.106)
Constant	-5.915	-0.862	-3.631	-0.263	-8.380	-0.937
Obs	3800	2430	3819	2430	3800	2430
Adj R^2	0.359	0.625	0.037	0.175	0.488	0.639

Panel B: Cash Flow Variance

	(1)	(2)	(3)	(4)	(5)	(6)
Self-origin (demand)	0.357*** (0.091)	0.549*** (0.090)			0.544*** (0.091)	0.558*** (0.091)
Across (demand)	0.112*** (0.023)	0.055** (0.027)			-0.059** (0.025)	0.011 (0.030)
Between (demand)	0.118*** (0.023)	0.066** (0.030)			0.373*** (0.027)	0.108*** (0.040)
Self-origin (supply)			0.030 (0.060)	-0.026 (0.073)	-0.141* (0.062)	0.084* (0.078)
Across (supply)			0.129*** (0.029)	0.199*** (0.039)	0.170*** (0.031)	0.091** (0.043)
Between (supply)			-0.074*** (0.023)	-0.156*** (0.032)	-0.347*** (0.028)	-0.157*** (0.041)
Size		-0.313*** (0.050)		-0.242*** (0.048)		-0.273*** (0.051)
Upstream centrality		0.185* (0.135)		0.327** (0.132)		0.090 (0.138)
Downstream centrality		3.213*** (0.406)		0.791*** (0.300)		3.587*** (0.413)
Durability		1.097*** (0.171)		1.104*** (0.171)		1.145*** (0.170)
Vertical position		3.495*** (0.193)		0.226 (0.164)		3.102*** (0.251)
Constant	-10.060	-1.041	-5.694	-0.256	-12.580	-1.162
Obs	3800	2430	3819	2430	3800	2430
Adj R^2	0.15	0.235	0.005	0.047	0.205	0.242

Notes: This table reports panel regressions of realized industry variance on a variety of characteristics including the average log variance of supply and demand shocks, log concentration across trade partners, log concentration between trade partners, log total employment (size), log centrality of the upstream and downstream propagation networks, durability of output, vertical position in the supply chain, and industry cluster and year fixed effects. I calculate network components as the average value of 1000 bootstrap samples that randomly drop 10% of estimated pairwise non-zero correlations. In Panel A, the dependent variable is the log variance of annualized monthly returns on an equal-weighted industry portfolio. In Panel B, the dependent variable is the log variance of total quarterly year-on-year industry sales growth. I obtain return data from CRSP and sales data from Compustat. Concentration between and across trade partners are calculated as in 1.4 and 1.5, where I calculate the variance-covariance matrix of supply and demand shocks using the product similarity distances from Hoberg and Phillips (2016). Following Ahern (2013), I compute industry centrality as the eigenvector centrality of upstream and downstream propagation adjacency matrices. I obtain durability classifications from Gomes et al. (2009) and calculate vertical position of each industry as in Antràs et al. (2012) and Gofman et al. (2020). ***, ** and * indicate significance at the 1%, 5%, and 10% levels, respectively. Standard error are clustered at the BEA 15 major industry group level. Sample is at an annual frequency from 1988 to 2018 for 479 BEA manufacturing industries.

1.12 Simulation Evidence

Proposition 1.2.3 claims that there is no matrix $\mathbf{W}$ with entries $w_{ij} \in (0,1)$ such that Assumption 1.2.2 and (1.1) are satisfied. This implies that $\Delta := ||(\mathbf{I}-\mathbf{W})\mathbf{\Sigma}_y(\mathbf{I}-\mathbf{W}) - \mathbf{\Sigma}_u||$ should always be different from zero when $\mathbf{\Sigma}_u$ is diagonal. I test $H_0 : \Delta = 0$ numerically as follows:

1. Fix dimension n, number of iterations S, and consider the following constrained optimization problem:

$$f_{min}(\mathbf{\Sigma}_u) = \min_{w_{ij,i\neq j}\in(0,1)} f(\mathbf{W};\mathbf{\Sigma}_u) = \min_{w_{ij,i\neq j}\in(0,1)} \left|\left|\mathbf{\Sigma}_u - (\mathbf{I}_n - \mathbf{W})^{-1}\mathbf{\Sigma}_u(\mathbf{I}_n - \mathbf{W}^\top)^{-1}\right|\right|_F, \tag{1.49}$$

2. Draw two random samples S_u^d and S_u^* for the variance-covariance matrix of residuals $\mathbf{\Sigma}_u$ such that S_u^d is diagonal and S_u^* is not. I sample the non-diagonal matrix as follows:

$$S_u^*(s) \sim_{iid} \mathcal{W}^{-1}(\mathbf{\Psi}_u, \nu),$$

where $\mathcal{W}^{-1}(\mathbf{\Psi}, \nu)$ denotes an Inverse-Wishart random variable with scale $\mathbf{\Psi}$ and degrees of freedom $\nu > n-1$. On the other hand, the diagonal matrix is given by:

$$S_u^d(s) = \text{diag}(X_1, ..., X_n), \quad X_i \sim_{iid} \Gamma^{-1}(\alpha_i, \beta_i),$$

where $\Gamma^{-1}(\alpha_i, \beta_i)$ denotes an Inverse-Gamma distribution such that $\alpha_i > 1$ and $\mathbb{E}[X_i] = \frac{\beta_i}{\alpha_i - 1} = [\Psi_u]_i i$ for all i. This ensures that the means of diagonal elements are the same across the two samples.

3. Solve equation (1.49) using off-the-shelf constrained quasi-Newton algorithms in both cases and construct S realizations of $f_s^* = f_{min}(S_u^*)$ and $f_s^d = f_{min}(S_u^d)$ for

$s = 1, ..., S$.

4. Construct the numerical p-value using

$$p = \frac{1}{S} \sum_{s=1}^{S} \mathbb{1}\{f_s^* \geq f_s^d\},$$

and repeat Step 3 for different values of n and $\mathbf{\Phi}_u$. Reported results in Table 1.12.

5. For optimal $W_{min}^{spec} = \arg\min f(W, S_u^{spec})$, consider the $n \times n$ elementwise difference $\tilde{\Delta}_s$ for each iteration s:

$$\tilde{\Delta}_s = f(W_{min}^d, S_u^d) - f(W_{min}^*, S_u^*)$$

Then I test the marginal hypotheses $H_0 : [\tilde{\Delta}_s]_{ij} = 0$ for all i, j. This test corresponds to the numerical t-statistic $t_{ij} = m_{ij}/SE_{ij}$, where m_{ij} and SE_{ij} are the mean and standard error of $[\tilde{\Delta}_s]_{ij}$, respectively. Reported results in Table 1.13.

Table 1.12: Simulation p-values

Specification	n=2	n=3	n=4	n=5
(1)	0.000	0.002	0.036	0.004
(2)	0.000	0.000	0.041	0.005
(3)	0.000	0.003	0.046	0.006
(4)	0.000	0.001	0.048	0.009
(5)	0.000	0.000	0.053	0.008
(6)	0.000	0.003	0.060	0.008
(7)	0.000	0.001	0.061	0.006
(8)	0.000	0.001	0.061	0.010
(9)	0.000	0.004	0.068	0.011
(10)	0.000	0.001	0.072	0.011

Notes: This table reports numerical p-values from the procedure described in Appendix 1.12 for different values of n and $\mathbf{\Phi}$. Number of iterations is $S = 1000$.

Table 1.13: Simulation t-statistics

n=2	n=3	n=4	n=5
121	-27	-54	-118
863	48	133	201
861	60	163	152
-504	48	2	-502
	-68	133	-152
	80	-247	201
	60	-26	-583
	80	190	172
	-91	163	205
		-26	61
		-255	152
		204	172
		2	-177
		190	174
		204	138
		-60	-502
			205
			174
			-89
			-334
			-152
			61
			138
			-334
			-75

Notes: This table reports elementwise t-statistics from the procedure described in Appendix 1.12 for different values of n and a fixed value $\mathbf{\Phi}$. Number of iterations is $S = 1000$.

1.13 Comovement in Industry Volatility

Recent research documents significant common variation in both market and fundamental volatility at the granular level.[36] Importantly, the factor structure in volatility is significant even after removing all common variation in returns and cash flow growth, which suggests that this is being driven by underlying sources of systematic risk and rather than an omitted set of returns or sales growth factors.[37] When economic units are connected via input-output networks, a shock to any given unit can generate systematic effects. Moreover, networks mechanically generate volatility comovement regardless of whether shocks to individual units are uncorrelated. When shocks are correlated, this comovement is even more pronounced.

My results so far establish a significant relationship between network concentration across and between customers and suppliers and realized variance. Consequently, comovement in supply chain concentration should also generate comovement in realized variance. Table 1.14 reports average loadings and R^2 values of univariate factor regressions for panel of industry network concentration and realized volatility measures. The main takeaway is that 20-30% of dynamic variation in input-output concentration and return and sales growth volatilities across industries of these variables can be explained by a single factor. Moreover, there is a significant degree of comovement between these common factors. That is, network concentration factors can explain up to 40% of time-series variation in both market and sales growth volatility factors. These results are robust to a variety of specifications.

[36]Herskovic et al. (2016) show that a single common factor explains around 30% of variation in log variance for the panel of CRSP stocks. Other work also documents common variation in option-implied volatilities (Engel and Figlewski (2015)), intra-daily returns (Barigozzi et al. (2014)), and dispersion in firm sales growth (Bloom et al. (2018)).

[37]Herskovic et al. (2016) verify that there is a factor structure even when the pairwise correlation between idiosyncratic return or sales growth residuals is statistically indistinguishable from zero. I also verify that the factor structure holds in residual returns after a non-parametric regression using deep feed-forward neural networks, which have favorable universal approximation properties (see e.g., Hornik et al. (1989)).

Table 1.14: Comovement in Industry Variance

Panel A: Loadings

Factor / Outcome	Across (d)	Between (d)	Across (s)	Between (s)	Var (mkt)	Var (cf)
Across (d)	0.788	1.576	0.364	0.507	0.410	0.574
Between (d)	0.584	0.735	0.394	0.567	0.206	0.383
Across (s)	0.701	1.185	0.700	0.518	0.156	0.497
Between (s)	0.455	0.422	0.401	0.736	0.157	0.310
Var (mkt)	0.091	0.124	0.016	0.035	0.880	0.185
Var (cf)	0.109	0.181	0.156	0.218	0.283	0.721

Panel B: R^2 (avg univariate)

Factor / Outcome	Across (d)	Between (d)	Across (s)	Between (s)	Var (mkt)	Var (cf)
Across (d)	0.201	0.258	0.217	0.318	0.170	0.079
Between (d)	0.189	0.318	0.257	0.400	0.174	0.087
Across (s)	0.164	0.274	0.229	0.344	0.168	0.082
Between (s)	0.189	0.322	0.265	0.347	0.165	0.086
Var (mkt)	0.042	0.046	0.041	0.042	0.343	0.068
Var (cf)	0.063	0.097	0.085	0.115	0.147	0.104

Panel C: R^2 (aggregate)

Factor / Outcome	Across (d)	Between (d)	Across (s)	Between (s)	Var (mkt)	Var (cf)
Across (d)	1	0.620	0.575	0.566	0.373	0.352
Between (d)		1	0.628	0.670	0.255	0.422
Across (s)			1	0.668	0.196	0.469
Between (s)				1	0.229	0.428
Var (mkt)					1	0.165
Var (cf)						1

Notes: This table reports the results of factor regressions for industry variance components. I calculate factors as the first principal component of an industry panel of the variable of interest. Panels A and B report the average loading and R^2 from the regressions $outcome_{it} = \alpha_i + \beta_i \cdot f_t + u_{it}$, respectively. Each column, row pair denotes a different outcome, factor pair. Panel C reports the R^2 of aggregate time-series regressions of factor pairs (i.e., $y_t = \alpha + \beta \cdot x_t + u_t$, where y_t is the column factor and x_t is the row factor). All variables are log-transformed. Sample is at an annual frequency from 1997 to 2019 for 66 non-government BEA industries.

1.14 Additional Firm-Level Results

Table 1.15: Predicting Firm Sales Growth

Variable	(1)	(2)	(3)
a_{t-1}	0.031**	0.029**	0.007**
	(0.002)	(0.001)	(0.000)
g_{t-1}	-0.016**	-0.018**	0.015
	(0.003)	(0.003)	(0.011)
$ROA_{i,t-1}$	0.080**	0.063**	0.064**
	(0.003)	(0.002)	(0.002)
$size_{i,t-1}$	-0.008**	-0.008**	-0.008**
	(0.000)	(0.000)	(0.000)
age_i	0.037**	0.050**	0.046**
	(0.004)	(0.004)	(0.003)
Constant	0.44	0.372	0.39
Obs	259,976	259,976	259,976
Adj R2	0.265	0.259	0.258

Notes: This table reports the regression results based on the model in (1.21). The dependent variable is year-on-year quarterly sales growth and the covariates are aggregate TFP growth from Fernald (2012a), procurement proxy from Briganti and Sellemi (2022), log return on assets, size (log market value), and age as year appears on the database. Each column reports results for different industry aggregations: (1) is 66 BEA non-government industries, (2) is 405 BEA non-government industries, and (3) is 15 major BEA non-government industries. Standard errors are clustered at the same granularity. $**$ and $*$ indicates significance at the 1% and 5% levels, resp. Sample is quarterly between 1997-2019 for 10,700 firms.

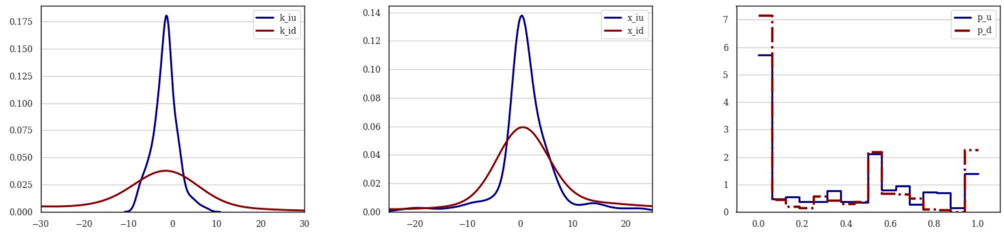

(a) k_i (b) x_i (c) $\bar{p}_i$

Figure 1.7: Distribution of Calibrated Propensities with No Network Connections

Notes: This figure plots the kernel density of calibrated parameters k_{iq} and x_{iq} as described in Section 1.5. Note that $\bar{p}_{iq} = 1/(1+\exp(-k_{iq}x_{iq}))$.

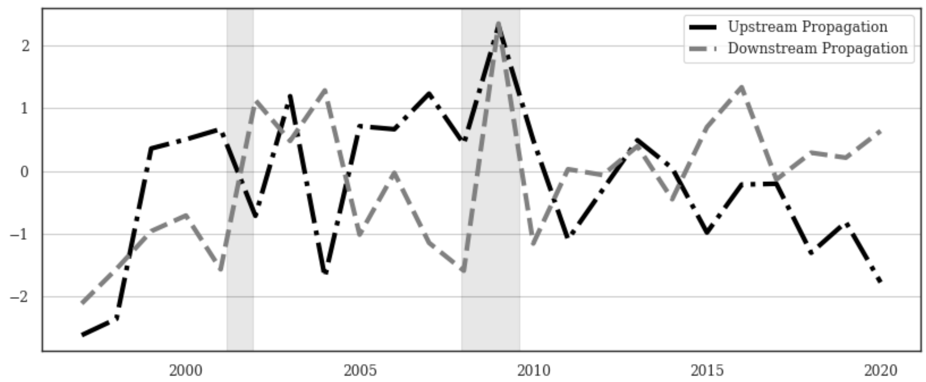

(a) Panel A: Propagation Factors

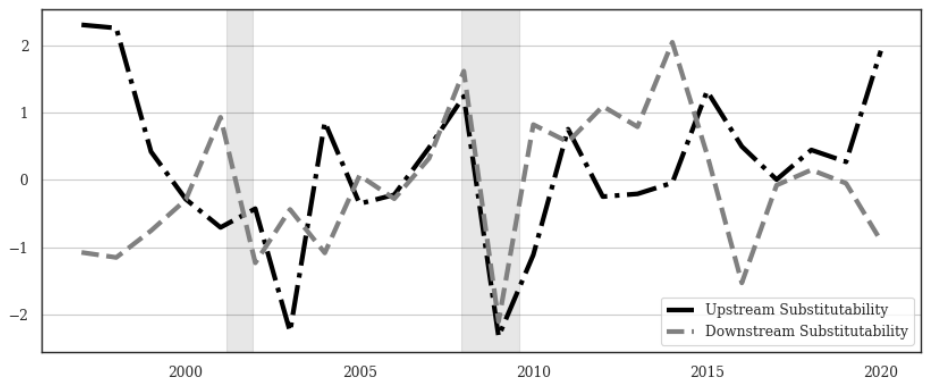

(b) Panel B: Innovations in Average Substitutability

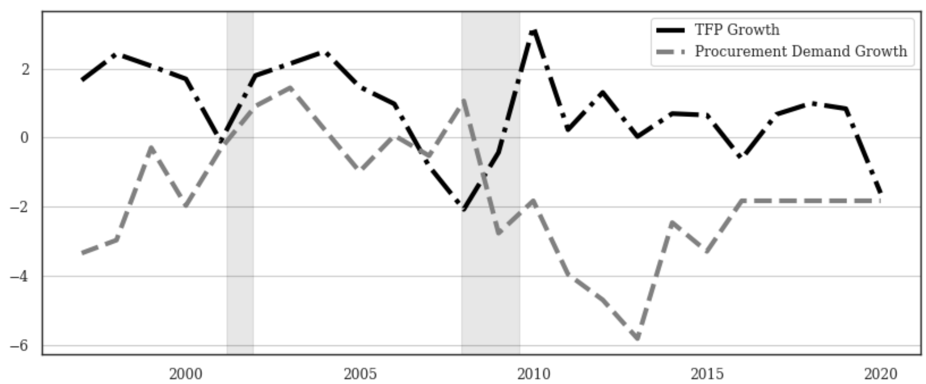

(c) Panel C: Productivity and Demand Growth

Figure 1.8: Network Propagation Risk Factors

Notes: This figure plots the time series of network propagation risk factors (Panel A), the cross-sectional average industry substitutability (Panel B), and aggregate TFP and procurement demand growth (Panel C). Shaded regions indicate NBER-dated recession periods.

Table 1.16: One-Way Sorted Portfolios on Network Propagation Factors (controlling for volatility factors)

Panel A: One-way sorts on upstream propagation beta (controlling for a_t, g_t, and σ_t^{civ})

	1 (Low)	2	3	4	5 (High)	H-L	t(H-L)	MR p-val
$\mathbb{E}[r]-r_f$	12.97	12.12	11.34	10.44	10.11	-2.86	-5.83	0.22
α_{capm}	-0.04	-0.21	-0.21	-0.28	-0.68	-0.64	-4.06	0.00
α_{ff3}	-0.10	-0.14	-0.17	-0.25	-0.39	-0.29	-0.83	0.01
Volatility (%)	13.37	14.81	14.71	19.1	12.64	-	-	-
Book-to-market	0.52	0.53	0.58	0.54	0.50	-	-	-
Market value ($bn)	12.18	6.21	18.07	6.14	15.75	-	-	-

Panel B: One-way sorts on downstream propagation beta (controlling for a_t, g_t, and σ_t^{civ})

	1 (Low)	2	3	4	5 (High)	H-L	t(H-L)	MR p-val
$\mathbb{E}[r]-r_f$	14.97	12.94	12.15	9.78	9.35	-5.61	-9.36	0.09
α_{capm}	0.03	-0.24	-0.29	-0.32	-0.39	-0.43	-1.88	0.00
α_{ff3}	-0.02	-0.10	-0.26	-0.29	-0.30	-0.29	-3.40	0.01
Volatility (%)	18.69	16.42	12.11	14.20	13.02	-	-	-
Book-to-market	0.59	0.55	0.51	0.51	0.52	-	-	-
Market value ($bn)	5.11	12.13	8.96	16.43	15.64	-	-	-

Panel C: One-way sorts on upstream propagation beta (controlling for a_t, g_t, and σ_t^{mkt})

	1 (Low)	2	3	4	5 (High)	H-L	t(H-L)	MR p-val
$\mathbb{E}[r]-r_f$	13.83	13.07	11.33	10.10	9.73	-4.09	-8.10	0.21
α_{capm}	0.01	-0.26	-0.28	-0.35	-0.45	-0.47	-6.84	0.04
α_{ff3}	-0.07	-0.17	-0.20	-0.29	-0.30	-0.23	-2.06	0.03
Volatility (%)	13.68	16.44	17.09	14.12	13.03	-	-	-
Book-to-market	0.51	0.54	0.55	0.54	0.53	-	-	-
Market value ($bn)	11.43	5.87	7.24	19.44	14.40	-	-	-

Panel D: One-way sorts on downstream propagation beta (controlling for a_t, g_t, and σ_t^{mkt})

	1 (Low)	2	3	4	5 (High)	H-L	t(H-L)	MR p-val
$\mathbb{E}[r]-r_f$	13.66	13.5	12.67	9.77	9.24	-4.42	-7.17	0.37
α_{capm}	0.07	-0.26	-0.32	-0.39	-0.41	-0.48	-1.81	0.00
α_{capm}	0.02	-0.17	-0.24	-0.25	-0.33	-0.35	-1.52	0.01
Volatility (%)	19.47	12.94	15.7	14.27	13.03	-	-	-
Book-to-market	0.63	0.54	0.44	0.53	0.52	-	-	-
Market value ($bn)	4.60	10.24	13.33	15.77	14.30	-	-	-

Notes: This table reports average excess returns and post-sample alphas in annual percentages for value-weighted portfolios sorted into quintiles on annual upstream and downstream propagation factors. Sample is between 1997-2021 for more than 10,000 stocks belonging to the BEA 66 non-government industry classifications. Panels A and B control for productivity growth and federal procurement demand growth, while Panels C and D have no controls. I also report average return volatility, book-to-market ratio and market value for each portfolio. To test for significant return spreads, I report t-statistics for the null hypothesis $H_0 : xr_5 = xr_1$, where xr_q is the average return of the q^{th} quintile single sorted portfolio. Moreover, I report p-values for the test $H_0 : xr_{q+1} < xr_q \forall q \leq 4$, calculated via bootstrap following Patton and Timmermann (2010).

Victor Sellemi is the sole author of Chapter 1.

Chapter 2 is coauthored with Edoardo Briganti. The dissertation author is one of the primary investigators of this material.

Chapter 3 is coauthored with Tjeerd de Vries. The dissertation author is one of the primary investigators of this material.

Chapter 2

Why Does GDP Move Before G? It's all in the Measurement

Abstract

We find that the early impact of defense news shocks on GDP is due to a rise in business inventories, as contractors ramp up production for new defense contracts. These contracts do not affect government spending (G) until payment-on-delivery, which occurs 2-3 quarters later. Novel data on defense procurement obligations reveals that contract awards Granger-cause shocks to G identified via Cholesky decomposition, but not defense news shocks. We show that Cholesky shocks to G miss early changes in inventories, and thus result in lower multiplier estimates relative to the narrative method.

2.1 Introduction

The fiscal policy literature has long aimed to quantify the effects of government spending (G) and its underlying transmission mechanism. To do so, researchers must first identify unpredictable government spending shocks that are exogenous to the business cycle. According to Ramey (2016), the two most commonly used approaches for identification are the Cholesky decomposition (Blanchard and Perotti (2002)) and the narrative method (Ramey (2011)). The Cholesky decomposition approach places G first in a vector autoregressive (VAR) model, relying on the assumption that G is predetermined at time t due to decision lags. Practically, this entails regressing G on its lags and on lags of other pertinent state variables, assuming that the resulting OLS residuals represent structural shocks (henceforth Cholesky shocks). By contrast, the narrative approach uses an external instrument (e.g., defense news shocks) which reflects the anticipated shifts in defense spending brought on by exogenous military events, and places this instrument first in a VAR. Both approaches are valid under the right assumptions. Yet, the Cholesky-based method estimates smaller multipliers than the narrative method (i.e., "multiplier-gap"), especially at small horizons. This paper provides an empirical explanation of the multiplier-gap.

We start from the key empirical finding that GDP increases immediately while G increases with a delay following narratively-identified shocks to government spending.[1] Since narrative shocks predict Cholesky-identified shocks to G, proponents of the narrative approach use this as evidence that Cholesky-identified shocks fail to account for anticipation effects of fiscal policy. For instance, Ramey (2011) shows that war-dates Granger-cause (or predict) Cholesky shocks, thus leading to an identification problem since those shocks capture military build-ups with a delay. Moreover, delaying war-dates in the VAR can reconcile resulting estimates from the two methods (i.e., "*it's all in the timing*").

[1]See Ramey and Shapiro (1998), Edelberg et al. (1999), Burnside et al. (2004), Eichenbaum and Fisher (2005), Ramey (2011), Barro and Redlick (2011), Ben Zeev and Pappa (2017) and Ramey and Zubairy (2018). Leeper et al. (2013)'s also suggests to control for anticipation effects to correctly identify fiscal shocks.

However, one question still remains. What causes GDP to move before G in the narrative approach? Ramey (2011) suggests that it is Ricardian behavior of agents to drive the anticipation effect of government spending. In particular, the existence of implementation lags during military build-ups leads to a time-mismatch between the agents' expectations of future G and the actual change in G. Since Ricardian agents respond to changes in the present discounted value of G and taxes, GDP responds even before any actual change in G. However, the strength of this mechanism is still a matter of debate among economists.[2]

We provide empirical evidence of an alternative mechanism. In particular, we show that an increase in business inventories accounts for the initial movement of GDP following a narrative shock. We trace back the inventory effect to an increase in newly awarded defense procurement contracts following a defense news shock. However, war-related contract awards and associated early-stage production occur several quarters before payment-on-delivery. Since government spending tracks payments, early-stage production is recorded in aggregate inventories until delivery. The differential response of aggregate inventories explains the difference in government spending multipliers calculated via the narrative method and the Cholesky decomposition (i.e., "*it's all in the measurement*").

We start by decomposing the increase in GDP after a defense news shock, using quarterly data from the National Income and Product Accounts (NIPA). At the aggregate level, we observe that G responds two quarters after the defense news shock, while GDP has a positive and significant response on impact and in the first quarter. The impact (horizon 0) response is entirely driven by durable consumption, but is not robust to the exclusion of the Korean war from the sample.[3] The horizon 1 response is entirely driven by a strong and robust increase in aggregate investment, and more specifically the business inventories component of investment. Even more specifically, we find from a panel of manufacturing industries that the increase in inventories after war events is driven exclusively by higher

[2]For instance, Monacelli and Perotti (2008), Galí et al. (2007) and Gabaix (2020) propose theoretical models which can dampen the strength of this mechanism. Coibion et al. (2020) find little survey evidence in support of a strong negative income effect.

[3]This is a well-known fact in the fiscal policy literature (see Perotti (2014) and Ramey (2016)).

real inventories in defense sectors. In other words, the response of inventories is a result of contractors ramping-up production.

We directly document the time delay between obligations and payments using our novel quarterly time-series of defense procurement spending and defense procurement obligations. We find that obligations precede payments (and G) by an average of 2-3 quarters. The time-mismatch is discussed in the Department of Commerce's Government Transaction Methodology Paper, which shows that the production of government contractors is not immediately reflected in government spending. Rather, G primarily tracks payments which occur after the delivery of the ordered items, and defense production takes time. To summarize, the recorded time-delay between payments and new orders provides an accounting origin of the positive response of inventories during a military build-up: it is the unpaid production-in-progress which does not yet show up in G.

To better capture shocks to obligated government funds, we order defense procurement obligation first in a VAR. We show that these shocks Granger-cause the Cholesky shocks of government spending. Shocks to obligations, however, do not predict defense news shocks. Intuitively, fluctuations in real government spending, as measured by NIPA, reflect changes in defense spending brought on by military events. The Cholesky shocks to NIPA government spending thus capture these fluctuations. The timing of these shocks, however, is delayed relative to the initial economic impact of a military event reflected in new government orders. As a result, shocks to defense procurement obligations predict the Cholesky shocks. On the other hand, defense news shocks are recorded at the start of a military build-up, when new contracts are awarded and contractors increase production. Thus, defense news shocks are not predictable by shocks to defense procurement obligations.

Finally, we show that narrative shocks lead to higher estimates of the fiscal multiplier than the Cholesky shocks and, on average, more than 84% of their difference (multiplier-gap) is explained by the differential response of inventories. In other words, whenever defense production is characterized by long time-to-build, and contractors are paid after-delivery, the Cholesky shocks will overlook the initial production by defense contractors that is recorded in inventories. Therefore, under these conditions, our findings support the robustness of the narrative method in accurately (i) identifying government spending shocks, and (ii) estimating fiscal multipliers. Under the assumption that obligated funds are predetermined, identification via Cholesky decomposition is still valid as long as the government spending variable is set to obligations, which better captures the timing of federal funds as soon as they are committed to be spent.

The idea that inventories absorb the time-to-build of defense contractors can be traced back to **?**'s analysis of the US economy during the Korean war. Ginsburg argues that changes in government spending have effects before the actual disbursement of money, as captured by G, and that these effects are temporarily reflected in inventories.[4] Therefore, researchers should take into account new government orders to fully understand the impact of government spending changes. To overcome this implementation lag problem, Leduc and Wilson (2013) study the effects of local fiscal policy using obligations rather than outlays.

Similarly, Brunet (2020) suggests that the National Income and Product Account "measures G too late in the process", and constructs an annual measure of funds appropriations by the Department of Defense, termed budget authority. Brunet finds that this measure leads G and uses it to estimate a fiscal multiplier between 1.3 and 1.6, which is higher than typical estimates from the national multiplier literature (see Ramey (2016)). Brunet attributes the difference to implementation-lags and time-to-build in the government spending process, which leads to increased production reflected in private inventory investment before government expenditures.

Our work contributes to this literature in a few ways. To the best of our knowledge, we are the first to study the aggregate and sectoral effects of fiscal shocks on invento-

[4] Extract from page 10 of their NBER book: "*It is apparent that a defense mobilization will provide a stimulus to economic expansion if government expenditures increase the aggregate demand for goods and services. However, the expansion need not await the actual growth of government expenditures. In the first place, some government expenditures for defense will lag behind the placement of orders. For a time, the increased production consequent on the orders will be reflected in private inventory investment rather than in government expenditures.*"

ries.[5] Although **?** also studies inventories, the analysis is restricted to the outbreak of the Korean War. Moreover, we focus on national government spending multipliers and relate them to aggregate obligations. This complements the cross-sectional analysis of Leduc and Wilson (2013), who use obligations to study the effects of state-level highway-construction expenditure, and estimate cross-sectional government spending multipliers.[6]

We also build on the work of Brunet (2020), who provides accounting evidence on the behavior of inventories during a military build-up. We verify this theory empirically using both an aggregate and sectoral analysis of inventories. Additionally, our novel quarterly measure of federal defense procurement obligations has several advantages relative to Brunet (2020)'s annual budget authority series. Firstly, our measure is available at the quarterly frequency rather than annual, which (i) considerably increases the sample size, (ii) allows for a more direct comparison with the other quarterly multiplier estimates from the literature, and (iii) allows us to understand the time-mismatch between contracts and payments at sub-annual frequencies.

With two quarterly series on defense procurement obligations and defense procurement spending, we are able to precisely quantify the time-mismatch between newly awarded contracts and payment to contractors. The focus on defense contracts illustrates the role of time-to-build in generating an accounting delay. Our results show that obligations precede payments (and G) by an average of 2-3 quarters, which could not have been detected with annual data. Finally, we directly relate this accounting delay to the anticipation effect measured by Ramey (2011), and use our findings to reconcile the difference in multiplier estimates obtained using narrative and Cholesky shocks.

The paper is organized as follows. Section 2.2 establishes the positive response of

[5] Researchers have historically overlooked the role of inventories in analyzing government spending shocks, likely due to the use of log-transformations in VAR models, which cannot handle negative inventory values. However, the adoption of other transformation of the data, such as the Gordon and Krenn (2010)'s transformation, does not require the adoption of logs and allows us to analyze the response of aggregate inventories to fiscal policy shocks.

[6] It is well-known that national and local multipliers are two different objects. In particular, the local multiplier is a rough lower bound of the deficit-financed, closed-economy, no-monetary-policy-response national multiplier (see Chodorow-Reich (2019)).

contractor inventories following a defense news shock. Section 2.3 carries out the sectoral level analysis of inventories. Section 2.4 studies the underlying economic and accounting mechanisms driving the response of inventories using novel procurement data. Section 2.5 explores implications of our results in estimating government spending multipliers. Section 3.6 concludes.

2.2 Response of Inventories to Fiscal Shocks

In this section, we decompose changes in the components of real output that are driven by news about future government spending rather than actual government spending. We find that the early response of GDP to defense news shocks is driven by a positive and robust response in business inventories.[7]

Our starting point is Ramey (2011), who finds that aggregate output reacts immediately to news about future war-related defense spending (defense news shocks), while government spending itself has a delayed response.[8] We replicate this result in the top panels of Figure 2.1. Note that GDP responds immediately, while G only responds starting from the second period, marked with a dashed red line.

In particular, we estimate the quarterly impulse response function (IRF) of some outcome y_t of interest (e.g., GDP) using lag-augmented local projections:[9]

$$y_{t+h} = \theta_h \cdot \text{Shock}_t + \beta \cdot \mathbf{X}_t + \varepsilon_{t+h} \tag{2.1}$$

where y_{t+h} is the outcome, Shock_t is the updated series of narratively identified defense

[7]Note that we use the term "inventories" to refer to "Aggregate Changes in Business Inventories", which is one component (along with fixed - residential plus non-residential - investment) of I in the decomposition $GDP = C + I + G + NX$.

[8]See similar results in Ramey and Shapiro (1998), Edelberg et al. (1999), Eichenbaum and Fisher (2005), Ben Zeev and Pappa (2017).

[9]See Jordà (2005) for local projections, LPs, and Montiel Olea and Plagborg-Møller (2020) for econometric details on lag-augmented LPs. Notice that the IRFs obtained via LPs are asymptotically equivalent to the IRFs estimated via VAR (Plagborg-Møller and Wolf (2020)). LPs are more precise in terms of bias-reduction than VAR, however, this comes at a great efficiency cost (Li et al. (2021)). We use LPs for their simplicity and to compare with the literature (e.g. Ramey and Zubairy (2018)).

news shocks from Ramey and Zubairy (2018), and $\mathbf{X}_t$ is a vector of four lags of shocks, government spending, consumption, investments, net-exports, hours worked by the private sector, the three-month Treasury Bill rate and a linear time trend. Following Ramey and Zubairy (2018), we divide all nominal variables by real potential output and the GDP price deflator.

