Inflation and Nominal Rigidity in Supply Chains*

Robert Minton† (Job Market Candidate) Brian Wheaton‡

September 8, 2022

Abstract

We study theoretically and empirically the propagation of price movements through supply chains, with a specific focus on oil and commodity price shocks. Full pass-through of commodity shocks to industries directly using the commodity occurs quickly. Pass-through to industries using the commodity indirectly – through their suppliers, suppliers' suppliers, etc. – is also full but occurs more slowly. We show that a core reason for this result is that price rigidity compounds through supply chains, a phenomenon capable of translating relative price flexibility at the industry level into price rigidity at the aggregate level. In a Calvo New Keynesian model augmented with myopia in expectations, we argue that price rigidity compounds even more in supply chains when firms are myopic about gradual upstream pass-through of shocks. Empirically, we find that, for large oil shocks, firms pass through anticipated marginal cost increases before they have occurred, consistent with rational expectations. For small oil shocks and shocks to other commodities, firms are more myopic, and rational expectations can be rejected. In an application of our model, we measure the component of aggregate inflation induced by indirect network effects of oil price movements, finding that these effects take years to manifest fully and are more than twice as large as the direct effect on consumer gas prices. Our measure explains 30% of the variation in PCE inflation and 12% of the variation in Core PCE inflation.

*We are grateful to Pol Antras, Robert Barro, Adrien Bilal, Raj Chetty, Gabe Chodorow-Reich, Jason Furman, Xavier Gabaix, Jordi Galí, Ed Glaeser, Larry Katz, Greg Mankiw, Andrei Shleifer, Jim Stock, Ludwig Straub, and John Taylor for excellent comments and advice. We also thank participants of the Macroeconomics, International Economics, and Labor/Public lunch seminars at Harvard, the Macroeconomics faculty lunch and Hoover Economic Policy seminar at Stanford, and the Institute for Humane Studies SuRF seminar.

†Harvard University, Economics Department. Contact: rminton@g.harvard.edu
‡UCLA, Anderson School of Management. Contact: brian.wheaton@anderson.ucla.edu

1 Introduction

In recent years, there has been a profusion in the use of network models and associated methods in macroeconomics. These papers, however, have typically focused on modeling and calibration, and consequently there is limited causal empirical evidence for these models' predictions - particularly predictions about prices. Yet researchers have increasing turned to input-output models as a way to amplify aggregate price rigidity above microeconomic price rigidity, a modeling choice with implications including the efficacy of monetary policy, the intensity of business cycles, and beyond.

Our contribution is to study the propagation of price movements through supply chains in reduced-form, semi-structural, and fully structural contexts, unifying network models and input-output data to uncover a coherent story about the determinants of price pass-through and the rigidities that mediate it. Applying our empirically-validated model allows us to measure the aggregate inflationary effects of commodity price movements once they have fully filtered through supply chains. We show in particular that the indirect network effects of oil price movements are responsible for a much higher fraction of the variation in aggregate inflation than previously thought: the indirect network effects are almost 2.5 times larger than the direct effects of oil on aggregate inflation through consumer gas purchases, but they take years to manifest fully in aggregate inflation.

Setting up a network New Keynesian model with Calvo pricing, we first detail the predictions of the model for price pass-through in the context of a production network. We then empirically study the extent, pace, and mechanism of the propagation of prices through the production network in a reduced-form causal setting – providing direct tests of the validity of network models in macroeconomics. We find substantial evidence that price pass-through of oil and commodity shocks throughout the production network is strongly significant, statistically indistinguishable from 100%, and relatively rapid. By and large, the extent and pace of pass-through in the data is statistically indistinguishable from our network-enriched Calvo New Keynesian model, appropriately calibrated,[1] suggesting it can be used for the purpose of predicting how shocks will propagate through industry prices along the supply chain.

Both theoretically and in our empirics, we show that the precise pace of pass-through is strongly influenced by a number of factors. In reduced-form empirical analyses, we show that industries far upstream in the supply chain – such as those using commodities directly in their production process – experience full pass-through of cost shocks within a month or two. On the other hand, downstream industries – such as consumer-facing industries using commodities only indirectly in their production process – take six to eight months or more to realize full pass-

[1] A key parameter in calibration is the average industry markup of price over marginal cost, which we estimate using a generalized method of moments procedure.

through. Indeed, we find evidence that nominal rigidities are compounded through the supply chain, a phenomenon whereby even minor industry-level price rigidities can yield substantial aggregate price rigidity.

In more concrete terms, we use data from the approximately 400-500 industry BEA input-output tables published every five years to compute the network share of oil (and other commodities) in the production process of each industry. This combines both direct and indirect exposure to oil. For example, the Petroleum Refineries industry has nearly 80% of its total costs in oil, and nearly all of these costs constitute direct purchase of crude oil – i.e., first-order exposure to crude oil. The Synthetic Rubber Manufacturing industry has approximately 20% of its total costs in oil, but nearly none of it constitutes direct purchase of crude oil. Most of this originates from purchases of refined oil from petroleum refineries (i.e., second-order exposure to crude oil), purchases of petrochemical products from the Petrochemical Manufacturing industry, which itself purchased most inputs from refineries (i.e., third-order exposure to crude oil), and beyond.

Using these measures of exposure, we first pool all variation in oil prices since 1997, and we show that pass-through of both direct and indirect exposure to oil price shocks is strongly significant and statistically indistinguishable from 100%. Pass-through of direct exposure occurs mostly on impact, whereas pass-through of indirect exposure occurs over a slower period of six to eight months or more. These results are replicated using as an instrument the Kanzig (2021) series of exogenous oil price shocks obtained through high-frequency identification of the effects of OPEC announcements. Results are also analogous when we investigate the effects of other commodity shocks beyond oil. As an additional robustness check, we study a few specific cases of large, plausibly-exogenous movements in the oil price – including the 1979 oil price spike driven by the Iranian Revolution, the 2014-15 oil price crash driven by the U.S. oil shale boom, and the 2020 COVID shock to oil prices. Once again, we find evidence of strongly-significant pass-through of 100%, with a greater lag for indirect exposure. These results provide evidence for network models' predictions that shocks propagate through prices in supply chains, albeit gradually.

Performing heterogeneity analysis, we show that industries with less flexible pricing – using data from Pasten, Schoenle, and Weber (2017) on the frequency of price adjustment by industry – experience less rapid pass-through. And, even conditional on this finding, industries further downstream experience significantly less rapid pass-through. On the other hand, we do not find consistent evidence that the sign or size of cost shocks are important sources of heterogeneity in the pace or fraction of the shock which is passed through.

In the framework of our Calvo New Keynesian model, we show that the compounding of rigidities through the supply chain is intensified in a setting with myopic firms, as opposed to a setting where rational expectations holds. Intuitively, with full rational expectations (and full attention), downstream sectors will tend to adjust their prices to account for anticipated future cost changes shortly after they observe upstream sectors (e.g., their suppliers' suppliers) being hit by cost shocks – even before the actual cost changes have worked their way down the supply chain to reach the downstream sectors. In a more myopic context, sectors will wait until the cost shocks have filtered down to them step-by-step before adjusting their prices, leading to a much higher degree of price rigidity. In semi-structural and structural empirical analyses, we find that the extent of myopia is quite context-dependent. We develop a generalized method of moments (GMM) procedure to estimate myopia from our observational data, and we find that industries tend to behave in line with rational expectations during episodes of large oil shocks while behaving more myopically during episodes of small oil shocks or shocks to other commodities. We verify these findings in semi-structural empirical analyses based on the Calvo New Keynesian model that are identified using different variation. These findings may be relevant for research employing an input-output network to amplify macroeconomic price rigidity, such as research on the efficacy of monetary policy or intensity of business cycles, and for research studying how firms respond to changes in their expectations, such as research on forward guidance.

We argue that our approach to production networks has a variety of applications. Amongst them, it allows computing revised measures of inflation that fully strip out the influence of specific industries. For example, official measures of core inflation – which simply remove the food and energy sectors from computations of inflation – do not fully purge the influence of energy from the resulting measure of inflation; it is still heavily intertwined through the production network. Our approach makes it possible to fully strip the influence of oil throughout the production network – directly and indirectly – from inflation, resulting in a new, network-corrected measure of core inflation. Using our GMM-calibrated Calvo New Keynesian model, we show that the indirect network effects of oil inflation are more than twice as large as the direct effects of oil inflation on consumer gas prices. We formalize the notion that core PCE inflation does not fully purge the effects of oil price movements by showing that it is predictable using our series measuring the indirect network effects of oil price movements. Moreover, network oil inflation explains 30% of the variation in aggregate PCE inflation. In an application, we show that the indirect network effects of oil price increases are an important component of the 2021-22 spike in inflation – responsible for a much greater fraction of that inflation than direct effects alone would suggest. Further, as of May 2022, we show that 2 percentage points of indirect network components of oil inflation has not yet passed through to aggregate inflation. Declines in oil prices starting in May 2022 would mediate this effect.

The remainder of this paper is organized as follows. Section 2 reviews the relevant literature on production networks. Section 3 sets up our models, beginning with a simple flexible-price framework in 3.1 and then turning to a network-enriched Calvo New Keynesian model in 3.2. We also provide a simple linear-network version of this model - useful for building intuition - in 3.3. In Section 4.1, we discuss various datasets used in our estimation. In Section 4.2, we walk through the results of our reduced-form empirical analysis, a variety of robustness checks, and heterogeneity analysis. In Section 4.3, we reconcile our empirical results with the model and then turn to our semi-structural and structural analysis of rational expectations and myopia. In Section 5, we focus on our application: stripping the full network effects of oil price movements from core inflation. In Section 6, we conclude by summarizing our findings and discussing some of their implications.

2 Literature Review

Our paper contributes to three strands of existing literature. First, we add to the recent but growing literature performing empirical tests of production network models. We do so by studying the extent and pace of price pass-through along supply chains in a causal manner. Second, we provide direct evidence of nominal rigidity amplification in supply chains, a conjecture which dates back to Gordon (1981) and underlies a growing recent literature. Finally, we contribute to the literature on forward-looking expectations and myopia, as we develop a way of estimating myopia from observational macroeconomic data.

While production network models have garnered substantial interest within macroeconomics over the past few years, their origins can be traced back many decades. Input-output tables mapping industry-to-industry linkages in the production process have their origins with Leontief (1936). Inspired by Quesnay (1758), who wrote about the interdependencies between economic classes, Leontief used data from a broad variety of government publications on specific industries or sectors to construct a 40-sector input-output representation of the United States economy circa 1919. Leontief's table was the first input-output representation of any economy, and it was the foundation of later tables produced by the BLS and the BEA (Kohli 2001). Over the subsequent decades, input-output modeling found wide application in the study of international trade flows, regional economic modeling, and socialist economic planning. Leontief later won the 1973 Nobel Prize in Economics for his work in developing input-output tables.

In the latter decades of the 20th century, input-output modeling declined in popularity and research activity within the economics literature. However, the work of Hulten (1978) is often regarded as seminal by more recent literature on production networks in that it studies the importance (or lack thereof) of network linkages in production for aggregate impacts of industry-specific shocks. Hulten's finding – sometimes called Hulten's Theorem – is that, to a first-order approximation, the macroeconomic impact of a total factor productivity shock in industry i depends only on i's sales share of gross domestic product. Consequently, for macro outcomes, one can ignore the details of industry-to-industry linkages, complementarities in production, and reallocation of factors across industries. Another seminal paper from this era is Long and Plosser (1983), which developed the multi-sector general equilibrium model drawn on by many recent papers on production networks. Hulten's Theorem applies in this context as well.

Far from serving as the final chapter on interest in production networks, these papers have been a foundation for a resurgent body of recent literature. Horvath (1998) presents a modified model inspired by but distinct from that of Long and Plosser. In this setting, Hulten's Theorem does not hold. Positive shocks to certain sectors are not equally offset by negative shocks in other sectors; interactions amongst producing sectors stymie the Law of Large Numbers from producing this result. Consequently, Horvath argued that sector-specific shocks can explain a substantial fraction of aggregate disturbances; as much as 80% of the volatility in GDP growth is due to sector-specific shocks in Horvath's findings. Horvath (2000), Acemoglu et al. (2012), and Baqaee and Farhi (2019) present additional modeling evidence of significant deviations from Hulten's Theorem. Furthermore, in calibration exercises, Foerster, Sarte, and Watson (2011), Carvalho and Gabaix (2013), and Atalay (2017) attribute half or more of aggregate volatility to sector-specific shocks.