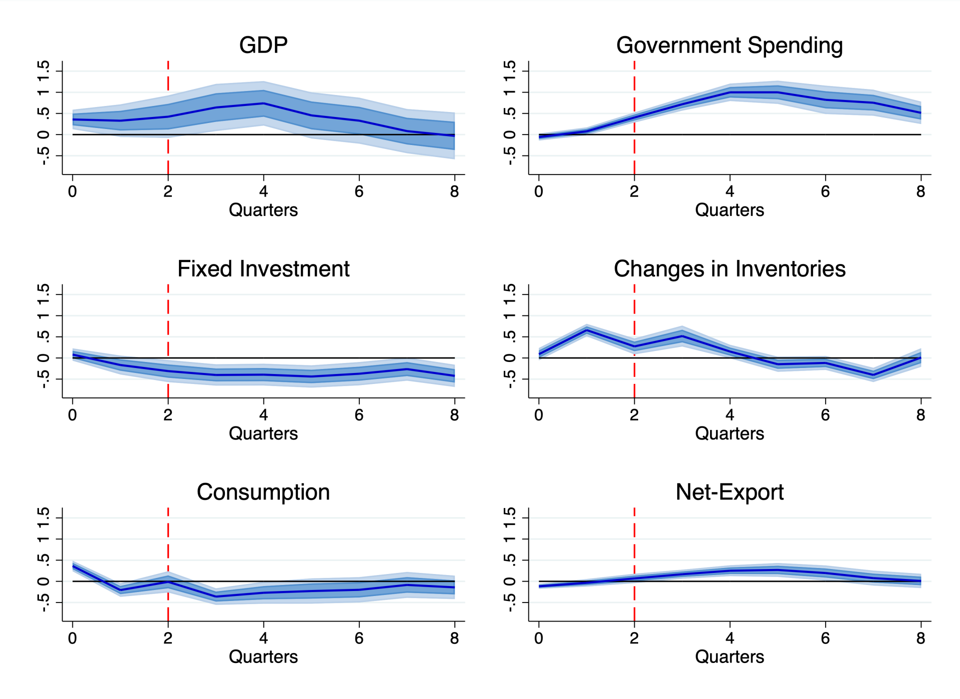

Figure 2.1: Response of GDP and its Components to a Defense News Shock

Notes: IRFs of GDP, G, Investment and Changes in Inventories to a defense news shock are obtained via lag-augmented local projections. Bands represent the 68% and 90% heteroskedasticity robust standard errors. Defense news shocks are obtained from the updated series in Ramey and Zubairy (2018). Sample goes from 1947Q1 to 2015Q4. Values in the Figures are normalized by the peak response of G.

To further investigate the underlying mechanism here, we decompose GDP and estimate the aggregate response of consumption, fixed investment, inventories, government spending, and net-exports to defense news shocks. Note that the IRF of GDP (top-left panel) can be obtained by summing up the ones of all its components.[10] The middle-left panel of Figure 2.1 shows the IRF of Fixed Investments, the middle-right panel the one of Inventories, the bottom-left panel the one of consumption and, finally, the bottom-right panel the one of net-export. Values are normalized by the peak response of G.

[10]This follows from (i) the linearity of the OLS estimator used in local projections and (ii) the way NIPA constructs GDP, as the sum of the components of final demand. See Online Appendix A for the formal proof.

Firstly, consumption at horizon 0 is almost 50% of the peak response of government spending and accounts for almost all of the impact response of GDP. However, it is a well-known fact in the fiscal policy literature that this response is driven by durable consumption at the onset of the Korean war.[11]

Secondly, the positive response of inventories at horizon 1 is equal to more than 50% of the peak response of G. Since we detect either negative or insignificant responses of fixed investment (middle-left panel), consumption (bottom-left panel) and net-export (bottom-right panel) at horizon 1, it is clear that the early increase in GDP relative to G following a defense news shock initially shows up as an increase in inventories.

To the best of our knowledge, we are the first to detect positive effects of inventories to defense news shocks and relate it to the anticipation effect of G detected in Ramey (2011).[12]

Robustness The positive response of inventories is robust to the exclusion of the Korean War (the largest military build-up after World War II) from the sample, indicating that the response of inventories is not driven by periods in which defense shocks dominate.[13]

Secondly, we find that the positive response of inventories is robust to the adoption of other types of fiscal shocks. In particular, we use the Cholesky shocks and shocks identified from a VAR which orders defense procurement obligations first, where defense procurement obligations capture the all universe of defense prime contract awards (we will discuss the construction of this variable in the next sections). We report all robustness checks in the Online Appendix B.

Next, we show in the panel of manufacturing industries that the aggregate response of inventories is driven by an increase in industries which heavily contract to the federal government.

[11] See **?**, **?**, Ramey (2016) and Binder and Brunet (2021). Consistently with the literature, we detect no significant effect of durables in samples which exclude the Korean war.

[12] **?** estimate the effect of shocks to G identified via Cholesky decomposition on a multitude of variables and also find a positive early response of inventories. They do not discuss this result in the paper.

[13] We believe that it is important to include the largest war events in the sample as they mimic natural experiments involving government spending. However, we are aware of potential confounding factors (see Perotti (2007), Fisher and Peters (2010), Perotti (2014) and Ramey (2016)).

2.3 Industry Analysis: Who is Responding?

Given the positive and robust aggregate response of inventories, we study heterogeneity in this response across industries in response to war events. We find that the positive response is driven by defense industries which increase inventories during a military build-up. To do so, we use monthly data from the Bureau of Economic Analysis (BEA) to construct a panel of real inventories for 18 manufacturing industries between January 1959 and December 1997.[14]

The production of defense goods is concentrated in the manufacturing sector (see e.g., Ramey and Shapiro (1998), Nekarda and Ramey (2011) and Cox et al. (2021)). However, the level of government involvement varies greatly among manufacturing sub-industries. For example, the "Other Transportation Equipment" sector has 34% of its sales directly from the government. Accounting for indirect sales via input-output connections, the sector's dependence on government purchases rises to 42% and 44% with first and second order downstream connections included (as done in Nekarda and Ramey (2011)). This heavy reliance on government purchases is unsurprising given that the sector includes sub-industries like Aircraft, Ship Building, Guided Missiles, and Space Vehicles. Conversely, the "Wood Products" sector has no sales to the government as it does not include any defense item producers.

Therefore, we construct a weight θ_i for each industry which captures the long-run average share of industry sales coming from government purchases. Using industry-by-industry input-output matrices, our weights include up to second-order downstream connections.[15] Then we estimate the following equation:

$$\text{Invt}_{i,t+h} = \lambda_{ih} + \alpha_h \cdot \text{War}_t + \beta_h \cdot \text{War}_t \cdot \theta_i + \sum_{p=1}^{12} \varphi_{ph} \cdot \text{Invt}_{i,t-p} + \varepsilon_{i,t+h} \tag{2.2}$$

[14] We thank Valerie Ramey for providing this data. Our data ends in 1997, however, most of the variation in defense spending comes from before the Nineties (Vietnam War and Soviet invasion of Afghanistan).

[15] We don't find that downstream linkages matter beyond the second order degree of connection. See Online Appendix C.2 for a detailed derivation of industry weights.

where $h = 0, 1, ..., 24$, Invt_{it} is total real inventories of industry i in month t, λ_{ih} is an industry fixed-effect, and War_t is war dates.[16] Consistent with Ramey and Shapiro (1998) and Eichenbaum and Fisher (2005), our war event variable is a weighted dummy with value 1 on March 1965 and 0.3 on January 1984 to indicate the start of the Vietnam War and Soviet invasion of Afghanistan, respectively.

We are interested in the estimands α_h and $\alpha_h + \beta_h$. The former is the response of inventories for those industries not connected to the government (i.e., $\theta_i = 0$). The latter is the response of industries which are highly connected to the government through government purchases (i.e., $\theta_i = 1$). If war dates have a differential positive effect on sectoral inventories which is proportional to the degree of connection to the government, we expect $\beta_h > 0$.[17]

Figure 2.2 shows a significant positive and long-term differential response ($\alpha_h + \beta_h$) of defense industries' inventories to war dates. On the other hand, the change in inventories for those industries who do not supply the government (α_h) is negative and close to zero. Therefore, all of the effect of war dates on inventories is explained by the degree of connection of each sector to the government.

Robustness We verify that this differential response of defense industries' inventories is not driven by their different sensitivity to the business-cycle. In particular, we replace War_t with monetary policy shocks constructed narratively by Romer and Romer (2004) and updated by Wieland and Yang (2020) and estimate the differential response ($\alpha_h + \beta_h$) to be statistically indistinguishable from zero. This confirms that the reaction of federal contractors to defense news shocks is driven by war-related forces and not the associated

[16] We use war dates instead of defense news shocks since the former can easily be converted into monthly frequency to match our inventories data.

[17] Our approach differs from traditional shift-share methods, such as those examined in Goldsmith-Pinkham et al. (2020) and Borusyak et al. (2022). Unlike those studies, which primarily focus on cross-sectional frameworks and require instrumental variables, we investigate the impact of an aggregate exogenous shock (i.e., war-dates) on sectoral inventories and its heterogeneous effects on defense industries, as captured by the interaction between the shock and industry weights. Moreover, since we use long-run averages for our industry weights and we account for any time-invariant fixed effects through industry fixed effects, we are not concerned about the potential endogeneity of our industry weights.

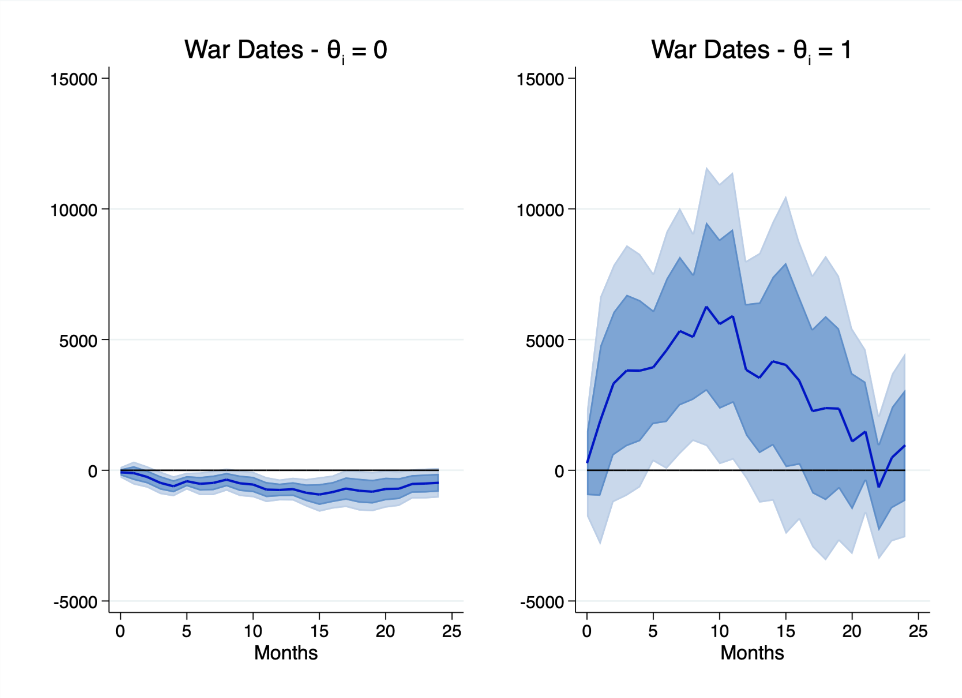

Figure 2.2: Response of Sectoral Inventories to War Events

Notes: Left panel shows estimates of α_h (response when $\theta_i = 0$), right panel reports estimates of $\alpha_h + \beta_h$ (response when $\theta_i = 1$). Weights are normalized by maximum weight (i.e. the one of Other Transportation Equipment Manufacturing). Since Real Inventories are trending, data is filtered using Hamilton (2018)'s filter (we set $h = 24$ and $p = 12$, that is two years lag plus one more year of lags). The unit of real inventories is millions of 2005 chained dollars. Sample goes from 1959-Jan to 1997-Dec and uses 18 sectors breakdown of Manufacturing. Confidence bands are 68% and 90%. Standard errors are obtained via Bootstrap (standard Stata routine for `xtreg`: we use `vce(boot)` and set the seed for replicability of results; Stata uses a non-parametric type of bootstrap which resamples data with replacement).

business-cycle fluctuations.[18]

Furthermore, we make sure that the differential response of defense industries during a military build-up is not driven by spurious correlation. In particular, we re-estimate Equation (2.2) using randomly re-shuffled weights as commonly done in the production network literature (e.g. see Ozdagli and Weber (2020)). Again, we estimate the differential response ($\alpha_h + \beta_h$) to be statistically indistinguishable from zero and we report the results of these robustness checks in the Online Appendix C.1.

[18] We thank Juan Herreño for suggesting this test.

2.4 Why Inventories and not G?

This section explains why the early stage production of defense industries during a military build-up is absorbed by inventories and not government spending (G). Briefly, part of the production process occurs between contract award and delivery, and contractors are paid after delivery. Since G is constructed primarily using payments, it measures production with delay (see also Brunet (2020)). To accurately track production as it happens, NIPA uses inventories to align the timing of production with the contract award and payment. Chapter 7 of NIPA's Handbook states:[19]

"*A general principle underlying NIPA accounting is that production should be recorded at the time it occurs. [...] The recording of movements of goods in inventory — materials and supplies, work-in-process, and finished goods — and from inventories to final sales provides the means to allocate production to the period in which it occurred.*"

The Procurement Process In the defense procurement process, obligations and spending are two distinct stages. The process starts with the award of a contract, which is when the government is legally bound to pay for goods/services. Although contractors are notified of contract opportunities before the award date through pre-award solicitations, these solicitations are typically posted in the same quarter as the award date and made available to contractors on a federally managed online database.[20]

After contracts are awarded, contractors launch a potentially long production process. In particular, contract-level data indicates that the mean and median duration of \$1 defense procurement contract are 4.2 and 5.4 years, respectively. We measure duration as the period of performance, or the number of days between award date and contract end (full delivery) date. We find that total defense procurement spending is dominated by few very large contracts with very long duration. Using the same data, Cox et al. (2021) report

[19] We thank Junyuan Chen and Valerie Ramey to bring up to our attention this meaningful passage.

[20] Generally, solicitations are posted on beta.sam.gov and are linked to the eventual contract award using the solicitation ID. Further discussion can be found in the Online Appendix D.5

a very short average contract duration. However, their estimated duration is not weighted by contract size. Weighting is necessary to find the duration of \$1 of spending and not the average duration of contracts. This difference matters, since most of procurement spending comes from few very large contracts. If we do not weigh by contract size, our results are consistent.[21]

Given that production takes a long time, when do associated payments actually occur? According to the Federal Acquisition Regulation (FAR), the canonical rule for payments to federal contractors from government agencies is *payments-after-delivery* (see FAR 32.904).[22] Finally, NIPA constructs G using mainly outlays, that is, payments to contractors (see Brunet (2020)).

Therefore, NIPA's accounting rules result in a delay in tracking defense production due to the time it takes to produce items. In the following sections, we create a measure of defense procurement spending and obligations to directly observe the time gap between the start of production (when the contract is awarded) and when NIPA records it (at delivery).

Construction of Defense Procurement Spending and Obligations We construct a novel database of defense procurement spending and obligations. Spending measures payments from federal agencies to contractors, while obligations measure the total value of federal funds as soon as they are contractually obligated to firms. To construct the spending

[21]We use defense contract data from the federal procurement data system (FPDS) from 2000 to 2020. FPDS encompasses every federal transaction at daily frequency. We report results in the Online Appendix D.2.

[22]Certain contracts are also subject to partial-delivery-payments. However, given the multiple year average duration of \$1 of procurement spending, we still observe several quarter-long delays in partial deliveries. We further clarify this point in the Online Appendix D.3.

series, we use the accounting identity discussed in Cox et al. (2021):[23]

$$\begin{aligned}(\text{Procurement Spending})_t &\approx (\text{Intermediate Goods \& Services Purchased})_t + \\ &+ (\text{Change in Government Fixed Assets})_t + \\ &- (\text{Investment R\&D})_t \\ &\approx (\text{Payment to Contractors})_t,\end{aligned}$$

where all variables are obtained from the National Income and Product Accounts (NIPA). The top panel of Figure 2.3 plots this measure of defense procurement spending along with the annual measure of procurement spending of Dupor and Guerrero (2017), aggregated over states. The two measures are virtually identical before 1984, but afterwards the Dupor and Guerrero (2017) series omits contract actions with value less than $25,000 and thus systematically underestimates our NIPA-based series. From 2000 onward, we also aggregate federal agency payments from the universe of procurement contracts, available in the Federal Procurement Data System (FPDS), and find that our measure is consistent.

To construct the obligations series, we aggregate the value of procurement contracts awarded by the Department of Defense (DoD) from the universe of procurement contracts recorded in the Federal Procurement Data System (FPDS). Since this data is only available from 2000 onward, we also collect historical information from the periodical *Business Conditions Digest* (henceforth BCD) which is available from January 1951 to November 1988. We use information from the contract and spending data to impute missing quarters and construct a quarterly time series of defense procurement obligations.[24]

Direct Evidence of Time Mismatch in Defense Procurement The bottom-left panel of Figure 2.3 plots spending and obligations from Jan 1951-Nov 1988, and the bottom-

[23] Further details on the accounting origin of procurement spending is discussed in the Online Appendix D.1.

[24] Many thanks to Valerie Ramey for providing the BCD data. We remand to our Online Appendix D.4 for extra details on the sources of contract level data and the construction of the series.

(a) Annual Defense Procurement Spending

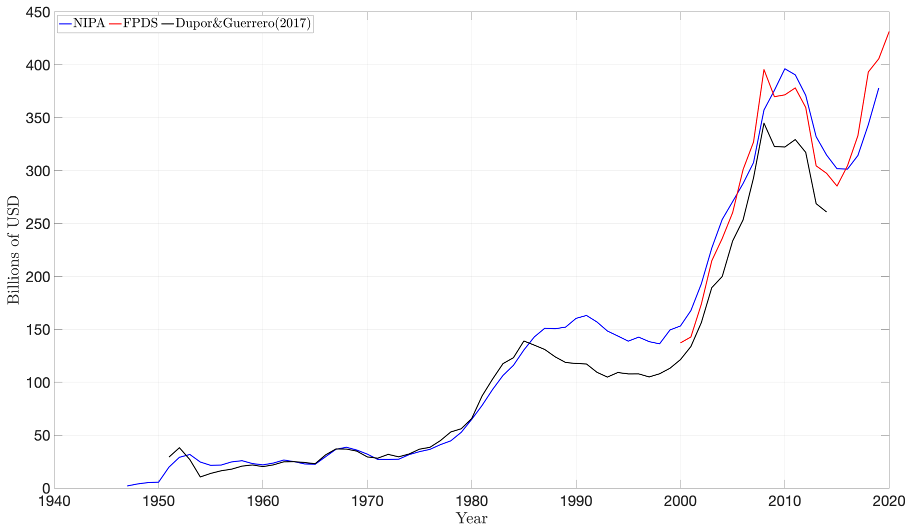

(b) Time Lag in "BCD"

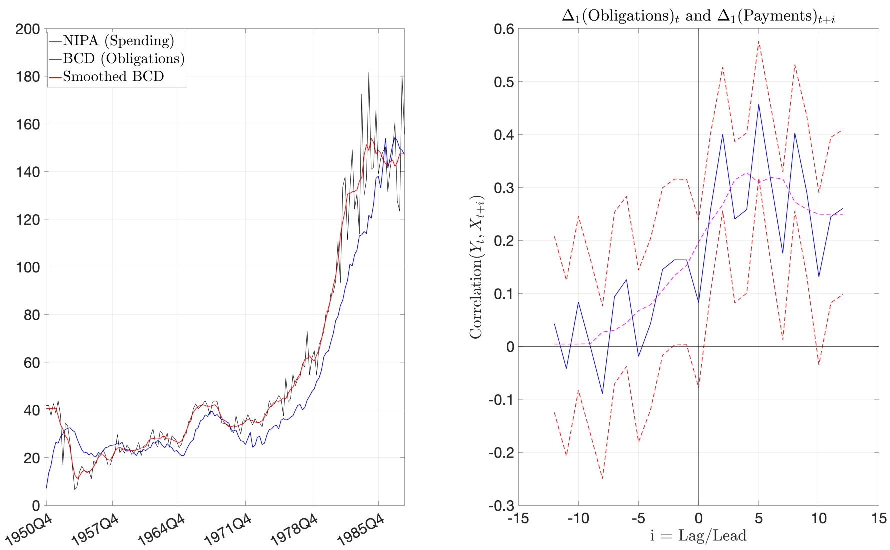

Figure 2.3: Federal Defense Procurement Obligations vs Spending

Notes: Top panel (a) compares different measures of defense procurement at annual frequency. The bottom-left panel (b) compares defense procurement spending (i.e. payments) as we construct it from NIPA data, to defense procurement obligations (i.e. awards) from "BCD". The bottom-right panel shows the lead-lag correlation map between the two: *Corr* $(\Delta_1(\text{Obligations})_t, \Delta_1(\text{Payments})_{t+i})$, where Δ_1 is the first difference operator. Sample: 1951Q1 to 1988Q4.

right panel reports the lead-lag correlation.[25] From the right panel, the average lead-lag correlation significantly peaks in the North-East quadrant of the map. This suggests that

[25]Lead-lag correlations are useful for studying relationships in time between variables. For example, Smets et al. (2019) use it to study the timing of propagation of inflation from upstream to downstream sectors.

changes in obligations are more highly correlated with delayed changes in spending rather than current changes in spending. The results replicate for more recent obligations data obtained from FPDS and when we look at quarterly year-to-year changes instead of simple changes. We report these robustness checks in the Online Appendix D.2. On average, we find that obligations lead spending by 2-3 quarters.

The payment (or government outlay) thus occurs several quarters after the defense contract award. This finding is consistent with the results of Leduc and Wilson (2013) and Brunet (2020) in the context of highway spending and the aggregate annual defense budget. Moreover, this is confirmed directly by the Department of Commerce's Government Transaction Methodology Paper:[26]

"*The largest timing difference is for national defense gross investment for relatively long-term production items, such as aircraft and missiles, for which the work in progress is considered as part of business inventories until the item is completed and delivered to the Government.*"

In other words, early-stage production associated with long procurement contracts is recorded at an aggregate level in inventories until the delayed payment-on-delivery. The value of completed and paid contract work is then moved from inventories to G. We can observe the delay between defense contract awards and payment directly from our data.

Finally, in the Online Appendix E, we distinguish between the response of defense contractors to actual contract awards and the anticipation of future contract awards. Firms may increase their inventories in preparation for future awards, whether to minimize adjustment costs or reduce delivery times (i.e., production smoothing). While we identify evidence of the latter, it is of lesser importance compared to the response to actual contract awards.

[26] Many thanks to Gillian Brunet for redirecting us to that document.

2.5 Implications for the Government Spending Multiplier

In this section, we argue that the Cholesky shocks to government spending as measured by NIPA do not capture early-stage production associated with newly awarded federal procurement contracts during a military build-up. This leads to lower multiplier estimates relative to the narrative method. We show that 84% of the difference in multipliers (multiplier gap) is driven by a differential early response of inventories following a defense news shock.

Shock Predictability Ramey (2011) shows that narrative shocks predict (Granger-cause) the Cholesky shocks, which implies that those shocks are missing part of the early response in GDP. To show that the missing early response is associated with early-stage production related to defense procurement contracts, we further show that shocks to defense procurement obligations Granger-cause the Cholesky shocks to G, while do not Granger-cause defense news shocks. We construct defense procurement obligation shocks by ordering defense procurement obligations first in a VAR.[27] In turn, we use two series of defense procurement obligations: one which goes from 1947Q1 to 1988Q4, which uses data from BCD ("*BCD series*") and one which uses information from defense procurement spending and FPDS to extend the BCD data up to 2015Q4 ("*extended series*"). Our full sample spans 1947Q1 to 2015Q4, and Table 2.7 summarizes the results.

[27] The variables employed here are identical to the ones utilized in Section 2.2.

Table 2.1: Predictability of Cholesky Shocks via Obligations

Predicted	Predictor	F	Pvalue	Korea
Cholesky Shocks	Obligation Shocks (Extended Series)	5.63	0.0%	Yes
Cholesky Shocks	Obligation Shocks (BCD Series)	3.45	0.1%	Yes
Cholesky Shocks	Obligation Shocks (Extended Series)	4.24	0.0%	No
Cholesky Shocks	Obligation Shocks (BCD Series)	2.41	1.9%	No
Obligation Shocks (Extended Series)	Cholesky Shocks	1.07	38.7%	Yes
Obligation Shocks (BCD Series)	Cholesky Shocks	0.57	84.2%	Yes
Obligation Shocks (Extended Series)	Cholesky Shocks	1.67	10.7%	No
Obligation Shocks (BCD Series)	Cholesky Shocks	1.12	35.31%	No
Defense News Shocks	Obligation Shocks (Extended Series)	0.73	66.1%	Yes
Defense News Shocks	Obligation Shocks (BCD Series)	0.75	64.4%	Yes
Defense News Shocks	Obligation Shocks (Extended Series)	0.32	95.7%	No
Defense News Shocks	Obligation Shocks (BCD Series)	0.59	78.7%	No

Notes: Granger Causality test is a Wald test on the 8 lags of the predictor while controlling for 4 lags of the predicted variable. In Appendix F, we report analogous results for Cholesky shocks to an index of Top 3 defense contractor excess returns, constructed as in Fisher and Peters (2010). We find no significant predictability in either direction for this index.

The top panel of Table 2.7 shows that shocks to defense procurement obligations predict the Cholesky shocks. On the other hand, the second panel shows a much weaker relationship in the other direction, especially when you omit the Korean War from the sample. Our results are consistent with Ramey (2011). The bottom panel shows that shocks to defense procurement obligations do not predict defense news shocks. This indicates that early economic effects of newly awarded contracts, which are missed by the Cholesky shocks to G, are captured using defense news shocks.

Government Spending Multipliers In most macroeconomic studies, researchers are interested in the economic effects of government spending from the moment funds are contractually obligated and contractors begin reacting. In this setting, the actual transfer of cash is not the main focus. Given our results from the previous section, we argue that the Cholesky shocks are capturing transfers of cash rather than obligation of funds.

We begin with an illustrative example of this problem around the outbreak of the Korean War in the top-panel of Figure 2.4.

(a) Illustration of the Delay - Korean War

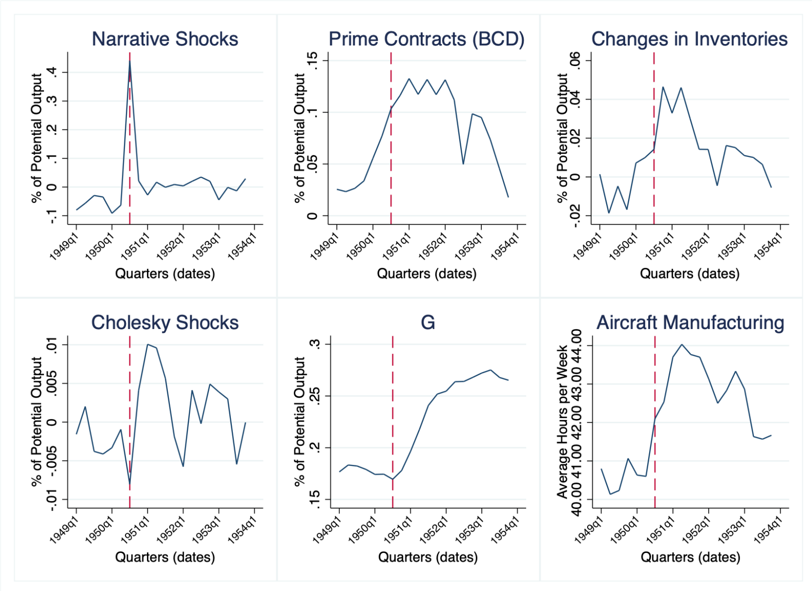

(b) Consequences of the Delay - Multiplier Underestimation

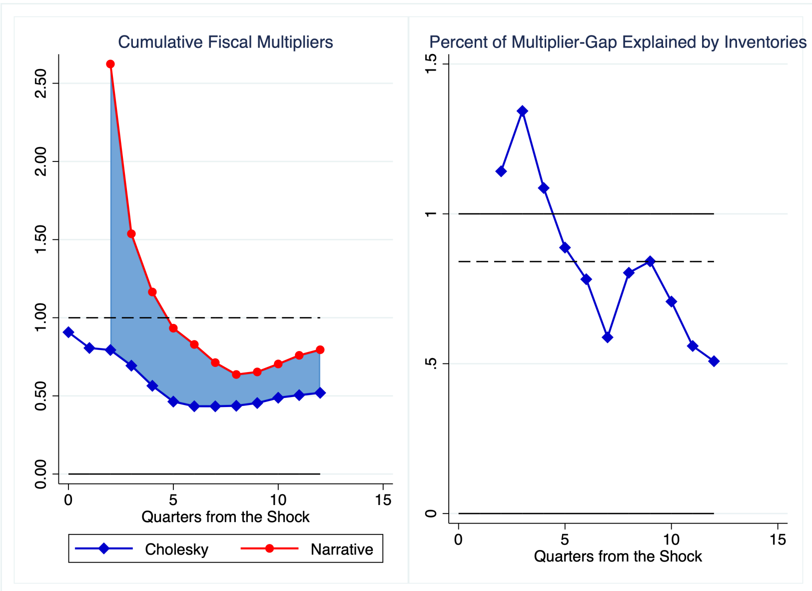

Figure 2.4: Illustration and Consequences of the Delay

Notes: Top panel (a) illustration of the delay during the Korean war. The bottom-left panel (b) compares the point estimates of the calculated fiscal multipliers from horizon 0 to 12 quarters. Sample: 1947Q1 to 2015Q4. The bottom-right panel shows the share of *multipliers-gap* explained by the differential response of inventories (dashed black line is the average of the response). Share is calculated only when the multiplier gap is finite and positive.

In the summer of 1950 (Q3), we observe a large defense news shock associated with the outbreak of the Korean War. However, the Cholesky shock to NIPA's measure of G does not spike until 2-3 quarters later. Unsurprisingly, G has a slow positive response. On the other hand, defense procurement obligations react almost immediately to the shock.

In other words, the DoD begins awarding defense procurement contracts at the onset of the war. We also observe quick increases in inventories starting from 1950Q4 as well as in defense production, proxied by average hours of production and non-supervisory workers in the aircraft industry.[28] Therefore, the Cholesky shocks fail to capture the initial production of defense industries in response to newly granted contracts at the onset of the Korean war. This is consistent with our previous Granger-causality test results.

We now show that this delay leads to the underestimation of the fiscal multiplier when using Cholesky decomposition as an identification method. In particular, we show that, on average, 84% the difference in fiscal multipliers estimated using the Cholesky and narrative methods is explained by a difference in capturing the early response of inventories.

Following Ramey (2016), we estimate cumulative fiscal multipliers using LP-IV with both Cholesky shocks to G and narratively identified defense news shocks. We use the following estimation equation:[29]

$$\sum_{h=0}^{H} y_{t+h} = \gamma_H + \hat{\mathcal{M}}(H) \cdot \underbrace{\sum_{h=0}^{H} g_{t+h}}_{\text{instrument with Shock}_t} + \text{lags}_t + \varepsilon_{t+h}, \tag{2.3}$$

where $\hat{\mathcal{M}}(H)$ is the cumulative government spending multiplier at horizon H, y_t is GDP at time t, g_t is government spending at time t, Shock_t is an exogenous instrument for cumulative government spending, and lags_t contains lagged values of the shock, government spending, consumption, investment, hours worked and 3 months T-Bill rate. We rescale nominal variables by potential output. The narrative method sets Shock_t equal to the defense news shock variable, while the Cholesky identification is equivalent to setting

[28] Production workers account for 82% of total private employment, on average (see Nekarda and Ramey (2020)). We choose the Aircraft industry since it specializes in defense production and we use average hours of production workers since total hours is a lagged measure of production (see Bils and Cho (1994) and Fernald (2012b)). We further clarify this point in the Online Appendix C.3. Furthermore, in the Online Appendix C.1 we show that this measure of defense production responds strongly and positively to both defense news shocks and defense procurement obligations.

[29] More technical details on LP-IV are available in Stock and Watson (2018).

Shock_t equal to G.

The bottom-left panel of Figure 2.4 shows that the Cholesky method delivers uniformly lower point estimates of the multiplier relative to the narrative method. To investigate how much of the multiplier gap can be explained by a differential response in inventories, we break down the multiplier in different components, each accruing to one of the components of GDP.

We start from the result discussed in Ramey (2016) and Stock and Watson (2018), that the one-step LP-IV approach delivers an estimate of the multiplier which is analytically equivalent to the one obtained following a two steps procedure consisting in (i) estimating the cumulative impulse response functions of GDP and G to a government spending shock via local projections and (ii) by taking their ratio:

$$\hat{\mathcal{M}}_{GDP}(H) = \frac{\sum_{h=0}^{H} \hat{\theta}_{GDP,h}}{\sum_{h=0}^{H} \hat{\theta}_{G,h}}, \quad \forall H = 0, 1, ...$$

where $\hat{\theta}_{GDP,h}$ and $\hat{\theta}_{G,h}$ are the estimated IRFs of G and GDP to a government spending shock. For instance, if we used defense news shocks, they would be equal to the estimated IRFs of GDP and G shown in the top-left and top-right panel of Figure 2.1. Furthermore, since IRF of GDP can be obtained by summing up the IRFs of each of its components, we can break down the fiscal multiplier as follows:

$$\underbrace{\frac{\sum_{h=0}^{H} \hat{\theta}_{GDP,h}}{\sum_{h=0}^{H} \hat{\theta}_{G,h}}}_{\hat{\mathcal{M}}_{GDP}(H)} = 1 + \underbrace{\frac{\sum_{h=0}^{H} \hat{\theta}_{C,h}}{\sum_{h=0}^{H} \hat{\theta}_{G,h}}}_{\hat{\mathcal{M}}_{\text{C}}(H)} + \underbrace{\frac{\sum_{h=0}^{H} \hat{\theta}_{I_{\text{Fixed}},h}}{\sum_{h=0}^{H} \hat{\theta}_{G,h}}}_{\hat{\mathcal{M}}_{I_{\text{Fixed}}}(H)} + \underbrace{\frac{\sum_{h=0}^{H} \hat{\theta}_{I_{\text{Invy}},h}}{\sum_{h=0}^{H} \hat{\theta}_{G,h}}}_{\hat{\mathcal{M}}_{I_{\text{Invy}}}(H)} + \underbrace{\frac{\sum_{h=0}^{H} \hat{\theta}_{NX,h}}{\sum_{h=0}^{H} \hat{\theta}_{G,h}}}_{\hat{\mathcal{M}}_{\text{NX}}(H)}$$

Notice that each component of the fiscal (GDP) multiplier corresponds to the ratio of the area under the IRF of the corresponding component of GDP and the area under the IRF of G. For instance, the inventory-multiplier obtained via defense news shocks, $\hat{\mathcal{M}}_{I_{\text{Invy}}}^{\text{News}}(H)$, is equal to the area under the IRF of inventories up to horizon H, shown in the middle-right panel of Figure 2.1, divided by the one of Government spending, plotted in the top-right panel of the same figure.

If we differentiate the above expression, and divide by the left-hand side, we have:

$$(\forall H) \quad 1 = \underbrace{\frac{d\hat{\mathcal{M}}_{I_{\text{Invy}}}(H)}{d\hat{\mathcal{M}}_{GDP}(H)}}_{:=\Delta\% I_{\text{Invy}}(H)} + \underbrace{\frac{d\hat{\mathcal{M}}_{\text{C}}(H)}{d\hat{\mathcal{M}}_{GDP}(H)} + \frac{d\hat{\mathcal{M}}_{I_{\text{Fixed}}}(H)}{d\hat{\mathcal{M}}_{GDP}(H)} + \frac{d\hat{\mathcal{M}}_{\text{NX}}(H)}{d\hat{\mathcal{M}}_{GDP}(H)}}_{\Delta\%\text{Other}(H)}$$

$$1 = \Delta\% I_{\text{Invy}}(H) + \Delta\%\text{Other}(H),$$

where $\Delta\% I_{\text{Invy}}(H)$ represents the share of the multiplier-gap, $d\hat{\mathcal{M}}_{GDP}(H)$, explained by differences in the response of inventories, $d\hat{\mathcal{M}}_{I_{\text{Invy}}}(H)$, while $d\Delta\%\text{Other}(H)$ refers to all the other components of GDP.

Therefore, we calculate and breakdown the fiscal multiplier using both defense news shocks (News) and Cholesky shocks (Chol), then we calculate the share of multiplier gap explained by inventories, as suggested by the previous expression:

$$\Delta\% I_{\text{Invy}} = \frac{\hat{\mathcal{M}}^{\text{News}}_{I_{\text{Invy}}}(H) - \hat{\mathcal{M}}^{\text{Chol.}}_{I_{\text{Invy}}}(H)}{\hat{\mathcal{M}}^{\text{News}}_{\text{GDP}}(H) - \hat{\mathcal{M}}^{\text{Chol.}}_{\text{GDP}}(H)}$$

which computes the proportion of the multiplier gap (denominator) arising from using the narrative and Cholesky methods, explained by differences in the inventory multiplier (numerator). The bottom-right panel of Figure 2.4 plots $\Delta\% I_{\text{Invy}}(H)$ up to horizon 8 (solid pink line) along with its average (dark dash line). On average, 84% of the multiplier gap can be explained by the differential response of inventories as captured by the shocks. In the Online Appendix F, we show that this result is robust to the exclusion of the Korean War.

To summarize, the identification of government spending shocks via Cholesky decomposition fails to fully capture early-stage defense production which is reflected in inventories, which results in underestimated multipliers. This is due to NIPA G's delayed tracking of defense production during military build-ups. Our Granger-causality test results are consistent with this intuition. This finding raises a major challenge in identifying government spending shocks through the Cholesky decomposition, provided there exists a long enough time-mismatch between orders and payments in the government spending process.

2.6 Conclusion

The National Income and Product Accounts (NIPA) tracks production by monitoring changes in inventories. During a military buildup, defense industries increase production in response to new procurement contracts, which results in a rise in inventories and GDP. Once the production of defense items, such as aircraft and missiles, is finished, they are delivered to the government and the contractors receive payment. This causes inventories to decrease and government spending (G) to increase as payments are recorded. The onset of a war results in GDP responding faster than G due to (1) accounting procedures and (2) the time required for production in the defense sector.

The findings of our study support the idea that the early rise in GDP relative to G after a defense news shock, as described by Ramey (2011), can be attributed to an increase in inventories. Our analysis of manufacturing sector data reveals that defense industries are responsible for the rise in inventories. By creating new quarterly time series that track defense procurement contract awards and payments, we were able to observe a 2-3 quarter gap between the two. This delay provides evidence for the existence of a time-to-build period for defense production.

Our study has three significant implications. Firstly, it provides a straightforward explanation for the early reaction of GDP compared to G in response to a defense news shock, which was previously believed to be due to households' Ricardian behavior (negative wealth effect). Secondly, the results indicate that shocks to defense procurement obligations predict Cholesky shocks to government spending, which is a major issue in the identification of macroeconomic shocks (as noted by Ramey (2016)). Lastly, the delay in these shocks leads to an under-estimation of the response of inventories which is responsible for 84%, on average, for the under-estimation of the fiscal multiplier estimated by the narrative method. The impact of shocks to defense procurement obligations on macroeconomic variables extends beyond the scope of this paper and remains a subject for future investigation.

Our findings highlight the significance of the early effects of G, as reflected in the increase in inventories. Policymakers and economists should take into account measurement delays in government spending when evaluating the impact of government purchases on the economy.

Appendices

2.7 Breaking Down the Response of GDP

In this section, we decompose the response of GDP to a defense news shock into its underlying components. To do so, we exploit the linearity of the OLS estimates which are used to construct the impulse response functions (IRFs) via local projection.

In particular, we first calculate the IRF of GDP to a defense news shock by regressing GDP on defense news shocks and four lags of investment, government spending, net-export, consumption total hours worked in the private sector, the 3-months T-Bill rate, defense news shocks and a linear time trend. We divide all nominal variables by nominal potential GDP (we take real potential GDP from Ramey and Zubairy (2018) and multiply it by the GDP price deflator). In particular, we group this set of lagged variables and the time trend into matrix X_t, and the IRF of GDP is the coefficient θ_h^{GDP} in the following linear equation:

$$\text{GDP}_{t+h} = \theta_h^{\text{GDP}} \cdot \text{News}_t + X_t \cdot \boldsymbol{\beta}^{\text{GDP}} + \varepsilon_{t+h} \quad h = 0, 1, ..., 8.$$

We report the estimated IRF of GDP in the left panel of Figure 1 in the main text. Repeating this procedure for all four components of GDP, we estimate the following set of linear equations:

$$\mathrm{G}_{t+h} = \theta_h^{\mathrm{G}} \cdot \mathrm{News}_t + X_t \cdot \boldsymbol{\beta}^{\mathrm{G}} + \varepsilon_{t+h}^{G} \quad h = 0, 1, ..., 8$$

$$\mathrm{C}_{t+h} = \theta_h^{\mathrm{C}} \cdot \mathrm{News}_t + X_t \cdot \boldsymbol{\beta}^{\mathrm{C}} + \varepsilon_{t+h}^{C} \quad h = 0, 1, ..., 8$$

$$\mathrm{I}_{t+h} = \theta_h^{\mathrm{I}} \cdot \mathrm{News}_t + X_t \cdot \boldsymbol{\beta}^{\mathrm{I}} + \varepsilon_{t+h}^{I} \quad h = 0, 1, ..., 8$$

$$\mathrm{NX}_{t+h} = \theta_h^{\mathrm{NX}} \cdot \mathrm{News}_t + X_t \cdot \boldsymbol{\beta}^{\mathrm{NX}} + \varepsilon_{t+h}^{NX} \quad h = 0, 1, ..., 8$$

Given that the decomposition of GDP is additive and all equations have the same set of controls X_t, it is easy to show that:

$$\hat{\theta}_h^{\mathrm{GDP}} = \hat{\theta}_h^{\mathrm{G}} + \hat{\theta}_h^{\mathrm{C}} + \hat{\theta}_h^{\mathrm{I}} + \hat{\theta}_h^{\mathrm{NX}} \quad \text{for all } h = 0, 1, ..., 8.$$

where the ˆ denotes an OLS estimate. Therefore, we decomponse IRF of GDP to a defense news shock into its four underlying components, reported in Figure 2.5.

Figure 2.5 shows that aggregate consumption at horizon 0 and aggregate investment at horizon 1 drive the early increase in GDP after a defense news shock.

Consumption and Investment We can further decompose the responses of consumption and investment to better understand what drives their early response. In particular, we apply the same methodology to estimate the IRFs of inventories and residential plus non-residential fixed investment (components of investment) to a defense news shock. Similarly, we estimate the IRFs of durable consumption and the sum of non-durable and service consumption. As before, we consider variables in nominal terms, divide by the GDP price deflator and multiply by real potential output (Gordon and Krenn (2010) transformation).

We report the IRFs of these four components of consumption and investments to a defense news shocks in Figure 2.6. We observe that the horizon 0 response of consumption largely shows up in durables while the horizon 1 response of investment is driven by inventories.

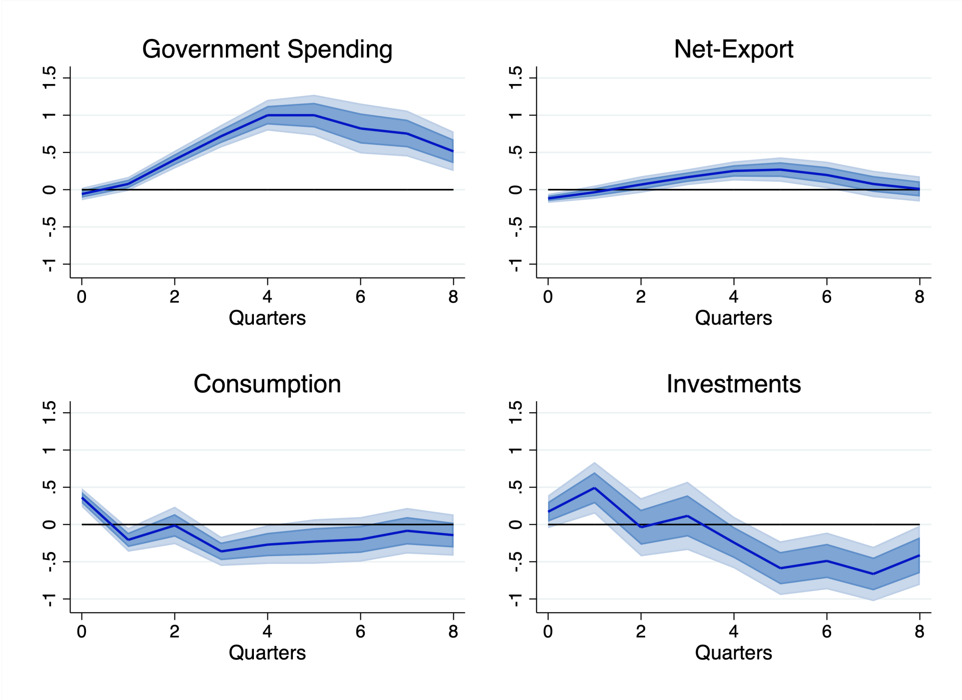

Figure 2.5: Response of GDP Components to a Defense News Shock

Notes: IRFs of GDP, G, Investment and Changes in Inventories to a defense news shock are obtained via lag-augmented local projections. Bands represent the 68% and 90% heteroskedasticity robust standard errors. Defense news shocks are obtained from the updated series in Ramey and Zubairy (2018). Sample goes from 1947Q1 to 2015Q4. Values in the Figures are normalized by the peak response of G.

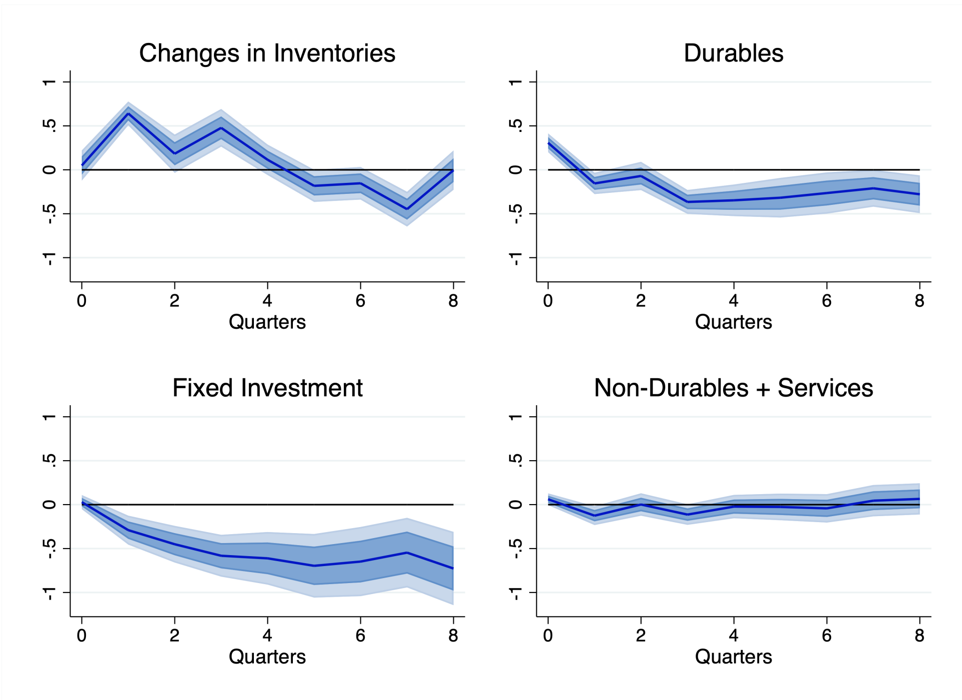

Figure 2.6: Response of Consumption and Investment to a Defense News Shock

Notes: See notes of Figure 2.5

2.8 Robustness - Section I of the Paper

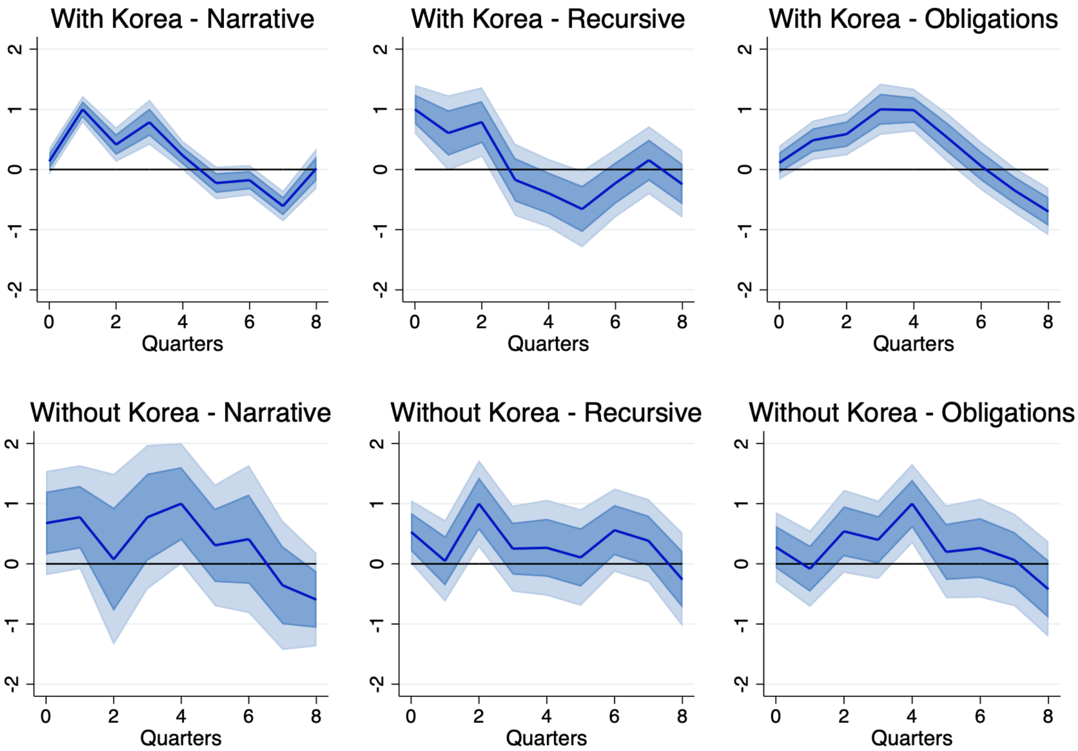

Figure 2.7: Response of Inventories - Robustness

Notes: Response of inventories to different fiscal shocks over two samples (with and without Korean war). All the rest is identical to notes of Figure 2.5.

In Figure 2.7, we verify that the positive response of inventories is robust to the inclusion of the Korean War in the sample period. In particular, we estimate IRFs of inventories via lag-augmented local projections with respect to three different fiscal shocks (narratively identified, recursively identified, and shocks to obligations) over two samples. The first sample includes the Korean war and goes from 1947Q1 to 2015Q4 (top row of Figure 2.7). The second sample runs from 1954Q1 to 2015Q4 and excludes the Korean war (bottom row of Figure 2.7).

For all results, we control for a linear time trend and four lags of government spending, consumption, investment, net-export, hours in the private sectors and 3-months

T-Bill rate. To implement the narrative method, we include defense news shocks and its four lags and estimate the IRF using the OLS coefficients associated with defense news shocks (first column of Figure 2.7). To implement the recursive method, we add contemporaneous government spending and obtain the IRF from its OLS coefficient (second column of Figure 2.7). Finally, we consider shocks to defense procurement obligations. We control for four lags of obligations using the series discussed in the main text of the paper, and estimate the IRF from the OLS coefficient on contemporaneous defense procurement obligations (third column of Figure 2.7).

Although excluding the Korean War from the sample leads to less precise estimates of the IRF, our results are still significant especially at early horizons. The difference in precision is not a surprising result since the Korean War represents the largest military build-up after WWII. As discussed in the paper, we support the idea of including the Korean war in the sample since wars represent natural experiments where G increases exogenously.

2.9 Details on Industry Level Analysis

In this section, we implement robustness checks for the industry-level analysis of inventories (see Appendix 2.9.1) and provide details on our construction of industry weights θ_i (see Appendix 2.9.2).

2.9.1 Robustness - Section II in the Paper

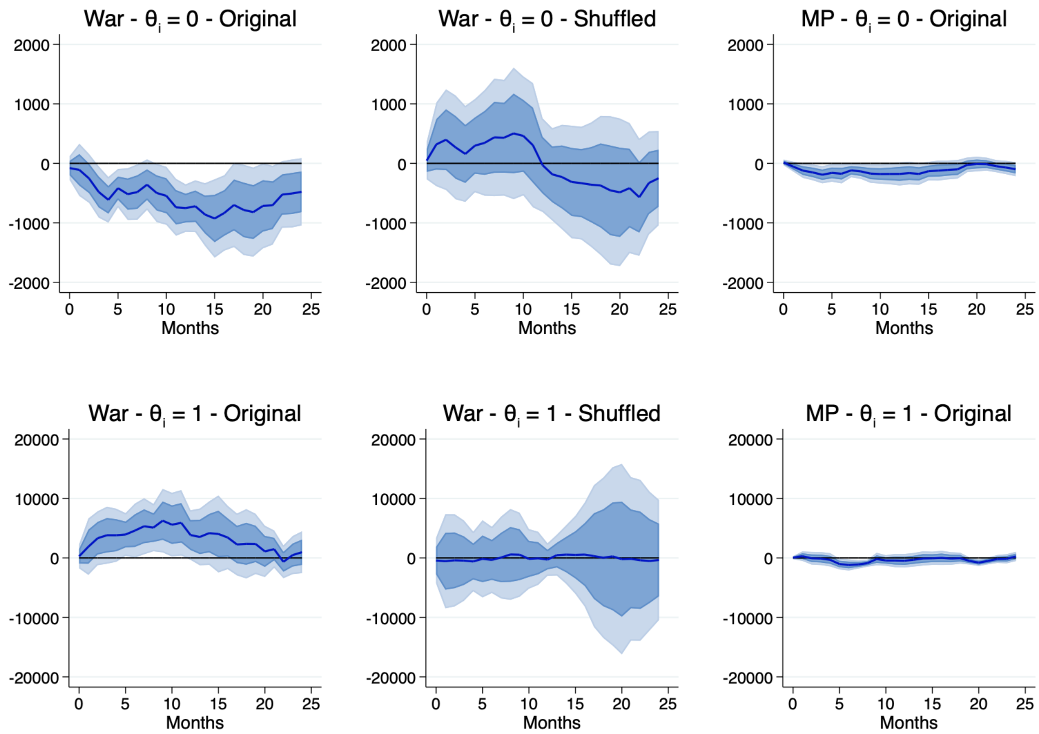

Figure 2.8: Response of Sectoral Inventories to War Events (Robustness).

Notes: Same as in Figure 2 of the paper.

Figure 2.8 shows the results of the robustness checks associated with Section II of the main text. The first column replicates the results reported in the paper, where our $Shock_t$ variable is war dates and industry weights (θ_i) are baseline weights constructed directly from the BEAs Make and Use tables. We report IRFs conditional on setting $\theta_i = 0$ (top panel) and $\theta_i = 1$ (bottom panel). Recall that setting $\theta_i = 0$ indicates the effect of a shock on a sector not connected to the government while setting $\theta_i = 1$ indicates the effect of such a shock on a sector which is fully connected to the government.

Additionally, in the middle panels we report the same results using shuffled weights. In this case, we randomly assign a weight θ_j to an industry i to verify that the result is not driven by the aggregate distribution of weights. Lastly, the right panels report the results when the weights are fixed at their empirical value, but where the shock is a monetary policy shock rather than a war date. The goal of this robustness check is to verify that the result is not driven by industry-level exposure to the business cycle. Notice that the inventory response of industries connected to the government $\theta_i = 1$ (bottom panels) vanishes for both robustness checks.

The Response of the Aircraft Industry Here we estimate the following lag-augmented local projection:

$$\bar{h}_{t+h}^{aircraft} = \beta_h \cdot Shock_t + lags_t + \varepsilon_{t+h}$$

where $\bar{h}_{t+h}^{aircraft}$ is average hours of production workers in the aircraft industry in quarter t+h, $Shock_t$ is either defense news shocks or defense procurement obligations, $lags_t$ is four lags of the dependent variable and four lags of the shock. We believe average hours of production workers in the aircraft industry is an excellent proxy for defense production (see Appendix 2.9.3). We report IRFs in Figure 2.9. We observe that defense production quickly ramps up in response to defense news or newly awarded procurement contracts.

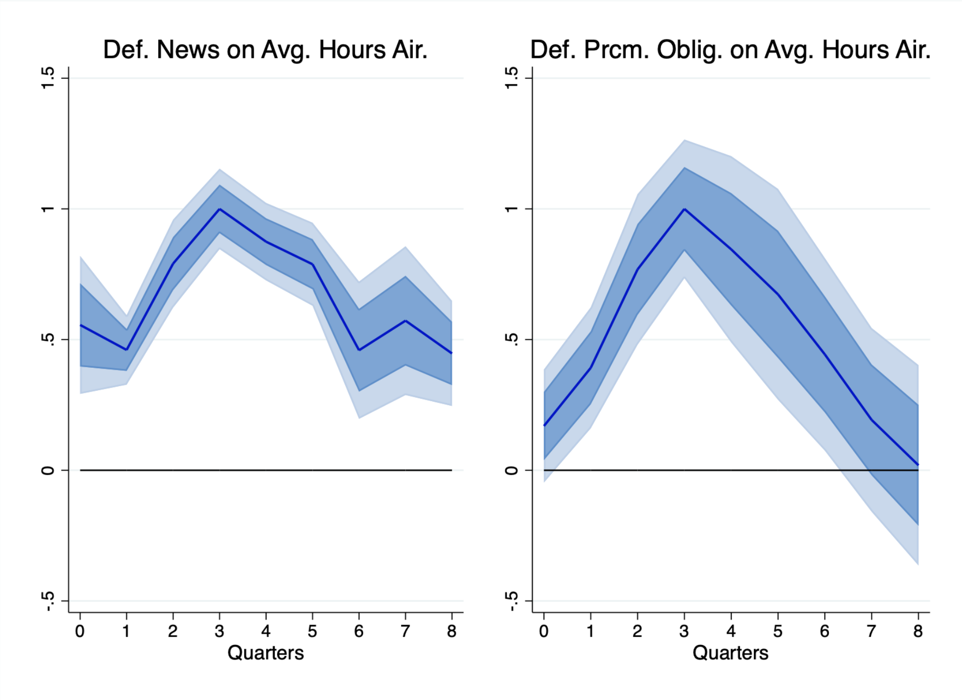

Figure 2.9: Effects of Military Build-ups on Defense Production

Notes: IRFs are obtained via lag-augmented local projections. Sample goes from 1947Q1 to 2002Q4 (sample stops in 2002 because data are no longer available). Data Source: BLS Discontinued Databases. Standard errors are heteroskedasticity robust. Confidence bands are 90% and 68%.

2.9.2 Construction of Industry Weights

To construct industry weights, we combine information from the Make and Use table with more than 60 non-government sectors between 1963 to 1996. Following Horowitz and Planting (2009), we derive direct requirement industry-by-industry matrices A_t and direct sales from the private sectors to the government. We use these two elements to construct our final industry weights as follows.

Government Direct Purchases. We construct a vector of government purchases (i.e., direct requirements) relative to industry output:

$$\underset{n\times 1}{\boldsymbol{\gamma}_{0,t}} = \begin{bmatrix} \dfrac{\text{SALES}_{1\to G,t}}{\text{SALES}_{1,t}} \\ \vdots \\ \dfrac{\text{SALES}_{n\to G,t}}{\text{SALES}_{n,t}}, \end{bmatrix}$$

where t denotes the year, n is the number of manufacturing sub-industries, G denotes the federal general government, and the 0 subscript in a vector's name refers to the order of included input-output connections (e.g., a zero subscript suggests that the vector only accounts for direct sales to the government). Moreover, $\text{SALES}_{i\to G,t}$ for a given sector i includes government gross investments, which show up as final uses in the Use tables. We report the time-average values of $\boldsymbol{\gamma}_{0,t}$ in the third column of Table 2.2.

Government Indirect Purchases Following Nekarda and Ramey (2011), we also include downstream input-output linkages to account for indirect sales to the government. In order to do so, we construct yearly $n \times n$ input-output matrices A_t in which (i,j)th element of matrix A_t is given by:

$$\frac{\text{SALES}_{i\to j,t}}{\text{SALES}_{i,t}}.$$

We then construct a vector of direct and first-order indirect sales shares as follows:

$$\boldsymbol{\gamma}_{1,t} = (I_n + A_t) \cdot \boldsymbol{\gamma}_{0,t}.$$

Notice that the ith element of $\boldsymbol{\gamma}_{1,t}$ is given by:

$$\gamma_{1,i,t} = \underbrace{\frac{\text{SALES}_{i\to G,t}}{\text{SALES}_{i,t}}}_{\text{Direct Sales}} + \underbrace{\sum_{j=1}^{n} \frac{\text{SALES}_{i\to j,t}}{\text{SALES}_{i,t}} \cdot \frac{\text{SALES}_{j\to G,t}}{\text{SALES}_{j,t}}}_{\text{Indirect Sales.}}$$

We report the time-average of $\boldsymbol{\gamma}_{1,t}$ in the fourth column of Table 2.2. Similarly, we construct direct, first and second order indirect sales to the government, shares of total output as:

$$\boldsymbol{\gamma}_{2,t} = \left(I_n + A_t + A_t^2\right) \cdot \boldsymbol{\gamma}_{0,t}.$$

We report the time-average values of $\boldsymbol{\gamma}_{2,t}$ in the fifth column of Table 2.2. Finally, we construct our industry weights θ_i as:

$$\theta_i := \frac{\mathbb{E}\left[\gamma_{2,i,t}\right]}{\max_i \mathbb{E}\left[\gamma_{2,i,t}\right]}$$

We report the weights in the last column of Table 2.2.

Table 2.2: Industry Weights

Sector	*Commodity Description*:	$\gamma_{0,i}$	$\gamma_{1,i}$	$\gamma_{2,i}$	θ_i
3364	Other transportation equipment	34.43%	42.00%	43.94%	1.00
334	Computer and electronic products	13.09%	17.04%	18.38%	0.42
323	Printing and related support activities	7.98%	9.35%	9.95%	0.23
332	Fabricated metal products	3.73%	4.78%	5.37%	0.12
3361	Motor vehicles, bodies and trailers, and parts	2.09%	3.70%	4.64%	0.11
339	Miscellaneous manufacturing	2.31%	3.80%	4.49%	0.10
333	Machinery	2.65%	3.84%	4.44%	0.10
335	Electrical equipment, appliances, and components	2.37%	3.66%	4.31%	0.10
325	Chemical products	1.91%	3.50%	4.27%	0.10
324	Petroleum and coal products	2.71%	3.50%	4.17%	0.09
326	Plastics and rubber products	1.13%	2.20%	2.89%	0.07
337	Furniture and related products	0.66%	1.63%	2.19%	0.05
331	Primary metals	0.54%	1.44%	2.06%	0.05
313	Textile mills and textile product mills	0.48%	1.31%	2.01%	0.05
315	Apparel and leather and allied products	0.57%	1.37%	1.98%	0.05
327	Nonmetallic mineral products	0.49%	1.35%	1.91%	0.04
322	Paper products	0.51%	1.25%	1.83%	0.04
311	Food and beverage and tobacco products	0.38%	1.16%	1.77%	0.04
321	Wood products	0.19%	0.91%	1.53%	0.03

Notes: Last column divides $\theta_{2,i}$ by the maximum value (i.e. the one of Other Transportation Equipment Manufacturing). In the paper, the weights θ_i that we use are the ones which include second order degree of connections, normalized (last column).

2.9.3 Tracking Defense Industrial Production

In Section II of the paper, we use use *Average Hours of Production Workers* of the *Aircraft industry* to keep track of the "*defense production machine*". We now explain the reasons behind that choice. Firstly, we plot the quarterly time series in Figure 2.10.

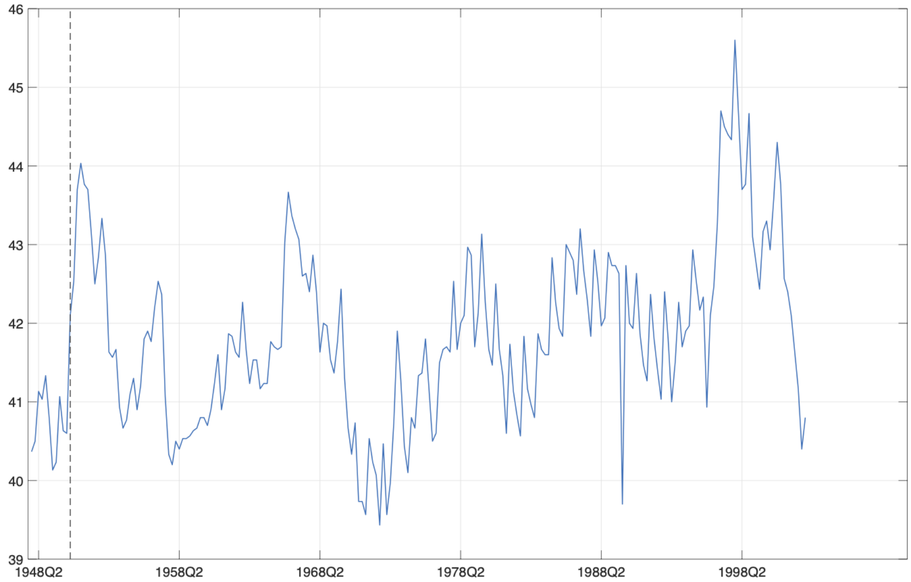

Figure 2.10: Average Hours of Production Workers in the Aircraft Industry

Aircraft Industry We choose the Aircraft industry for two reasons: (i) great data availability (monthly data from BLS discontinued series starting from 1939) and (ii) high dependency on government purchases (see Nekarda and Ramey (2011)).

Hours-per-Worker In general, there are no direct measures of industrial output. In the case of the aircraft industry, we do not observe the exact number of aircraft produced every month nor their percentage of completion. However, we have three variables which can proxy for industrial production: (i) average weekly hours of production workers, (ii) number of production workers (i.e., employment) and (iii) their product, namely total hours worked. The first one is a measure of intensity of production, while the other two are stock variables measuring the extensive margin of production.

In order to understand which one is more suitable to measure changes in production, we consider as an illustrative example the outbreak of the Korean War. During this period, defense manufacturers foresee a period of high demand of weapons by the government and adjust production accordingly. The first sensible thing is to increase production, given the predetermined level of capital and labor inputs. For instance, increasing production requires extra use of electricity to operate machinery in the assembly lines as well as a higher number of shifts with longer duration for production workers. By consequence, hours per worker increase immediately. Over time, contractors expand production by widening their stock of capital and workers, thus overcoming problems related to capital immobility (see e.g., Ramey and Shapiro (1998)) and labor market frictions. As contractors expand their production facilities and hire new production workers, intensity of production returns back to normal.

This example highlights two facts. Firstly, intensity of production of manufacturing industries is a good indicator of switches in the production regime. Secondly, intensity of production leads employment and other stock variables which tend to move more slowly. This intuition is consistent with Bils and Cho (1994), who find that hours per worker lead employment and the business cycle. Moreover, they emphasize how hours-per-worker co-moves with other relevant but unobserved measures of intensity of production.[30] Along these lines, Fernald (2012b) suggest to use hours of production workers to proxy other unobserved measures of intensity of production. According to them, a cost-minimizing firm operates on all margins simultaneously, both observed (i.e., hours per worker) and unobserved (i.e., labor effort and workweek of capital).

In what follows, we show that (i) hours-per-worker in the aircraft industry leads employment and (ii) employment drives the dynamics of total hours, overshadowing very informative lumpy changes in hours-per-worker. In light of all this, *hours-per-worker is the most suitable variable to timely measure changes in production*.

[30] They find that (i) "looms hours" are strongly related to hours per worker in the US textile industry and (ii) electricity use of manufacturing industries and hours worked per week co-moves. On the contrary, they find that the relationship between their measures of capital utilization and the number of production workers is much weaker.

Hour per Worker, Employment and Total Hours in the Aircraft Industry Figure 2.11 shows in its top-left panel the lead-lag correlation map between changes in average hours of production workers and changes in the number of production workers in the Aircraft industry. Clockwise from the top-right panel we show the time series of average hours of production workers ($\bar{h}_t$), number of production workers (e_t) and total hours of production workers ($\bar{h}_t \cdot e_t$) around the onset of the Korean war for the Aircraft industry (i.e., 1950Q3), respectively.

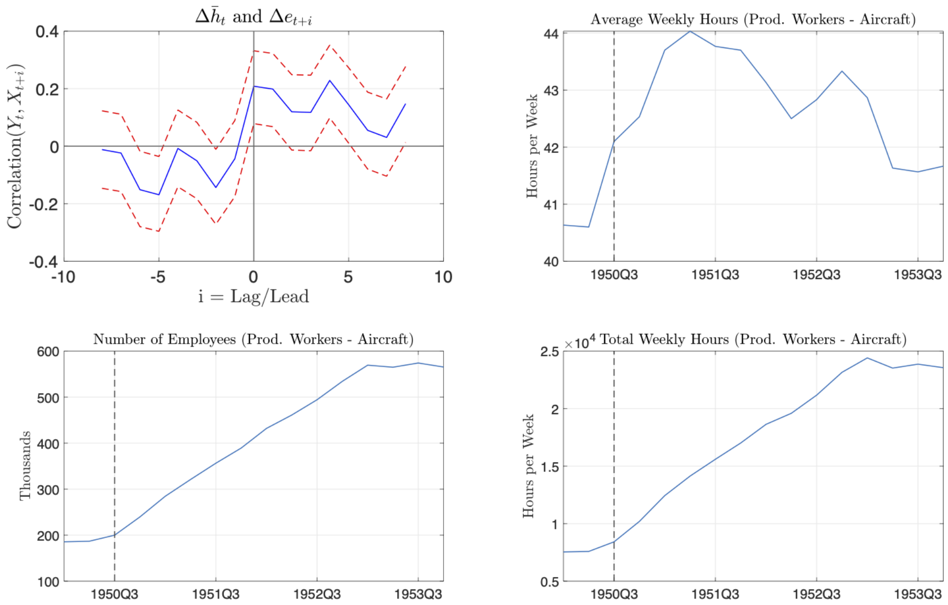

Figure 2.11: Average Hours of Production Workers Vs Production Workers - Aircraft Industry

Firstly, from the lead-lag correlation map, we observe that average hours of production workers lead employment. This is consistent with the findings of Bils and Cho (1994). Secondly, notice that the dynamics of total hours is dominated by employment, and not by average hours per worker. Therefore, if we gauge industrial production by simply looking at the dynamics of total hours, we would conclude that the response of the Aircraft industry at the outbreak of the Korean war was mild and slow. On the contrary, average hours per production worker anticipated the peak response of employment and total hours of production, signaling that defense production had already fired up at the onset of the war.