Bartelme and Gorodnichenko (2015) and Caliendo, Parro, and Tsyvinski (2017) extend this logic internationally, finding that sector- and country-specific distortions have meaningful implications for global macroeconomic output. Pasten, Schoenle, and Weber (2017) find that industry heterogeneity in price rigidity amplifies aggregate fluctuations. Baqaee (2018) finds that industry-level market structure is responsible for such an amplifying effect. Acemoglu, Ozdaglar, and Tahbaz-Salehi (2017) further document the importance of industry heterogeneity.

All of the aforementioned papers, however, are theoretical in nature – focused on modeling and calibration exercises. A much smaller set of papers has tested some empirical implications of network models. Barrot and Sauvagnat (2016) study the propagation of firm-specific shocks from natural disasters in the United States, focusing on propagation to immediate suppliers and consumers and finding statistically-significant transmission of shocks. Boehm, Flaeen, and Pandalai-Nayar (2019) and Carvalho et al. (2021) focus on the 2011 Tohoku Earthquake, studying how its effects on output propagate upward and downward through supply chains, with the result being that a non-trivial part of the drag on Japanese real GDP growth from the disaster was due to network propagation effects.

Acemoglu, Akcigit, and Kerr (2016) study the effects on industry-level output of a variety of supply and demand shocks (Chinese import shocks, government spending changes, TFP growth, and foreign-industry patenting) propagating through the production network. The authors find important network effects – dwarfing the own-sector effects – of all four types of shocks. However, their empirical specification assumes away potential pre-trends (and effects on impact), and their data is annual, making it difficult to apply a causal interpretation or track propagation through the supply chain at a higher-frequency – virtually all of the cumulative effect has already occurred after the first year.

We contribute to this empirical segment of the literature on production networks. Unlike most existing empirical work on the topic, we are interested in studying the propagation of *price* changes through production networks. Studying prices over other outcomes provides substantial empirical power because good measures of highly disaggregated industry price indices are available at a monthly frequency from the Bureau of Labor Statistics' Producer Price Index database. Furthermore, we aim to use identification techniques and approaches necessary to conduct a causal, reduced-form test of production network models with techniques standard in the modern applied-econometrics literature.

The two papers which relate closest to our work are Auer, Levchenko, and Saure (2019) and Smets, Tielens, and Van Hove (2018). Auer et al. study international input-output linkages, presenting evidence that global input-output linkages contribute to the synchronization of PPI inflation across countries. Smets et al. empirically analyzes the network price pass-through patterns of estimated micro-level shocks in a structural Bayesian framework. While these papers have qualitatively similar findings to ours on the dynamics of network price pass-through, our methodologies vary substantially: (1) we use a frequentist framework to track the industry price pass-through of observed upstream shocks, e.g. to the oil price, while Smets et al. track the pass-through of price shocks estimated through the model; and (2) our empirical framework features monthly data on prices along with the 400- to 500-sector benchmark input-output tables, while the model-driven approach of Smets. et al. relies computationally on analysis of a 35-sector model with quarterly price data. Smets et al. find that the data favors their models with inflation propagation through supply chains over models with no such propagation. We use a Calvo New Keynesian model that assumes propagation of price shocks through supply chains and show that this model can match the pass-through dynamics of upstream cost shocks to industry prices throughout the network.

Another subset of the literature on production networks to which we make a particular contribution is work on the compounding of nominal rigidities through the supply chain. This idea extends back at least to Gordon (1981), who called attention to "the role of the input-output table in translating prompt price adjustment at the individual level to gradual price adjustment at the aggregate level." Deep supply chains with a long sequence of links could translate short lags at the firm- or industry-level into lengthy lags between monetary or commodity cost shocks and their incorporation into the aggregate price level. Blanchard (1983) formalized these ideas in a model featuring a linear production network. Building on this work, Basu (1995) develops a model arguing that this compounding of nominal rigidities magnifies productivity fluctuations and thereby contributes to the intensity of business cycles. Basu found some empirical validation for the model—changes in industries' intermediate input shares of output covary with changes in the relative price of intermediate inputs to output—but this result is not causally identified and can follow from complementarities in production or from prices being imperfectly flexible, without this rigidity compounding through the supply chain. We test the hypothesis that nominal rigidities compound in supply chains more directly by tracking propagation deep into supply chains and in a causal setting, further controlling for the fact that high intermediate input shares can be associated with very fast price adjustment (as, e.g., in the case of petroleum refineries).

More recent work has studied the policy implications of the compounding of nominal rigidities through the production network, with Rubbo (2020) and La'O and Tahbaz-Salehi (2022) investigating the implications for monetary policy. They find that optimal monetary policy in the context of production networks targets an alternative price index that more greatly weights certain sectors with greater influence in the production network (such as those that are more upstream, larger, or stickier, in the case of La'O and Tahbaz-Salehi). Other papers discuss the importance of price rigidity compounding in supply chains, including Carvalho (2006) and Nakamura and Steinsson (2010). With the exception of the described section in Basu (1995), all of the analyses in this strand of the literature are based on theory and calibrations; our contribution is to show that the compounding of nominal rigidities actually does occur in the data.

Finally, we contribute to the literature on forward-looking expectations and myopia. As illustrated by Carlstom, Fuerst, and Paustian (2012) and discussed in detail by Del Negro, Giannoni, and Patterson (2013), in a workhorse New Keynesian model (such as Smets and Wouters 2007), central bank promises with regard to policies that will be undertaken in the (sometimes-distant) future can have unreasonably large effects on present-day inflation and output. This is referred to as the "forward guidance puzzle." Gabaix (2020) discusses myopia - the notion that agents are not perfectly forward-looking in all contexts - as a solution to the forward guidance puzzle. We show a way of estimating myopia from observational data on the macroeconomy (as opposed to microeconomic experiments) and, in so doing, provide empirical evidence of its importance.

3 Model

Our modeling seeks to guide all the empirical analysis we perform in our study of price rigidity. We begin with a standard, Dixit-Stiglitz setup with flexible pricing in a network economy. The flexible pricing solutions allow us to define and develop intuition for the "network exposures" of industries to commodity price movements, even if those industries do not directly purchase those commodities for use in production—if, that is, they are instead exposed to commodity price movements only indirectly through their suppliers' use (or their suppliers' suppliers' use, etc.).

We then develop measures for how price rigidity compounds in the production network if firms in each sector can only change their prices with some probability, following the long tradition in macroeconomics of using Calvo pricing to study price rigidity. This model most closely follows the setup of Rubbo (2020). We extend this model by allowing firms to be myopic about the pass-through of upstream shocks to suppliers' prices, following the setup of Gabaix (2020) and nesting the case of rational expectations. This model guides our semi-structural empirical analysis as well as providing the framework for estimating key pass-through parameters using the generalized method of moments in our later structural estimation.

Finally, we use a simple, linear network to develop intuition about the role of myopia in supply chain propagation of shocks. A linear network consists of a single supply chain in which each firm uses only labor and inputs from the previous link in the supply chain—with the final link in the supply chain ultimately selling to consumers. The linear network model admits convenient, closed-form solutions in continuous time that allow us to characterize how myopia affects the speed of shock pass-through.

3.1 Flexible Pricing

The setup is standard Dixit-Stiglitz at the industry level. There is a continuum of firms $j \in [0,1]$ in each industry $i \in \{1, ..., I\}$. Intermediate and final demand consume a constant elasticity of substitution bundle of firm output in each industry, which we denote by

$$Y_{i,t} = \left(\int_0^1 X_{i,j,t}^{\frac{\sigma_i - 1}{\sigma_i}} \, dj \right)^{\frac{\sigma_i}{\sigma_i - 1}} .$$

The elasticity of substitution σ_i is constrained to be greater than 1 for a well-defined monopoly profit maximization problem. The optimal choice of each $X_{i,j,t}$ by intermediate and final demand minimizes the cost of using the bundle $Y_{i,t}$, leading to the demand for $X_{i,j,t}$:

$$X_{i,j,t} = Y_{i,t} \left(\frac{P_{i,t}}{P_{i,j,t}} \right)^{\sigma_i}$$

with

$$P_{i,t} \equiv \left(\int_0^1 P_{i,j,t}^{1-\sigma_i} \right)^{\frac{1}{1-\sigma_i}} .$$

Total factor productivity $A_{i,t}$ (hereafter, TFP) is exogenous and common to all firms in an industry. The production process F_i is also common to all firms within an industry and is constant returns to scale. Each firm j in industry i may produce using a bundle of intermediate inputs from all modeled industries, denoted $\boldsymbol{X_{i,j,t}} = (X_{i,j,t}^1, ..., X_{i,j,t}^I)$ and labor $L_{i,j,t}$. They also may use a commodity from an unmodeled, commodity-producing industry $Z_{i,j,t}$ sold at price $P_{Z,t}$. We can think about Z for now as oil, supplied on a global market. It follows that marginal cost does not vary across firms and depends on the industry-specific wage, $W_{i,t}$, the vector of modeled industry prices $\boldsymbol{P}_t = (P_{1,t}, ..., P_{I,t})$, the commodity price $P_{Z,t}$, and the level of TFP, $A_{i,t}$.

Noting that, without dynamics, there is no current use for the time variable, we define industry cost shares in each input as

$$s_i^L = \frac{W_i L_i}{C_i}, \quad s_i^{X_i} = \frac{P_i X_j^i}{C_i}, \quad s_i^Z = \frac{P_Z Z_i}{C_i}$$

where C_i is the sum of industry i's expenditures on labor and intermediate inputs. The intermediate input shares can be arranged into a row vector $(\boldsymbol{s}_i)' = (s_i^{X_1}, ..., s_i^{X_I})$ and stacked across industries to form what is commonly called the input-output matrix,

$$\boldsymbol{\Phi} = \begin{pmatrix} \boldsymbol{s}_1' \\ \vdots \\ \boldsymbol{s}_I' \end{pmatrix}.$$

Similarly, we gather labor and commodity shares into the column vectors $\boldsymbol{s}^L = (s_1^L, ..., s_I^L)'$ and $\boldsymbol{s}^Z = (s_1^Z, ..., s_I^Z)'$.

Finally, we allow markups in industry i to vary from the optimal monopoly markup by a factor $\mu_{i,t}$. Though we will not use this parameter in the flexible pricing model, it will allow us to endogenize commodity production later and consider, for example, OPEC guidance on oil prices as moving the markups in the U.S. oil sector rather than modeling OPEC announcements as a shift in oil sector TFP. The profit maximization problem for a firm j in industry i is therefore

$$P_{i,t} = \max_{P_{i,t}} X_{i,j,t} \left(\frac{P_{i,t}}{\mu_{i,t}} - MC_{i,j,t} \right),$$

subject to

$$X_{i,j,t} = Y_{i,t} \left(\frac{P_{i,t}}{P_{i,j,t}} \right)^{\sigma_i}.$$

From this setup, we can derive the following:

Proposition 1. *The optimal industry price is*

$$P_{i,t} = \mu_{i,t} \frac{\sigma_i}{\sigma_i - 1} MC_{i,t},$$

and

$$d \ln \boldsymbol{P}_t = (\boldsymbol{I} - \boldsymbol{\Phi})^{-1} \left(d \ln \boldsymbol{\mu}_t - d \ln \boldsymbol{A}_t + \left(\boldsymbol{s}^L\right)' \boldsymbol{I} d \ln \boldsymbol{W}_t + \boldsymbol{s}^Z d \ln P_{Z,t} \right).$$

Moreover, the movements in the exogenous parameters $A_{i,t}$ and $\mu_{i,t}$ can be measured as

$$d \ln \boldsymbol{\mu}_t - d \ln \boldsymbol{A}_t = d \ln \boldsymbol{P}_t - \left(\boldsymbol{\Phi} d \ln \boldsymbol{P}_t + \left(\boldsymbol{s}^L\right)' \boldsymbol{I} d \ln \boldsymbol{W}_t + \boldsymbol{s}^Z d \ln P_{Z,t} \right),$$

the residual price movement after deducting changes in industries' marginal costs.