We further clarify what is happening by breaking down the change in total hours into two components, one which accrues to changes in hours worked (intensive margin) and one which accrues to changes in number of workers (extensive margin):

$$H_t = \bar{h}_t \cdot e_t$$

where z_t is a defense news shock. We break down the dynamic response of total hours to the Korean War using the previous expression:

$$(H_{1950Q3+h} - H_{1950Q2}) = \underbrace{\left(\bar{h}_{1950Q3+h} - \bar{h}_{1950Q2}\right) \cdot e_{1050Q3+h}}_{\text{Intensive Margin}} + \underbrace{\left(e_{1950Q3+h} - e_{1950Q2}\right) \cdot \bar{h}_{1050Q3+h}}_{\text{Extensive Margin}}$$

with $h = 0, 1, ..., H$. We show this breakdown in Table 2.3:

Table 2.3: Breakdown Total Hours - Korean War

Date	$\bar{h}_t$	e_t	H_t	$H_{1950Q3+h} - H_{1950Q2}$	Int. Margin	Ext. Margin	Int. Margin (%)	Ext. Margin (%)
1950Q2	40.60	186.83	7585.43	0.00	0.0	0.0	-	-
1950Q3	42.10	200.00	8420.00	834.57	300.0	554.3	35.9%	66.4%
1950Q4	42.53	239.70	10195.24	2609.81	463.4	2248.6	17.8%	86.2%
1951Q1	43.70	284.57	12435.56	4850.13	882.2	4270.9	18.2%	88.1%
1951Q2	44.03	321.00	14134.70	6549.27	1102.1	5907.8	16.8%	90.2%
1951Q3	43.77	356.37	15596.98	8011.55	1128.5	7419.9	14.1%	92.6%
1951Q4	43.70	389.27	17010.95	9425.52	1206.7	8846.3	12.8%	93.9%
1952Q1	43.13	432.00	18633.60	11048.17	1094.4	10574.9	9.9%	95.7%
1952Q2	42.50	461.07	19595.33	12009.90	876.0	11654.9	7.3%	97.0%
1952Q3	42.83	494.30	21172.52	13587.08	1103.9	13169.8	8.1%	96.9%
1952Q4	43.33	534.37	23155.89	15570.46	1460.6	15059.8	9.4%	96.7%
1953Q1	42.87	569.43	24409.71	16824.28	1290.7	16400.8	7.7%	97.5%

Notice that the dynamic of Total hours, H_t is dominated by the extensive margin. Therefore, using total hours would overshadow the early change in hours-per-worker, which is a clear signal that contractors were already responding to the shock in the third quarter of 1950.

Delay in the FED's Defense Industrial Production Index Notice that the Board of Governors of Federal Reserve System constructs a monthly real index of industrial production of manufacturing equipment in defense industries.[31]

The Fed makes clear that such defense production index is *mainly obtained from BLS data on production-hours (i.e., total hours). Hours are then used to infer output.* However, we have just seen that the dynamics of total hours worked are delayed relative to average hours worked. In fact, we now show that hours-per-worker in the Aircraft industry leads defense production as measured by the Fed.

In particular, we study the lead-lag correlation map between each labor margin and defense procurement obligations, production, and spending. Figure 2.12 plots the results.

[31]Data is available from 1947 to present at monthly and quarterly frequency, both seasonally adjusted and not. It can be downloaded at this link. Detailed information on the Real Index of Industrial Production of Manufacturing Equipment in Defense sector is available at this link. In particular, the underlying industries used for the construction of the series are discussed in these two tables: (i) market structure (Equipment); (ii) Industry Group (defense and space).

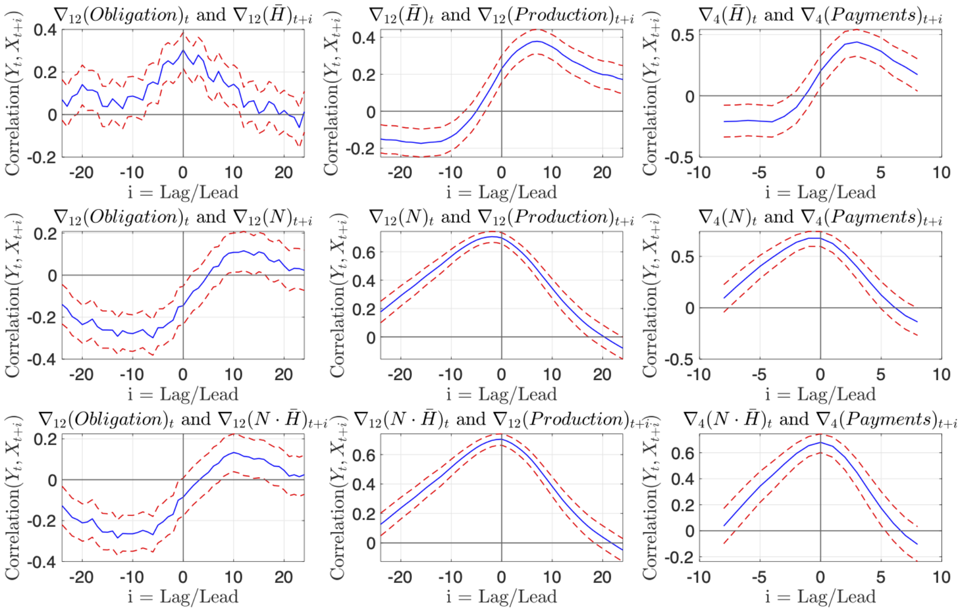

Figure 2.12: Lead-Lag Correlation Graph - Defense Industrial Production

Notes: Defense procurement spending is constructed as discussed in the paper and therefore tracks payments to contractors (sample from 1947Q1). Defense procurement obligations come from the original series from Business Condition Digest, discussed in Appendix 2.10.2 and track new contract awards (monthly data from 1951M1 to 1988M11). Defense Production is the monthly seasonally adjusted index constructed by the Fed (data available from 1947M1). Hours and employment data come from the BLS discontinued data series on production workers data (available from 1939M1 to 2003M12).

Firstly, looking at the first row, average hours of production workers in the Aircraft industry (intensive margin) appear to: (i) co-move with obligations, (ii) lead industrial output by 8 months (2 quarters), and (iii) lead payments by 4 quarters.

From the second row, we notice that the number of production workers (extensive margin) appear to: (i) lag behind obligations (the delay is about 3 quarters), (ii) co-move with the production index and (iii) co-move with payments.

Finally, the third row shows that total hours of production workers co-move with industrial production as measured by the Fed. This confirms the fact that the Federal Reserve adopts total hours to construct the defense production variable. Moreover, the maps of total hours and employment are basically identical, confirming our previous finding that the dynamics of employment drive movements in total hours.

To summarize, we show that the Fed measures defense production using total hours of production workers. However, the dynamics of total hours is dominated by employment, which is a delayed measure of production and overlooks the ability of producers to ramp-up production by using more intensively their input of production (i.e., capital utilization and average hours worked). Specifically, the Fed's measure lags behind defense procurement obligations but co-moves with spending. This confirms that the Fed's production index is subject to the same delays which characterize employment. In light of this, we believe that using average hours of production workers in the Aircraft industry is best suited for capturing real-time changes in defense production.

2.10 Details on Defense Procurement in the Data

In this section, we outline the details about measurement of defense procurement spending.

Section 2.10.1 clarifies the accounting origin in the NIPA of outlays which refer to the purchase of goods. Section 2.10.2 shows how we calculate the 2 to 3 quarters delay between defense procurement obligations and spending. Section 2.10.3 uses contract level data from the 2000 to rationalize the existence of a time delay and address the issue of partial delivery payments. Section 2.10.4 illustrates how we construct the quarterly time series of defense procurement obligations. Section 2.10.5 uses data from the 2000 on contracts' opportunities (i.e., contract level *solicitations*) to show that it is unlikely for contracts awards to be anticipated by more than one quarter.

2.10.1 Accounting Origin of Procurement in the NIPA

In this section, we provide further details on the accounting origin of public procurement contracts in the NIPA tables. Figure 2.13 summarizes the accounting of G, according to Chapter 9 of Bea (2017), which explains how the NIPA record all the entries of G. It highlights in red the two entries which contain public procurement spending: (i) Intermediate Goods and Services and (ii) Investment in Fixed Assets.

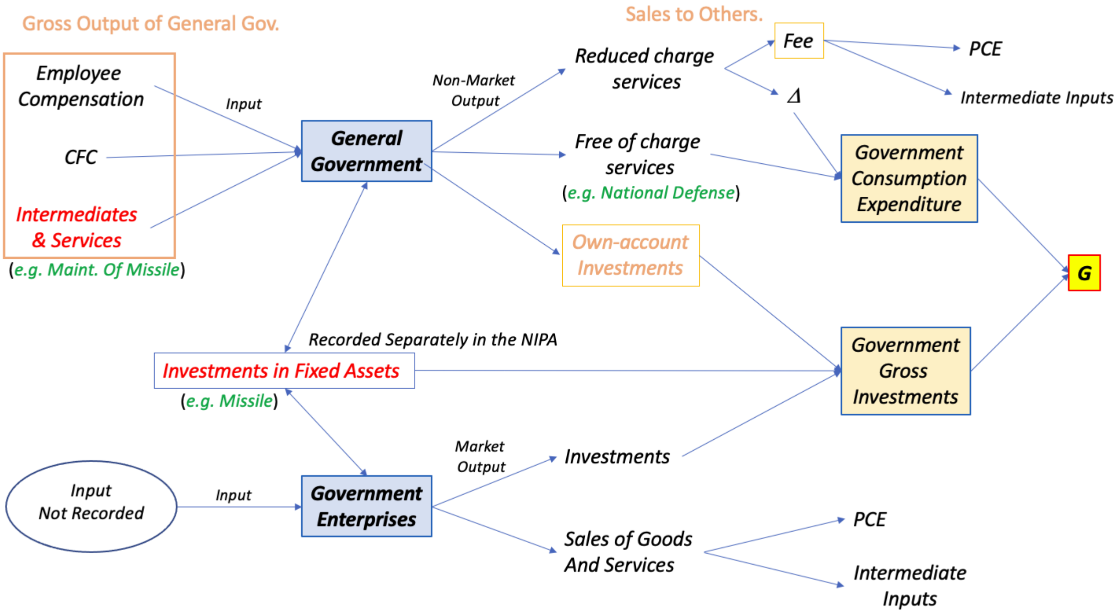

Figure 2.13: Accounting of G - Summary

Notes: CFC means "Cost of Fixed Capital" and measures depreciation of government assets. PCE means Personal Consumption Expenditure, the NIPA measure of Consumption which absorbs reduced charge services from the government (e.g. tuition fees from public universities). Own Account Investment is own resources reinvested in the public capital stock.

First of all, notice that G is made of two components, consumption and investments:

$$G = \underbrace{\text{Government Consumption Expenditure}}_{G^C} + \underbrace{\text{Government Gross Investments}}_{G^I}$$

Government Consumption Expenditure Government consumption originates from the gross output of the government after deducting (i) Sales to Other Sectors and (ii) Own-Account Investments:

$$G^C = \underbrace{\text{Compensations} + \text{CFC} + \text{Intermediates and Services Purchased}}_{\text{Gross Output of General Government}} - \ldots$$
$$- \text{Own Account Investments} - \text{Sales to Others}$$

When a general government entity (e.g., DoD) decides to purchase goods and/or services, they are mainly accounted as Intermediates, which eventually end up in G as government consumption.

Government Gross Investments The government also makes three types of investments. Firstly, the General Government makes Own-Account-Investments, which are deducted from the gross output of general government, in order to account them as investments. Secondly, both the General Government and Government Enterprises make investments in fixed assets. Investment in Fixed Assets contain other purchases from the private sector.

Example 2.10.1 (Purchasing a Missile). To clarify this point, consider the case of the government purchasing a new set of guided missiles:

1. The missile is accounted as Equipment in the Federal Defense category of Change in Government Fixed Assets and therefore contributes to G as part of Government Gross Investments.

2. Installation, Maintenance, Quality Control and other services related to the Missile are accounted as Intermediate Goods and Services Purchased (input of production).

3. The missiles and the related services are used to produce a non-market output of production, namely, national defense.

The production of the missile shows up in business inventories as long as the contractor supplying the missile delivers it to the government. Once delivered, inventories decrease and government investment increase. Notice that the reduction in inventories and the corresponding increase in G is a zero-sum game which does not increase GDP (recall that GDP in the US is constructed as the sum of final demand). GDP increases while production takes place and is recorded as inventories. In absence of time-to-build, inventories do not increase and the purchase of the item by the government directly shows up in G. For instance, this is the case of the Installation, Maintenance, Quality Control and other services related to the missile purchased by the Government.

Figure 2.14 provides an example of official accounting table of G, namely NIPA Table 3.10.5A, taken from Bea (2017).

Finally, to clarify the timing, we provide a visual representation of the process in Figure 2.15.

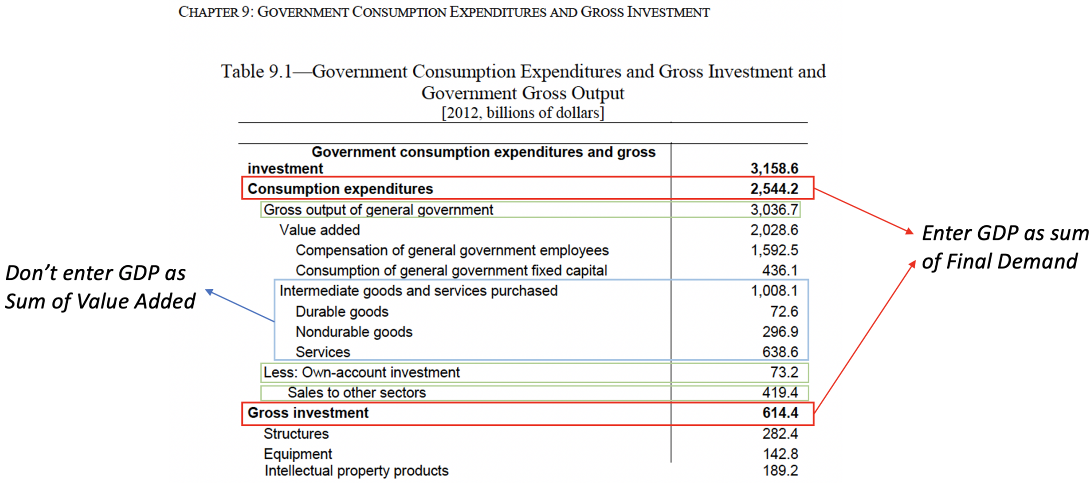

CHAPTER 9: GOVERNMENT CONSUMPTION EXPENDITURES AND GROSS INVESTMENT

Table 9.1—Government Consumption Expenditures and Gross Investment and Government Gross Output
[2012, billions of dollars]

Government consumption expenditures and gross investment	**3,158.6**
Consumption expenditures	**2,544.2**
Gross output of general government	3,036.7
Value added	2,028.6
Compensation of general government employees	1,592.5
Consumption of general government fixed capital	436.1
Intermediate goods and services purchased	1,008.1
Durable goods	72.6
Nondurable goods	296.9
Services	638.6
Less: Own-account investment	73.2
Sales to other sectors	419.4
Gross investment	**614.4**
Structures	282.4
Equipment	142.8
Intellectual property products	189.2

Figure 2.14: NIPA Table 3.10.5A - Example

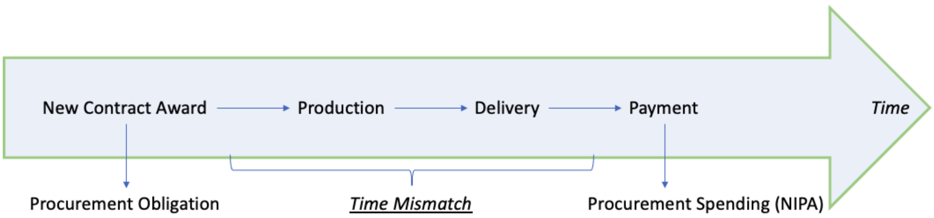

Figure 2.15: Timeline of Procurement Contracts

Notes: The procurement timeline follows information from the Federal Acquisition Regulation (FAR) and the BEA's Concepts and Methods of NIPA.

2.10.2 Time Mismatch Between Obligations and Payments

In Section 2.10.2, we show how we construct a proxy for defense procurement spending using data from the NIPA. We now show how we construct the defense procurement obligation proxy. Recall that obligations arise when the DoD awards new contracts while spending reflect government outlays, that is, payment to contractors. We observe obligations through two data sources, discussed below.

Business Condition Digest The periodical *Business Conditions Digest*, available on Fraser at this link, provided Business Cycle Indicators, among which a list of Defense Indicators. The original source of the data was the Department of Defense, Office of the

Assistant Secretary of Defense (seasonal adjustment implemented by BEA). In particular, we use Series 525, "*Defense Prime Contracts Awards*". This series was firstly collected by Valerie Ramey for her papers: Ramey (1989) and Ramey (1991). We are grateful to her for providing the data. The periodical was issued monthly from October 1961 until March 1990. However data is available from January 1951 to November 1988.

Business Condition Digest was discontinued in March 1990, and data on prime contracts is no longer recorded starting December 1988 (all year 1989 is missing). Most business indicators on *Business Condition Digest* were moved to another monthly periodical, namely the *Survey of Current Business*. Prime award contracts (series 525) was preserved and moved to Appendix C on *Business Cycle Indicators* (section 2.4: other important economic measures/government activities). Data is available in the form of scanned versions of the *Survey of Current Business* at this link. For some reason, data starting from 1991 does report values of prime contract awards for months in the fourth quarter (Q4) of every year. We believe this is a systematic omission, which results in less reliable data for this time period. Finally, due to reorganization of resources at the BEA, the *Business Cycle Indicators* section was discontinued, and prime award contracts were no longer disclosed to public, following the joint November-December 1995 issue. Therefore data is not available after this date.

To summarize, we obtain reliable monthly data on prime contract awards from January 1951 to November 1988. We plot defense prime contract awards versus defense procurement spending in the top-left panel of Figure 5 of the paper. Notice that in order to match the quarterly frequency of procurement spending, obligation data is aggregated by quarters. Moreover, since NIPA data are annualized (their quarterly averages return their yearly values), we do the same for obligation data to allow for a closer comparison between the two series.

We observe from the graph that obligations lead spending. Notice that obligation data tends to be more noisy than spending data. The main reason for this is that large contracts are often awarded and then terminated a few months later for convenience, or due to litigation with a losing offeror (this is also highlighted in Auerbach et al. (2020)). Moreover, obligations are more lumpy than payments which get smoothed over the duration of a contract. In order to account for this, we use a simple MA smoother (red line in the graph). We then provide a quantitative assessment of the delay by looking at the lead-lag correlation map between the growth rates of smoothed obligations and the growth rates of spending (see top-right panel).[32]

Overall, we find a positive correlation between the two series which increases when obligations are delayed (top-right quadrant of the lead-lag correlation map). In particular, correlation spikes when obligations are delayed by 2, 5 and 8 quarters. Results are robust to a different approach which looks at the lead-lag correlation between year-to-year quarterly changes ($\nabla_4\, x_t = (1-L^4)x_t = x_t - x_{t-4}$) of original -i.e., not smoothed - obligations and spending. In this case, the spikes happen at 2 and 5 quarters.

Federal Procurement Data System Next Generation On September 26th 2006, the Federal Funding and Accountability Act is passed by congress as a first step towards a more transparent procurement system, which allows full disclosure of information involving federal contracts. The transparency effort by FFATA culminates in 2019 with the opening of a public website, USASpending.gov, which discloses information on all federal procurement contracts.[33] Data from USASPending.gov is pulled from FPDS-NG, the Federal Procurement Data System Next Generation, which actually includes the whole universe of procurement contracts. FPDS is the system used by government contracting officers to officially input data on awarded contracts in the government-wide system. Data from FPDS can be downloaded from USASpending.gov. The data spans 2000Q4 to the present with a caveat: contract data awarded before the beginning of the construction of the database could have been lost or not recorded. We collect data on all defense procurement contracts on FPDS between 2000Q4-2020Q3.

We again compare obligations and spending in Figure 2.16. The top-right panel

[32] We look at growth rates ($\Delta_1\, x_t = (1-L)^1 x_t = x_t - x_{t-1}$) to cope with the non-stationarity of the series.
[33] More information on the history of USASpending.gov can be found here.

plots again the lead-lag correlation between the growth rates of (smoothed) obligations and the growth rates of spending. The highest correlation is recorded when obligations are delayed by 1 quarter. Once again, the results are robust to looking at the lead-lag correlation of year-to-year quarterly changes between original obligations and spending. In this case, the peak occurs from 0 to 2 quarters.

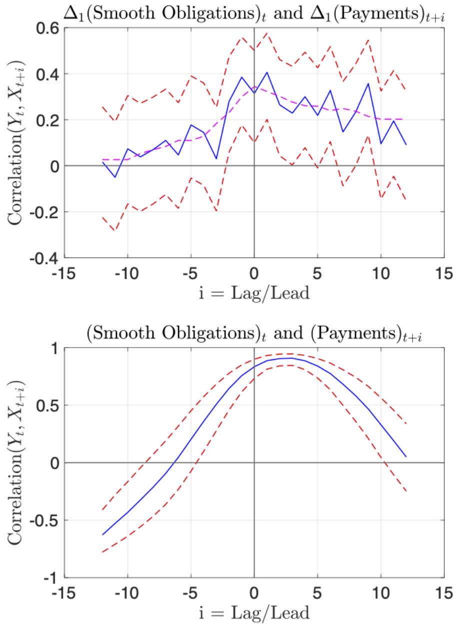

Figure 2.16: Accounting Mismatch - January 2000 onward

Notes: Figure 2.16's notes: the FPDS measure of obligation (blue line) is constructed by: (i) summing the daily data to obtain quarterly data; (ii) convert to annual units using NIPA's procedure; (iii) seasonally adjust data data using Brockwell and Davis (1991)'s method (the Matlab code can be found here.

Before the signing of the FFATA in 2006, obligations seem under-reported relative to spending, thus inducing a downward bias in estimates of the accounting time-mismatch. We have two possible explanations for this counter-intuitive result, either this is a consequence of missing contract modifications awarded before the introduction of FPDS, or those modifications could have been classified before FFATA made most of contracts available to the public. In fact, FPDS-obligation data catches up and starts leading spending after the signing of the FFATA. Moreover, we show in Figure 2.17 that the share of large contracts (top 1%, 5% and 10%) out of all procurement spending stabilizes after 2006, indicating that large classified contracts are not showing up in FPDS.

We take this into account and we repeat the analysis only on those quarters following the signing of the FFATA (bottom panel of Figure 2.16). We observe a single clear spike in the lead-lag correlation, which indicates that obligations are delayed by 3 quarters relative to payments.[34]

Summary of Time Mismatch We summarize the time delay between obligations and spending in Table 2.4

Table 2.4: Summary of Time Mismatch Between Spending and Obligations

Period	*Data Source*	*Correlation Spike Delay (Quarters)* Δ_1	∇_4
1951M1 to 1988M11	BCD	2-5	2-5
2000M10 to 2020M9	FPDS	1	2
2006M1 to 2020M9	FPDS	3	-

These results suggest that the accounting delay between beginning of production (award date) and the first payment (outlay) is on average between 2 to 3 quarters. Notice also that the time delay seems to shorten over time, when we use FPDS data.

Overall, our results are consistent with anecdotal evidence that government payments happen once every 180 days.[35]

2.10.3 Rationalizing the Time Mismatch

In this section, we rationalize the existence of an aggregate time-mismatch between defense procurement obligations and spending. In particular, we provide both theoretical and empirical micro-level evidence of the time mismatch.

Duration of Defense Procurement Contracts Firstly, a necessary condition for the existence of an accounting mismatch is the long duration of contracts. If contracts were

[34]The peculiar non-trending sinusoidal-wave shape of the data referring this period allows us to directly look at the correlation between the two series in levels. The super-positioning of waves which happen when we shift one series back and forth in time, allows to observe a single clear spike which refers to the exact period when the two series overlap. The correlation spikes when obligations are delayed by 3 quarters.

[35]We confirm this timeline in discussions with several federal procurement contractors.

less than 90 days in duration, then payments would be processed in the same quarter as the award date.

We use FPDS data pulled from USASpending.gov from 2000Q4 to 2020Q3 to construct the distribution of duration of defense government contracts. In this context, contracts have two main types: (i) single transaction and (ii) multiple transaction.[36] We calculate the duration of a single transaction contract from the award date to the end of work. The award date almost always indicates the start of work associated with a contract. To calculate the duration of multiple transaction contracts, we take the oldest contract modification end date and subtract from it the "new-action" award date.[37] Table 2.5 shows contract durations without distinguishing between single and multiple transaction contracts.

Table 2.5: (Log)Duration of Defense Contracts

Stats		*Unweighted*		*Weighted (by Obligation)*	
		Duration (days)	*Log-Duration*	*Duration (days)*	*Log-Duration*
Percentiles	1%	0	0	0	0
	5%	0	0	46	3.85
	10%	0	0	193	5.27
	25%	3	1.39	514	6.24
	50%	20	3.04	1519	7.33
	75%	126	4.84	2962	7.99
	90%	377	5.93	4844	8.49
	95%	794	6.68	5464	8.61
	99%	2584	7.86	6887	8.84
Mean		173.03	3.09	1988.02	6.94
Std.		485.32	2.14	1746.81	1.57
Min.		0.00	0	0.00	0
Max.		7300.00	8.89	7300.00	8.89

Table 2.5: defense contracts are identified by reporting DoD as funding/awarding agency. Data is taken from FPDS, all defense contracts from 2000Q1 to 2020Q1. Sample is restricted to those contracts for which the entire history of transactions (from the first new contract to the last modification) are available. Number of contracts in the sample is about 17 millions. Almost 5 thousands contracts with duration more than 20 years (7,300 days) are eliminated from the sample.

[36]Transactions which refer to the same contracts are pooled together through a unique contract identifier field in FPDS.

[37]This is possible because FPDS reports both the beginning and the end of the PoP (Period of Performance).

The median contract duration is 20 days and 90% of contracts have duration less than one year. These results are in line with the findings of Cox et al. (2021) and suggest that contracts have a short duration.[38] However, this measure does not take into account the size of contracts, as larger contracts might have longer duration. The right columns of Table 2.5 report the distribution of the contracts' (log)duration, weighted by the total obligation amount. The weighted distribution can be interpreted as the duration distribution of a $1 of spending in defense procurement. The following remark characterizes the mean and median of this distribution.

Remark [Median/Mean Duration of $1] The median duration of $1 of defense procurement spending is 4.16 years. The mean duration of $1 of defense procurement spending is 5.44 years.

Notice that after weighting, the shape of the distribution drastically changes. This suggests that procurement spending is characterized by a small number of large and long-duration contracts. We confirm this in Figure 2.17, which plots the share of total spending of the largest 1%, 5% and 10% of contracts. We find that the largest 10% of contracts account for 95% of total spending, on average. Similarly, the top 1% of contracts accounts for roughly 80% of total spending on average.

To summarize the results of this section: (i) large contracts make the bulk of defense spending and (ii) large contracts have long duration.

Partial Delivery Payments Now, we want to rationalize the observed aggregate time delay. We do so by assuming there exists a "representative large contract" which follows a specific delayed payment schedule consistent with partial delivery payments.

Firstly, consider an example of a top 5% defense contract from FPDS. For instance, on December 22nd, 2015, the Department of Defense (DoD) awards a new "multi-transactions" contract to L-3 Communications Corporation.[39] The contract has a duration

[38] They use a sample from 2001 to 2018 and find a median duration of defense contracts of 26 days and 90% of them have duration less than a year.

[39] See contract here.

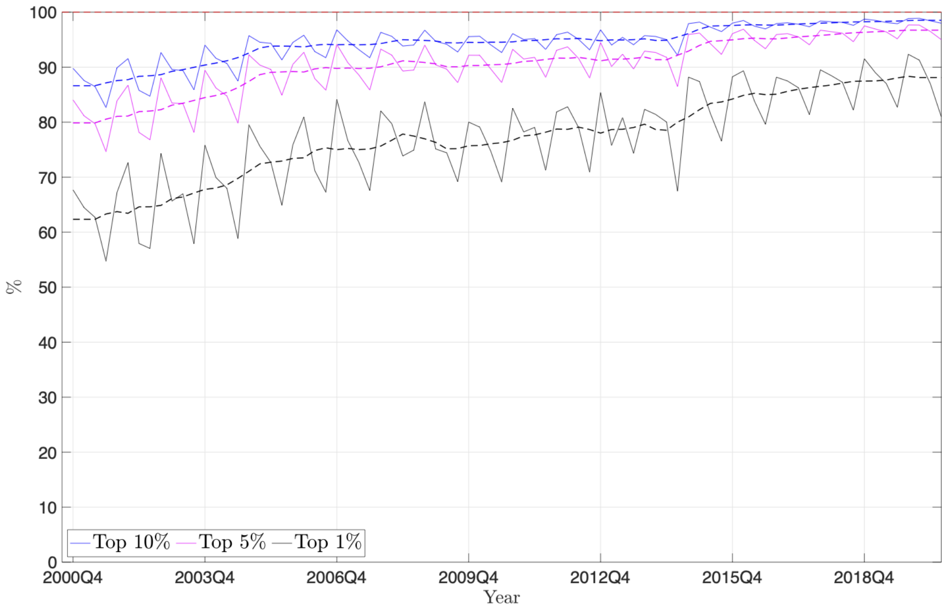

Figure 2.17: Large Contracts Share of Total Procurement Spending

of two years and involves the reparation and maintenance of some aircraft components and accessories.[40] At the time of obligation, this contract has several components, denoted child contracts, and 24 contract modifications. Each modification represents a new child contract with its own duration.[41]

In the top panel of Figure 2.18, we show on the left axis the amount of dollars obligated every quarter by this contract, and on the right axis the number of (child) contracts signed every quarter. The bottom panel shows the corresponding payment schedule which assumes that payments are disbursed once every 180 days, by uniformly spreading the initial amount of obligated funds over a contract duration.[42]

For instance, the first new child contract, signed in December 2015, lasts 375 days and obligates almost $3 million by the DoD. The payment schedule assumes that the contractors start producing the parts to be replaced immediately with partial delivery and partial reimbursement after 180 days, Therefore, the contractor is paid $1.5 million in June

[40] Duration is measured as the number of days between the Period of Performance (PoP) end date and the PoP start date.

[41] Modifications can have two types: (i) uni-lateral (e.g., administrative actions which obligate new funds for the specific contract) or (ii) bi-lateral (e.g., change to the original orders or additional work).

[42] This assumption is also made in Auerbach et al. (2020).

2016. Finally, in December 2016, the period of performance ends and the DoD pays to the contractor the remaining half of obligated funds.

Figure 2.18: Example of a Contract's Obligation and Payment Schedule

Notice that payments look like a delayed version of obligations for this particular contract. The choice of 180 days delay between payments is consistent with our estimates for the average time mismatch between defense obligations and payments found earlier. The assumption of uniform production and payments is standard and consistent with the work of Auerbach et al. (2020). In general, contractors are often incentivized to distribute production associated with an obligation throughout the whole duration of the contract.[43] In the data, cost-overruns and delays are common (see e.g., Gonzalez-Lira et al. (2021)).

Therefore, consider a representative contract with a structure similar to the one just analyzed: few new child contracts followed by several modifications. Overall, the contract lasts 48 months - consistent with the median weighted duration of defense contracts (see Table 2.5) and is characterized by payments disbursed once every 6 months (for a total of 8 payments). If we denote by P_t the total payments to contractors at time t and by O_t the amount of obligations, it is easy to show that the mapping between spending and

[43]Consider a simple firm optimization problem with convex adjustment costs.

obligations is given by the following equation:

$$P_t = \frac{1}{8} \cdot \sum_{j=1}^{8} O_{t-6\cdot j}. \tag{2.4}$$

We take the obligation data from BCD and feed it into Equation (2.4) to construct a time series of simulated payments. The left panel of Figure 2.19 plots BCD defense obligations data and the so constructed payments.

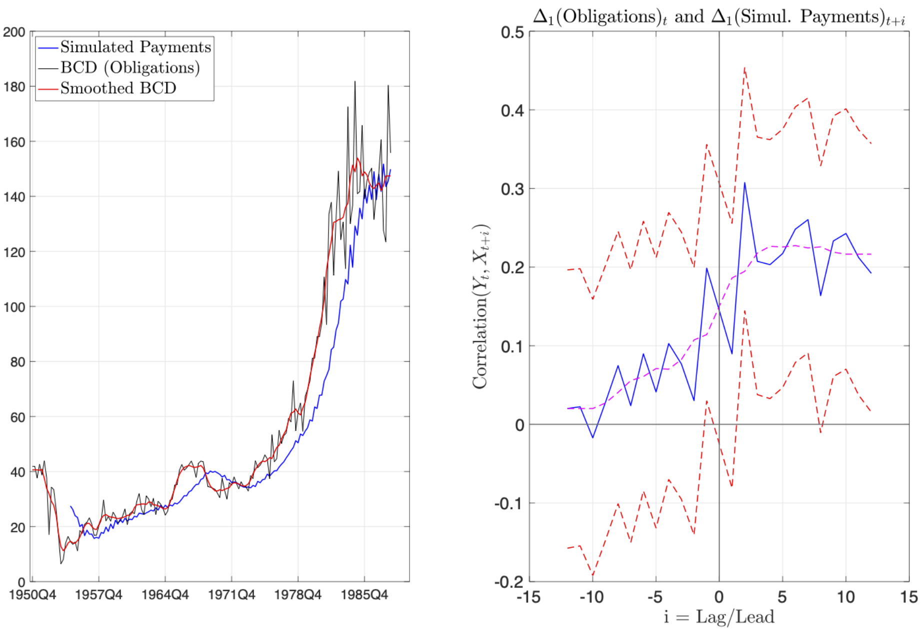

Figure 2.19: BCD Obligations and Simulated Payments

Despite the simplicity of the payments data generating process given by Equation (2.4), the simulated payments data approximate quite well the actual ones shown in Figure 5 of the paper. Similarly, the right panel shows the lead-lag correlation map between the growth rates of (smoothed) obligations and simulated payments. Notice that the results are very similar to the ones obtained using real spending data.[44]

[44] We highlight that by allowing time varying number of payments (here 8) and payments delays (here 6 months), we can improve by far the approximation to the actual data. Here, we preferred to keep things simple in the interest of brevity and clarity.

2.10.4 Construction of Quarterly Defense Obligation

We show here how we construct the time series of defense procurement obligations.

We face two main challenges: (i) we have obligations data only from 1951 to 1989 and from 2000 onward; (ii) obligations are very lumpy because contracts also get cancelled and we want to focus on obligations which turn into actual production.

i. We turn BCD and FPDS monthly data into quarterly annualized data (sum monthly observations within a quarter and multiply by 4).

ii. We apply the standard Brockwell and Davis (1991) filter to seasonally adjust the data.

iii. We construct a time trend which takes value of 1 in 1947Q1. Denote it by t.

iv. We predict obligations using 4 leads and lags and contemporaneous defense procurement spending, as well as time trends t and t^2.

v. We construct obligations from 1951Q1 to 1988Q4 using the predicted values from the previous regression. We use the estimated coefficients and the values of defense procurement spending from 1947 to 1951 to extrapolate obligations for those years.

vi. We predict obligations from year 2006 onward (FFATA introduction) in the same way. We use the predicted values to be our new series of obligations for those years. Since defense procurement (smoothed) obligations and spending overlap from 2000 to 2006, we use actual defense procurement spending for those years.

vii. From 1989 to 2000 we use defense procurement spending to proxy obligations which turn into actual production.