Proposition 1 establishes that industry prices are moved only by changes in desired markups, TFP, wages, or the commodity price. It follows that the pass-through of a change in the commodity price to a first-order, $d \ln P_{Z,t}$, if wages, TFP, and desired markups do not adjust is

$$d \ln \boldsymbol{P}_t = (\boldsymbol{I} - \boldsymbol{\Phi})^{-1} \boldsymbol{s}^Z d \ln P_{Z,t}.$$

Using the geometric sum formula for matrices, we can decompose industries' exposure to the commodity price movement:

$$(\boldsymbol{I} - \boldsymbol{\Phi})^{-1}\boldsymbol{s}^Z = \boldsymbol{s}^Z + \boldsymbol{\Phi}\boldsymbol{s}^Z + \boldsymbol{\Phi}^2\boldsymbol{s}^Z + ...$$

From this decomposition, we define the exposure of an industry to the commodity price at order k as

$$\text{NetworkExposure}_{i,k} = \left[\boldsymbol{\Phi}^k\boldsymbol{s}^Z\right]_i.$$

It is clear that $\text{NetworkExposure}_{i,0}$ is an industry's *first-order* exposure to the commodity through its own purchases of the commodity. If 10% of an industry's labor and intermediate costs are comprised of purchases of the commodity, then $\text{NetworkExposure}_{i,0} = .1$. Through simple matrix multiplication, it also follows that $\text{NetworkExposure}_{i,1}$ is an industry's exposure to the commodity through its suppliers' costs. Suppose only one of an industry's suppliers buys the commodity directly, and this supplier has a direct cost share of .8. Then if 20% of an industry's cost is comprised of purchases from this supplier, $\text{NetworkExposure}_{i,1} = .2 \times .8 = .16$. Connecting our definition of network exposure to the data, Figure 1 shows multiple orders of a selection of industries' network exposures to oil. We use the Bureau of Economic Analysis's benchmark input output table from 2012 to perform these calculations, as described in the data section.

3.2 Calvo New Keynesian

To study how price rigidity affects the dynamics of industry price responses to commodity price movements, we turn to a Calvo New Keynesian model allowing for myopia, as in Gabaix (2020). The setup is totally standard, following the textbook treatment in Gali (2010) at the industry level. Each firm $j \in [0, 1]$ in industry $i \in \{1, ..., I\}$ is permitted to change prices with some probability $(1 - \theta_i)$ in each period. The optimal reset price, $P^*_{i,j,t}$ that the firm sets when it gets the opportunity to change its output price maximizes expected discounted profits for as long as that price is expected to remain the firm's market price. Denote the stochastic discount factor, the relevant discount rate for firms, between periods t and $t+k$ by $SDF_{t,t+k}$. As before, the optimal reset price will not vary across firms within an industry, and it solves

$$\max_{P^*_{i,j,t}} \sum_{k=0}^{\infty} \theta_i^k \tilde{\mathrm{E}}_t \left[SDF_{t,t+k} X_{i,j,t} \left(\frac{P^*_{i,j,t}}{\mu_{i,t+k}} - MC_{i,t+k} \right) \right],$$

subject to the sequence of demand constraints

$$X_{i,j,t} = Y_{i,t} \left(\frac{P_{i,t}}{P^*_{i,j,t}} \right)^{\sigma_i}.$$

Figure 1 – Network Oil Shares for a Selection of Industries

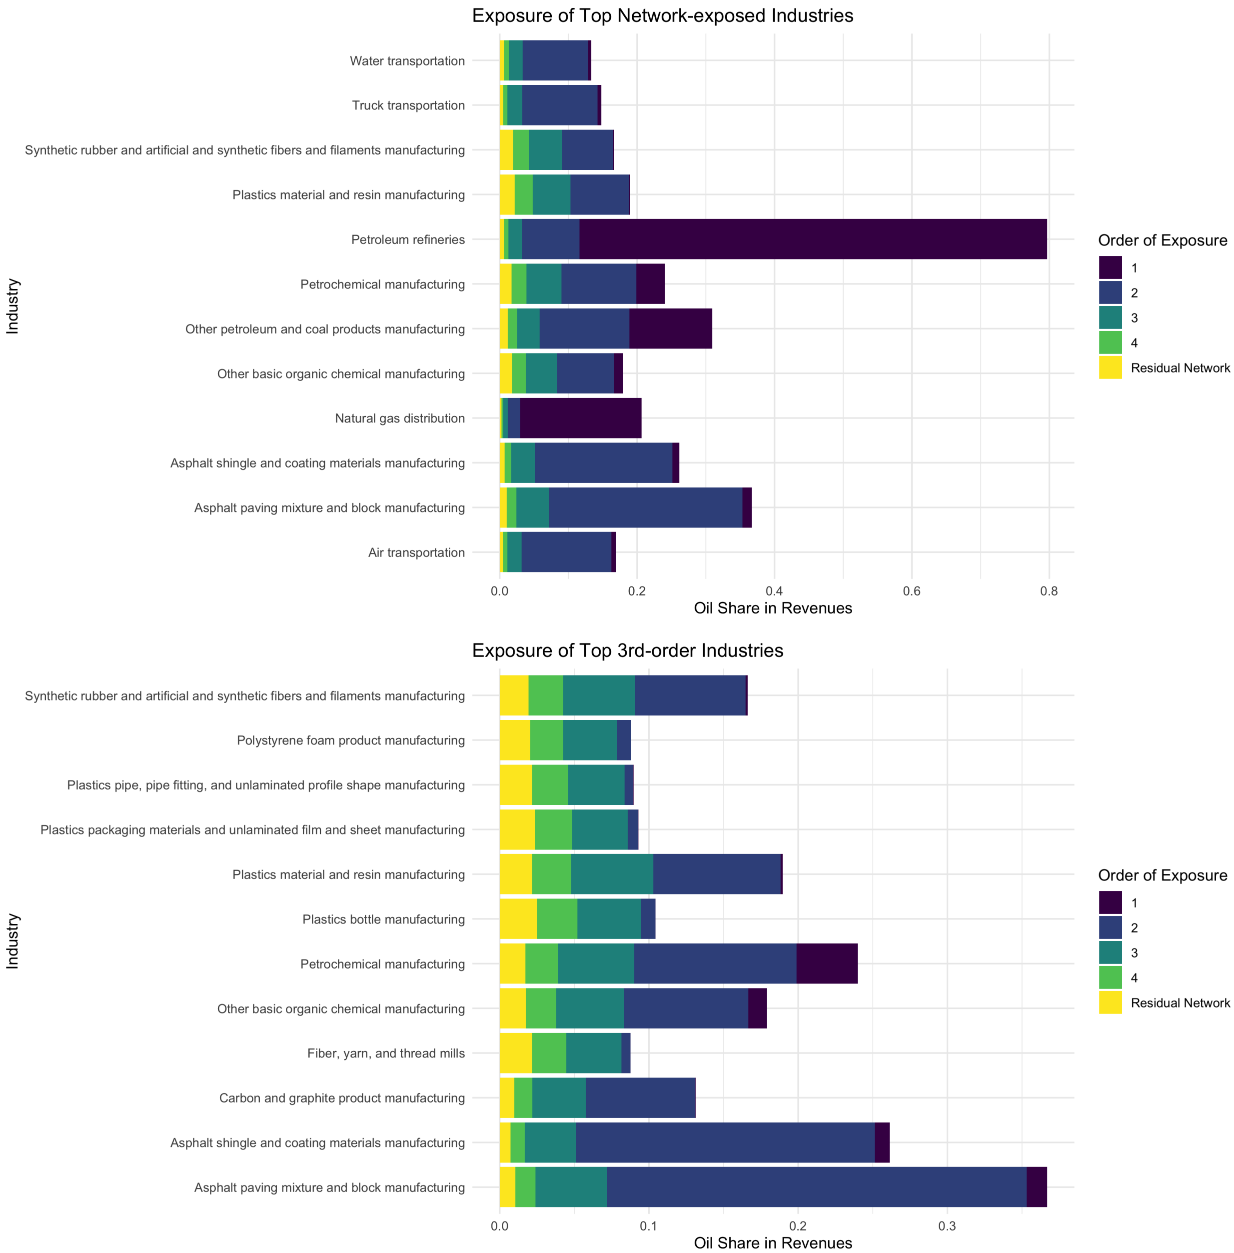

Note: This figure displays the total network share of oil in each industry's revenues for a selection of industries. Panel 1 displays the twelve industries with the highest network oil shares. Panel 2 displays the twelve industries with the highest third-order oil cost shares (i.e., indirect exposure to oil through suppliers' suppliers). Exposure of the natural gas distribution sector occurs because we are formally plotting industries' network exposures to the "oil and gas extraction" sector, the most disaggregated oil extraction sector available in the input-output data.

Now, $\tilde{\mathrm{E}}_t$ is the potentially myopic expectations operator given the period t information set and will be defined in a moment. It follows from random selection of which firms get the opportunity to change prices within industries and the definition of our earlier price index that

$$P_{i,t} = \left(\theta_i P_{i,t-1}^{1-\sigma_i} + (1-\theta_i)\left(P_{i,t}^*\right)^{1-\sigma_i}\right)^{\frac{1}{1-\sigma_i}}.$$

Now, rather than considering derivatives of log variables, which was sufficient in the flexible pricing model, we consider deviations of log variables from steady state, which we will denote by a hat above a lower-case variant of a variable. For example, the deviation of an industry's log price from steady state will be denoted $\hat{p}_{i,t}$. When taking the potentially myopic expectation of the deviation of a random variable from steady state, e.g., $\hat{p}_{i,t+k}$, with $k \geq 0$, the operator is defined as

$$\tilde{\mathrm{E}}_t[\hat{p}_{i,t+k}] = m_f^k \mathrm{E}_t[\hat{p}_{i,t+k}],$$

with $m_f \in (0,1]$ ($m_f = 1$ denoting rational expectations) and E_t being the rational expectations operator under the period t information set. Finally we define the notation $\mathrm{diag}(\cdot)$ to stack its argument in order across industries in a diagonal matrix. For example,

$$\mathrm{diag}(\theta_i) = \boldsymbol{I}\boldsymbol{\theta} = \boldsymbol{I}(\theta_1, ..., \theta_I)'.$$

We characterize the solution as follows:

Proposition 2. *The log deviation of optimal reset prices in an industry follows*

$$\hat{p}_{i,t}^* = (1-\theta_i\beta m_f)\sum_{k=0}^{\infty}(\theta_i\beta m_f)^k \mathrm{E}_t[\widehat{mc}_{i,t+k} + \hat{\mu}_{i,t+k}]. \tag{1}$$

Furthermore, the model solution for industry prices is the solution to a second-order expectational difference equation:

$$\begin{aligned}\mathrm{E}_t[\hat{\boldsymbol{p}}_{t+1}] = &\left(\mathrm{diag}\left(\frac{1+\theta_i^2\beta m_f}{\theta_i\beta m_f}\right) - \mathrm{diag}\left(\frac{(1-\theta_i)(1-\theta_i\beta m_f)}{\theta_i\beta m_f}\right)\boldsymbol{\Phi}\right)\hat{\boldsymbol{p}}_t - \frac{1}{\beta m_f}\hat{\boldsymbol{p}}_{t-1} \\ &- \mathrm{diag}\left(\frac{(1-\theta_i)(1-\theta_i\beta m_f)}{\theta_i\beta m_f}\right)(\boldsymbol{s}^Z\hat{p}_{Z,t} + \mathrm{diag}(s_i^L)\hat{\boldsymbol{w}}_t - \hat{\boldsymbol{a}}_t + \hat{\boldsymbol{\mu}}_t).\end{aligned}$$

It readily follows that if we have a one-time, fully persistent shock to the commodity price at time 0, and wages, TFP, and desired markups do not move, long-run industry prices converge to

$$\hat{\boldsymbol{p}}_\infty = (\boldsymbol{I}-\boldsymbol{\Phi})^{-1}\boldsymbol{s}^Z\hat{p}_{Z,0},$$

which was the solution for the on-impact effect of commodity shocks in the flexible price model.

Further, we can use equation 1 of Proposition 2 to see how myopia affects pass-through, a point we will develop further in our next subsection focusing on a linear network. Under rational expectations, industries understand that their marginal costs will increase gradually as upstream suppliers slowly pass through the shock at least in part according to their industry-specific frequency of price adjustment. They take this into account in their optimal reset price, meaning that firms that get the opportunity to adjust prices will partially pass through marginal cost increases that have not yet been realized. For equation 1, it is not problematic if we consider the case of complete myopia, $m_f = 0$, and in this case the equation collapses to

$$\hat{p}^*_{i,t} = \widehat{mc}_{i,t} + \hat{\mu}_{i,t}.$$

Under complete myopia, firms will only pass through changes in marginal cost that they observe. They are completely myopic with respect to the gradual pass-through of their upstream suppliers to the commodity price increase.

Further, it follows from equation 1 that the dynamics of pass-through for a hypothetical industry that is only directly exposed to the commodity price increase, i.e. that has no indirect network exposure to the commodity through its suppliers or suppliers' suppliers, etc., is unaffected by the distinction between myopia and rational expectations when the shock is fully persistent. Intuitively, there is no need for them to consider how gradually its suppliers are responding to the commodity price increase, as it only purchases the commodity directly. Myopia still plays a role if the shock is not fully persistent, since even this hypothetical industry will neglect that the price will revert to the steady state in the future.

We show in the appendix how, in a continuous time approximation, industry price movements can be decomposed into components due to direct (or first-order) exposure to oil and indirect (or second-order and higher) exposure to oil. Calibrating the frequencies of price adjustment to those measured in Pasten et al. (2017), which we describe in the data section, using rational expectations, and using a standard discount factor, we show the results of the model for a subset of industries in Figure 2.

There are several implications of the model worth differentiating in Figure 2. First and most directly by assumption, only some share of firms in each sector changes prices in each period, so pass-through is mechanically slower than in the flexible pricing case. Second, if two industries have the same frequency of price adjustment, pass-through will still be slower in the industry whose suppliers (throughout the network) have less frequent price adjustment. Finally, if firms are less myopic—more forward-looking about the gradual pass-through of their upstream suppliers—they will pass through the shock faster.

We seek to develop empirical measures of how an industry's "downstreamness" from the commodity shock should slow pass-through, as predicted by the model, as distinct from how het-

Figure 2 – Selected Results of the New Keynesian Model

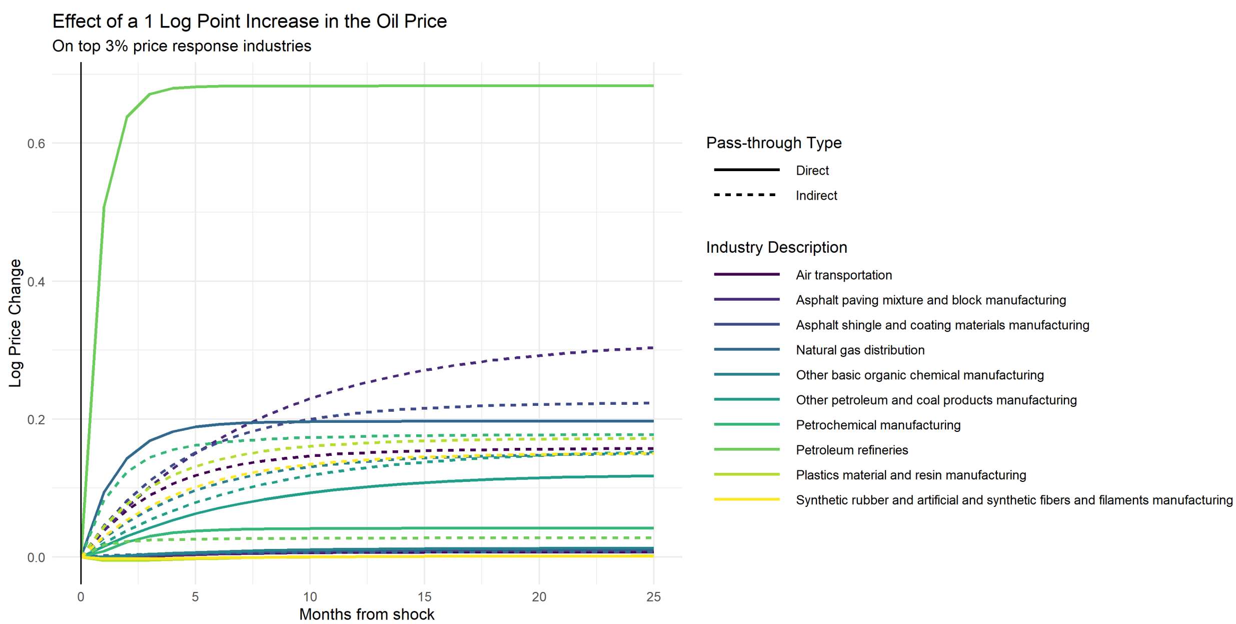

Note: We plot the effect of a log point oil price increase on the industries with the top 3% highest long-run price changes. Pass-through in some industries is fast, while in others it is gradual. Pass-through due to indirect exposure can be large. Pass-through in the natural gas sector is due to the shock originating in the "oil and gas extraction" sector, the most disaggregated sector available in the BEA's input output tables.

erogeneity in price adjustment frequency by industry should affect pass-through dynamics. To do this, we first measure how many months it takes each industry to reach 50% pass-through under the models with rational expectations and myopia separately. This is straightforward to compute from the model's solution. Formally, the fraction of long-run pass-through of a fully persistent commodity shock in sector i at time $t = 0$ can be denoted $b_{i,t}(m_f)$, with $\hat{p}_{i,t} = b_{i,t}(m_f)\hat{p}_{i,\infty}$. The time required to reach a fraction X of long-run pass-through, denoted $t_i(X; m_f)$, is implicitly defined as $b_{i,t_i(X;m_f)}(m_f) = X$. The half life under rational expectations is $t_i(.5; 1)$, while under myopia it is $\lim_{m_f \to 0} t_i(.5; m_f)$.

We compare these pass-through "half-lives" to industry pass-through half-lives in a model where each industry operates in isolation and is hit directly with the long-run marginal cost shock that will ultimately be realized as a result of the shock—formally, $\lim_{m_f \to \infty} t_i(.5; m_f)$, the time it would take to achieve 50% pass-through if industries were hyper forward-looking and focused only the long-run change in their marginal costs. We call these half-lives the "no network half-lives" since we have removed all of the time it would normally take upstream suppliers to pass through the commodity price shock to an industry's marginal costs.

An industry's downstreamness from a shock is then defined as the excess half life required under the myopic or rational expectations solution relative to the solution where the industry is affected immediately by its long-run marginal cost shock. The units are months: if, under rational expectations in the network, an industry takes 4 more months to reach 50% of long-run pass-through relative to the case where it is hit by its long-run marginal cost shock immediately, its downstreamness is 4 months. If, under myopia, the same industry requires 6 more months to reach 50% of long-run pass-through relative to the no-network half-life, its excess downstreamness due to myopia is 2 months. Formally, downstreamness under rational expectations is $t_i(.5; 1) - \lim_{m_f \to \infty} t_i(.5; m_f)$, and excess downstreamness due to myopia is $\lim_{m_f \to 0} t_i(.5; m_f) - t_i(.5; 1)$.

In the next subsection on the linear network, we show that excess downstreamness due to myopia increases in downstreamness under rational expectations; intuitively, the more that suppliers' gradual response to the shock slows down industry pass-through, the more that myopia can amplify the industry's response time. In Figure 3, Panel A, we show the half-lives of pass-through under rational expectations and myopia in the network, relative to the no-network half-life when industries are hit immediately with the long-run shock to their marginal cost. Half-lives are about 35% longer under rational expectations relative to the no-network case and just over two times longer under complete myopia. In Panel B, we show that excess downstreamness due to myopia is increasing in downstreamness due to rational expectations: an industry that takes 5 months longer to achieve 50% pass-through under rational expectations relative to the no-network case on average takes 7.5 additional months to achieve 50% pass-through under complete myopia.

Figure 3 – Half-lives of Oil Pass-through by Industry

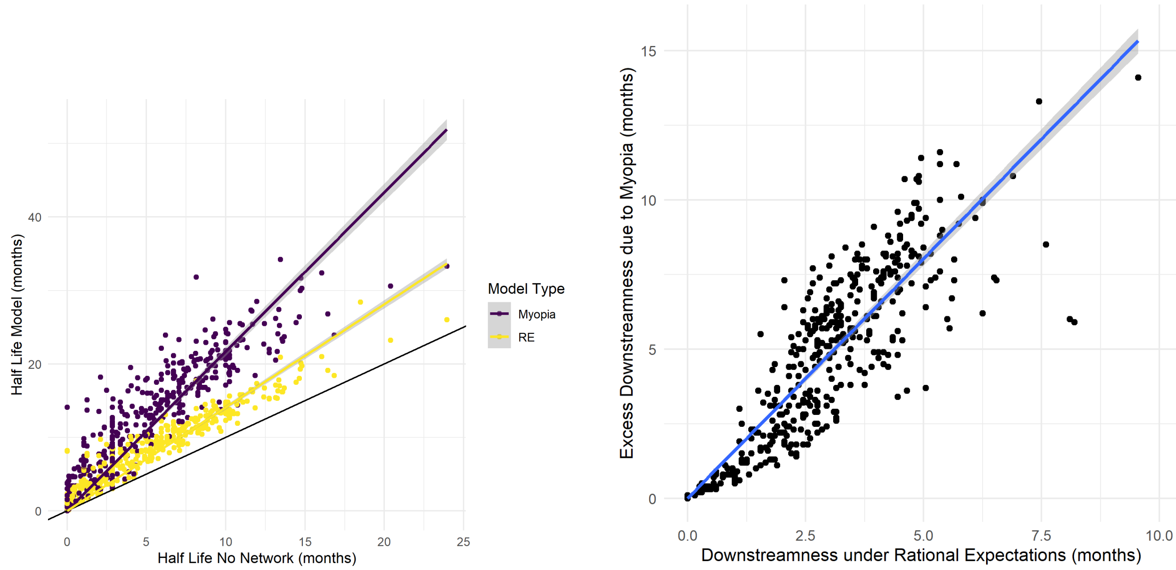

Note: In the left panel, we plot the time it takes each industry to reach 50% of long-run pass-through in response to a persistent oil price increase under the myopic and rational expectations network New Keynesian models against the half-life of pass-through if there were no network. In the right panel, we plot how much longer it takes industries to reach 50% pass-through under the myopic model as a function of their downstreamness from the shock under rational expectations.

3.3 Linear Network

In this section, we develop intuition on how expectations influence the dynamics of commodity shock pass-through with closed-form solutions in a linear network. A linear network of length N is comprised of industries $n \in \{1, ..., N\}$, each populated by a continuum of firms that only use labor and an input from industry $n-1$. Industry 1 uses labor and an exogenous supplied commodity with price denoted $P_{0,t}$. Industry n's price is denoted $P_{n,t}$. This setup is nested in our general setup from the previous subsection.