Figure 2.20 plots the so constructed defense procurement obligations variable (pink dash line) along with defense procurement spending (blue line) and original defense procurement obligations (dark solid line).

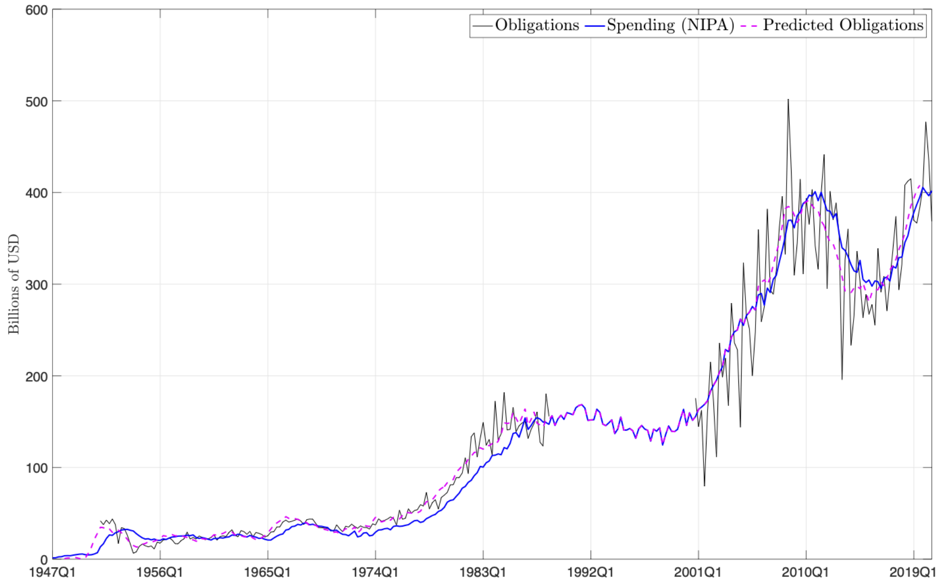

Figure 2.20: Quarterly Defense Obligation and Spending

2.10.5 What Goes On Before Contract Awards?

Although there is still uncertainty about the contract award when a pre-award notice is posted, firms might still take action in anticipation of the award. This might occur if a firm wants to become more competitive in the bidding process or predicts a contract win with a high probability. In addition, some pre-award notices justify the lack of competition in a sole-sourced contract proposal. In this case, the contractor might even predict a contract award with full certainty.

We argue from the data that any anticipatory behavior is likely to occur at a frequency higher than the frequency of aggregate analysis in this work. In other words, almost all information about contract opportunities is revealed to contractors within the quarter of the contract award. We summarize this finding by notice type in Table 2.6 and plot the distribution of pre-award notice lags in Figure 2.21.

Table 2.6: Average Lag Between Pre-Award Notices and Award Date

Notice Type	*Avg Lag in Days*	*Proportion of Notices*
Justification / Fair Opportunity	87	1.2%
Other	54	62.5%
Special Notice	41	2.1%
Pre-Solicitation	28	14.6%
Sources Sought	21	4.1%
Solicitation/Contract Solicitation	16	15.5%
TOTAL	43	

Notes: Based on matched notices between FPDS and Contract Opportunities.

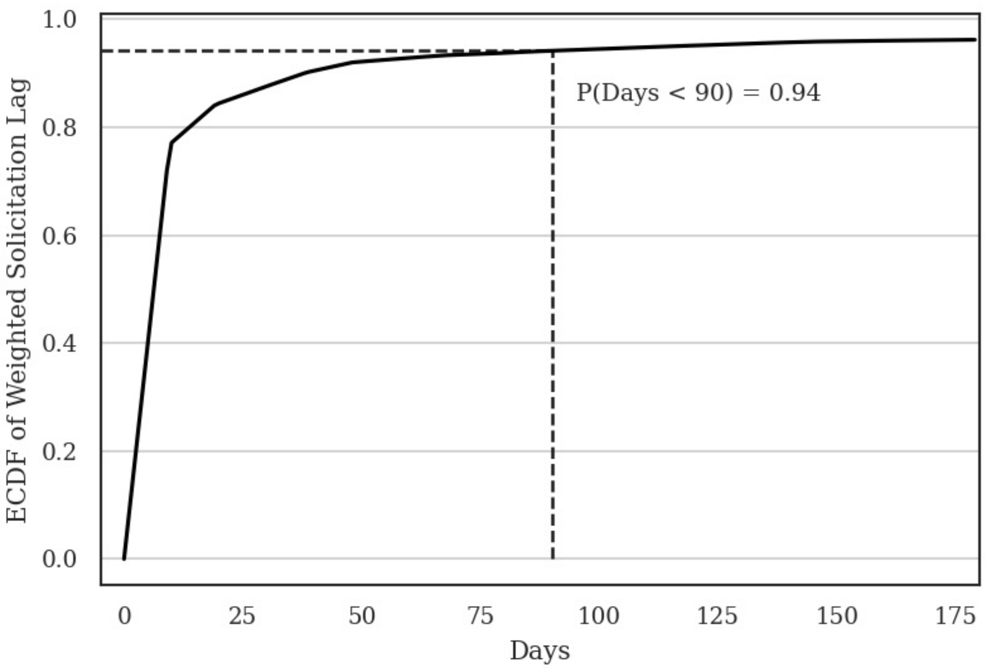

Figure 2.21: Empirical CDF of (Weighted) Solicitation Lag

Notes: Weighted duration between Contract Solicitation and Award dates measured in days. Dark dashed line represents 1 quarter (90 days). The Empirical CDF is estimated using Gaussian Kernel Density.

Detailed Description of Solicitation Process: Although public procurement contracts are awarded at a highly decentralized level (i.e., by over 69 federal agencies, 209 sub-agencies), all contracting officers are required to abide by the guidelines proposed in the Federal Acquisition Regulation (FAR). The FAR is a set of principles and procedures intended to organize and guide the procurement process across all federal agencies. In this section, we focus on the publicizing requirements associated with procurement contracts, depicted in Figure 2.22.

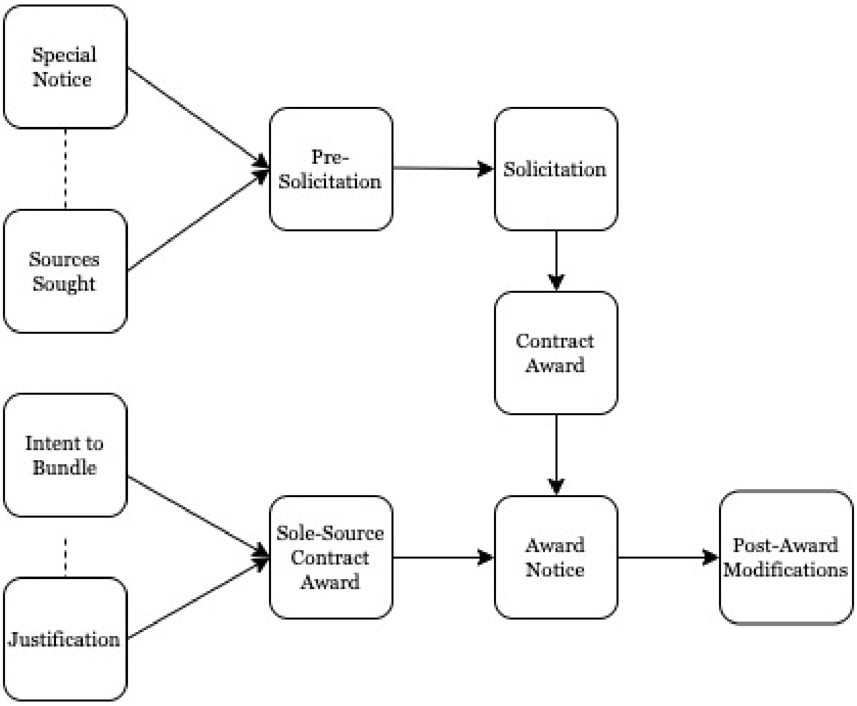

Figure 2.22: Timeline of the Procurement Process

Notes: Notice prior to contract award step occur on average within the quarter. Source: beta.sam.gov daily files.

In particular, FAR Part 5 (*Publicizing Contract Actions*) requires that contractors publicize contract opportunities with the goal of increasing competition, broadening industry participation, and assisting small businesses in obtaining contracts. Since October 1, 2001, contract actions with an expected value of over $25,000 must be publicized in an online and easy-to-access government platform, which we refer to as *Contract Opportunities*. Contract actions below the threshold might still be posted to increase visibility. On the other hand, FAR allows for exemptions to the requirement above the threshold when the posting might "compromise national security" or when the posting is "not in the government's interest". The result is that many contracts which are awarded are never solicited. When the regulation applies, Contract Opportunity notices are posted publicly at beta.sam.gov and include award notices such as solicitations, pre-solicitations, or other pre-award and post-award actions.

We describe the types of contract notices below.[45]

Special Notice Agencies use Special Notices to announce important pre-award events such as business fairs, long-run procurement forecasts, or pre-award conferences and meetings. Special Notices might also refer to "Requests for Information" (RFI) or draft solicitations.

Sources Sought Agencies post Sources Sought Notices in order to seek possible sources for a project. As discussed in FAR 7.3, the Sources Sought notice is not a solicitation for work or a request for proposal. Agencies typically use Sources Sought notices to collect industry feedback on key contracting strategy decisions and to perform market research on firm capabilities.

Pre-Solicitation Agencies post a pre-solicitation to notify vendors that a solicitation may follow. Potential vendors might then express interest in the contract by adding themselves to the Interested Vendors List. Government agencies use pre-solicitations to determine the number of qualified vendors to perform the desired work. Contracting officers can also use pre-solicitations to gather information on interested suppliers and determine if a set-aside for a small business might be applicable.

[45] Gonzalez-Lira et al. (2021) also provides a useful description and analysis of the publicizing requirements for Federal Procurement and the effects of information diffusion via public notices. We thank Andres Gonzales-Lira for directing us to the General Services Administration Technical Documentation for the FedBizOpps (FBO) website, whose information is now migrated to Contract Opportunities.

Intent to Bundle Requirements Agencies post "Intent to Bundle Requirements (DoD-Funded)" (ITB) whenever awarding actions are funded solely by the DoD. ITB supports the requirements in Section 820 of the Fiscal Year (FY) 2010 National Defense Authorization Act (NDAA) for contracting officers to post a notice of intent to use contract bundling procedures 30 days prior to releasing a solicitation or placing an order - if a solicitation is not required.

Solicitation Agencies post a solicitation to clearly define government requirements for a potential contract so that businesses can submit competitive bids. A "Request for Proposal" (RFP) is the most common type of solicitation used by federal agencies. The solicitation also sets conditions and requirements for contractor proposals and includes the government's plan for evaluating submissions for potential award.

Combined Synopsis/Solicitation Agencies post a combined synopsis/solicitation when a contract is open for bids from eligible vendors. The Synopsis/Solicitation includes specifications for the product or service requested and a due date for the proposal, as well as the bidding procedures associated with the solicitation.

Award Notice Agencies post an award notice when they award a contract in response to a solicitation. Federal agencies may choose to upload a notice of the award to make aware other interested vendors of the winning bid. Note that the requirement guidelines for posting the award notice vary based on the agency and the solicitation.

Justification Agencies are required to post a justification in order to obtain approval to award a contract without posting a solicitation as required by FAR 41 U.S.C. 253(c) and 10 U.S.C. 2304(c). Under certain conditions, agencies are authorized for contracting without full and open competition. The Department of Defense, Coast Guard, and National Aeronautics and Space Administration are subject to 10 U.S.C. 2304(c). Other executive agencies are subject to 41 U.S.C. 253(c). Contracting without providing for full and open competition or full and open competition after exclusion of sources is a violation, unless permitted by one of the exceptions in FAR 6.302.

Sale of Surplus Property Agencies post a sale of surplus property notice when they wish to sell federal real estate properties for public use. These properties are typically made available for public use to state and local governments, regional agencies, or nonprofit organizations to state and local governments. Public uses for properties are those that are accessible to and can be shared by all members of a community, and include community centers, schools and colleges, parks, municipal buildings and many more.

2.11 Time-to-build or Production Smoothing?

We decompose the early response of inventories to a defense news shock into time-to-build and production smoothing. We already have contract level evidence of a long time-to-build, but at the onset of a military build-up, contractors should also presumably change their expectations about future government demand. Even if contractors lack resources to forecast government demand, federal agencies are required by the FAR to provide procurement forecasts each quarter.[46] If contractors anticipate winning future contracts, they might decide to increase production today to smooth convex adjustment costs or reduce future delivery times. We do not take a stance on the exact mechanism here. We consider a recent example of this type of behavior.

Example 2.11.1 (Lockheed Martin in 2022). In the context of an ongoing military conflict between Russia and Ukraine, new military tests in North Korea, and escalating tension in relationship between China and Taiwan, US-based contractors have modified expectations about future defense spending. In particular, the largest American defense contractor, Lockheed Martin, decided in October 2022 to increase production of HIMARS (High Mobility Artillery Rocket System). When asked about this decision, CEO Jim Taiclet responded as follows:[47]

"*The company has met with its long lead supply chain and spent $65 million — which will eventually be paid back by the US government — to fund parts in advance, shortening the time needed to manufacture the rocket system. That was without a contract or any other memo or whatnot back from the government. We just went ahead and did that because we expected it to happen. So those parts are already being manufactured now*".

In order to measure production smoothing, we first provide a formal definition.

Definition 2.11.1 (Production Smoothing of Defense Industries)**.** We define production

[46] See e.g., Agency Recurring Procurement Forecasts.

[47] Find the associated article on *Breaking Defense*. here.

smoothing $\Delta(h)$ as the effect of a defense news shock on inventories, orthogonal to shocks to newly awarded contracts (i.e., defense procurement obligations). In particular, production smoothing is the impulse response of inventories to a defense news shock conditional on zero shocks to defense procurement obligations (i.e. orthogonalized IRF):

$$\Delta(h) = \mathbb{E}_t[\text{Invt}_{t+h}|Z_t = 1, \varepsilon_t^O = 0] - \mathbb{E}_t[\text{Invt}_{t+h}|Z_t = 0, \varepsilon_t^O = 0], \tag{2.5}$$

where Invt_t is changes in aggregate inventories as from the NIPA, Z_t is a defense news shock, and ε_t^O is a shock to defense procurement obligations.

We estimate production smoothing using the following tri-variate VAR using quarterly data from 1948Q1 to 2015Q4:

$$\begin{bmatrix} 1 & 0 & 0 \\ -\alpha & 1 & 0 \\ -\beta_{\text{News}} & -\beta_{\text{Oblg}} & 1 \end{bmatrix} \cdot \underbrace{\begin{bmatrix} Z_t \\ O_t \\ \text{Invt}_t \end{bmatrix}}_{\boldsymbol{X}_t^3} = \underset{3\times 3}{\boldsymbol{B}_3(L)} \cdot \boldsymbol{X}_t^3 + \boldsymbol{\varepsilon}_{3,t}$$

where $\mathbf{B}_3(L)$ is a polynomial in the lag operator. The parameter α captures the contemporaneous effect of a defense news shock on obligations, while β_{News} and β_{Oblg} capture the contemporaneous effect of shocks to news and obligations on inventories.

By including our aggregate series for defense procurement obligations O_t, we are able to calculate the effect of defense news shocks on inventories which is independent of the effect of shocks to newly awarded contracts. Figure 2.23 shows the impulse response function to a defense news shock estimated using the above tri-variate VAR as well as the total response of inventories estimated in a bi-variate VAR without obligations.

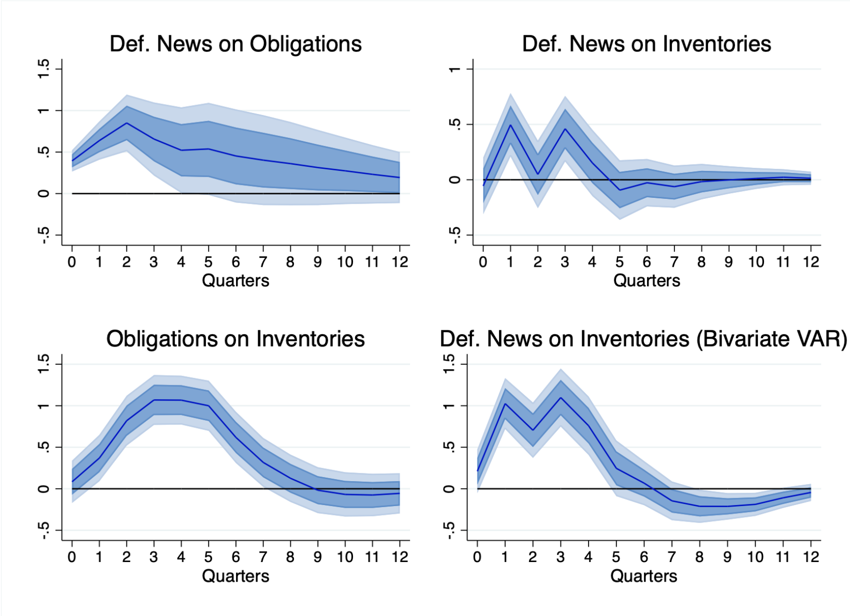

Figure 2.23: Orthogonalized VAR Impulse Response Functions

Notes: Variables are divided by potential GDP and include a linear time trend. Sample goes from 1948Q1 to 2015Q4. Confidence Bands are 68% and 90%. Values are rescaled by the peak response of Inventories to a defense news shock from the bivariate VAR which excludes defense procurement obligations.

The top-left panel of Figure 2.23 shows the positive response of defense procurement obligations to a defense news shock. This indicates that new contracts start being awarded as soon as a defense news shock occurs. This confounds the effects of news (i.e., anticipation) with the effects of newly awarded contracts which show up in G with delay. In the bottom-left panel of the figure, we show the effect of shock to obligations ε_t^O, on inventories, orthogonal to defense news. The effect is positive and significant. Additionally, the top-right panel reports production smoothing, or the response of inventories to a defense news shock which is orthogonal to newly awarded contracts. The positive and significant estimates of $\Delta_(h)$ at horizons 1 and 3 suggest that production smoothing plays a role in the response of inventories. The bottom-right panel shows the IRF of inventories to a defense news shock without including obligations in the VAR, i.e. bivariate VAR.

For interpretability, we rescale the IRFs by the peak response of inventories to a defense news shock in the bivariate VAR occurring at horizon 1. Since the horizon 1 response of inventories to a defense news shock in the tri-variate VAR is slightly more than 0.4 it means that roughly 40% of the response of inventories at horizon 1 comes from production smoothing, while the residual part (gap between bottom-right and top-right responses) originates from the effects of newly awarded contracts, i.e. time-to-build production. Intuitively, this can be seen by shrinking the tri-variate VAR into a bivariate one by plugging obligations into the equation of inventories:

$$\begin{bmatrix} 1 & 0 \\ -\left(\beta_{\text{News}} + \alpha \cdot \beta_{\text{Oblg}}\right) & 1 \end{bmatrix} \cdot \underbrace{\begin{bmatrix} Z_t \\ \text{Invt}_t \end{bmatrix}}_{\boldsymbol{X}_t^2} = \underset{2\times 2}{\boldsymbol{B}_2(L)} \cdot \boldsymbol{X}_t^2 + \boldsymbol{\varepsilon}_{2,t}$$

Notice that the impact effect of a defense news shock on inventories is the combination of production smoothing (β_{News}) and the effect of a shock to new contracts on inventories triggered by the news ($\alpha \cdot \beta_{\text{Oblg}}$). Basically, without controlling for new contracts, defense news shock capture both production smoothing and the time-to-build, while augmenting the VAR with new contracts allows us to tell-apart the two effects.

2.12 Robustness - Section IV in the paper

Firstly, we construct an index of cumulative excess returns similar to the Top3 index constructed in Fisher and Peters (2010). The variable is shown in Figure 2.24 along with red lines denoting the Ramey-Shapiro episodes.

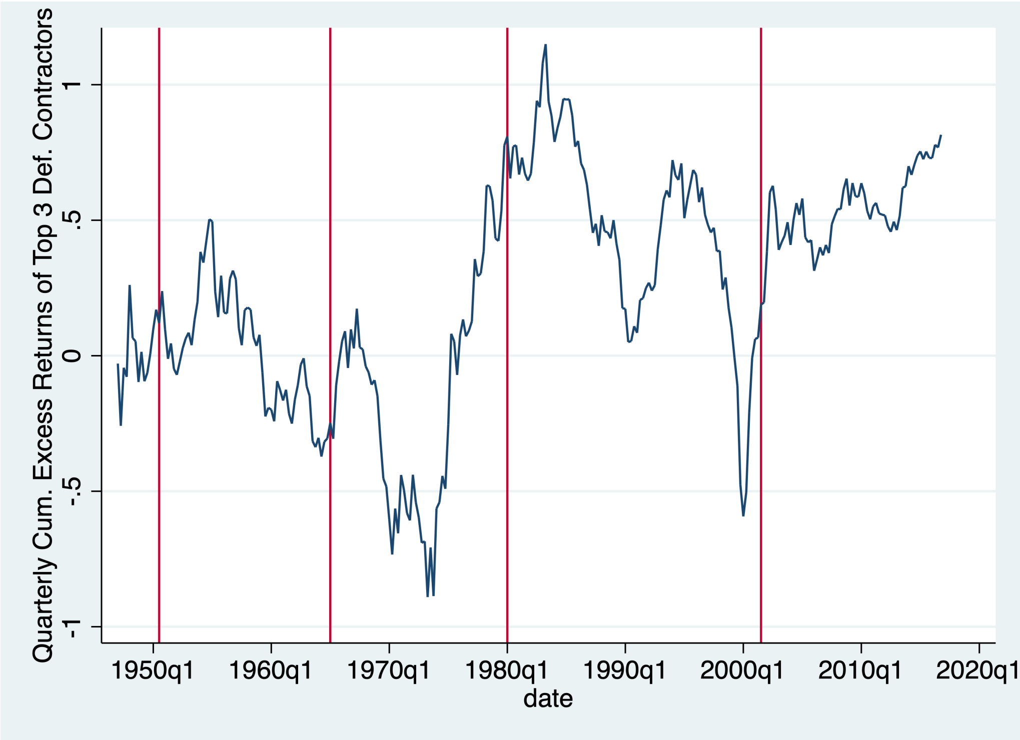

Figure 2.24: Top3 Defense Contractors Cumulative Excess Stock Returns Index

Notes: Red solid lines are the Ramey-Shapiro episodes.

Similarly to Fisher and Peters (2010), the Top3 index onlr responds to the Vietnam war and 9/11, but not the Carter-Reagan military build-up nor the Korean war.

We construct shocks to this variable by ordering it first in the same VAR used in Section I in the paper. Furthermore, we complement the Granger Causality test in the paper by using these new shocks. Results are shown in table below.

Table 2.7: Predictability of Recursive Shocks via Obligations

Predicted	Predictor	F	Pvalue	Korea
Recursive Shocks	Top3	0.26	97.84%	Yes
Obligation Shocks	Top3	1.25	26.81%	Yes
Defense News Shocks	Top3	0.42	90.67%	Yes
Recursive Shocks	Top3	0.63	74.93%	No
Obligation Shocks	Top3	0.88	53.53%	No
Defense News Shocks	Top3	0.62	76.22%	No
Top3	Recursive Shocks	1.00	43.57%	Yes
Top3	Obligation Shocks	0.49	86.54%	Yes
Top3	Defense News Shocks	0.89	52.84%	Yes
Top3	Recursive Shocks	0.94	48.09%	No
Top3	Obligation Shocks	0.39	92.70%	No
Top3	Defense News Shocks	0.46	88.44%	No

Notes: Granger Causality test is a Wald test on the 8 lags of the predictor while controlling for 4 lags of the predicted variable.

It is clear from Table 2.7 that we find no significant predictability in either direction for the Top3 index.

Secondly, we replicate the bottom panel of Figure 4 in the paper, by excluding the Korean war from the sample. Results are shown in Figure 2.25.

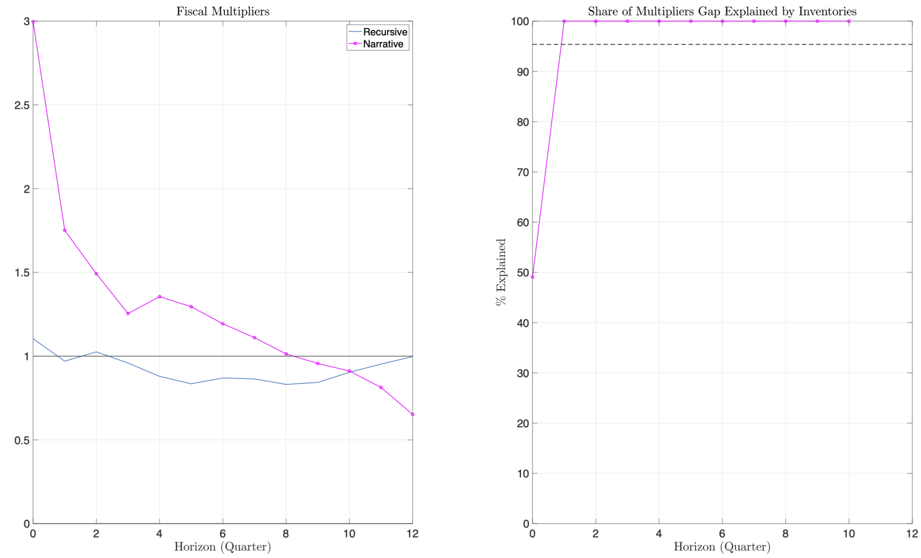

Figure 2.25: Cumulative Fiscal Multipliers and Multiplier-Gap (Robustness)

Notes: Sample goes from 1954Q1 to 2015Q4. All the rest is identical to Figure 7 in the Paper.

2.13 Comparison of Multipliers with Brunet (2020)

Finally, we estimate multipliers using our new quarterly measure of defense procurement obligations. We compare our results to the recent estimates from Brunet (2020), who also uses a measure of government spending which is better aligned with the timing of obligated funds. In particular, Brunet (2020) estimates multipliers by regressing cumulative changes in GDP on cumulative changes of *Budget Authority*, which tracks defense spending when it is authorized, before funds are dispersed from the Treasury:

$$\sum_{k=0}^{H} \frac{GDP_{t+h} - GDP_{t-1}}{GDP_{t-1}} = \mathcal{M}(H) \cdot \sum_{k=0}^{H} \frac{BA_{t+h} - BA_{t-1}}{GDP_{t-1}} + lags + \varepsilon_{t+h}$$

where BA_t is Budget Authority in year t. We report these estimates in Table 2.8 for two post WWII samples which either include or do not include the Korean War.

Table 2.8: Multiplier Comparison

Sample	Horizon (Years)	0	1	2	3	4
Post WWII Sample	Budget Authority	1.76	1.51	1.30	1.28	
		(4.08)	(2.73)	(1.63)	(1.29)	
	Def. Proc. Oblig.	4.84	3.96	1.61	1.27	1.17
		(1.71)	(3.01)	(3.94)	(3.18)	(2.61)
Post Korean War	Budget Authority	1.83	1.84	1.72	1.67	
		(4.54)	(3.56)	(2.25)	(1.66)	
	Def. Proc. Oblig.	1.44	1.7	1.15	1.09	1.14
		(2.85)	(2.46)	(1.99)	(2.04)	(2.09)

Notes: t-statistics reported in parentheses below the multipliers' point estimates. Budget Authority samples go from 1948 to 2016 and from 1955 to 2016 (annual frequency). Samples using defense procurement obligations go from 1948Q1 to 2015Q4 and from 1954Q1 to 2015Q4 (quarterly frequency).

On the other hand, we estimate multipliers as suggested in Ramey (2016), instrumenting the cumulative change in the NIPA-measured government spending, G, with defense procurement obligations (one step LP-IV).[48] We plot our estimates of cumulative fiscal multipliers in Figure 2.26 for the two sample periods from Brunet (2020).

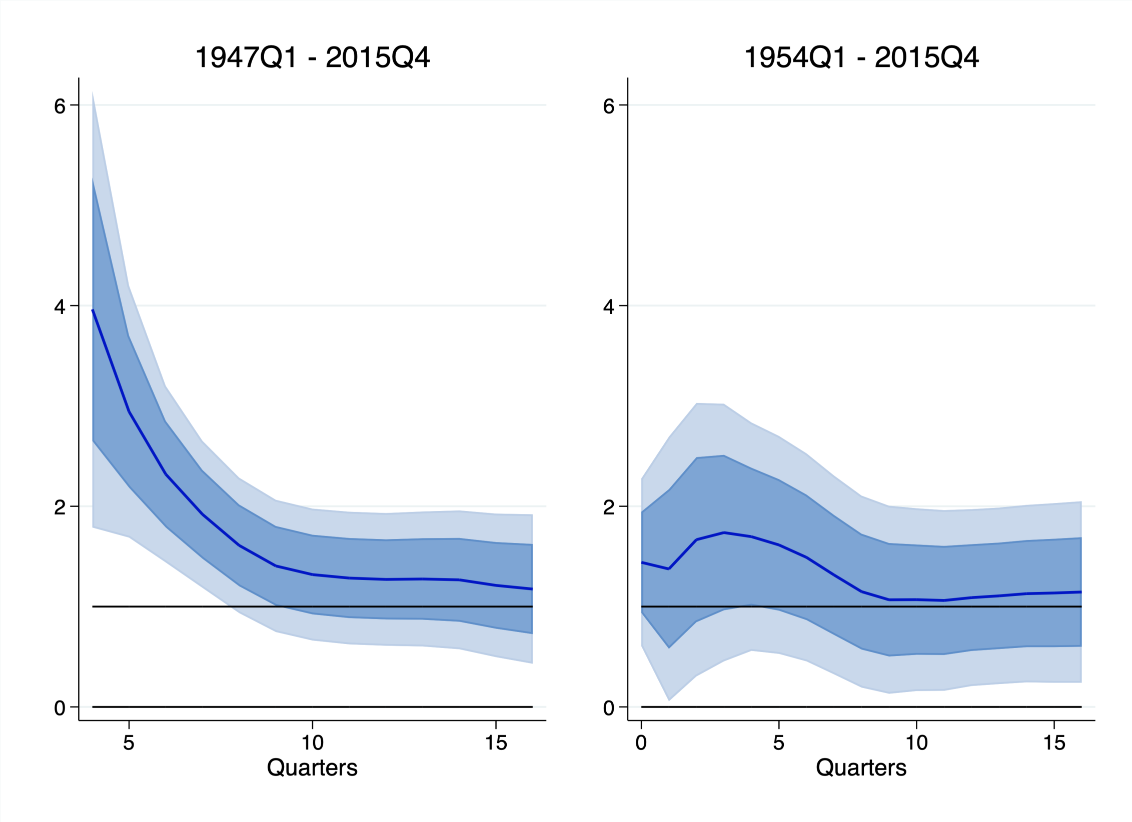

Figure 2.26: Cumulative Fiscal Multipliers via Shocks to Defense Procurement Obligations

Notes: Standard errors are two-stage-least-squares robust standard errors and bands are the 68% and 90% confidence levels. Dark horizontal lines referring to the values of zero and one. Data are quarterly.

Notice how in both sample periods, the multipliers are higher at short horizons and smaller at longer horizons. This is a consequence of anticipation effects: GDP increases even before G moves.[49] We also report the point estimates of the multipliers in Table 2.8 for different years. This is done to facilitate the comparison with the results of Brunet (2020).

In the top panel of the table, it is clear that analysis on both Post WWII samples deliver similar results, particularly at the three-year horizon. On the contrary, results slightly differ from each other when the Korean war is excluded from the sample. In fact,

[48] Recall that this is equivalent in population to ordering defense procurement obligations first in a VAR.

[49] When the Korean war is in the sample, anticipation is so strong that the multiplier is infinite until horizon 4.

our point estimates are smaller than Brunet (2020), even if they both remain above one. Finally, our multipliers tend to be more statistically significant at longer horizons, which is likely due to the fact that our analysis is carried out at quarterly frequency rather than an annual frequency.

Despite minor discrepancies between our results and those of Brunet (2020), our obligations-based method also delivers point estimates for the multiplier which are greater than one. Nevertheless, we note three important differences in our methodology, without taking a stand on the relative effects of each. First of all, defense procurement obligations is a quarterly variable which captures the whole universe of newly awarded defense procurement contracts, while Budget Authority is an annual variable which captures authorizations for defense-spending and is broader than procurement spending. Secondly, our LP-IV multiplier is interpretable as the ratio of the IRFs of GDP and G following a shock to defense procurement obligations, and is therefore a spending multiplier. On the contrary, Brunet (2020) regresses cumulative changes of GDP on cumulative changes of Budget Authority. Since Budget Authority does not map directly to NIPA defense spending (i.e., changes in Budget Authority are not necessarily changes in NIPA G), their estimates cannot be directly interpreted as a spending multiplier. Thirdly, contemporaneous changes in non-defense spending are not captured in the Budget Authority measure.

Victor Sellemi is the sole author of Chapter 1.

Chapter 2 is coauthored with Edoardo Briganti. The dissertation author is one of the primary investigators of this material.

Chapter 3 is coauthored with Tjeerd de Vries. The dissertation author is one of the primary investigators of this material.

Chapter 3

Deep Learning and Long-Run Risk

Abstract

In dynamic asset pricing, stochastic discount factor (SDF) processes summarize the relationship between equilibrium asset prices and underlying economic conditions. SDFs can be factorized into permanent and transitory components, where the permanent component captures pricing at long payoff horizons. Hansen and Scheinkman (2009) show that the permanent-transitory decomposition can be cast as the unique solution to a Perron-Frobenius eigenvalue problem, for which analytic solutions are only available for a limited array of examples. Moreover, standard numerical approaches are not equipped to handle this general class of problems due to the curse of dimensionality and lack of well-developed boundary conditions or parametric restrictions. We develop a novel algorithm for solving this class of eigenvalue problems in a very general class of asset pricing models without boundary conditions or parametric assumptions on the eigenfunctions. We demonstrate the accuracy of the algorithm in the context of several workhorse structural asset-pricing models, and argue that our approach applies to models which feature a very general class of Lévy processes in arbitrarily high dimensions.

3.1 Introduction

Asset prices encode dynamic information about magnitude and duration of exposure to risk and uncertainty in future cash flow growth. To capture these trade-offs in arbitrage-free markets, dynamic asset pricing models summarize the relationship between equilibrium asset prices and underlying economic conditions (state variables) using stochastic processes called stochastic discount factors (SDFs). To study the limiting behavior of asset prices as the investment horizon gets very large, Alvarez and Jermann (2005) and Hansen and Scheinkman (2009) develop a factorization of SDF processes into permanent and transitory components. The permanent component characterizes prices over long investment horizons, while the transitory component is related to the return of a discount bond with (infinitely) long maturity. Our focus is long-run permanent-transitory components of SDFs in continuous-time asset pricing models.