We make the assumption that industry wages, markups, and TFP do not move in response to a shock to the commodity price $P_{0,t}$, nor do firms' expectations of their future values. Further, all industries have the same frequency of price adjustment, $(1-\theta)$, and the same cost share in the intermediate input, s. From these assumptions, the log deviation of an industry's marginal cost from steady state simplifies from the general case,

$$\widehat{mc}_{i,t} = -\hat{a}_{i,t} + \hat{\mu}_{i,t} + s_i^L \hat{w}_{i,t} + (\boldsymbol{s}_i^X)' \hat{\boldsymbol{p}}_t + s_i^Z \hat{p}_{Z,t}$$

to

$$\widehat{mc}_{n,t} = s\hat{p}_{n-1,t}, \quad \hat{p}_{0,t} = \hat{p}_{Z,t}.$$

Now, the industry price index is still given by

$$\hat{p}_{n,t} = \theta\hat{p}_{n,t-1} + (1-\theta)\hat{p}^*_{n,t},$$

with the optimal reset price still given by equation 1 of Proposition 2, but now with changes in markups eliminated:

$$\hat{p}^*_{n,t} = (1-\theta\beta m_f)\sum_{k=0}^{\infty}(\theta\beta m_f)^k \mathrm{E}_t[\widehat{mc}_{n,t+k}].$$

We can eliminate marginal cost, yielding the following system to solve:

$$\hat{p}_{n,t} = \theta\hat{p}_{n,t-1} + (1-\theta)\hat{p}^*_{n,t} \tag{2}$$

$$\hat{p}^*_{n,t} = (1-\theta\beta m_f)\sum_{k=0}^{\infty}(\theta\beta m_f)^k \mathrm{E}_t[s\hat{p}_{n-1,t+k}]. \tag{3}$$

From now on, we drop the hats on all variables to simplify notation. First, consider long-run pass-through of a persistent unit shock to the commodity price, i.e. $\hat{p}_{0,t} = 1$ for all t. We have $p_{n,\infty} = p^*_{n,\infty}$ from equation 2, and it follows that $p_{n,\infty} = s^n$ from equation 3. Therefore, if $s = .5$, so that 50% of costs are in intermediate inputs, $p_{1,\infty} = .5$, $p_{2,\infty} = .25$, and so on. Intuitively, 25% of the second sector's network costs are in oil, while 75% are in labor (25% in its supplier's labor, and 50% in its own labor).

We now solve for the impulse response of prices in sector n to a fully persistent commodity shock at time 0. There is a closed form solution when we set $m_f = 0$, the case of complete myopia. Our solution will be for the object determining the fraction of long-run pass-through a sector has achieved: $b_{n,t}$, where

$$p_{n,t} = b_{n,t}p_{n,\infty}.$$

Finally, denote by $\phi = -\ln\theta$ the rate of price adjustment. Then

Proposition 3. *The fraction of long-run pass-through in sector n at time t for a persistent commodity shock occurring at time 0 is*

$$b_{n,t} \approx 1 - e^{-\phi t}\sum_{i=0}^{n-1}\frac{(\phi t)^i}{i!}, \tag{4}$$

where the equation holds with equality in continuous time. Denote by $t_n(X)$ the time at which $b_{n,t} = 1 - e^{-X}$, for $X \geq 0$. For large X,

$$t_n(X) \approx \frac{X}{\phi} + (n-1)\frac{\ln X}{\phi} + o\left(n\frac{\ln X}{\phi}\right), \tag{5}$$

where the equation again holds with equality in continuous time.

This result is useful for several reasons. First, one can simply plot impulse response functions for each sector n using equation 4, without having to solve a difference or differential equation.[2] Second, we have a characterization of how impulse response functions vary with n. The time required for an industry to achieve a given level of pass-through differs with n according to the factor $\ln X/\phi$. Note that, for $X = -\ln 0.5$, $t_n^*(X)$ is the half-life of pass-through we defined in the subsection on our general model under the assumption of complete myopia.

Now, we turn to the effect of forward-lookingness. While we will not have a closed-form solution in this case, we can still characterize $t_n^*(X)$. First, define the discount rate $\rho = -\ln\beta$ and the myopic rate $\tilde{m}_f = -\ln m_f$, and gather our three continuous time parameters into $\psi = \phi + \rho + \tilde{m}_f$. Then

Proposition 4. *Let $b_{n,t}^F$ denote the continuous time pass-through solution $p_{n,t} = b_{n,t}^F p_{n,\infty}$ when agents can be partially forward looking. Denote by $t_n^F(X)$ the time at which $b_{n,t}^F = 1 - e^{-X}$, for $X \geq 0$. Then*

$$t_n^F(X) = \frac{X}{\phi} + (n-1)\frac{\ln\left(X\frac{\psi}{\psi+\phi}\right)}{\phi} + o\left(n\frac{\ln\left(X\frac{\psi}{\psi+\phi}\right)}{\phi}\right).$$

We see now that the spacing between sectors is smaller, mediated by a factor $\frac{\ln\left(X\frac{\psi}{\psi+\phi}\right)}{\phi}$. Note that myopia is increasing in $\tilde{m}_f$, and $\tilde{m}_f = 0$ is rational expectations. Therefore, pass-through could scale with n to a degree $\frac{\ln\left(X\frac{\rho+\phi}{\rho+2\phi}\right)}{\phi}$ under rational expectations or as much as $\frac{\ln(X)}{\phi}$ with complete myopia.

Now, we visualize how much longer pass-through takes to occur in each sector in the supply chain. We set $\rho = .04$, in accord with $\beta = .96$, and we set $\phi = .7$, in accord with $\theta = .5$, or half of the firms in each industry getting the chance to change prices each period. Finally, we will consider rational expectations, a somewhat myopic case, and the fully myopic case. Figure 4 shows our results. As we stated in our discussion of Proposition 2, we see that $b_{1,t}$ does not vary with the degree of forward-lookingness. The rate of pass-through for every other sector, however, does vary, and the degree to which these rates vary is increasing in n, just as predicted under Proposition 4. This variation can be substantial. Take the most extreme case, $n = 7$, in the myopic model compared to the rational expectations model. By month four from the shock, there has been almost no pass-through under the myopic model, while, under rational expectations, pass-through is already around 50%.

[2]The mapping between discrete and continuous time is straightforward. Recall that the rate in continuous time is a constant, but in discrete time the rate varies with whether we are considering days, weeks, months, etc. We typically measure θ at a monthly frequency, and we will denote this measurement by θ_1. Then the continuous time rate is $\phi = -\ln\theta_1$. Generally, however, the rate used for a different time interval in discrete time would be $\theta_T = e^{-\phi T} = \theta_1^T$. So the daily rate of price adjustment sets $T = 1/30$ to get $\theta_{1/30} = \theta_1^{1/30}$. The discrete and continuous time solutions match up very well, even in the fully general model, for $T = 1/90$, or eight hours.

Figure 4 – The Influence of Expectations on Pass-through Dynamics

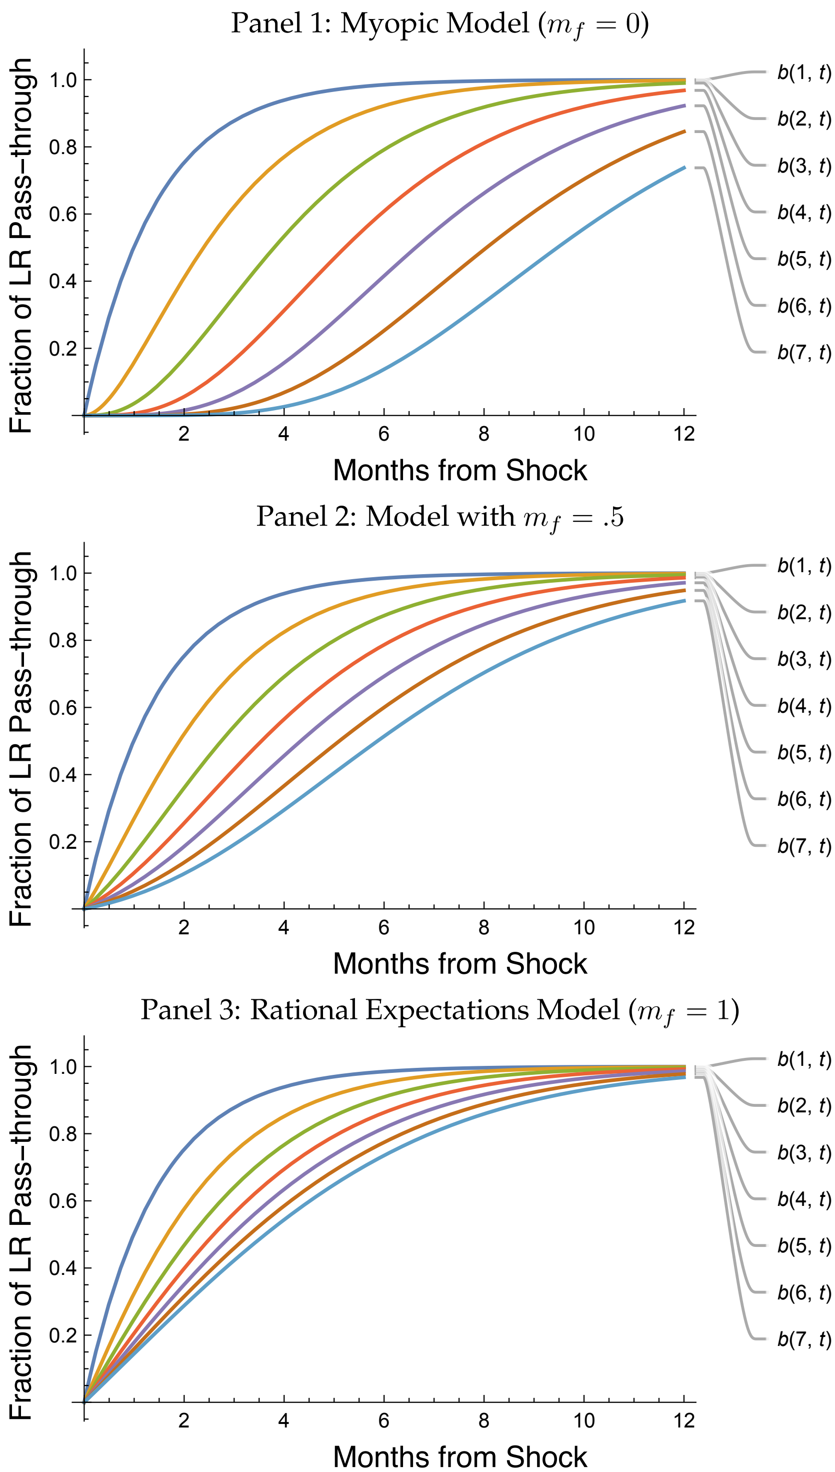

Note: We display the time it takes each of seven sectors in the linear network to pass through a persistent shock to the upstream price. More downstream sectors take more time to pass through the shock, and the amount of additional time is increasing in the extent of myopia.

4 Empirics

In order to test whether full pass-through of commodity shocks deep into the production network - as suggested by the preceding frameworks - actually exists in practice, we turn to the data. We begin by using the standard data source measuring the network dependencies of the macroeconomy – the input-output tables compiled by Bureau of Economic Analysis (BEA) – along with a regression setup directly informed by the flexible-price model. We initially focus on the effects of oil price changes in a reduced-form setting, but we branch out to perform a variety of other analyses using other sources of variation and regression setups.

4.1 Data

Since 1939, the Bureau of Economic Analysis (BEA) has periodically compiled input-output tables of the US economy, representing the interdependencies between industries. Specifically, the tables display the extent to which the output a given sector is used as an input by every other sector in the economy – or consumed by final demand from individuals or the government. Beginning in 1967, the level of granularity of the published tables was considerably increased. Since then, every five years, the BEA has published a detailed input-output table approximately 400 sectors in size. For our analysis, we make use of all input-output tables from 1977 onward. For a case study on the 1979 oil price shock, we will also use the input-output table published for 1977. We outline our processing of the input-output tables in Appendix B.

Panel 1 of Figure 1 plots the twelve industries with the highest network oil share; it also decomposes the total network oil share into first-order (i.e. direct) exposure to crude oil, second-order (i.e., through suppliers) exposure to crude oil, third-order (i.e., through suppliers' suppliers) exposure to crude oil, and beyond. For example, the Petroleum Refineries industry has nearly 80% of its total costs in oil, and nearly all of these costs constitute direct purchase of crude oil. The Asphalt Paving industry has nearly 50% of its total costs in oil, but almost none of these costs are direct purchase of crude oil; primarily, the industry purchases refined oil from refineries, who themselves bought crude oil (i.e., second-order exposure). Panel 2 of Figure 1 plots the twelve industries with the highest third-order oil share. Some sectors with high third-order exposure to oil – such as Petrochemical Manufacturing – also have high first-order and second-order exposure, whereas others – such as Polystyrene Manufacturing – have very little first- or second-order exposure. In short, there is considerable variation across sectors in both the network oil share itself and the breakdown of the network share into different orders of exposure.

Our second major source of data is Bureau of Labor Statistics (BLS), which publishes monthly data on prices (the Producer Price Index, or PPI) by industry at a great many different levels of granularity. The BLS compiles these data from price micro data it collects on an extensive variety of individual goods. By weighting these individual price series, the BLS generates its industry-level PPI data – a procedure described in detail in BLS (2011). While some industry PPIs extend back to the 1940s, it is sparse in early decades; most sectors have no price data. Coverage improves considerably over time, with the bulk of the expansion occurring in the 1960s and 1970s. The BLS measures industry prices in a given period by taking the average price of all transactions occurring in that industry and period - inclusive of both transactions on the spot market and transactions at (previously-agreed) contracted prices. This data is consequently able to provide an accurate picture of the true pace of price pass-through.

The industries in the BEA input-output tables are identified by BEA codes, whereas the industries in the BLS PPI data are primarily identified by SIC codes prior to 1997 and by NAICS codes after that date.[3] The BEA released correspondences between SIC codes and BEA codes with each input-output table through 1992, and correspondences between NAICS codes and BEA codes are publicly-available through the BEA for the 1997 table onward. By utilizing these various correspondences, it is possible to merge the BLS industry price data with the BEA input-output tables.

Additionally, we obtain data compiled by Pasten, Schoenle, and Weber (2017) on the frequency of price adjustment.[4] The authors gained access to the BLS price micro data and computed the average frequency of price changes by industry.

4.2 Reduced-form Analysis

4.2.1 Regression Specifications

Our regression strategy is directly informed by the flexible-price model. We run two main versions of this regression – one for our regressions using pooled variation, another for our individual case studies. Beginning with the regression equation used pooled variation, in the framework of Proposition 1, we measure

$$\Delta(\boldsymbol{\mu}_t - \boldsymbol{a}_t) = \Delta \boldsymbol{p}_t - \boldsymbol{\Phi}_t^X \Delta \boldsymbol{p}_t$$

as the price change that remains after residualizing the cost-share weighted sum of input cost changes in the US industry producing that input.[5] Note that we can compute these shocks for all industries, though we will generally use shocks only for certain commodities whose prices are thought to be less dependent on movements in U.S. demand. For our regressions on oil alone, we note that using our measured shocks as treatments and shocking the oil price directly yield essentially the same regression estimates below. The resulting regression specification is given

[3]Some SIC PPI data continued to be reported – with decreasing coverage – through the mid-2000s. And the NAICS PPI data has been extended backward into the early 1990s and before for a limited set of industries.

[4]We kindly thank the authors for sharing their data with us.

[5]Our results are robust to additionally residualizing changes in seasonally-adjusted industry wages.

by

$$\Delta p_{i,t} = \lambda_t + \sum_{j=-6}^{12} \beta_j \left[(\boldsymbol{I} - \boldsymbol{\Phi}_t)^{-1} \Delta(\boldsymbol{\mu}_{t-j} - \boldsymbol{a}_{t-j}) \right]_i + \epsilon_{i,t} \tag{6}$$

where $p_{i,t}$ denotes the price of industry *i* at time *t*, λ_t is a time fixed-effect, $[(I - \boldsymbol{\Phi}_t\ )^{-\mathbf{1}}]_i$ represents the network cost shares of industry i in all commodities, and $\epsilon_{i,t}$ is an error term. To shock just oil, we set the composite shock $(\boldsymbol{\mu}_{t-j} - \boldsymbol{a}_{t-j})$ to 0 for all other industries, and to shock non-oil commodities we set $(\boldsymbol{\mu}_{t-j} - \boldsymbol{a}_{t-j})$ to 0 for the oil sector and all other non commodity-producing sectors. Note that $\beta_t = 1$ corresponds to full pass-through, as the right-hand-side variable of interest corresponds to the size of the cost shock experienced by industry *i*. Note that the flexible price model predicts specifically that $\beta_0 = 1$. Furthermore, we run an instrumental variables (IV) version of this regression which instruments oil price changes with Kanzig's (2021) series of exogenous oil price shocks.

It is also possible to separately analyze the price pass-through of direct (first-order) exposure and indirect (higher-order) exposure to cost shocks with a slightly modified version of the preceding specification.

$$\Delta p_{i,t} = \lambda_t + \sum_{j=-6}^{12} \gamma_j [\boldsymbol{\Phi}_t \Delta(\boldsymbol{\mu}_{t-j} - \boldsymbol{a}_{t-j})]_i + \sum_{j=-6}^{12} \theta_j [((I - \boldsymbol{\Phi}_t)^{-1} - \boldsymbol{\Phi}_t)\Delta(\boldsymbol{\mu}_{t-j} - \boldsymbol{a}_{t-j})]_i + \epsilon_{i,t} \tag{7}$$

Here, the network share of industry i's costs that are due to commodity/commodities is decomposed into the direct and indirect components of the network cost share. We exclude all industries that are directly shocked from the regression. The coefficients γ_t correspond to the extent of pass-through of direct exposure to cost shocks; the coefficients θ_t correspond to the extent of pass-through of indirect exposure to cost shocks. This specification makes it possible to investigate whether any price pass-through is driven entirely by direct exposure or whether cost shocks mediated through potentially-complex network structures are equally passed into prices.

Turning away from the pooled variation, we slightly modify the initial regression specification to make it more suitable for case studies wherein price variation is driven by large narrative shocks. In a variant of an event-study difference-in-differences specification, we regress the price change in an industry on a time fixed-effect and the network share of the industry's costs that are due to the commodity or commodities of interest. Specifically,

$$\sum_{j=0}^{t} \Delta p_{i,j} = \lambda_t + \beta_t [(I - \boldsymbol{\Phi}_t)^{-\mathbf{1}} \boldsymbol{\Phi}_t^Z]_i + \epsilon_{i,t} \tag{8}$$

where $p_{i,t}$ denotes the price of industry *i* at time *t*, $\Sigma_{j=0}^{t} \Delta p_{i,j}$ denotes the cumulative change in the price of industry *i* from some designated base period 0 through period *t*, λ_t is a time fixed-effect, and $[(I - \boldsymbol{\Phi}_t^X\ )^{-\mathbf{1}}\ \boldsymbol{\Phi}_t^Z\ ]_i$ represents the network cost share of commodity or commodities **Z** in industry i, and $\epsilon_{i,t}$.

To provide an intuitive description of this specification, note that full pass-through of a commodity cost shock into prices means something very different for an industry with 10 percent of its costs originating from that commodity than it does for an industry with 50 percent of its costs originating from that commodity. The former industry should experience an increase in its prices one-tenth the increase in the commodity price, whereas the latter industry should experience an increase in its prices one-half the increase in the commodity price. This is what the network share on the right-hand-side of the regression equation adjusts for. Concretely, if there is a cumulative $\sum \Delta p_{Z,t}$ log-point increase in the price of our commodity of interest, full pass-through would imply a coefficient value of $\beta_t = \sum \Delta p_{Z,t}$. This regression specification is well-suited for case studies, as it allows us to plot the values of β_t for each time period t against the cumulative increase $\sum \Delta p_{Z,t}$ in the commodity price itself.

As before, it is possible to decompose the right-hand-side variable of interest into direct and indirect exposure to cost shocks in order to study these separately:

$$\sum_{j=0}^{t} \Delta p_{i,j} = \lambda_t + \gamma_t[\mathbf{\Phi}_t^Z]_i + \theta_t[(I - \mathbf{\Phi}_t)^{-\mathbf{1}}\mathbf{\Phi}_t^Z - \mathbf{\Phi}_t^Z]_i + \epsilon_{i,t} \tag{9}$$

We note that in all of our regressions, the shocked industries will be excluded from the regression analysis. For example, when we consider shocks to the oil price, we will only be interested in how that shock propagates to non-oil industry prices. Including oil in the regression introduces a source of mechanical dependence in the analysis; most obviously, in the commodity price approach, including the oil sector would mean regressing a change in the oil price on the oil sector's network cost share in oil, multiplied by the change in the oil price. Similarly, when pass-through of shocks to multiple commodities is assessed jointly, all of the industries producing these commodities are excluded from the regression analysis. We are only interested in how those commodity price movements affect non-commodity industry prices.

Finally, we provide some intuition for interpreting our regression coefficients in light of our inclusion of time fixed effects. Consider regression specification (9), our case study specification separating first-order from higher-order exposure. The coefficient θ_t measures a *relative* effect, answering the following question: holding first-order exposure to oil constant, how much more do industries with high indirect exposure to oil change their prices relative to industries with low indirect exposure to oil? This relative measurement emerges because the time fixed effect purges any national effects on all industries' prices, such as those related to inflation, the Federal Reserve's response to oil price movements, or oil price movements' effects on inflation expectations. Formally, oil price increases may cause price increases even in industries with no network exposure to oil (though we will see later that they do not), but this effect will be missed in our estimates to the degree it affects all industry prices equally in each time period. This aspect of the measurement is desirable for us because it allows us to focus specifically on the network model's predictions about relative oil price pass-through across industries.

4.2.2 Main Results

We begin by using all variation in oil prices from 1997 onward, running the regression specification given by Equation (6). The 1997 BEA input-output table is the first table with BEA codes based on the NAICS classification, and most all BLS PPI series have become available in NAICS format by 1997 as well.[6]

The red coefficients in Panel 1 of Figure 5 plot the results of this specification month-by-month. There is no evidence of any pre-trend in the months prior to impact of the shock. Then, in the month of impact, roughly 50% of the shock is passed through into prices. Over the course of the next few months, pass-through increases until reaching 100%.

There are, of course, a wide variety of commodities other than oil which are also of substantial importance in US supply chains. Consequently, we pool all commodities apart from oil and compute the network share in all other commodities by industry. Using all price variation non-oil commodities since 1997, we then run a modified version of the previous regression regression. The blue coefficients plot the results of this specification using non-oil commodity price variation instead of oil price specification. The results are nearly identical in both the time pattern and extent of pass-through.

The bottom panel includes both the oil price variation and non-oil commodity price variation as separate terms in the same regression to deal with any potential omitted variable bias. The results scarcely change relative to the top panel, suggesting omitted variable bias is not a serious concern here.

Figure 6 plots the results of the regression specification given by Equation (7), decomposing total network exposure to oil price variation into direct and indirect exposure. Panel 1 again focuses on all oil price variation. Here, the red coefficients correspond to pass-through of direct exposure to oil shocks, whereas the blue coefficients correspond to pass-through of indirect exposure through the network to oil shocks (i.e., total network exposure minus direct exposure). The results reveal no evidence of substantial pre-trends. That is, in the months prior to an oil price movement, coefficients are not substantially positive or negative. In month 0, on impact of the shock, a high degree of direct pass-through occurs (approximately 75%). One month after impact of the shock, additional direct pass-through occurs (approximately 25%). At this point, after just a couple months, full pass-through has already occurred. This contrasts with the pattern of indirect

[6]Both the BLS and the BEA recommend against attempting to merge NAICS codes with the older, pre-1997 SIC codes, as the underlying industries the codes describe – even at the most granular level – are fundamentally not comparable in many cases.