Alvarez and Jermann (2005) and Bakshi and Chabi-Yo (2012) demonstrate that both permanent and transitory components of SDFs must be non-trivial in order to match key features of historical returns data. Moreover, Qin and Linetsky (2017) show that the permanent-transitory decomposition is a fundamental feature of arbitrage-free asset pricing models in very general semi-martingale environments. However, quasi-analytical solutions for the permanent and transitory components are only available for a very limited subset of examples, which is a challenge to verifying that the long-run predictions of asset pricing models are consistent with very salient features the data.[1] In this work, we develop a flexible non-parametric numerical approach for computing the long-run components in very general set of asset pricing models.

Our starting point is the seminal work of Hansen and Scheinkman (2009), who show that the permanent and transitory components of SDFs can be constructed as the principal

[1]Christensen (2017) develops a non-parametric empirical framework for estimating the permanent component of SDF processes from historical data. Such empirical benchmarks can be used to compare predictions of asset pricing models to estimated properties. Alvarez and Jermann (2005) and Bakshi and Chabi-Yo (2012) also develop bounds on the various moments of the permanent and transitory components as a function of asset returns.

solution to a corresponding Perron-Frobenius eigenvalue problem. The eigenfunction captures dependence of the price of long-horizon payoffs on the underlying economic state, while the principal eigenvalue determines the average yield on asymptotically long-horizon payoffs. In other words, the eigenfunction of the SDF operator fully characterizes the probability measure relevant for pricing assets of asymptotically long maturity.[2] However, analytic solutions to this eigenvalue problem are rare, and standard numerical approaches either suffer from the curse of dimensionality or require strict parametric assumptions.

To overcome these limitations, we build on recent developments in the intersection of computational physics and deep learning (e.g., E and Yu (2017), Han et al. (2020)). We start by augmenting a time dimension and transforming the stationary eigenvalue problem into a parabolic partial differential equation (PDE). We conveniently construct the PDE such that its stationary solutions correspond to eigenpairs of the SDF operator. Our object of interest is the principal eigenpair, which characterizes its permanent-transitory components. However, the PDE transformation is still numerically intractable, especially in high dimensions. We therefore leverage the Feynman-Kac Theorem to map the PDE to a more tractable stochastic differential equation (SDE) counterpart with the same stationary solutions. Like the PDE, the SDE process depends on both the eigenvalue and eigenfunction.

Since the eigenvalue and eigenfunction are unknown to us ex-ante, we initialize them randomly. We draw from a uniform distribution for the eigenvalue, and from a flexible class of deep feed-forward neural networks (DNNs) for the eigenfunction.[3] We then optimize both our eigenvalue and neural network approximation of the eigenfunction via gradient descent, by minimizing a squared error fixed-point loss function which is equal to zero at stationary solutions of the SDE. We add a penalty to avoid convergence to the trivial solution. Moreover, since the fixed-point loss function is non-convex, we validate

[2] See Hansen (2012), Backus et al. (2014), Borovička et al. (2016), Qin and Linetsky (2017) for additional theoretical results.

[3] Related research in economics also leverages the flexibility and computational efficiency of deep neural networks to solve Hamilton-Jacobi-Bellman equations, options pricing problems, and heterogeneous-agent macroeconomic models (see e.g., Han et al. (2018), Germain et al. (2021), Han et al. (2022)).

the optimization with different initial conditions to ensure that we are truly converging to the stationary solution corresponding to the principal eigenpair, which can then be used to construct the permanent and transitory SDF components.[4]

When the dimension of the (Markov) state variable is low enough, we might opt for classic numerical approaches such as finite differences, finite elements, or the spectral method to approximate the eigensolutions. However, besides the common requirement of specified boundary conditions, such approaches suffer from the curse of dimensionality as number of degrees of freedom grows exponentially with the dimension of the problem. We find that performance of these benchmark deterministic approaches deteriorates in multivariate asset pricing models with more than a single state variable. Since the dimension of the Markov state corresponds to the number of variables required to summarize the state of the economy, most realistic asset pricing models tend to have at least two dimensions (e.g., mean and volatility of aggregate growth rates). To overcome the curse of dimensionality, our proposed algorithm leverages more tractable stochastic elements. Moreover, to avoid making parametric assumptions and for robustness to arbitrary high dimensionality, we leverage flexible neural network approximations of the eigenfunction.[5]

We demonstrate the accuracy of our approach in the selected set of (low-dimensional) workhorse structural macro-finance models whose permanent-transitory decomposition has a known closed form solution. Moreover, according to recent research, these models are a good starting point for our analysis since they showcase challenges that many standard asset pricing models have in explaining key features of historical returns data. The first

[4]Under non-convex loss, we cannot guarantee global convergence to the true solution. However, Hansen and Scheinkman (2009) and Qin and Linetsky (2016) prove that the principal eigenvalue associated with a wide class of multiplicative SDFs is strictly positive. As long as the principal eigenvalue of the SDF operator is bounded away from zero, the true solution is guaranteed to at least be a local minimum. Given the (recent) advances in computational efficiency of gradient-based optimization (e.g., Kingma and Ba (2017)), we are able to easily validate the optimization under various perturbations to the algorithm initialization.

[5]The two alternative stochastic approaches (common in quantum and statistical physics) for solving high-dimensional eigenvalue problems are variational and diffusion Monte Carlo (see e.g., Ceperley et al. (1977), Blankenbecler et al. (1981), Zhang et al. (1997), Foulkes et al. (2001), Needs et al. (2009)), but typically require researchers to specify a parametric form for the solution eigenfunction and/or two-sided boundary conditions. See Hornik et al. (1989) for one of the earliest proofs of the universal approximation property of Neural Networks. See Bauer and Kohler (2019) for theoretical and empirical evidence that neural networks perform well in the high dimensional setting.

two examples are special cases of the Breeden (1979) model and the Kreps and Porteus (1978) model with Duffie and Epstein (1992) stochastic differential utility. Backus et al. (2014) shows that these standard consumption-based models correspond to SDFs which cannot generate large enough permanent components in line with empirical measurements without also generating unrealistically large spreads between long- and short-term yields. Our third example, the Bansal and Yaron (2004) long-run risk model, tries to amplify long-run risk premia by combining intertemporal preferences and predictable components in consumption and consumption volatility. We find that this model does not fully resolve the previous critique.[6]

Given its flexibility and scalability, our proposed approach is a useful model selection and analysis tool for addressing these limitations using asset pricing models which do not have solutions for the long-run risk components. One important set of models without well-developed permanent-transitory components introduce jump processes in the underlying state dynamics to help explain key features of the data. For instance, Bansal and Shaliastovich (2011) propose a model with jumps in asset prices to capture large return fluctuations observed in the data. Similarly, Shaliastovich and Tauchen (2011) motivate the importance of incorporating very general Lévy processes with a potentially infinite number of jumps to explain higher order moments of consumption and dividend growth rates. Although these new models are able to better capture some important features of the data, it remains unclear on if jumps can directly resolve the critique raised by Alvarez and Jermann (2005) and Backus et al. (2014) that asset pricing models undershoot empirically observed permanent-transitory SDF components. We provide evidence that incorporating jump processes in the state dynamics amplifies the long-run risk components in the Breeden (1979) and Kreps and Porteus (1978) models. Building on recent developments in numerical solutions to PDEs from Gnoatto et al. (2022), we argue that our solution algorithm can be extended to approximate the permanent-transitory components of asset

[6]The other issue, pointed out by Bakshi and Chabi-Yo (2012), is that the long-run risk model cannot replicate the empirically observed positive covariance between the permanent and transitory components of SDFs.

pricing models whose underlying state dynamics follow very general Lévy processes, including processes which feature infinite jumps.

The paper is structured as follows. Section 3.2 introduces eigenvalue problems in the asset pricing setting. Section 3.3 presents and develops our solution algorithm. Section 3.4 demonstrates the accuracy of our algorithm in settings with known solutions. Section 3.5 develops the extension of the algorithm to more general processes. Section 3.6 concludes.

3.2 Long-Run Asset Prices as an Eigenvalue Problem

3.2.1 Setup

Our setting is the continuous-time Markovian environment developed in Hansen and Scheinkman (2009) (HS henceforth).[7] Consider an economy whose state is summarized by a d-dimensional continuous-time Markov state process $(X_t)_{t\geq 0}$ with completed natural filtration $(\mathcal{F}_t)_{t\geq 0}$ generated by its history.[8] In particular, the process X_t is a vector-valued diffusion capturing dynamic economic information relevant for valuation, such as aggregate growth rates or stochastic volatility. Under the no-arbitrage assumptions, there exists a strictly positive stochastic discount factor (SDF) process $(S_t)_{t\geq 0}$ that is $\mathcal{F}_t$-adapted, and can price assets with payoffs at time t (i.e., contingent t-claims). In particular, the current (time zero) price, $\mathcal{P}_0$, of a claim with payoff Π_t at time t can be written as follows:

$$\mathcal{P}_0 = \mathbb{E}\left[S_t \Pi_t | \mathcal{F}_0\right], \tag{3.1}$$

[7] Hansen (2012), and Borovička et al. (2016), Qin and Linetsky (2017) also develop this theoretical framework.

[8] More formally, $(X_t)_{t\geq 0}$ is a time-homogeneous, strictly stationary, and ergodic Markov process, defined on probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and taking values in $\mathcal{X} \subseteq \mathbb{R}^d$.

where the $(\Pi_t)_{t\geq 0}$ is also an $\mathcal{F}_t$-adapted random process, and the expectation operator $\mathbb{E}[.]$ is taken with respect to investor beliefs.[9] Note that S_t must satisfy $S_0 = 1$ under the no-arbitrage assumption. To price the payoff Π_t at intermediate times $\tau \leq t$, we can re-index the pricing equation in (3.1) and write:

$$\mathcal{P}_\tau = \mathbb{E}\left[S_{t-\tau}\Pi_t|\mathcal{F}_\tau\right]. \tag{3.2}$$

The law of one price guarantees that a date zero claim to the date τ purchase price $\mathcal{P}_\tau$ is equivalent to a date zero claim to Π_t, or alternatively, a date zero claim to payoff Π_t/S where $S = S_\tau$ is the realized value of the SDF process at time τ.[10] We can leverage this property to infer future prices based on a ratio date zero stochastic discount factors, and rewrite the price in (3.2) as follows:

$$\mathcal{P}_\tau = \mathbb{E}\left[\frac{S_t}{S_\tau}\Pi_t\middle|\mathcal{F}_\tau\right]. \tag{3.3}$$

Suppose the random payoff Π_t can be written as a function of the current Markov state (i.e., $\Pi_t = \phi(X_t)$), then the Markov restriction in asset pricing models ensures that current prices only depend on the current Markov state. To capture, this we define the pricing operator $\mathbb{S} = (\mathbb{S})_{t\geq 0}$ which maps random future payoffs to time zero prices as a function of time zero state x:

$$\mathbb{S}_t\phi(x) := \mathbb{E}\left[S_t\phi(X_t)|X_0 = x\right]. \tag{3.4}$$

We can price the same payoff at intermediate dates $\tau \leq t$ based on the representation in (3.3).[11] In this specification, the ratio S_t/S_τ depends only the evolution of the Markov state

[9]See e.g., Hansen and Renault (2010) for more details on the probability measure underlying this expectation operator.

[10]The law of one price implies that $\mathcal{P}_0 = \mathbb{E}[S_\tau P_\tau|\mathcal{F}_0]$. We can use this to verify that P_τ is the time zero price of an asset that pays off Π_t/S_τ at time t as follows: $\mathbb{E}[S_t(\Pi_t/S_\tau)|\mathcal{F}_0, S_\tau] = \mathbb{E}[\mathbb{E}[S_t\Pi_t|\mathcal{F}_0]/S_\tau|\mathcal{F}_\tau] = \mathbb{E}[P_0/S_\tau|\mathcal{F}_\tau] = \mathbb{E}[\mathbb{E}[S_\tau P_\tau|\mathcal{F}_0]/S_\tau|\mathcal{F}_\tau] = P_\tau$.

[11]We can also use the operator $\mathbb{S}_{t-\tau}$ to price this payoff. Note that $\mathbb{S}_0 = \mathbb{I}$ and $\mathbb{S}_t = \mathbb{S}_{t-\tau}\mathbb{S}_\tau$ ensures that $\mathbb{S}$ satisfies the semi-group property, a property remarked in early work by Garman (1985).

process between τ and t, and can be written as follows:

$$\frac{S_t}{S_\tau} = S_{t-\tau}(\theta_\tau), \tag{3.5}$$

where θ_τ is an operator that shifts the Markov process forward by τ time units (i.e., $X_\tau(\theta_t) = X_{t+\tau}$).[12] The property in (3.5) along with the boundary condition $S_0 = 1$ define a semigroup family of multiplicative pricing functionals. HS prove that this family of SDF operators permit the following multiplicative factorization:

$$S_t = \exp(\lambda t)\hat{S}_t \frac{e(X_0)}{e(X_t)}, \tag{3.6}$$

where $\hat{S}_t$ is a martingale whose log has stationary increments and $e(.)$ is a positive function.[13] Intuitively, the scalar λ can be thought of as a deterministic growth rate, while the ratio of positive random variables $(e(X_0)/e(X_t))$ is the transitory contribution of the Markov state. Although this decomposition is not necessarily unique, HS prove that these is a unique pair of values $(\lambda, e(.))$ under which X_t remains stationary and ergodic (stochastically stable). This unique solution characterizes pricing for payoffs with (infinitely) long horizons. To see this, recall from HS that the martingale $\hat{S}_t$ defines a new probability measure via $\hat{P}(A) = \mathbb{E}(\mathbb{1}_A \hat{S}_t)$ for all $A \in \mathcal{F}_t$. Taking expectations under this twisted probability measure $\hat{P}$ (denoted similarly by $\hat{\mathbb{E}}$), we can write the time-zero price of a long-term discount bond by:

$$p_0 = \mathbb{E}[S_t | X_0 = x] = \exp(\lambda t) \cdot \hat{\mathbb{E}}\left[\frac{e(x)}{e(X_t)}\middle| X_0 = x\right]. \tag{3.7}$$

[12] $\theta : \Omega \to \Omega$ maps current states to future states.

[13] Qin and Linetsky (2017) extend this decomposition to very general semi-martingale environments.

Rearranging the above expression and taking limits as the payoff horizon tends to infinity ($t \to \infty$), we can write:

$$\lim_{t\to\infty} \mathbb{E}[S_t \exp(-\lambda t)] = \lim_{t\to\infty} \hat{\mathbb{E}}\left[\frac{1}{e(X_t)}\right] e(x), \tag{3.8}$$

where the left-hand side is the value of a discount bond with an infinitely long payoff horizon, and the right-hand side, particularly $e(x)$, captures the dependence of long-run prices on the current state x.[14]

HS show that we can rewrite $e(x)$ and $\lambda > 0$ as the principal eigenfunction and eigenvalue of the valuation operator $\mathbb{S}$. By definition, we rewrite the pricing equation from (3.7) as follows:

$$\mathbb{S}_t e(x) = \mathbb{E}[S_t e(X_t)|X_0 = x] = \exp(\lambda t) e(x). \tag{3.9}$$

Next, we divide both sides by t and take the limit as $t \downarrow 0$ to obtain the time-invariant eigenvalue problem associated with long-run prices:

$$\mathbb{S}e = \lambda e, \tag{3.10}$$

where $e : \mathbb{R}^d \to \mathbb{R}_+$. In later sections, we provide the few examples of asset pricing models with closed form expressions for e and λ. However, as highlighted by HS, there is no analytical solution method for solving this general class of problems. Moreover, standard deterministic numerical approaches tend to be very error-prone or completely intractable at values of $d > 1$. In the following section, we propose a novel algorithm for solving (3.10) with arbitrarily high dimension d. Before doing so, we provide some concrete examples of of the permanent-transitory decomposition discussed above.

[14]This limit is only possible since X_t is stochastically stable under $\hat{P}$.

3.2.2 Examples from Structural Asset Pricing

In this section, we consider the few continuous-time structural macro-finance models with closed form solutions for the permanent-transitory SDF components. These decompositions have been well-developed in papers such as Hansen and Scheinkman (2009), Hansen (2012), and Borovička et al. (2016). In these cases, the closed-form solutions to (3.10) are obtained using a "guess-and-check" approach which was only viable due to the low dimensionality $d = 2$. As far as we know, Models with higher dimensionality or a vastly different structural formulation have not been solved analytically in the literature.

For each of these models, we write down the eigenvalue problem in (3.10) and its analytic solution. In Section 3.4, we verify that our proposed algorithmic approach is able to accurately converge to the known solutions under standard calibration.[15]

Breeden Model Our first example is a continuous-time case of the Breeden (1979) consumption-based asset pricing model.[16] Assume that the Markov state has two components $X_t = (X_t^o, X_t^f)$, where X_t^o is a real-valued Ornstein-Uhlenbeck process and X_t^f is a positive Feller square root process:

$$
\begin{aligned}
dX_t^o &= \xi_o(\bar{x}_o - X_t^o)dt + \sigma_o dW_t^o \\
dX_t^f &= \xi_f(\bar{x}_f - X_t^f)dt + \sqrt{X_t^f}\sigma_f dW_t^f,
\end{aligned}
$$

where $(\xi_i, \bar{x}_i)_{i\in\{o,f\}} > 0$ and $\sigma_o > 0$ are positive scalars, and $2\xi_f\bar{x}_f \geq \sigma_f^2$ to ensure that X_t^f is positive and stationary. In this model, the process X_t^o captures aggregate growth rates, while X_t^f captures aggregate market volatility. Suppose that the equilibrium dynamics of log consumption are as follows:

$$
d\log(C_t) = X_t^o dt + \vartheta_o dW_t^o + \sqrt{X_t^f}\vartheta_f dW_t^f, \tag{3.11}
$$

[15] See Appendix 3.8, for model calibrations

[16] For more details, see also example 3.8 in Hansen and Scheinkman (2009).

where $\vartheta_o, \vartheta_f > 0$ and investors have power utility:

$$U((C_t)_{t\geq 0}) = \mathbb{E}_0 \int_0^\infty \exp(-bt) \frac{C_t^{1-a}}{1-a} dt,$$

for constants $a, b > 0$. Then the implied stochastic discount factor thus has the following dynamics:

$$d\log S_t = (-aX_t^o - b)dt - a\vartheta_o dW_t^o - a\sqrt{X_t^f}\vartheta_f dW_t^f.$$

Kreps-Porteus Model We also consider a special case of the Kreps and Porteus (1978) with Duffie and Epstein (1992) stochastic discounted utility, which leads to the same long-run risk formulation.[17] In this setup with time separable logarithmic utility, log-consumption is related to the investor continuation process W which satisfies:

$$\frac{dW_t}{dt} = b(W_t - \log(C_t)).$$

HS prove that W_t takes the following form:

$$W_t = \frac{1}{1-a} \exp\left[(1-a)(w_f X_t^f + w_0 X_t^o + c_t + \bar{w})\right].$$

Moreover, HS solve this model assuming that consumption follows the same dynamics as in (3.11), and show that the SDF is the product of two multiplicative functionals (i.e., $\log S_t = A_t^B + A_t^w$), where the first functional A_t^B takes the exact same form as the Breeden model:

$$A_t^B = -\int_0^t X_s^o ds - bt - \int_0^t \sqrt{X_s^f}\vartheta_f dB_s^f - \int_0^t \vartheta_o dB_s^o.$$

[17] See e.g., Duffie and Epstein (1992), Schroder and Skiadas (1999) and Hansen and Scheinkman (2009) for details.

The second term, exponential of A_t^w, is a martingale constructed from the forward-looking continuation value process implied by the recursive utility function:

$$A_t^w = (1-a)\int_0^t \sqrt{X_s^f}(\vartheta_s^f + w_f\sigma_f)dB_s^f + (1-a)\int_0^t (\vartheta_o + w_o\sigma_o)dB_s^o$$
$$-\frac{(1-a)^2}{2}\int_0^t X_s^f(\vartheta_f + w_f + \sigma_f)^2 ds - \frac{(1-a)^2(\vartheta_o + w_o\sigma_o)^2}{2}t.$$

For both the Breeden and Kreps-Porteus models mentioned above, HS show that the principal long-run eigenfunction $e : \mathbb{R}^2 \to \mathbb{R}_+$ implied by (3.10) takes the form:

$$e(x) = \exp(c_o x_o + c_f x_f), \tag{3.12}$$

where the constants c_o and c_f are given by:

$$c_o = \frac{-a}{\xi_o}, \qquad c_f = \frac{(\xi_f + \vartheta_f a\sigma_f) \pm \sqrt{(\xi_f + \vartheta_f a\sigma_f)^2 - \sigma_f^2(\vartheta_f a)^2}}{\sigma_f^2}, \tag{3.13}$$

and corresponding long-run eigenvalue is given by:

$$\lambda = -b + \frac{1}{2}(a\vartheta_o)^2 + c_f\xi_f\bar{x}_f + c_o(\xi_o\bar{x}_o - a\vartheta_o\sigma_o) + \frac{1}{2}c_o^2\sigma_o^2. \tag{3.14}$$

To summarize, the pricing operator corresponding to (3.10) is given by:

$$\mathbb{S} = \sigma\gamma\nabla + 0.5|\gamma|^2 + \beta, \tag{3.15}$$

where $\gamma(x) = (-\vartheta_f a\sqrt{x_f}, -\vartheta_o a)$ and $\beta(x) = -ax_o - b$.

The Long-Run Risk Model As another example, we consider the structural asset pricing model from Bansal and Yaron (2004), which features growth rate predictability and stochastic volatility in aggregate consumption. Borovička et al. (2016) solve for this model analytically and show that it implies a large martingale component, and is thus a good benchmark model for our numerical approach.[18] Assume that the economy is driven by a (2×1) vector of state variables $(X_t)_{t\geq 0}$, where the terms represent the predictable component of the growth rate and stochastic volatility. The state variables evolve according to the Itô process

$$dX_t = \mu_x(X_t)dt + \sigma_x(X_t)dW_t,$$

where $(W_t)_{t\geq 0}$ is a 3-dimensional standard Brownian motion. Following Hansen and Scheinkman (2009), we parameterize the model by

$$\mu(x) = \bar{\mu}(x-\iota), \qquad \sigma(x) = \sqrt{x_2}\bar{\sigma},$$

where

$$\bar{\mu} = \begin{bmatrix} \bar{\mu}_{11} & \bar{\mu}_{12} \\ 0 & \bar{\mu}_{22} \end{bmatrix}, \quad \bar{\sigma} = \begin{bmatrix} \bar{\sigma}_1 \\ \bar{\sigma}_2 \end{bmatrix}, \quad \iota = \begin{bmatrix} \iota_1 \\ \iota_2 \end{bmatrix},$$

and $\bar{\sigma}_1, \bar{\sigma}_2$ are both (1×3) row vectors and ι is the vector of means in a stationary distribution. The dynamics imply that X_t is a mean-reverting Markov process. In addition, the stochastic discount factor process $(S_t)_{t\geq 0}$ evolves according to

$$d\log S_t = -\delta dt - d\log C_t + d\log H_t,$$

where δ is a time discount factor, C_t is aggregate consumption and H_t is a continuous-time martingale constructed from the forward-looking continuation value of a representative agent with recursive homothetic preferences (see Borovička et al. (2016) for further details). Borovička et al. (2016) show that the SDF inherits the functional form

$$d\log S_t = \beta_s(X_t)dt + \alpha_s(X_t)dW_t,$$

[18]This model is first developed in Hansen et al. (2007). Our benchmark results rely on the calibrated consumption dynamics from the original paper.

where

$$\beta_s(x) = \beta_{s,0} + \beta_{s,1} \cdot (x - \iota)$$
$$\alpha_s(x) = \sqrt{x_2}\bar{\alpha}_s.$$

The scalars $\beta_{s,0}, \bar{\alpha}_s$ and vector $\beta_{s,1}$ can be constructed from the parameters driving the consumption process (Borovička et al., 2016, Appendix D). The resulting valuation operator can be written as follows:

$$\mathbb{S}e = \left(\beta_s(x) + \frac{1}{2}|\alpha_s(x)|^2\right)e + \nabla e \cdot (\mu(x) + \sigma(x)\alpha_s(x)) + \frac{1}{2}\mathrm{Tr}\left(\mathrm{Hess}(e)\sigma(x)\sigma(x)^\top\right) \tag{3.16}$$

Borovička et al. (2016) show that when the principal eigenfunction is constrained to be positive, there exists a unique eigensolution corresponding to the smallest eigenvalue. In particular, this solution, uniquely characterizes the measure which renders the $(X_t)_{t\geq 0}$ process stationary after a change of measure. The solutions to this eigenvalue problem are well-defined and are characterized as follows:

$$\hat{e}(x) = \exp\left(\bar{e}_1 x_1 + \bar{e}_2 x_2\right), \tag{3.17}$$

where our coefficients are determined by a system of equations (see Appendix 3.7.1 for a detailed derivation). Since these models can be solved analytically, we can use them to validate our proposed methodology before applying more broadly to asset pricing models that do not have well-known solutions. We discuss our proposed approach in the next section.

3.3 Methodology

3.3.1 Setup and Objective Function

In this section, we develop a non-parametric stochastic approach to solving the eigenvalue problem in (3.10). To do so, we start by recasting the problem in terms of a partial differential equation. Recall that the linear pricing operator $\mathbb{S} : C^2(\mathbb{R}^d) \to C(\mathbb{R})$ depends on the underlying dynamics of the Markov state process. The general form is written as follows:

$$\mathbb{S}e(x) = f(x) \cdot e(x) + \nabla e(x) \cdot b(x) + \frac{1}{2} \operatorname{Tr}\left(\sigma(x)\sigma(x)^\top \operatorname{Hess} e(x)\right), \tag{3.18}$$

where $f : \mathbb{R}^d \to \mathbb{R}$ is the stationary component, $b : \mathbb{R}^d \to \mathbb{R}^d$ is the drift component, and $\sigma\sigma^\top : \mathbb{R}^d \to \mathbb{M}_{d\times d}(\mathbb{R}^d)$ is the volatility component.[19] Without loss of generality, we restrict the eigenfunction $e \in C^2(\mathbb{R}^d)$ to a compact domain $\Omega \subset \mathbb{R}^d$. Applying the Feynman-Kac Formula[20] , we augment a time-dimension and consider the following backward parabolic parabolic partial differential equation (PDE) in the interval $t \in [0,T]$:

$$\begin{cases} \partial_t u(x,t) = -\mathbb{S}u(x,t) + \lambda u(x,t) & t \in [0,T], x \in \Omega \\ u(T,x) = e(x) & x \in \Omega \end{cases} \tag{3.19}$$

We can write solutions of (3.19) in terms of a backward propagation operator $(\mathcal{P}_t^\lambda)_t$ such that $\mathcal{P}_t^\lambda e(x) = u(T-t,x)$.[21] This formulation guarantees that if $e(.)$ is a stationary solution of (3.19) satisfying $\mathcal{P}_T^\lambda = u(T,x) = e(x)$, then (λ, e) must be an eigenpair of $\mathbb{S}$ (i.e., $\mathbb{S}e = \lambda e$). In other words, eigenpairs of $\mathbb{S}$ are local minima of the fixed-point squared error loss function $||\mathcal{P}_T^\lambda e - e||^2$. Moreover, the results of HS guarantee that the smallest positive eigenvalue λ and its corresponding eigenfunction uniquely characterize

[19] We assume that the volatility component $\sigma(\cdot)\sigma(\cdot)^\top$ is uniformly elliptic.

[20] See e.g., Theorem 8.2.1 in Oksendal (1992).

[21] Note that $\mathcal{P}$ forms a semigroup. Han et al. (2020) argue that such a PDE formulation is a continuous time analog of the power method for discrete matrix eigenvalue problems.

the permanent-transitory decomposition in (3.6).

Like the eigenvalue problem in (3.10), this PDE is difficult or impossible to solve in closed form and deterministic numerical approaches are subject to the curse of dimensionality. Stochastic numerical approaches such as variational and diffusion Monte Carlo (see e.g., Ceperley et al. (1977), Blankenbecler et al. (1981), Zhang et al. (1997), Foulkes et al. (2001), Needs et al. (2009)) are popular options for solving this problem in domains with two sided boundary conditions and strict parametric restrictions on the eigenfunction (e.g., trigonometric). Our asset pricing applications do not fit this narrow structure.

Recent work such as Han et al. (2019) and Pfau et al. (2020) introduce neural networks in the variational Monte Carlo approach to overcome these parametric limitations, but still in a very specific use case. Along these lines, recent work from E and Yu (2017) and Han et al. (2020) develop more general frameworks which combines the stochastic elements of Monte Carlo approaches with the flexibility of deep neural networks.[22] We adapt this class of approaches for the asset pricing setting.[23]. Intuitively, although differential equations are difficult to solve in high dimensions, we leverage the fact that stochastic processes are fairly easy to simulate (See e.g., Cochard (2019)). We consider the following Itô process:

$$X_t = X_0 + \int_0^t \mu(X_s)\,\mathrm{d}s + \int_0^t \sigma(X_s)\,\mathrm{d}W_s, \tag{3.20}$$

where W_t is a standard Brownian motion. Applying Itô's Lemma, the solution to (3.19) satisfies:

$$\begin{aligned} u(t,X_t) = u(0,X_t) + \int_0^t \big(\lambda u(s,X_s) - f(X_s)u(s,X_s) - (b(X_s) - \mu(X_s))\nabla u(s,X_s)\big)\,\mathrm{d}s \\ + \int_0^t \nabla u(s,X_s)^\top \sigma(X_s)\,\mathrm{d}W_s. \end{aligned} \tag{3.21}$$

Since the eigenfunction $e(x)$ and eigenvalue λ are a priori unknown, we cannot directly simulate the process in (3.19). Universal approximation properties of neural networks (see

[22]Hornik et al. (1989) demonstrates the universal approximation properties of neural networks. See e.g., Bauer and Kohler (2019) for more discussion on overcoming the curse of dimensionality with DNNs.

[23]Table 3.3 compares advantages and disadvantages of our proposed approach relative to other numerical benchmarks from the literature

e.g., Hornik et al. (1989)) implies that we can approximate eigenfunction arbitrarily closely with a fully-connected feed-forward neural network with a finite number of parameters $\theta \in \Theta \subset [-1,1]^N, N < \infty$, without loss of generality. We denote this approximation by $\mathcal{N}_e(x;\theta)$.[24] At a high level, we propose an iterative algorithm which starts by randomly initializing the eigenvalue $\hat{\lambda}^{(0)} \sim \text{Uniform}[0,\bar{\lambda}]$ and neural network parameters $\hat{\theta}^{(0)} \sim \text{Uniform}[-1,1]^N$ (see e.g., Glorot and Bengio (2010)) for uniform distributions.[25] We simulate approximate sample paths from (3.5) and iteratively update estimated parameters $(\hat{\theta}^{(s)})_{s=1,2,...,S}$ and $\hat{\lambda}^{(s)}_{s=1,2,...,S}$ for iterations $s = 1,2,...,S$ using gradient descent (see e.g., Kingma and Ba (2017)). To converge to a valid eigenpair of $\mathbb{S}$, our objective is to minimize the following fixed-point squared error loss:

$$\min_{\hat{\theta}} \left|\left|\mathcal{N}_e(x;\hat{\theta}) - u(T,X_T|\hat{\theta})\right|\right|_2^2, \tag{3.22}$$

where the simulated process $u(t,x;\hat{\theta})$ also depends on the estimated neural network (eigenfunction) parameters. The loss function in (3.22) attains a local minimum at stationary solutions of the propagator $\mathcal{P}_T(\hat{\lambda})$. Although this optimization problem is non-convex, we know that there exists an initialization $(\hat{\lambda}^{(0)},\hat{\theta}^{(0)})$ such that this iteration we converges to the true principal stationary solution of $\mathcal{P}_T^{\lambda}$.

3.3.2 Solution Algorithm

Our proposed algorithm is implemented as a numerical analogue to the the fixed point optimization defined in (3.22). For numerical tractability, we consider a discrete partition of the interval T into N sub-intervals with $\Delta t = \frac{T}{N}$ and $t_n = \frac{nT}{N}$ $(n = 0,1,\ldots,N)$. We apply an Euler-Maruyama (E-M) scheme (see e.g., Lamba et al. (2006)) to discretize

[24] In practice, we compute the gradient of $\mathcal{N}_e(x;\theta)$ with respect to θ numerically. For very high dimensions, numerical differentiation may not scale well, and thus we can initialize a second neural network to approximate the gradient.

[25] Since λ corresponds to an infinitely long bond yield, we know it can be bounded by 1 ($\bar{\lambda} < 1$), although we can further refine this bound when considering specific problems. In practice, we consider several initial conditions with extensive cross-validation to ensure our solution is not sensitive to the choice of $\bar{\lambda}$.

the Itô processes X_t and $u(t,X_t)$ as follows:

$$\begin{aligned}
\mathcal{X}_{t_{n+1}} &= \mathcal{X}_{t_n} + \mu(\mathcal{X}_{t_n})\Delta t + \sigma(\mathcal{X}_{t_n})\Delta W_n \\
\mathcal{U}_{t_{n+1}} &= \mathcal{U}_{t_n} + \big(\lambda \mathcal{U}_{t_n} - \mathcal{N}_{\sigma^\top \nabla e}^\top(\mathcal{X}_{t_n})\alpha(\mathcal{X}_{t_n}) - f(\mathcal{X}_{t_n})\mathcal{U}_{t_n}\big)\Delta t + \mathcal{N}_{\sigma^\top \nabla e}(\mathcal{X}_{t_n})^\top \Delta W_n,
\end{aligned} \tag{3.23}$$

where α is defined by $b = \mu + \sigma\alpha$. Note that that the two equations have the same realization of Brownian motion. Here, we denote by $\mathcal{X}_{t_n}$ and $\mathcal{U}_{t_n}$ the discrete analogues of the continuous processes X_t and $u(t,X_t)$, respectively. $\Delta W_n \sim N(0, \frac{T}{N} I_d)$ are independent samples of Brownian increments. For a tractable initialization, we uniformly sample $\mathcal{X}_0 \sim \mathrm{Unif}(\Omega)$. At each iteration of the algorithm $s = 1,2,...,S$, we sample K trajectories and denote each trajectory by $\mathcal{X}_{t_n}^{(k)}$ and $\mathcal{U}_{t_n}^{(k,s)}$ based on the estimated parameters $(\hat{\theta}^{(s)}, \hat{\lambda}^{(s)})$. To avoid the trivial (zero) solution, we define the following normalization factor:

$$\widehat{Z}(\mathcal{N}_e, \mathcal{X}_0; \hat{\theta}) = \left(\frac{1}{K} \sum_{k=1}^{K} \mathcal{N}_e(\mathcal{X}_0^{(k)}; \hat{\theta})^2 \right)^{\frac{1}{2}}. \tag{3.24}$$

Here, the notation $\mathcal{X}_0$ denotes the collection of $\{\mathcal{X}_0^{(k)}\}_{k=1}^K$. Then the initialization for $\mathcal{U}$ in the E-M scheme (3.23) is

$$\mathcal{U}_0^{(k)} = \frac{\mathcal{N}_e(\mathcal{X}_0^{(k)}; \hat{\theta}, \hat{\lambda}}{\widehat{Z}(\mathcal{N}_e, \mathcal{X}_0; \hat{\theta}, \hat{\lambda})}. \tag{3.25}$$

To update estimated neural network parameters $\hat{\theta}$ and estimated eigenvalue $\hat{\lambda}$, we minimize the discrete analogue of (3.22) based on the simulated processes of $\mathcal{X}_T^{(k)}$ and the estimated eigenfunction:

$$\widehat{L}(\theta, \lambda) = \frac{1}{K} \sum_{k=1}^{K} \left(\frac{\mathcal{N}_e(\mathcal{X}_T^{(k)}, \theta)}{\widehat{Z}(\mathcal{N}_e, \mathcal{X}_0, \theta)} - \mathcal{U}_T^{(k)}(\theta, \lambda) \right)^2. \tag{3.26}$$

In practice, introduce an absolute value in the final layer of $\mathcal{N}_e$ to ensure non-negativity of the solution. Furthermore, we optimize (3.26) using stochastic gradient descent as in Kingma and Ba (2017) with early stopping as in Shen et al. (2022). We consider multiple initialization of our neural network approximation (sampled as in Glorot and Bengio (2010)), and we use cross-validation to verify that the converged solution is robust to the choice of algorithm hyperparameters. We provide pseudocode in Algorithm 1 below.

Algorithm 1: Valuation Eigensolver

Input: Valuation operator $\mathbb{S}$, number of time intervals N, terminal time T, number of iterations

Hyperparameters: batch size K, loss hyperparameters γ, learning rate, network architecture

Output: Function approximation of smallest non-negative eigenfunction $\mathcal{N}_e(x)$ and corresponding eigenvalue λ.

Randomly initialize eigenvalue λ, neural network $\mathcal{N}_e$, and $\mathcal{N}_{\sigma^\top \Delta_e}$

while $i <$ *max iterations* ***or*** $\Delta loss < \varepsilon$: **do**

- Sample K initial points X_0 and K discretized Brownian motion paths
- Simulate discrete U process
- Compute gradient of the empirical loss (3.26) with respect to neural network parameters and eigenvalue
- Update parameters using Adam optimizer

end

Return $(\mathcal{N}_e, \lambda)$

3.4 Numerical Results

In this section, we validate the performance of our proposed Valuation Eigensolver in approximating the long-run components in the models from Section 3.2.2. Under baseline algorithm hyperparameters and asset pricing model calibrations detailed in Appendix 3.8, we can obtain closed form expressions for the long-run eigenfunction and eigenvalue of the two sets of models. We measure the accuracy of our eigenfunction approximation in terms of root mean squared approximation error (RMSE):

$$\widehat{RMSE} = \left[\frac{1}{K}\sum_{k=1}^{K}\left(\frac{\mathcal{N}_e(X_k;\hat{\theta})}{\hat{Z}(\mathcal{N}_e;\hat{\theta})} - e(X_k)\right)^2\right]^{1/2}, \tag{3.27}$$

where $\mathcal{N}_e(.;\hat{\theta})$ is our neural network approximation with normalization $\hat{Z}(\mathcal{N}_e;\hat{\theta})$ and learned parameters $\hat{\theta}$, $e(.)$ is the true function implied by the model, and $(X_k)_{k=1}^K$ is a set of points sampled from a uniform distribution on the domain of the eigenfunction. To ensure that our learned approximation is stable at all points in the domain, we also report the maximum approximation error:

$$\hat{L}^\infty = \max_k \left| \frac{\mathcal{N}_e(X_k;\hat{\theta})}{\hat{Z}(\mathcal{N}_e;\hat{\theta})} - e(X_k) \right|. \tag{3.28}$$

For the eigenvalue approximation, we report accuracy in terms of absolute errors ($\hat{\Delta}_\lambda = |\hat{\lambda} - \lambda|$). We compare performance of our proposed algorithm to two benchmark alternatives, one deterministic and one stochastic. The deterministic benchmark is a computationally optimized version of finite elements called stabilized multi-scale finite element (MSFE) (see e.g., Pichler et al. (2013)), which requires manually specified boundary conditions. Our other benchmark is the stochastic Deep-Ritz approach developed in E and Yu (2017), which also requires us to specify at least a one-sides boundary condition to converge. We are not aware of any other approaches which can be applied off-the-shelf in this class of problems.

Table 3.1 reports the error metrics for the proposed approach relative to the above benchmarks. The Valuation Eigensolver achieves a better eigenfunction approximation without any specified boundary conditions, achieving a worst-case approximation error of 2% across both model specifications, compared to about 4% in the two competing approaches. Similarly, the Valuation Eigensolver approximates the eigenvalue within 1% for both the Breeden/Kreps-Porteus and Bansal-Yaron specifications, with an absolute approximation error of 0.02 and 0.03 bps, respectively. Table 3.2 reports approximated vs. actual model-implied annualized eigenvalue, with values of 0.2% and 0.4% per year for the Breeden/Kreps-Porteus and Bansal Yaron models, respectively.

We further verify the robustness of our converged solution. Figure 3.1 plots the density of the approximated eigenfunctions on a uniform sample from the domain.

Notice that our approximation is stable across the entire domain of values, including the boundaries, even without explicitly specifying boundary conditions in any step of the solution algorithm. Additionally, Figures 3.6 and 3.4 verify that the converged solution is stable in both dimensions by comparing the approximated and actual eigenfunction densities when one dimension is fixed at its stationary mean. In Appendix 3.9, we also verify the robustness of the convergence under multiple specifications of initial value and model parameters.

Table 3.1: Eigenfunction and Eigenvalue Approximation Errors

A. Breeden/Kreps-Porteus Models

	$\hat{L}^2$	$\hat{L}^\infty$	$\hat{\Delta}_\lambda$ (bps)
Valuation Eigensolver	0.005017	0.018248	0.014061
DeepRitz	0.009963	0.044081	0.019319
MSFE	0.009159	0.039357	0.019422

B. Bansal-Yaron Long-Run Risk Model

	$\hat{L}^2$	$\hat{L}^\infty$	$\hat{\Delta}_\lambda$ (bps)
Valuation Eigensolver	0.004987	0.020841	0.031218
DeepRitz	0.010094	0.037994	0.049837
MSFE	0.008968	0.034615	0.044381

Notes: We report approximation errors of our proposed Eigensolver relative to MSFE and DeepRitz. Panel A reports results for the Breeden and Kreps-Porteus model, while Panel B reports results for the Bansal-Yaron model. The first two columns report the $\widehat{RMSE}$ and $\hat{L}^\infty$ of the estimated eigenfunction on a set of 10,000 points sampled uniformly from the state domain Ω, where $\Omega = [0, 0.02] \times [0.08, 0.72]$ for the Breeden model and $\Omega = [-0.006, 0.006] \times [0.6, 1.6]$ for Bansal-Yaron model. Both models are calibrated to the monthly frequency. The final column reports the absolute approximation error of the eigenvalue, whose units are reported in basis points (bps).

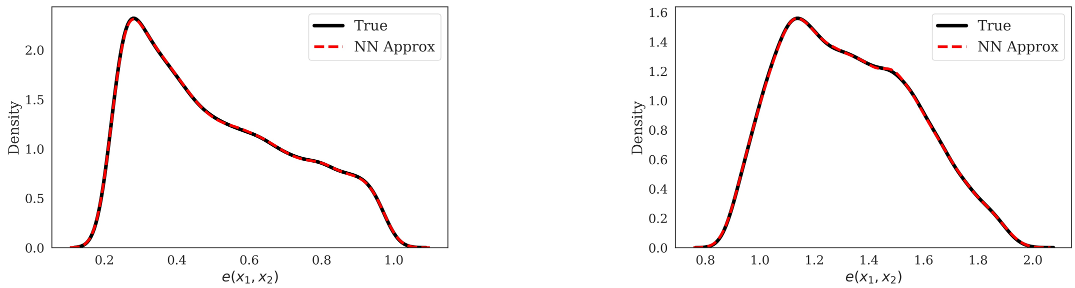

(a) Breeden / Kreps-Porteus (b) Bansal-Yaron

Figure 3.1: Approximated vs. Actual Eigenfunction Density

Notes: We plot the density of the actual model-implied and approximated eigenfunctions, evaluated on a set of 10,000 points sampled uniformly from the state domain Ω, where $\Omega = [0, 0.02] \times [0.08, 0.72]$ for the Breeden/Kreps-Porteus model and $\Omega = [-0.006, 0.006] \times [0.6, 1.6]$ for Bansal-Yaron model. Both models are calibrated to the monthly frequency.

Table 3.2: Approximated vs. Actual Long-Run Yields

	$\hat{\lambda}$	λ
Breeden/Kreps-Porteus	0.00189	0.00187
Bansal-Yaron	0.00375	0.00379

Notes: We report the annualized approximated vs. actual eigenvalues $\hat{\lambda}$ and λ for both the Breeden / Kreps-Porteus model and the Bansal-Yaron model.

3.5 Extension to Models with Jumps

In the previous section, we demonstrate the accuracy of the Valuation Eigensolver in solving for the long-run components of the SDF in two common structural macro-finance models. In particular, the approximated eigenvalue corresponds to the estimated yield of an infinite horizon bond, the permanent component of the SDF. Christensen (2017) provides the first empirical estimates of this quantity under various assumptions about investor discounting and risk-aversion, and obtains a lower bound of around 1% per quarter. This is much larger than our model implied approximations of about 0.2% and 0.4% per year. Our findings are thus consistent with the critique raised Alvarez and Jermann (2005) and Backus et al. (2014) that consumption-based asset pricing models tend to lead to SDFs which underestimate the magnitude of the permanent component observed in the data. Our

Valuation Eigensolver provides a path forward to reconciling model predictions with the data, namely, by solving for these components in models whose underlying state dynamics have jumps.

Recall the aggregate growth process in the Breeden/Kreps-Porteus specification, which we originally assumed to follow an Ornstein-Uhlenbeck process. Following Duffie et al. (2000), we can augment an affine jump component as follows:

$$dX_t^o = \xi_o(\bar{x}_o - X_t^o)dt + \sigma_0 dW_t^o + dZ_t, \tag{3.29}$$

where Z_t is a pure jump process whose jumps have a fixed probability distribution ν and arrive with intensity $\Lambda(x) = \bar{\omega}_1 x_f + \bar{\omega}_2$ with $\bar{\omega}_1, \bar{\omega}_2 \geq 0$, and can be modelled by the function $\kappa(y,x) = \bar{\kappa}(y_o - x_o)$ such that $\int_{\mathbb{R}}^d \exp\bar{\kappa}(z)dv(z) < \infty$. The implied SDF still takes the multiplicative form (i.e., $S_t = \exp(A_t)$), but in this case A_t includes another term:

$$A_t = \bar{\beta}t + \int_0^t \beta_o X_s^o ds + \int_0^t \beta_f X_s^f ds + \int_0^t \gamma_o dW_s^o + \int_0^t \sqrt{X_s^f}\gamma_f dW_s^f + \sum_{s \leq t} \bar{\kappa}(X_s^o - X_{s-}^o), \tag{3.30}$$

Hansen and Scheinkman (2009) show that this also adds an additional terms to the eigenvalue, as follows:

$$\lambda = \bar{\beta} + \frac{\gamma_o^2}{2} + c_f\xi_f\bar{x}_f + c_o(\xi_o\bar{x}_o + \gamma_o\sigma_o) + c_o^2\frac{\sigma_o^2}{2} + \bar{\omega}_2 \int \left[\exp\frac{\beta_0}{\xi_o}z - 1\right]\exp[\bar{\kappa}(z)]d\nu(z). \tag{3.31}$$

This provides a path forward for us to amplify the model predictions of the long-run component to be more consistent with the empirical estimates provided in Christensen (2017). Although we are still able to obtain a closed-form solution in this case, Hansen and Scheinkman (2009) emphasize the importance of developing a general computational approach that can be applied outside the limited array of such examples that can be solved. Before we demonstrate that our approach continues to work in this affine jump setting, we also argue that this is enough to generalize our approach to a very general class of Lévy process. Infinite jump processes cannot be simulated exactly on a pre-specific grid of time increments (see e.g., Fournier (2009)). To address this, we rely on recent results from Gnoatto et al. (2022), who show that we can approximate an infinite jump diffusion with a diffusion that has finitely many jumps. The residual small jumps can either be truncated or approximated by increasing the volatility of the diffusion component.

First, note that the updated SDF in (3.30) implies the following operator:

$$
\begin{aligned}
\mathbb{S}e(x) = \frac{1}{2}\operatorname{Tr}\left(\operatorname{Hess}(e)\sigma(x)\sigma(x)^{\top}\right) + \nabla e \cdot (\mu(x) + \sigma(x)\alpha_s(x)) \\
+ \left(\beta_s(x) + \frac{1}{2}|\alpha_s(x)|^2\right) e + \Lambda(x)\int (e(x+z) - e(x))\exp(\bar{\kappa}(z))\,\mathrm{d}\nu(z),
\end{aligned}
\tag{3.32}
$$

where the last term above comes from the jump. Cohen and Rosiński (2007) show that we can approximate this integration accurately as long as $\int_{\mathbb{R}} |z|^2 dv(z) < \infty$. For a jump diffusion with infinitely many jumps, our approach remains the same but with additional approximation error. To see this, we take $\varepsilon \in (0,1]$ and decompose the Poisson measure as the sum of two components, one corresponding to big jumps and one corresponding to small jumps:

$$
v(dz) = \underbrace{\mathbb{1}_{|z|>\varepsilon} v(dz)}_{v^{\varepsilon}(dz)} + \underbrace{\mathbb{1}_{|z|\leq\varepsilon} v(dz)}_{v_{\varepsilon}(dz)}. \tag{3.33}
$$

We then approximate the big jump component v^{ε} with a finite jump affine process, as above. Jum (2015) bounds the approximation error induced by excluding small jumps and shows that it can be made arbitrarily small by decreasing $\varepsilon \downarrow 0$.[26]

We leverage this approximation extend the Valuation Eigensolver to general Lévy processes with finite or infinite jumps. We continue with the extended Breeden/Kreps-Porteus specification from (3.29). We start by simulating the Markov state $X_t = (X_t^o, X_t^f)$, accounting for the new jump term in X_t^o. Applying Itô's Lemma, we rewrite to capture the

[26]See Proposition 4.1.2.

additional integration term over the jump measure:

$$u(t,X_t) = u(0,X_t) + \int_0^t \nabla u(s,X_s)^\top \sigma(X_s)\,\mathrm{d}W_s + \int_0^t \Big[(\lambda - \beta_s(X_s) - \frac{1}{2}|\alpha_s(X_s)|^2)u(s,X_s)$$
$$- \lambda(X_s)\int (u(s,X_s+z) - u(s,X_s))\exp(\bar{\kappa}(z))\,\mathrm{d}\nu(z)\Big]\,\mathrm{d}s. \tag{3.34}$$

Recall that (3.34) is the solution to the backward parabolic PDE (3.19). Let $N(t_n)$ denote the number of jumps in the state process occurring in the interval $[0,t_n]$. In this example, $N(t_n)$ is zero in the second dimension (i.e., $\kappa(0,x) = 0$ for any x). We then update the Euler-Maruyama discrete representation to include the additional integration term:

$$\mathcal{X}_{t_{n+1}} = \mathcal{X}_{t_n} + \mu(\mathcal{X}_{t_n})\Delta t + \sigma(\mathcal{X}_{t_n})\Delta W_n + \sum_{i=N(t_n)+1}^{N(t_{n+1})} \bar{\kappa}(\mathcal{X}_{t_n}, z_i) - \Delta t \int_{\mathbb{R}} \bar{\kappa}(\mathcal{X}_{t_n}, z)\nu(dz)$$
$$\mathcal{U}_{t_{n+1}} = \mathcal{U}_{t_n} + \left(\eta\,\mathcal{U}_{t_n} - (\beta(\mathcal{X}_{t_n}) + \frac{1}{2}|\alpha(\mathcal{X}_{t_n})|^2)\,\mathcal{U}_{t_n}\right)\Delta t + \mathcal{N}_{\sigma^\top\nabla e}(\mathcal{X}_{t_n})^\top \Delta W_n$$
$$- \lambda(\mathcal{X}_{t_n}) I(\mathcal{X}_{t_n}; \mathcal{N}_e)\Delta t,$$

where $I(\mathcal{X}_{t_n}; \mathcal{N}_e)$ denote the integration with respect to measure $\nu(dz)$ of our neural network approximation evaluated at $\mathcal{N}_e$ at $\mathcal{X}_{t_n}$. In low dimensions, one option is to use a simple Monte Carlo integration to compute $I(\mathcal{X}_{t_n}; \mathcal{N}_e)$, but this can be computationally taxing as part of the training process. As an alternative, we follow Gnoatto et al. (2022) to approximate the entire integration term with a second family of neural networks $\mathcal{N}_I(\mathcal{X}_{t_n}; \theta_I)$:

$$\mathcal{N}_I(\mathcal{X}_{t_n}; \theta_I) \approx I(\mathcal{X}_{t_n}; \mathcal{N}_e) = \int_{\mathbb{R}} \mathcal{U}_n(\mathcal{X}_n + \bar{\kappa}(\mathcal{X}_n, z)) - \mathcal{U}_n(\mathcal{X}_n)\nu(dz), \tag{3.35}$$

which is trained as part of the learning process according to the fixed point loss in (3.22). We demonstrate the accuracy of the adapted version of the algorithm in approximating the solution from the above example.

We start by assuming that $\bar{\kappa}(z) = \bar{\kappa} \in \mathbb{R}$ is a constant function, and that the jump size distribution $\nu(dz) = \varphi(z)d(z)$ is Gaussian. Panel A in Figure 3.2 plots actual vs. approximated eigenvalues from (3.31) for different values of $\bar{\omega}_2 \geq 0$ and $\kappa \in \mathbb{R}$, while

Panel B plots the actual vs. approximated eigenfunction. In the example with jumps, the Eigensolver achieves an average eigenvalue approximation error of 2.6%, and eigenfunction approximation errors of $\hat{L}^2 = 3\%$ and $\hat{L}^\infty = 5\%$. This approximation error is slightly higher relative to the no-jump case, likely due to additional discretization error from approximating the integration term in (3.35). Panel A in Figure 3.2 also demonstrates that adequate parametrization of the additional jump term can reconcile the Breeden/Kreps-Porteus specification with the empirical estimates in Christensen (2017). See Appendix 3.9 for additional robustness checks on the algorithm convergence.

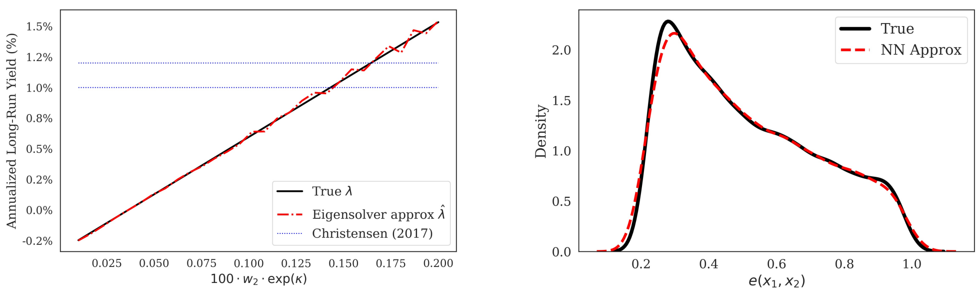

(a) Long-Run Yield

(b) Eigenfunction

Figure 3.2: Approximated vs. Actual Eigensolutions of Breeden/Kreps-Porteus model with jumps

Notes: Panel A plots the actual vs. Eigensolver approximation of λ from the Breeden/Kreps-Porteus model with jumps, (3.31) for different values of $\varpi_2 \cdot \exp \bar{\kappa}$, scaled by 100. Panel B plots the density of the actual model-implied and approximated eigenfunctions, evaluated on a set of 10,000 points sampled uniformly from the state domain Ω, where $\Omega = [0, 0.02] \times [0.08, 0.72]$. We fix $\bar{w}_1 = 0.001$ such that the eigenfunction boundary conditions are the same as the no jump case. We keep the same calibration from 3.8. To compute the actual eigenvalue, we approximate the addition integral term in (3.31) numerically numerical quadrature as in Piessens et al. (1983). The dotted blue lines in Panel A denote the range of semi-parametric empirical estimates for λ obtained in Christensen (2017).

3.6 Conclusion

In this paper, we propose a novel algorithm to compute the permanent-transitory components of stochastic discount factors in continuous time asset pricing models. We demonstrate the accuracy of our approximation and argue that the approach can be extended to a very general class of asset pricing models beyond those in which quasi-analytical solutions are available, including models which feature infinite jumps in the state dynamics. Future work could better develop the theoretical conditions under which our fixed point loss is minimized at the true value. Moreover, although our focus has been on the dominant eigenpair of valuation operators, our approach could likely be extended to solve for more refined characterizations of the long-run approximation which include non-dominant eigenpairs as well.

Appendices

3.7 Derivations

3.7.1 Bansal-Yaron Eigenfunction Coefficients from Section 3.2.2

We solve for the coefficients in (3.17) by providing the system of equations needed to pin down the unknowns in the model. The PDE that needs to be solved is given by

$$\left(\beta_s + \frac{1}{2}{\alpha_s}^2\right)\hat{e} + \hat{e}_x \cdot (\mu + \sigma\alpha_s) + \frac{1}{2}\mathrm{Tr}(\hat{e}_{xx}\sigma\sigma') - \eta e = 0.$$

Plugging the guess (3.17) into the PDE renders

$$\left(\beta_{s,0} + \bar{\beta}_1 \cdot (x - \iota) + \frac{1}{2}{\alpha_s}^2 - \eta\right)\hat{e} + \hat{e}\begin{bmatrix}\bar{e}_1\\ \bar{e}_2\end{bmatrix} \cdot \left(\begin{bmatrix}\bar{\mu}_{11}(x_1 - \iota_1) + \bar{\mu}_{12}(x_2 - \iota_2)\\ \mu_{22}(x_2 - \iota_2)\end{bmatrix} + x_2\begin{bmatrix}\bar{\sigma}_1\alpha_s\\ \bar{\sigma}_2\alpha_s\end{bmatrix}\right) \tag{3.36}$$

$$+\frac{1}{2}\hat{e}\left(\bar{e}_1^2\sigma_1\sigma_1' + 2\bar{e}_1\bar{e}_2\sigma_1\sigma_2' + e_2^2\sigma_2\sigma_2'\right) = 0. \tag{3.37}$$

Notice that the coefficients on x_1, x_2 all have to be zero, in addition to solving the system. For example in the equation above we get the coefficient of x_1

$$\bar{\beta}_{s,11}(x_1 - \iota_1) + \bar{e}_1\bar{\mu}_{11}(x_1 - \iota).$$

This gives rise to the parameter restriction

$$0 = \bar{\beta}_{s,11} + \bar{e}_1 \bar{\mu}_{11}.$$

The other restrictions are given by

$$\eta = \bar{\beta}_{s,0} - \bar{\beta}_{s,11}\iota_1 - \bar{\beta}_{s,12}\iota_2 - \bar{e}_1(\bar{\mu}_{11}\iota_1 + \bar{\mu}_{12}\iota_2) - \bar{e}_2\bar{\mu}_{22}\iota_2$$
$$0 = \bar{\beta}_{s,12} + \frac{1}{2}\alpha_s{}^2 + \bar{e}_1(\bar{\mu}_{12} + \bar{\sigma}_1\bar{\alpha}_s) + \frac{1}{2}\bar{e}_1^2\bar{\sigma}_1^2 + \bar{e}_2(\bar{\mu}_{22} + \bar{\sigma}_2\bar{\alpha}_s + \bar{e}_1\bar{\sigma}_1\bar{\sigma}_2') + \frac{1}{2}(\bar{e}_2)^2\bar{\sigma}_2^2.$$

The last equation is quadratic in $\bar{e}_2$ and we pick the solution that leads to the smallest eigenvalue $\hat{\eta}$.

3.7.2 Breeden/Kreps-Porteus Model with Jumps

Suppose that we augment the state variable X^o with an affine jump component as in Duffie et al. (2000). The process is no longer an Ornstein-Uhlenbeck process and instead satisfies:

$$dX_t^o = \xi_o(\bar{x}_o - X_t^o)dt + \sigma_o dB_t^o + dZ_t,$$

where Z is a pure jump process, whose jumps have compensator $\nu : \mathbb{R} \to [0,1]$ and arrive with intensity $\Lambda(x) = \bar{\omega}_1 x_f + \bar{\omega}_2$ with $\bar{\omega}_1 \geq 0$ and $\bar{\omega}_2 \geq 0$. The no-jump operator $\mathbb{S}$ now adds another term:

$$\begin{aligned}\mathbb{S}e(x) = {} & \frac{1}{2}\operatorname{Tr}\left(\operatorname{Hess}(e)\sigma(x)\sigma(x)^\top\right) + \nabla e \cdot (\mu(x) + \sigma(x)\alpha_s(x)) \\ & + \left(\beta_s(x) + \frac{1}{2}|\alpha_s(x)|^2\right)e + \underbrace{\Lambda(x)\int (e(x+z) - e(x))\exp(\bar{\kappa}(z))\,\mathrm{d}\nu(z)}_{\text{Jump term}}.\end{aligned}$$

For an eigenfunction of the form $e(x) = \exp(c_f x_f + c_o x_o)$, the additional jump term becomes:

$$(\bar{\omega}_1 x_f + \bar{\omega}_2)\exp(c_f x_f + c_o x_o)\int_{\mathbb{R}}[\exp(c_o z) - 1]\exp(\bar{\kappa}(z))d\nu(z).$$

As in the no-jump example we can write $c_o = \beta_o/\xi_o$, while c_f must satisfy the following equation:

$$\beta_f + \gamma_f^2/2 + c_f(\gamma_f\sigma_f - \xi_f) + c_f^2\sigma_f^2/2 + \bar{\omega}_1\int_{\mathbb{R}}\left(\exp(\beta_o z/\xi_o) - 1\right)\exp(\bar{\kappa})d\nu(z) = 0,$$

and the eigenvalue must satisfy:

$$\lambda = \bar{\beta} + \frac{\gamma_o^2}{2} + c_f\xi_f\bar{x}_f + c_o(\xi_o\bar{x}_o + \gamma_o\sigma_o) + c_o^2\frac{\sigma_o^2}{2} + \bar{\omega}_2\int \exp\frac{\beta_0}{\xi_o}z - 1\exp[\bar{\kappa}(z)]d\nu(z).$$

3.8 Algorithm Details

3.8.1 Comparison to Alternatives

Table 3.3: Comparison of Valuation Eigensolver to Alternatives

Algorithm	**Eigensolver***	DeepRitz	MSFE	VMC	DMC
Parametric Form	None	None	Smooth	Trig	Trig
Boundary Conditions	None	Required	Required	Required	Required
Nonlinear Operators	Yes	Yes	Yes	No	No
Dimensionality	High	High	$d \leq 2$	High	High
Theoretical Guarantees	None	None	Yes, optimal rate	No	No
Solves for Derivatives	Yes	No	No	Yes	No
Handles Jumps	Yes	No	No	No	No
Stochastic	Stochastic	Stochastic	Deterministic	Stochastic	Stochastic
Works for Asset Pricing	Yes	Yes, if BCs	Yes, if $d \leq 2$ + BCs	No	No

Notes: We compare our proposed Eigensolver to competing algorithms, including DeepRitz from E and Yu (2017), Multi-Scale Finite Elements from Pichler et al. (2013), the Variational Monte Carlo (VMC) approach from Han et al. (2019), and the Diffusion Monte Carlo Approach from Pfau et al. (2020).

3.8.2 Implementation Details and Hyperparameters

Eigensolver requires Pytorch v2.1.0 and Python 3.8. We implement gradient descent via Pytorch class Adam with manual learning rate decay and Pytorch Autograd class for automatic differentiation with gradient clipping. We implement iteration on a single NVIDIA V100 SMX2 GPU node with 284 GB DDR4 DRAM. We also introduce Wake/Sleep training for eigenvalue updating with a 3-step iteration cycle. Table 3.4 summarizes the implementation details, which are the same across all three models.

Table 3.4: Eigensolver Hyperparameters

Model	Bansal-Yaron	Breeden No Jumps	Breeden with Jumps
Hidden Dimension	40	40	40
Number of Layers	2	2	2
Learning Rate	[1e-2, 5e-4, 1e-4]	[1e-2, 5e-4, 1e-4]	[1e-2, 5e-4, 1e-4]
Learning Rate Boundaries	[300, 600]	[300, 600]	[300, 600]
Batch Size	1024	1024	1024
Number of Iterations	1000	1000	1000
Total Training Time (minutes)	0.05	0.1	0.19
Eigenvalue Initialization Range	[0,0.5]	[0,0.5]	[0,0.5]

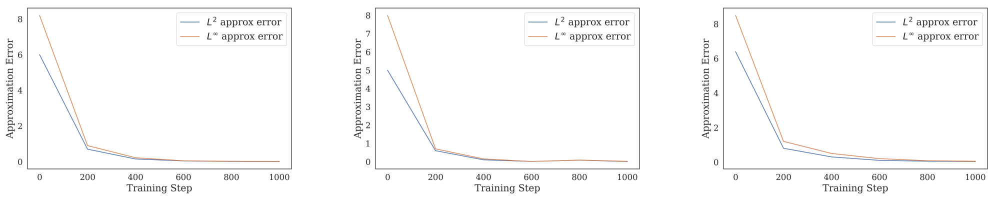

(a) Breeden (b) Bansal-Yaron (c) Breeden with jumps

Figure 3.3: Eigenfunction Training Curves

Notes: This figure plots the eigenvalue during training and the L^2/L^∞ errors of the eigenfunction every 200 iterations of the algorithm. We start both figures at iteration one.

3.8.3 Baseline Calibration

Table 3.5: Baseline Calibrated Parameters for Long-Run Risk Model

Parameter	**Calibration**
$\bar{\mu}_{11}$	-0.021
$\bar{\mu}_{12}$	0
$\bar{\mu}_{21}$	0
$\bar{\mu}_{22}$	-0.013
$\bar{\sigma}_1$	$[0, 0.00034, 0]$
$\bar{\sigma}_2$	$[0, 0, -0.038]$
ι_1	0
ι_2	1
$\bar{\beta}_{c0}$	0.0015
$[\bar{\beta}_{c1}, \bar{\beta}_{c2}]^\top$	$[1, 0]^\top$
$\bar{\alpha}_c$	$[0.0078, 0, 0]^\top$
δ	0.002
γ	10

Notes: We set baseline parameters for the long-run risk model based on the monthly calibration provided in Borovička et al. (2016) (see e.g., Figure 1 in that paper).

Table 3.6: Baseline Calibrated Parameters for Breeden and Kreps-Porteus Models

Parameter	Value
ξ_f	0.2
ξ_o	0.2
$\bar{x}_f$	0.2
$\bar{x}_o$	0.02
σ_f	-0.04
σ_o	0.2
a	5
b	$-\log(0.99)$
ϑ_f	1
ϑ_o	1

Notes: We calibrate the parameters in the Breeden and Kreps-Porteus to match aggregate asset returns data. Note that we must satisfy the restrictions $\xi_i, \bar{x}_i > 0$ for $i \in \{o, f\}$, $\sigma_o > 0$, and $2\xi_f \bar{x}_f \geq \sigma_f^2$ to ensure that X_t^f is positive and stationary. The corresponding parameters in the eigenfunction expression are $c_f = 0.4$ and $c_o = -2.11$.

3.9 Convergence Robustness

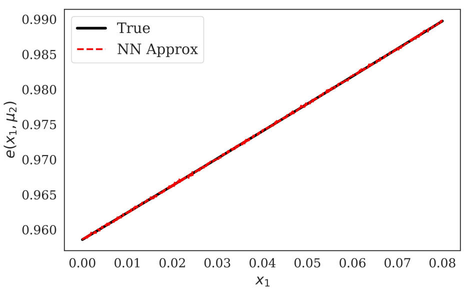

(a) $e(x_1, \bar{x}_2)$

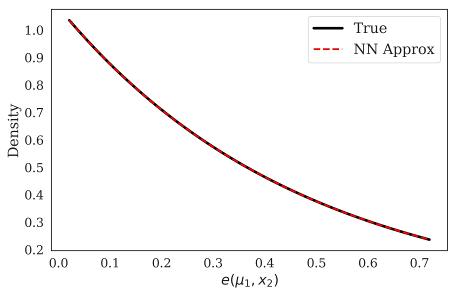

(b) $e(\bar{x}_1, x_2)$

Figure 3.4: Breeden/Kreps-Porteus Approximated vs. Actual Eigenfunction Density in 1D

Notes: In this figure, we plot the eigenfunction in 1-dimension, holding one state variable fixed at its stationary mean and varying the other state variable across its domain.

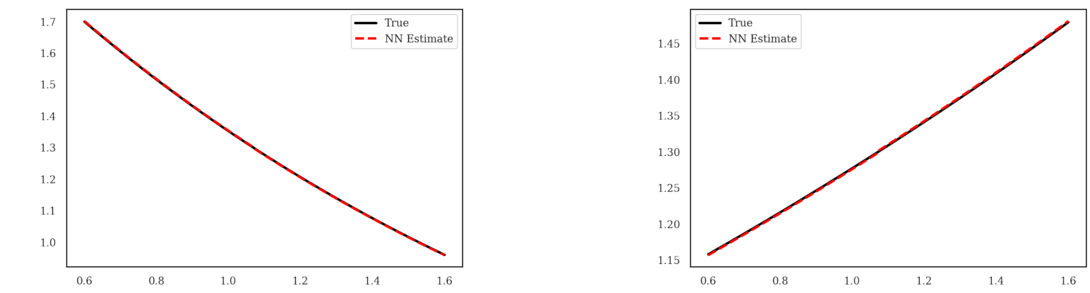

(a) $e(x_1, \bar{x}_2)$ (b) $e(\bar{x}_1, x_2)$

Figure 3.5: Bansal-Yaron Approximated vs. Actual Eigenfunction Density in 1D

Notes: In this figure, we plot the eigenfunction in 1-dimension, holding one state variable fixed at its stationary mean and varying the other state variable across its domain.

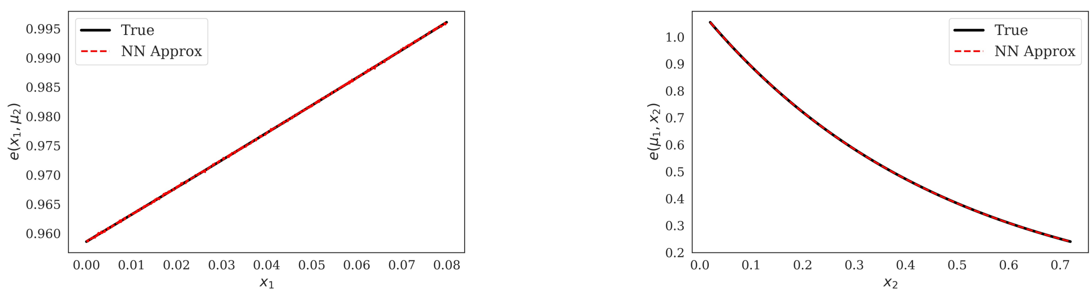

(a) $e(x_1, \bar{x}_2)$ (b) $e(\bar{x}_1, x_2)$

Figure 3.6: Breeden/Kreps-Porteus (with jumps) Approximated vs. Actual Eigenfunction Density in 1D

In this figure, we plot the eigenfunction in 1-dimension, holding one state variable fixed at its stationary mean and varying the other state variable across its domain.