Figure 5 – Month-by-Month pass-through of Oil & Non-Oil Commodity Price Changes

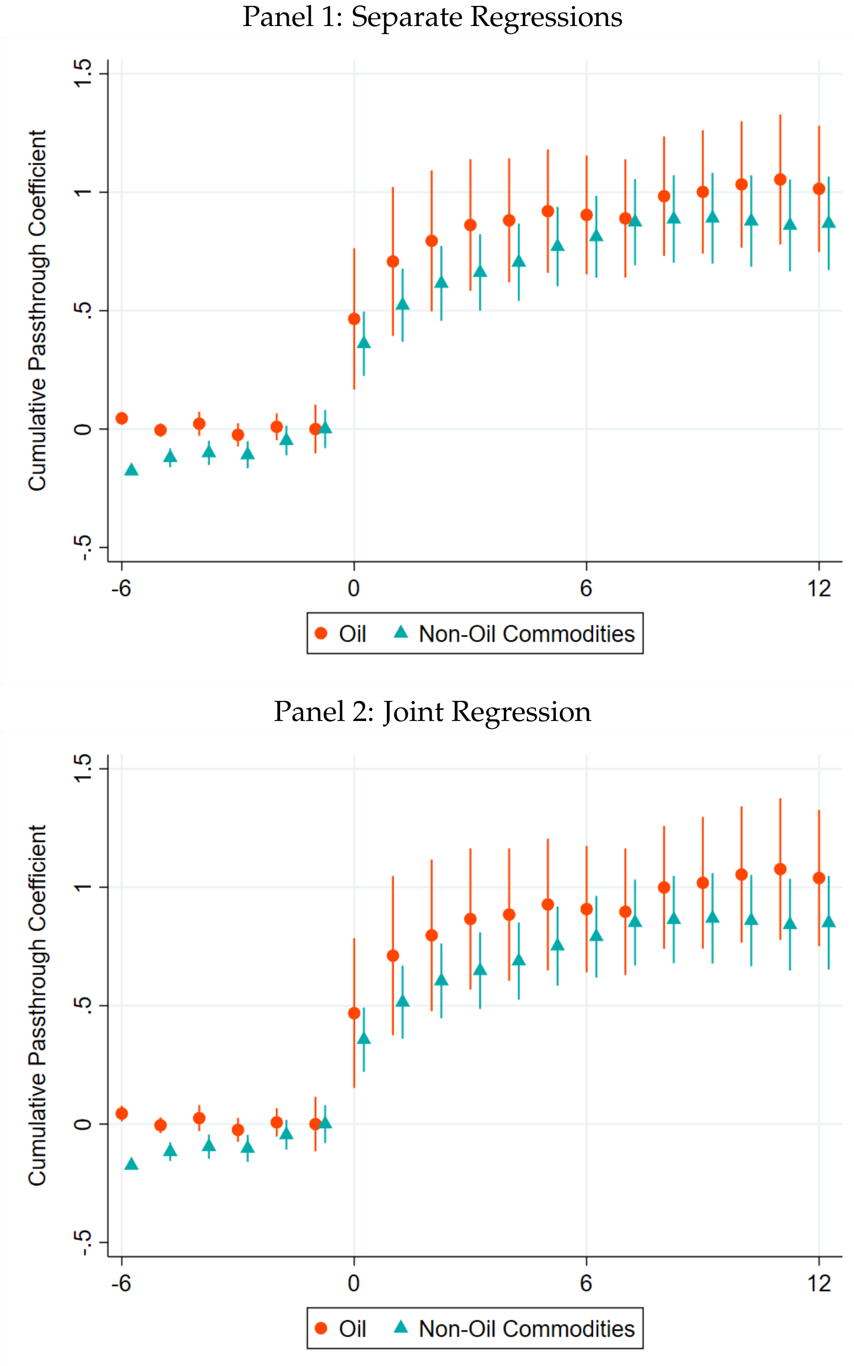

Note: The top panel plots two separate regression specifications, both corresponding to Equation (6). Red coefficients plot monthly price pass-through of network exposure to oil price changes; blue coefficients plot monthly price pass-through of network exposure to non-oil commodity price changes. The bottom panel includes these two separate sets of regressors in the same regression - one for which the shocked commodity is oil and the other for which the shocked commodities are non-oil commodities. The coefficients plot monthly cumulative price pass-through of total network exposure to crude oil price shocks.

pass-through, of which very little occurs on impact. Instead, pass-through phases in slowly over the course of a half-year or so.

Using all variation in oil prices and non-oil commodity prices helps demonstrate that full pass-through is not merely unique to a specific context. However, it may raise endogeneity concerns if, for example, oil price changes correlate with other variables that also disproportionately affect prices in sectors with a high network oil share. In particular, our flexible-price framework implies that TFP and wage changes are in the error term of these regressions. By using more exogenous sources of oil price variation that are unlikely to be correlated with TFP or domestic wage changes, we aim to minimize concerns relating to omitted-variable-bias. Consequently, Panel 2 of Figure 6 turns to the oil shock series of Kanzig (2021). The shock series is formed through high-frequency identification of the effects of OPEC announcements on oil prices. We use these shocks in a two-stage least-squares instrumental variables version of the regressions in the previous section – instrumenting the change in the oil price with Kanzig's shock series. Using this variation, pass-through is again 100%, but the dynamics are slightly different - virtually all of it occurs on impact.

In Appendix B, we show that these results are robust to a variety of modifications and alternative approaches. In B.1, we show that the same patterns are evident in a binscatter analysis. In B.2, we recompute measures of our shock variable to account for wage changes. In B.4, we control for gas and electricity exposure - two commodities likely to be close substitutes for oil in some contexts. In B.5, we use alternative cost shares that exclude payments to capital from the share denominator. In B.6, we study the propagation of non-commodity shocks (i.e., sectors further downstream than commodities with NAICS codes beginning in 23 or higher). In B.7, we conduct permutation tests on our main specifications as an alternative robust method of generating p-values from within-sample.

4.2.3 Case Study Results

As an additional approach to isolating plausibly-exogenous variation in oil prices, we examine a few case studies - major movements in oil prices known from the historical record to have been unanticipated. We begin our case studies with the 1979 oil crisis. BLS PPI data for the Petroleum Refineries industry – one of the most crucial links in the oil supply chain – does not exist for the earlier 1970s, nor does PPI data for many other oil-consuming sectors; consequently, the 1979 shock is the earliest we are able to analyze reliably. The shock occurred as a result of the 1979 Iranian Revolution. In the aftermath of the 1973 oil shock, Iran had increased its oil production in order to dampen the loss of oil exports from Arab nations to the West. Consequently, Iran became one of the most important oil exporters to Western economies. With the overthrow of Shah Mohammad Reza Pahlavi and the reconstitution of Iran as an Islamic Republic under Ayatollah Khomeini, Iranian oil production underwent a massive decline and, even after a partial recovery,

Figure 6 – Month-by-Month pass-through of Oil Price Changes: Direct vs. Indirect

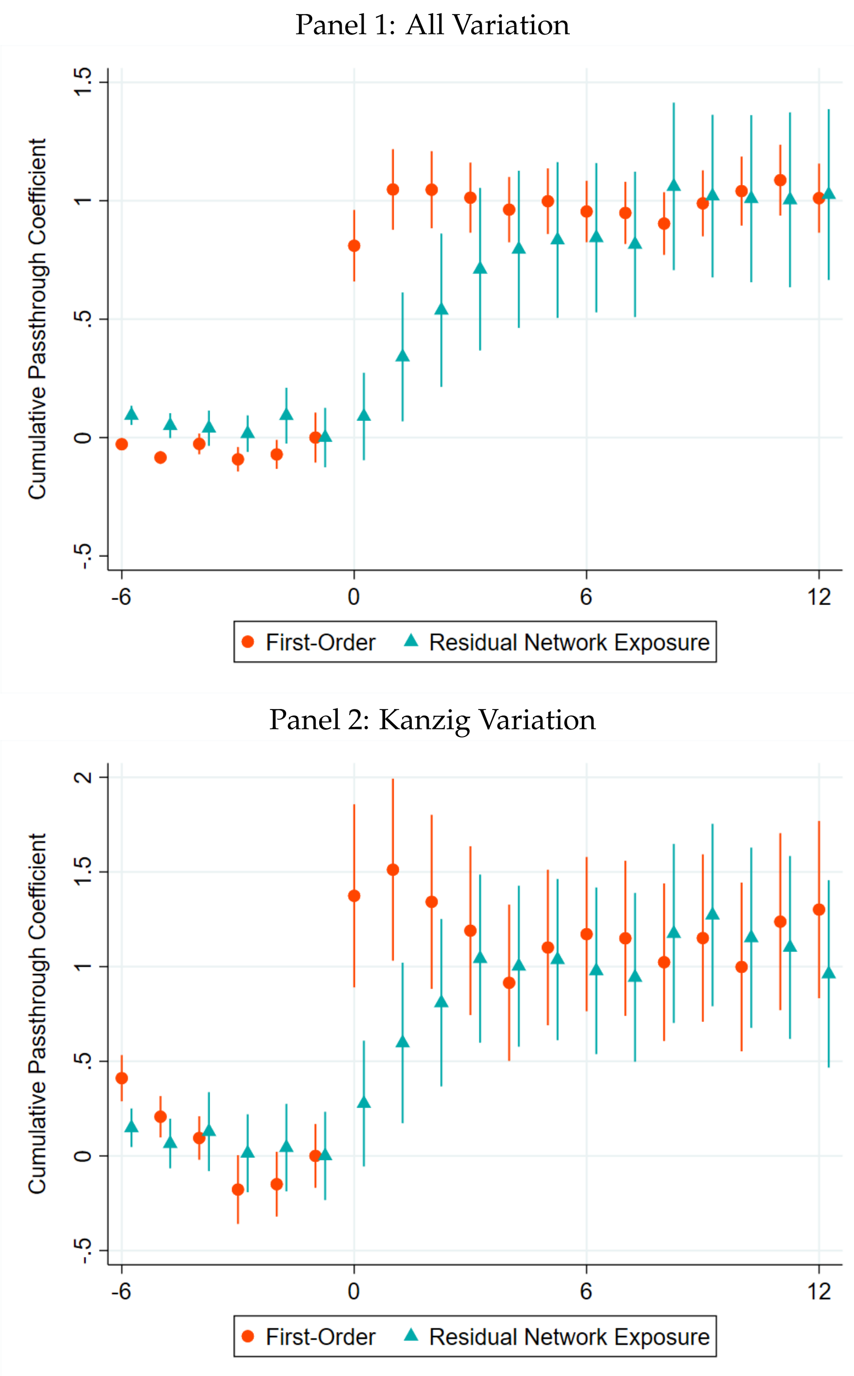

Note: Regression specifications correspond to Equation (7). Red coefficients plot monthly price pass-through of first-order exposure to crude oil price shocks; blue coefficients plot monthly price pass-through of second- and higher-order network exposure to crude oil price shocks.

exports to Western nations remained relatively low. This higher level of oil prices was thus mostly maintained until OPEC increased production in the mid-1980s.

Panel 1 of Figure 7 displays the regression coefficients originating from applying Equation (7) to the 1977-82 period surrounding the 1979 oil shock. The black line plots the monthly average spot price of West Texas Intermediate (WTI) crude oil. The red coefficients plot the extent of pass-through for direct exposure to oil, whereas the blue coefficients plot the extent of pass-through for indirect exposure to oil (i.e., network exposure minus direct exposure). As can be seen in the figure, by the end of the period of our case study, the coefficients are statistically indistinguishable from the black line representing the WTI spot price. In other words, both direct exposure and indirect exposure through the production network to the 1979 oil shock is fully passed through into industry prices. While standard error bars are wider on indirect pass-through relative to direct pass-through, point estimates are similar. There does appear to be some evidence of a slight lag in indirect pass-through.

We next turn to another case study: the 2014 oil shale boom. Despite some relief from the all-time peak in oil prices of nearly $150/barrel that occurred in 2008, oil prices remained near all-time highs throughout the early 2010s. They averaged over $90 a barrel between 2011 and 2014. These high prices coupled with the low-interest-rate regime in the aftermath of the Great Recession created a strong incentive for U.S. companies to invest in exploration and extraction of a source of oil theretofore untapped due to its comparative expense: shale oil. As shale oil extraction ramped up, US oil production expanded considerably in 2014-15, and OPEC announced that it would continue pumping oil at the same volumes to maintain marketshare – and, some have argued, to drive the shale oil producers out of business. This led to a considerable drop in oil prices over 2014-15 to a lower level that was mostly maintained for several years thereafter.

Panel 2 of Figure 7 displays the regression coefficients originating from applying Equation (9) to the 2012-17 period surrounding the 2014-15 oil shale boom. As before, the black line plots the monthly average spot price of West Texas Intermediate (WTI) crude oil, the red dots plot direct pass-through coefficients, and the blue dots plot indirect pass-through coefficients. The takeaway is the same: by the end of period of the case study, full pass-through of the shock into prices has occurred – in the case of both direct and indirect exposure to the shock. Once again, there is some evidence of a lag in the pass-through of indirect exposure to the shock; whereas the direct pass-through coefficients track the price of WTI crude very precisely, the indirect pass-through coefficients do not follow every small monthly variation but rather trace out a smoother curve that converges to the same point over time.