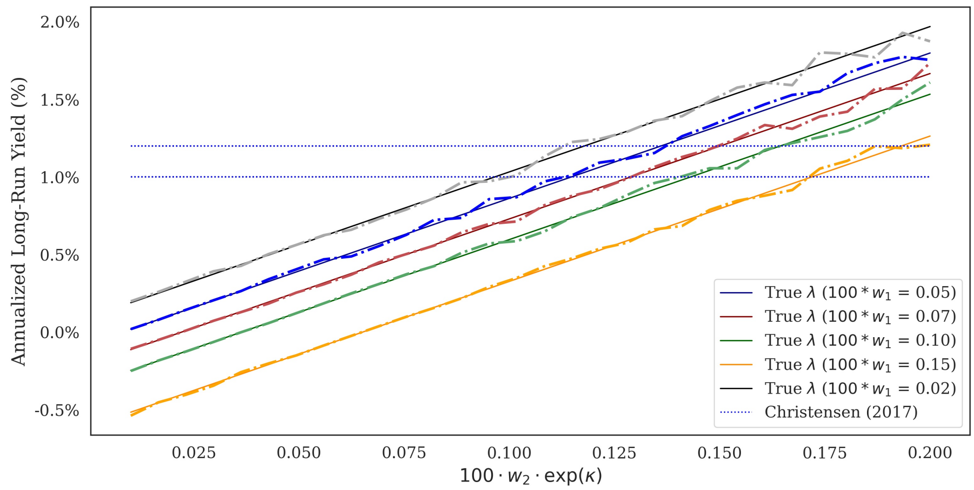

Figure 3.7: Approximated vs. Actual Eigenvalues of Breeden/Kreps-Porteus model with jumps

Notes: We plot the actual vs. Eigensolver approximation of λ from the Breeden/Kreps-Porteus model with jumps, (3.31) for different values of $\bar{\omega}_2 \cdot \exp \bar{\kappa}$, scaled by 100 and different values of $\bar{\omega}_1$. We keep the same calibration from 3.8. To compute the actual eigenvalue, we approximate the addition integral term in (3.31) numerically numerical quadrature as in Piessens et al. (1983). The dotted blue lines in Panel A denote the range of semi-parametric empirical estimates for λ obtained in Christensen (2017).

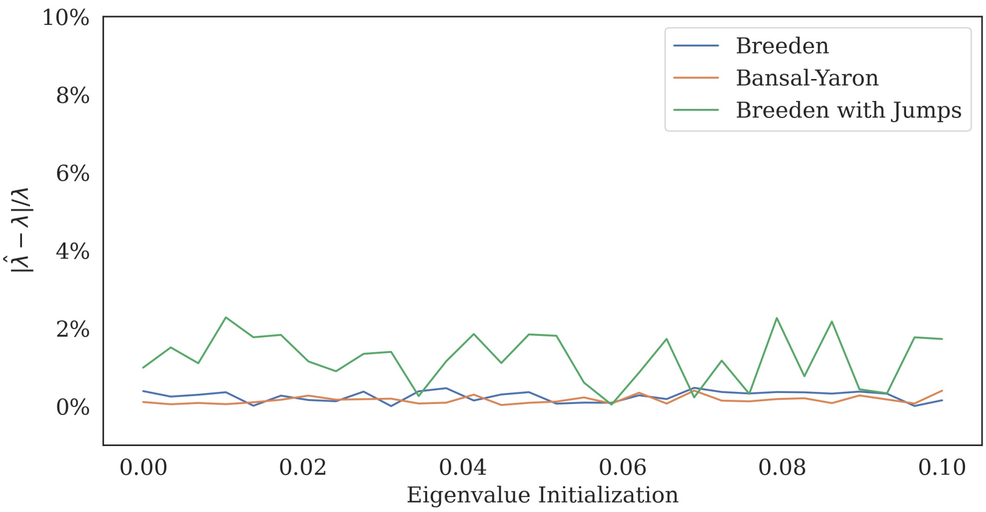

Figure 3.8: Eigenvalue Approximation Error under Various Initializations

Notes: We plot the Eigensolver approximation of λ from all three model specifications under various initial values of λ when training.

Victor Sellemi is the sole author of Chapter 1.

Chapter 2 is coauthored with Edoardo Briganti. The dissertation author is one of the primary investigators of this material.

Chapter 3 is coauthored with Tjeerd de Vries. The dissertation author is one of the primary investigators of this material.

Bibliography

D. Acemoglu, V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi. The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016, 2012a. doi: https://doi.org/10.3982/ECTA9623. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA9623.

D. Acemoglu, V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi. The Network Origins of Aggregate Fluctuations. *Econometrica*, 80(5):1977–2016, Mar. 2012b. ISSN 1468-0262. doi: 10.3982/ecta9623. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA9623.

D. Acemoglu, U. Akcigit, and W. Kerr. Networks and the macroeconomy: An empirical exploration. *NBER Macroeconomics Annual*, 30(1):273–335, 2016a. URL https://EconPapers.repec.org/RePEc:ucp:macann:doi:10.1086/685961.

D. Acemoglu, U. Akcigit, and W. Kerr. Networks and the Macroeconomy: An Empirical Exploration. *NBER Macroeconomics Annual*, 30(2015):63, 2016b. doi: https://doi.org/10.1086/685961.

K. R. Ahern. Network centrality and the cross section of stock returns. *Economics of Networks eJournal*, 2013.

F. Alvarez and U. J. Jermann. Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth. *Econometrica*, 73(6):1977–2016, November 2005. URL https://ideas.repec.org/a/ecm/emetrp/v73y2005i6p1977-2016.html.

A. Ang, R. J. Hodrick, Y. Xing, and X. Zhang. The cross-section of volatility and expected returns. *The Journal of Finance*, 61(1):259–299, 2006. doi: https://doi.org/10.1111/j.1540-6261.2006.00836.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.2006.00836.x.

P. Antràs, D. Chor, T. Fally, and R. Hillberry. Measuring the upstreamness of production and trade flows. *American Economic Review*, 102(3):412–416, May 2012. doi: 10.1257/aer.102.3.412. URL https://doi.org/10.1257/aer.102.3.412.

A. J. Auerbach, Y. Gorodnichenko, and D. Murphy. Local Fiscal Multipliers and Fiscal Spillovers in the USA. *IMF Economic Review*, 68(1):195–229, Mar. 2020. ISSN 2041-4161, 2041-417x. doi: 10.1057/s41308-019-00102-3. URL http://link.springer.com/10.1057/s41308-019-00102-3.

D. H. Autor, D. Dorn, and G. H. Hanson. The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review*, 103(6):2121–68, October 2013. doi: 10.1257/aer.103.6.2121. URL https://www.aeaweb.org/articles?id=10.1257/aer.103.6.2121.

D. Backus, M. Chernov, and S. Zin. Sources of entropy in representative agent models. *The Journal of Finance*, 69(1):51–99, 2014.

G. Bakshi and F. Chabi-Yo. Variance bounds on the permanent and transitory components of stochastic discount factors. *Journal of Financial Economics*, 105(1):191–208, 2012. ISSN 0304-405x. doi: https://doi.org/10.1016/j.jfineco.2012.01.003. URL https://www.sciencedirect.com/science/article/pii/S0304405X12000049.

R. Bansal and I. Shaliastovich. Learning and asset-price jumps. *The Review of Financial Studies*, 24(8):2738–2780, 2011. ISSN 08939454, 14657368. URL http://www.jstor.org/stable/20869322.

R. Bansal and A. Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance*, 59(4):1481–1509, 2004. doi: 10.1111/j.1540-6261.

2004.00670.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.2004.00670.x.

D. R. Baqaee. Cascading Failures in Production Networks. *Econometrica*, 86(5): 1819–1838, 2018. ISSN 0012-9682. doi: 10.3982/ecta15280. URL https://www.econometricsociety.org/doi/10.3982/ECTA15280.

D. R. Baqaee and E. Farhi. The macroeconomic impact of microeconomic shocks: Beyond hulten's theorem. *Econometrica*, 87(4):1155–1203, 2019. doi: https://doi.org/10.3982/ECTA15202. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA15202.

R. J. Barro and C. J. Redlick. Macroeconomic Effects From Government Purchases and Taxes *. *The Quarterly Journal of Economics*, 126(1):51–102, Feb. 2011. ISSN 0033-5533, 1531-4650. doi: 10.1093/qje/qjq002. URL https://academic.oup.com/qje/article-lookup/doi/10.1093/qje/qjq002.

J.-N. Barrot and J. Sauvagnat. Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks *. *The Quarterly Journal of Economics*, 131(3):1543–1592, 05 2016. ISSN 0033-5533. doi: 10.1093/qje/qjw018. URL https://doi.org/10.1093/qje/qjw018.

B. Bauer and M. Kohler. On deep learning as a remedy for the curse of dimensionality in nonparametric regression. *The Annals of Statistics*, 47(4):2261 – 2285, 2019. doi: 10.1214/18-aos1747. URL https://doi.org/10.1214/18-AOS1747.

Bea. Concepts and Methods of the U.S. National Income and Products Account, Nov. 2017.

A. Becker, W. Gray, and J. Marvakov. Nber-ces manufacturing industry database: Technical notes. *National Bureau of Economic Research*, 2016. URL https://www.nber.org/research/data/nber-ces-manufacturing-industry-database.

N. Ben Zeev and E. Pappa. Chronicle of A War Foretold: The Macroeconomic Effects of Anticipated Defence Spending Shocks. *The Economic Journal*, 127(603):1568–1597, Aug. 2017. ISSN 0013-0133, 1468-0297. doi: 10.1111/ecoj.12349. URL https://academic.oup.com/ej/article/127/603/1568-1597/5068886.

M. Bils and J.-O. Cho. Cyclical Factor Utilization. *Journal of Monetary Economics*, 33: 319–354, 1994.

C. C. Binder and G. Brunet. Inflation expectations and consumption: Evidence from 1951. *Economic Inquiry*, page ecin.13051, Nov. 2021. ISSN 0095-2583, 1465-7295. doi: 10.1111/ecin.13051. URL https://onlinelibrary.wiley.com/doi/10.1111/ecin.13051.

O. Blanchard and R. Perotti. An empirical characterization of the dynamic effects of changes in government spending and taxes on output. *The Quarterly Journal of Economics*, 117(4):1329–1368, Nov. 2002.

R. Blankenbecler, D. J. Scalapino, and R. L. Sugar. Monte carlo calculations of coupled boson-fermion systems. i. *Phys. Rev. D*, 24:2278–2286, Oct 1981. doi: 10.1103/PhysRevD.24.2278. URL https://link.aps.org/doi/10.1103/PhysRevD.24.2278.

N. Bloom and M. Shankerman. Identifying technology spillovers and product market rivalry. *Econometrica*, 81(4):1347–1393, 2013. doi: 10.3982/ecta9466. URL https://doi.org/10.3982/ecta9466.

J. Borovička, L. P. Hansen, and J. A. Scheinkman. Misspecified recovery. *The Journal of Finance*, 71(6):2493–2544, 2016.

K. Borusyak, P. Hull, and X. Jaravel. Quasi-Experimental Shift-Share Research Designs. *The Review of Economic Studies*, 89(1):181–213, Jan. 2022. ISSN 0034-6527, 1467-937x. doi: 10.1093/restud/rdab030. URL https://academic.oup.com/restud/article/89/1/181/6294942.

D. T. Breeden. An intertemporal asset pricing model with stochastic consumption and investment opportunities. *Journal of Financial Economics*, 7(3):265–296, 1979. URL https://EconPapers.repec.org/RePEc:eee:jfinec:v:7:y:1979:i:3:p:265-296.

E. Breza, A. G. Chandrasekhar, T. H. McCormick, and M. Pan. Using aggregated relational data to feasibly identify network structure without network data. *American Economic Review*, 110(8):2454–84, August 2020. doi: 10.1257/aer.20170861. URL https://www.aeaweb.org/articles?id=10.1257/aer.20170861.

E. Briganti and V. Sellemi. Who anticipates government spending? evidence from defense procurement. *Working Paper*, 2022.

P. J. Brockwell and R. A. Davis. *Time Series: Theory and Methods*. Springer Series in Statistics. Springer New York, New York, NY, 1991.

G. Brunet. When Does Government Spending Matter? Evidence from a New Measure of Spending. *Working Paper*, page 41, May 2020.

C. Burnside, M. Eichenbaum, and J. D. Fisher. Fiscal shocks and their consequences. *Journal of Economic Theory*, 115(1):89–117, Mar. 2004. ISSN 00220531. doi: 10.1016/s0022-0531(03)00252-7. URL https://linkinghub.elsevier.com/retrieve/pii/S0022053103002527.

D. Ceperley, G. V. Chester, and M. H. Kalos. Monte carlo simulation of a many-fermion study. *Phys. Rev. B*, 16:3081–3099, Oct 1977. doi: 10.1103/PhysRevB.16.3081. URL https://link.aps.org/doi/10.1103/PhysRevB.16.3081.

G. Chodorow-Reich. Geographic Cross-Sectional Fiscal Spending Multipliers: What Have We Learned? *American Economic Journal: Economic Policy*, 11(2):1–34, May 2019. ISSN 1945-7731, 1945-774x. doi: 10.1257/pol.20160465. URL https://pubs.aeaweb.org/doi/10.1257/pol.20160465.

T. M. Christensen. Nonparametric stochastic discount factor decomposition. *Econometrica*, 85(5):1501–1536, 2017.

G.-M. Cochard. *Introduction to stochastic processes and simulation*. Standards Information Network, Oct. 2019.

S. Cohen and J. Rosiński. Gaussian approximation of multivariate lévy processes with applications to simulation of tempered stable processes. *Bernoulli*, 13(1):195–210, 2007. ISSN 13507265. URL http://www.jstor.org/stable/25464871.

O. Coibion, Y. Gorodnichenko, and M. Weber. Fiscal Policy and Households' inflation Expectations: Evidence from a Randomized Control Trial. *Woriking Paper*, Dec. 2020.

L. Cox, J. J. Muller, E. Pasten, R. Schoenle, and M. Weber. BIg G. *Working Paper*, Jan. 2021.

D. Duffie and L. G. Epstein. Stochastic differential utility. *Econometrica*, 60(2):353–394, 1992. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/2951600.

D. Duffie, J. Pan, and K. Singleton. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, 68(6):1343–1376, 2000.

B. Dupor and R. Guerrero. Local and aggregate fiscal policy multipliers. *Journal of Monetary Economics*, 92:16–30, Dec. 2017. ISSN 03043932. doi: 10.1016/j.jmoneco.2017.07.007. URL https://linkinghub.elsevier.com/retrieve/pii/S0304393217300855.

W. E and B. Yu. The deep ritz method: A deep learning-based numerical algorithm for solving variational problems. *CoRR*, abs/1710.00211, 2017. URL http://arxiv.org/abs/1710.00211.

W. Edelberg, M. Eichenbaum, and J. Fisher. Understanding the Effects of a Shock to Government Purchases. *Review of Economic Dynamics*, 2(1):166–206, Jan. 1999.

M. Eichenbaum and J. Fisher. Fiscal Policy in the Aftermath of 9/11. *Journal of Money, Credit, and Banking*, 37(1):1–22, 2005. ISSN 1538-4616. doi: 10.1353/mcb.2005.

0005. URL http://muse.jhu.edu/content/crossref/journals/journal%5Fof%5Fmoney%5Fcredit%5Fand%5Fbanking/v037/37.1eichenbaum.pdf.

R. Engle. Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, 20(3):339–350, 2002. ISSN 07350015. URL http://www.jstor.org/stable/1392121.

J. G. Fernald. A quarterly, utilization-adjusted series on total factor productivity. Working Paper Series 2012-19, Federal Reserve Bank of San Francisco, 2012a. URL https://ideas.repec.org/p/fip/fedfwp/2012-19.html.

J. G. Fernald. A Quarterly, Utilization-Adjusted Series on Total Factor Productivity. *Federal Reserve Bank of San Francisco, Working Paper Series*, pages 01–28, Apr. 2012b. doi: 10.24148/wp2012-19. URL https://www.frbsf.org/economic-research/publications/working-papers/2012/19/.

J. A. Fioruci, R. S. Ehlers, and M. G. A. Filho. Bayesian multivariate GARCH models with dynamic correlations and asymmetric error distributions. *Journal of Applied Statistics*, 41(2):320–331, Oct. 2013. doi: 10.1080/02664763.2013.839635. URL https://doi.org/10.1080/02664763.2013.839635.

J. D. Fisher and R. Peters. Using Stock Returns to Identify Government Spending Shocks. *The Economic Journal*, 120(544):414–436, May 2010. ISSN 0013-0133, 1468-0297. doi: 10.1111/j.1468-0297.2010.02355.x. URL https://academic.oup.com/ej/article/120/544/414-436/5089580.

W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal. Quantum monte carlo simulations of solids. *Rev. Mod. Phys.*, 73:33–83, Jan 2001. doi: 10.1103/RevModPhys.73.33. URL https://link.aps.org/doi/10.1103/RevModPhys.73.33.

N. Fournier. Simulation and approximation of levy-driven stochastic differential equations, 2009.

X. Gabaix. The granular origins of aggregate fluctuations. *Econometrica*, 79:733–772, 2011.

X. Gabaix. A Behavioral New Keynesian Model. *American Economic Review*, 110 (8):2271–2327, Aug. 2020. ISSN 0002-8282. doi: 10.1257/aer.20162005. URL https://pubs.aeaweb.org/doi/10.1257/aer.20162005.

J. Galí, J. D. López-Salido, and J. Vallés. Understanding the Effects of Government Spending on Consumption. *Journal of the European Economic Association*, 5(1): 227–270, Mar. 2007.

M. B. Garman. Towards a semigroup pricing theory. *The Journal of Finance*, 40(3): 847–861, 1985. ISSN 00221082, 15406261. URL http://www.jstor.org/stable/2327810.

M. Germain, H. Pham, and X. Warin. Neural networks-based algorithms for stochastic control and pdes in finance, 2021.

X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In Y. W. Teh and M. Titterington, editors, *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9, pages 249–256, Chia Laguna Resort, Sardinia, Italy, 13–15 May 2010. Pmlr. URL http://proceedings.mlr.press/v9/glorot10a.html.

A. Gnoatto, M. Patacca, and A. Picarelli. A deep solver for bsdes with jumps, 2022.

M. Gofman, G. Segal, and Y. Wu. Production Networks and Stock Returns: The Role of Vertical Creative Destruction. *The Review of Financial Studies*, 33(12):5856–5905, 03 2020. ISSN 0893-9454. doi: 10.1093/rfs/hhaa034. URL https://doi.org/10.1093/rfs/hhaa034.

P. Goldsmith-Pinkham, I. Sorkin, and H. Swift. Bartik Instruments: What, When, Why, and How. *American Economic Review*, 110(8):2586–2624, Aug. 2020. ISSN 0002-8282. doi: 10.1257/aer.20181047. URL https://pubs.aeaweb.org/doi/10.1257/aer.20181047.

J. F. Gomes, L. Kogan, and M. Yogo. Durability of output and expected stock returns. *Journal of Political Economy*, 117(5):941–986, Oct. 2009. doi: 10.1086/648882. URL https://doi.org/10.1086/648882.

A. Gonzalez-Lira, R. Carril, and M. S. Walker. Competition under Incomplete Contracts and the Design of Procurement Policies. *Working Paper*, page 104, Jan. 2021.

R. Gordon and R. Krenn. The End of the Great Depression 1939-41: Policy Contributions and Fiscal Multipliers. page w16380, Sept. 2010. doi: 10.3386/w16380. URL http://www.nber.org/papers/w16380.pdf.

J. D. Hamilton. Why You Should Never Use the Hodrick-Prescott Filter. *The Review of Economics and Statistics*, 100(5):831–843, Dec. 2018. ISSN 0034-6535, 1530-9142. doi: 10.1162/rest_a_00706. URL https://direct.mit.edu/rest/article/100/5/831-843/58479.

J. Han, A. Jentzen, and W. E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018. ISSN 0027-8424. doi: 10.1073/pnas.1718942115. URL https://www.pnas.org/content/115/34/8505.

J. Han, L. Zhang, and W. E. Solving many-electron schrödinger equation using deep neural networks. *Journal of Computational Physics*, 399:108929, Dec. 2019. ISSN 0021-9991. doi: 10.1016/j.jcp.2019.108929. URL http://dx.doi.org/10.1016/j.jcp.2019.108929.

J. Han, J. Lu, and M. Zhou. Solving high-dimensional eigenvalue problems using deep neural networks: A diffusion monte carlo like approach. *ArXiv*, abs/2002.02600, 2020.

J. Han, Y. Yang, and W. E. Deepham: A global solution method for heterogeneous agent models with aggregate shocks, 2022.

L. P. Hansen. Dynamic valuation decomposition within stochastic economies. *Econometrica*, 80(3):911–967, 2012.

L. P. Hansen and E. Renault. *Pricing Kernels*. John Wiley Sons, Ltd, 2010. ISBN 9780470061602. doi: https://doi.org/10.1002/9780470061602.eqf19009. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/9780470061602.eqf19009.

L. P. Hansen and J. A. Scheinkman. Long-term risk: An operator approach. *Econometrica*, 77(1):177–234, 2009. doi: 10.3982/ecta6761. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA6761.

L. P. Hansen, J. Heaton, J. Lee, and N. Roussanov. Intertemporal substitution and risk aversion. *Handbook of econometrics*, 6:3967–4056, 2007.

B. Herskovic. Networks in production: Asset pricing implications. *The Journal of Finance*, 73(4):1785–1818, 2018a. doi: https://doi.org/10.1111/jofi.12684. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/jofi.12684.

B. Herskovic. Networks in Production: Asset Pricing Implications: Networks in Production: Asset Pricing Implications. *The Journal of Finance*, 73(4):1785–1818, Aug. 2018b. ISSN 00221082. doi: 10.1111/jofi.12684. URL http://doi.wiley.com/10.1111/jofi.12684.

B. Herskovic, B. Kelly, H. Lustig, and S. Van Nieuwerburgh. The common factor in idiosyncratic volatility: Quantitative asset pricing implications. *Journal of Financial Economics*, 119(2):249–283, 2016. ISSN 0304-405x. doi: https://doi.org/10.1016/j.jfineco.2015.09.010. URL https://www.sciencedirect.com/science/article/pii/S0304405X15001774.

B. Herskovic, B. Kelly, H. Lustig, and S. Van Nieuwerburgh. Firm volatility in granular networks. *Journal of Political Economy*, 0(0):000–000, 2020a. doi: 10.1086/710345. URL https://doi.org/10.1086/710345.

B. Herskovic, B. Kelly, H. Lustig, and S. Van Nieuwerburgh. Firm Volatility in Granular Networks. *Journal of Political Economy*, 128(11):4097–4162, Nov. 2020b. ISSN 0022-3808, 1537-534x. doi: 10.1086/710345. URL https://www.journals.uchicago.edu/doi/10.1086/710345.

G. Hoberg and G. Phillips. Text-based network industries and endogenous product differentiation. *Journal of Political Economy*, 124(5):1423–1465, 2016. doi: 10.1086/688176. URL https://doi.org/10.1086/688176.

K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. *Neural Netw.*, 2(5):359–366, July 1989. ISSN 0893-6080.

K. Horowitz and M. Planting. *Concepts and Methods of the U.S. Input-Output Accounts*. Apr. 2009.

C. J. Hottman, S. J. Redding, and D. E. Weinstein. Quantifying the sources of firm heterogeneity. *The Quarterly Journal of Economics*, 131(3):1291–1364, Mar. 2016. doi: 10.1093/qje/qjw012. URL https://doi.org/10.1093/qje/qjw012.

S. Janson. Normal Convergence by Higher Semiinvariants with Applications to Sums of Dependent Random Variables and Random Graphs. *The Annals of Probability*, 16 (1):305 – 312, 1988. doi: 10.1214/aop/1176991903. URL https://doi.org/10.1214/aop/1176991903.

Jordà. Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, 95(1):161–182, Feb. 2005. ISSN 0002-8282. doi: 10.1257/0002828053828518. URL https://pubs.aeaweb.org/doi/10.1257/0002828053828518.

E. Jum. Numerical approximation of stochastic differential equations driven by levy motion with infinitely many jumps. *PhD Thesis, University of Tennessee*, 13(1):195–210, 2015. ISSN 13507265. URL http://www.jstor.org/stable/25464871.

D. P. Kingma and J. Ba. Adam: A method for stochastic optimization, 2017.

F. Kramarz, J. Martin, and I. Mejean. Volatility in the small and in the large: The lack of diversification in international trade. *Journal of International Economics*, 122: 103276, Jan. 2020. ISSN 00221996. doi: 10.1016/j.jinteco.2019.103276. URL https://linkinghub.elsevier.com/retrieve/pii/S0022199618301296.

D. M. Kreps and E. L. Porteus. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica*, 46(1):185–200, 1978. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/1913656.

M. S. Kruttli, , B. R. Tran, and S. W. W. and. Pricing poseidon: Extreme weather uncertainty and firm return dynamics. *Finance and Economics Discussion Series*, 2019(054), July 2019. doi: 10.17016/feds.2019.054. URL https://doi.org/10.17016/feds.2019.054.

H. Lamba, J. C. Mattingly, and A. M. Stuart. An adaptive Euler–Maruyama scheme for SDEs: convergence and stability. *IMA Journal of Numerical Analysis*, 27(3): 479–506, 12 2006. ISSN 0272-4979. doi: 10.1093/imanum/drl032. URL https://doi.org/10.1093/imanum/drl032.

S. Leduc and D. Wilson. Roads to Prosperity or Bridges to Nowhere? Theory and Evidence on the Impact of Public Infrastructure Investment. *NBER Macroeconomics Annual*, 27 (1):89–142, 2013.

E. M. Leeper, T. B. Walker, and S.-C. S. Yang. Fiscal Foresight and Information Flows. *Econometrica*, 81(3):1115–1145, 2013. ISSN 0012-9682. doi: 10.3982/ecta8337. URL http://doi.wiley.com/10.3982/ECTA8337.

D. Li, M. Plagborg-Møller, and C. K. Wolf. Local Projections vs. VARs: Lessons From Thousands of DGPs. *Working Paper*, page 44, Mar. 2021.

T. H. McCormick and T. Zheng. Latent surface models for networks using aggregated relational data. *Journal of the American Statistical Association*, 110(512):1684–1695, 2015. doi: 10.1080/01621459.2014.991395. URL https://doi.org/10.1080/01621459.2014.991395.

A. Mian and A. Sufi. What explains the 2007-2009 drop in employment? *Econometrica*, 82(6):2197–2223, Nov. 2014. doi: 10.3982/ecta10451. URL https://doi.org/10.3982/ecta10451.

T. Monacelli and R. Perotti. Fiscal Policy, Wealth Effects, and Markups. Technical Report w14584, National Bureau of Economic Research, Cambridge, MA, Dec. 2008. URL http://www.nber.org/papers/w14584.pdf.

J. L. Montiel Olea and M. Plagborg-Møller. Local Projection Inference is Simpler and More Robust Than You Think. *Working Paper*, page 45, Dec. 2020.

R. J. Needs, M. D. Towler, N. D. Drummond, and P. L. Ríos. Continuum variational and diffusion quantum monte carlo calculations. *Journal of Physics: Condensed Matter*, 22(2):023201, dec 2009. doi: 10.1088/0953-8984/22/2/023201. URL https://dx.doi.org/10.1088/0953-8984/22/2/023201.

C. Nekarda and V. Ramey. Industry Evidence on the Effects of Government Spending. *American Economic Journal: Macroeconomics*, 3(1):36–59, Jan. 2011. ISSN 1945-7707, 1945-7715. doi: 10.1257/mac.3.1.36. URL https://pubs.aeaweb.org/doi/10.1257/mac.3.1.36.

C. Nekarda and V. Ramey. The Cyclical Behavior of the Price-Cost Markup. *Journal of Money, Credit and Banking*, 52(S2):319–353, Dec. 2020. ISSN 0022-2879, 1538-4616. doi: 10.1111/jmcb.12755. URL https://onlinelibrary.wiley.com/doi/10.1111/jmcb.12755.

E. Oberfield. A Theory of Input-Output Architecture. *Econometrica*, 86(2):559–589, 2018. ISSN 0012-9682. doi: 10.3982/ecta10731. URL https://www.econometricsociety.org/doi/10.3982/ECTA10731.

B. Oksendal. *Stochastic Differential Equations (3rd Ed.): An Introduction with Applications*. Springer-Verlag, Berlin, Heidelberg, 1992. ISBN 3387533354.

A. Ozdagli and M. Weber. Monetary policy through production networks: Evidence from the stock market. Working Paper 23424, National Bureau of Economic Research, May 2017. URL http://www.nber.org/papers/w23424.

A. Ozdagli and M. Weber. Monetary Policy through Production Networks: Evidence from the Stock Market. page 79, May 2020.

A. Patton and A. Timmermann. Monotonicity in asset returns: New tests with applications to the term structure, the capm, and portfolio sorts. *Journal of Financial Economics*, 98 (3):605–625, 2010. URL https://EconPapers.repec.org/RePEc:eee:jfinec:v:98:y:2010:i:3:p:605-625.

R. Perotti. In Search of the Transmission Mechanism of Fiscal Policy [with Comments and Discussion]. *NBER Macroeconomics Annual*, 22:169–249, Jan. 2007. ISSN 0889-3365, 1537-2642. doi: 10.1086/ma.22.25554966. URL https://www.journals.uchicago.edu/doi/10.1086/ma.22.25554966.

R. Perotti. Defense Government Spending Is Contractionary, Civilian Government Spending Is Expansionary. *NBER Working Paper*, page w20179, May 2014. doi: 10.3386/w20179. URL http://www.nber.org/papers/w20179.pdf.

D. Pfau, J. S. Spencer, A. G. D. G. Matthews, and W. M. C. Foulkes. Ab initio solution of the many-electron schrödinger equation with deep neural networks. *Phys. Rev. Res.*, 2:033429, Sep 2020. doi: 10.1103/PhysRevResearch.2.033429. URL https://link.aps.org/doi/10.1103/PhysRevResearch.2.033429.

L. Pichler, A. Masud, and L. A. Bergman. *Numerical Solution of the Fokker–Planck Equation by Finite Difference and Finite Element Methods—A Comparative Study*, pages 69–85. Springer Netherlands, Dordrecht, 2013. ISBN 978-94-007-5134-7. doi: 10.1007/978-94-007-5134-7_5. URL https://doi.org/10.1007/978-94-007-5134-7%5F5.

R. Piessens, E. de Doncker-Kapenga, C. W. Überhuber, and D. K. Kahaner. *Quadpack*. Springer Berlin Heidelberg, 1983. ISBN 9783642617867. doi: 10.1007/978-3-642-61786-7. URL http://dx.doi.org/10.1007/978-3-642-61786-7.

M. Plagborg-Møller and C. K. Wolf. Local Projections and VARs Estimate the Same Impulse Responses. *Econometrica (forthcoming)*, page 36, Oct. 2020.

L. Qin and V. Linetsky. Positive eigenfunctions of markovian pricing operators: Hansen-scheinkman factorization, ross recovery, and long-term pricing. *Operations Research*, 64(1):99–117, 2016. doi: 10.1287/opre.2015.1449. URL https://doi.org/10.1287/opre.2015.1449.

L. Qin and V. Linetsky. Long-term risk: A martingale approach. *Econometrica*, 85(1):299–312, 2017.

V. Ramey. Inventories as Factors of Production and Economic Fluctuations. *American Economic Review*, 79(3):338–354, June 1989.

V. Ramey. Nonconvex Costs and the Behavior of Inventories. *Journal of Political Economy*, 99(2):306–334, Apr. 1991. ISSN 0022-3808, 1537-534x. doi: 10.1086/261752. URL https://www.journals.uchicago.edu/doi/10.1086/261752.

V. Ramey. Identifying Government Spending Shocks: It's all in the Timing*. *The Quarterly Journal of Economics*, 126(1):1–50, Feb. 2011. ISSN 0033-5533, 1531-4650. doi: 10.1093/qje/qjq008. URL https://academic.oup.com/qje/article-lookup/doi/10.1093/qje/qjq008.

V. Ramey. Macroeconomic Shocks and Their Propagation. In *Handbook of macroeconomics*, volume 2, pages 71–162. Elsevier, 2016.

V. Ramey and D. Shapiro. Costly capital reallocation and the effects of government spending. *Carnegie-Rochester Conference on Public Policy*, 48(1998):145–194, 1998.

V. Ramey and S. Zubairy. Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data. *journal of political economy*, page 52, Mar. 2018.

C. Ramìrez. Inter-firm relationships and asset prices. *Finance and Economics Discussion Series 2017-014. Washington: Board of Governors of the Federal Reserve System*, 2017.

C. D. Romer and D. H. Romer. A New Measure of Monetary Shocks: Derivation and Implications. *The American Economic Review*, 94(4):64, 2004.

M. Schroder and C. Skiadas. Optimal consumption and portfolio selection with stochastic differential utility. *Journal of Economic Theory*, 89(1):68–126, 1999. ISSN 0022-0531. doi: https://doi.org/10.1006/jeth.1999.2558. URL https://www.sciencedirect.com/science/article/pii/S0022053199925589.

I. Shaliastovich and G. Tauchen. Pricing of the time-change risks. *Journal of Economic Dynamics and Control*, 35(6):843–858, 2011. ISSN 0165-1889. doi: https://doi.org/10.1016/j.jedc.2011.01.003. URL https://www.sciencedirect.com/science/article/pii/S0165188911000042.

J. Shea. Complementarities and comovements. *Journal of Money, Credit and Banking*, 34(2):412–433, 2002. ISSN 00222879, 15384616. URL http://www.jstor.org/stable/3270695.

R. Shen, L. Gao, and Y.-A. Ma. On optimal early stopping: Over-informative versus under-informative parametrization, 2022.

F. Smets, J. Tielens, and J. Van Hove. Pipeline Pressures and Sectoral Inflation Dynamics. *SSRN Electronic Journal*, 2019. ISSN 1556-5068. doi: 10.2139/ssrn.3346371. URL https://www.ssrn.com/abstract=3346371.

J. H. Stock and M. W. Watson. Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments. *The Economic Journal*, 128(610): 917–948, May 2018. ISSN 0013-0133, 1468-0297. doi: 10.1111/ecoj.12593. URL https://academic.oup.com/ej/article/128/610/917-948/5069563.

C. Syverson. Product substitutability and productivity dispersion. *Review of Economics and Statistics*, 86(2):534–550, May 2004. doi: 10.1162/003465304323031094. URL https://doi.org/10.1162/003465304323031094.

M. Taschereau-Dumouchel. Cascades and Fluctuations in an Economy with an Endogenous Production Network. *Review of Economic Studies*, page 62, 2020.

S. Tuzel and M. B. Zhang. Local risk, local factors, and asset prices. *The Journal of Finance*, 72(1):325–370, 2017.

J. F. Wieland and M. Yang. Financial Dampening. *Journal of Money, Credit and Banking*, 52(1):79–113, Feb. 2020. ISSN 0022-2879, 1538-4616. doi: 10.1111/jmcb.12681. URL https://onlinelibrary.wiley.com/doi/10.1111/jmcb.12681.

S. Zhang, J. Carlson, and J. E. Gubernatis. Constrained path monte carlo method for fermion ground states. *Phys. Rev. B*, 55:7464–7477, Mar 1997. doi: 10.1103/PhysRevB.55.7464. URL https://link.aps.org/doi/10.1103/PhysRevB.55.7464.