The final case study on which we focus is the 2020 COVID shock. During 2018 and 2019, the price oil averaged approximately $60 per barrel. In the early months of 2020, as it became apparent that COVID was becoming a global pandemic and that many nations would respond with large-scale shutdowns of economic activity in order to control the spread of the disease,

Figure 7 – Case Studies

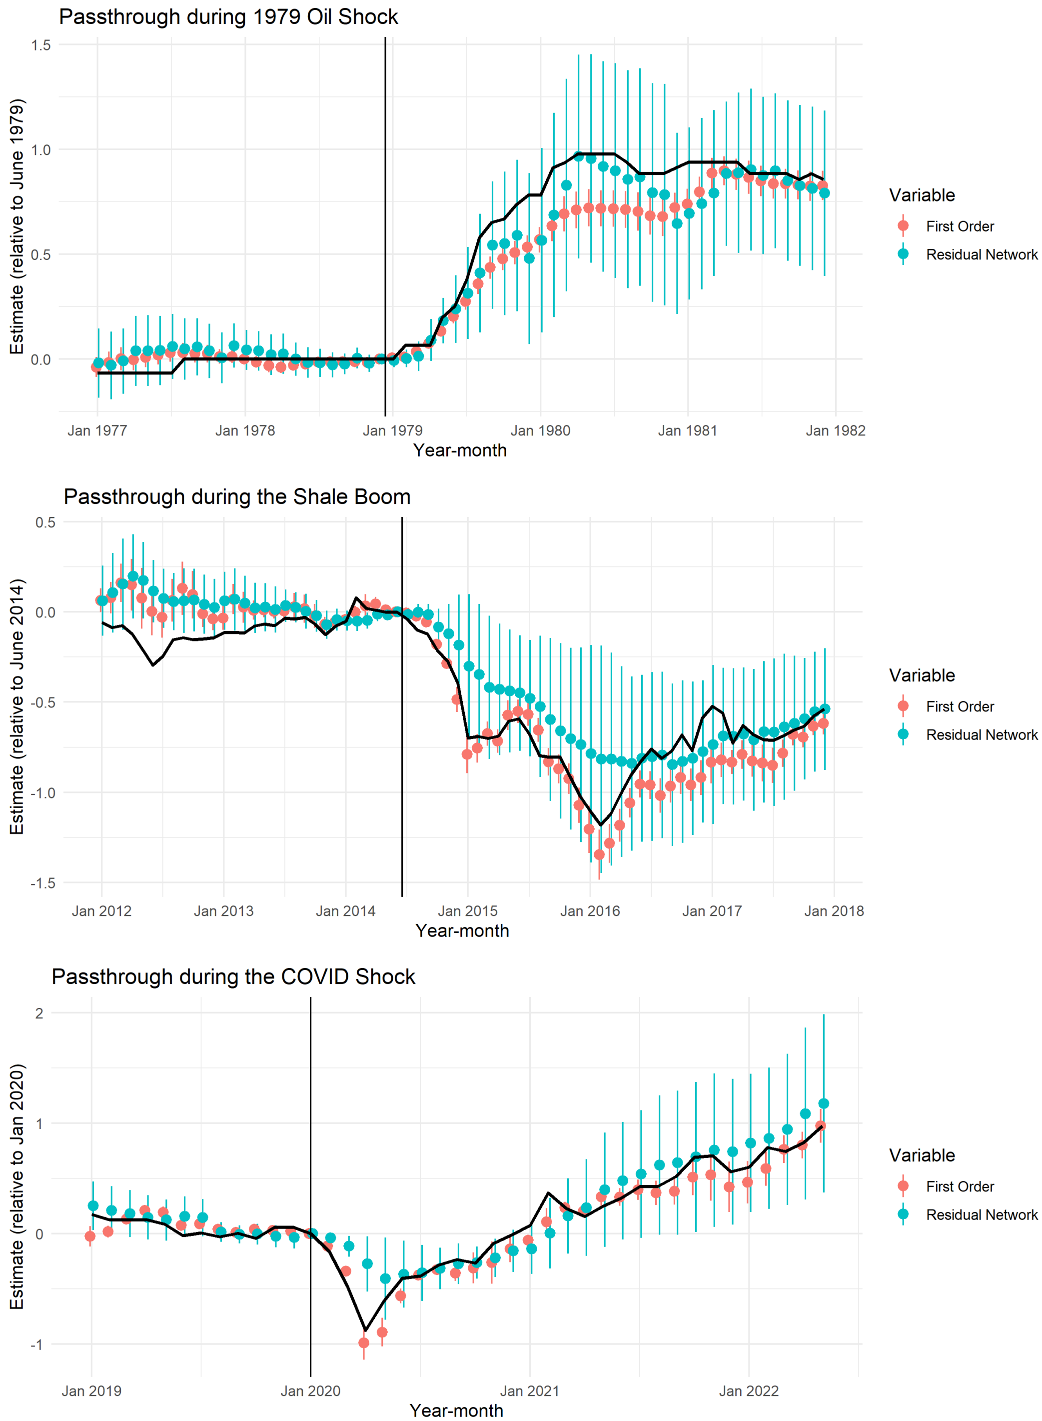

Note: Regression specification corresponds to Equation (9). Black line plots West Texas Intermediate (WTI) crude oil price, red coefficients plot cumulative industry price pass-through of direct/first-order exposure to crude oil price changes, and blue coefficients plot cumulative industry price pass-through of indirect/residual network exposure to crude oil price changes.

the price of oil plummeted, averaging \$20 per barrel in April and May of that year. Prices even briefly turned negative as producers scrambled to take production offline as soon as possible. However, the recovery from the COVID recession proved to be quicker than many anticipated, and demand for oil quickly rebounded as much production remained offline. Consequently, prices began to rebound, reaching pre-COVID levels by early 2021 and exceeding \$100 per barrel by early 2022.

Panel 3 of Figure 7 displays the regression coefficients originating from applying Equation (9) to the 2019-22 period surrounding the 2020 COVID shock. Again, the black line plots the monthly average spot price of WTI crude oil, the red dots plot direct pass-through coefficients, and the blue dots plot indirect pass-through coefficients. Yet again, by the end of the period of the case study, full pass-through has occurred in the case of both direct and indirect exposure to the shock, and there is evidence of a lag in the pass-through of indirect exposure - the indirect pass-through coefficients are smoothed over the depths of the downturn in March, April, and May of 2020.

4.2.4 Heterogeneity

We investigate some key dimensions of heterogeneity by interacting variables of interest with our price shocks. Specifically,

$$
\begin{aligned}
\Delta p_{i,t} = \lambda_t + heterogeneity_i + \sum_{j=-6}^{12} \beta_j \left[(\boldsymbol{I} - \boldsymbol{\Phi}_t)^{-1} \Delta(\boldsymbol{\mu}_{t-j} - \boldsymbol{a}_{t-j}) \right]_i \\
+ \sum_{j=-6}^{12} \beta_j \left[(\boldsymbol{I} - \boldsymbol{\Phi}_t)^{-1} \Delta(\boldsymbol{\mu}_{t-j} - \boldsymbol{a}_{t-j}) \right]_i * heterogeneity_i + \epsilon_{i,t}
\end{aligned}
\tag{10}
$$

We first confirm that our measure of the frequency of price adjustment is indeed predictive of the extent of price pass-through. In Table 1, we display the results of regression specifications with the no-network half-life of price adjustment as an interaction term in columns (1), (2), and (3). Whether we study all variation in oil prices, the Kanzig IV variation in oil prices, or variation in non-oil commodity prices, the interaction term is strongly significant. A higher half-life (i.e., lower price adjustment frequency) is associated with a lower level of price pass-through over a one-month horizon.

Table 1 – Heterogeneity Analysis

Dependent Variable: ΔPPI	(1) Oil	(2) Kanzig	(3) Non-Oil	(4) Oil	(5) Kanzig	(6) Non-Oil	(7) Oil	(8) Kanzig	(9) Non-Oil
Cost Change	0.641***	0.912***	0.349***	0.654***	0.937***	0.365***	0.380***	0.071	0.452**
	(0.063)	(0.241)	(0.097)	(0.051)	(0.227)	(0.097)	(0.116)	(0.224)	(0.190)
Cost Change*Price Adjust. Half-Life	-0.128***	-0.147***	-0.051***	-0.110***	-0.122***	-0.044***	-0.108***	-0.084**	-0.059***
	(0.017)	(0.038)	(0.015)	(0.016)	(0.033)	(0.013)	(0.019)	(0.034)	(0.017)
Cost Change*Downstreamness				-0.188***	-0.273**	-0.070			
				(0.050)	(0.109)	(0.055)			
Cost Change*SD[Marginal Cost]							5.152***	17.468***	-3.566
							(1.702)	(3.174)	(7.445)
R-Squared	0.1105	0.1099	0.0726	0.1115	0.1105	0.0728	0.1112	0.1062	0.0728
Observations	89,429	79,511	1,582,489	89,429	79,511	1,582,489	89,429	79,511	1,582,489

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The outcome variable of this table is the month-over-month change in industry prices. Columns (1) through (3) investigate heterogeneity on the (no-network) half-life of price adjustment, finding that a higher half-life (i.e., lower price adjustment frequency) is associated with a lower level of price pass-through. Columns (4) through (6) add downstreamness as another heterogeneity variable, finding that it matters above and beyond price adjustment frequency alone; downstream industries have lower levels of price pass-through over a one-month horizon. Columns (7) through (9) study the standard deviation of an industry's marginal costs as an interaction term. Because the half-life of price adjustment remains significant, said half-life measures more than just the volatility of an industry's marginal costs.

Columns (4), (5), and (6) repeat this exercise, adding downstreamness as another interaction term. In the case of all oil price variation and Kanzig IV variation, the downstreamness interaction term is negative and statistically significant, and in the case of the non-commodity price variation, it remains negative but falls short of significance at traditional levels, with a p-value of 0.2. Industries further downstream have lower levels of price pass-through over a one-month horizon. Because we have controlled for the no-network half-life of price adjustment, this suggests it is not merely the case that downstream industries experience less pass-through because they adjust prices more slowly. Downstreamness measures something distinct; it does matter, even conditional on frequency of price adjustment.

Figure 8 plots the impulse response function corresponding to the specification in column (4). It is apparent that the further downstream firms have delayed pass-through relative to upstream ones. After a year has passed, however, full pass-through has been realized by both upstream and downstream firms, and downstreamness loses its predictive power; the value of the interaction term becomes indistinguishable from zero.

One concern about the frequency of price adjustment measure is whether it merely measures the frequency with which industries are hit by marginal cost shocks, rather than how the industries would behave *if* they were hit by a marginal cost shock. We address this conjecture in columns (7), (8), and (9) by adding another interaction term: the standard deviation of an industry's marginal costs. The interaction is positive and statistically-significant in the case of all oil price variation and the Kanzig IV variation. And the frequency of price adjustment remains strongly significant in all three cases. This suggests that frequency of price adjustment is not merely measuring the extent to which industries have volatile marginal costs. Other information is contained in the measure.

In Appendix B.1, we investigate some additional dimensions of potential heterogeneity - size and sign of shock - finding little to no evidence of heterogeneity in these dimensions.

4.3 Analysis using Model-derived Dynamics

The New Keynesian model provides impulse response functions (IRFs) measuring how much each industry changes prices in response to a commodity price shock. In this section, we develop regression specifications to test how close these IRFs are to the industry pass-through patterns we find in the data. In addition to overall model fit, we seek to test whether firms are forward-looking about changes in expected future marginal costs coming from gradual pass-through of shocks by their suppliers. To do this, we decompose the New Keynesian IRFs into the component due to pass-through under the model under full myopia and the component due to forward-lookingness under rational expectations.

Figure 8 – Month-by-Month pass-through of Oil Price Changes: Upstream vs. Downstream Industries

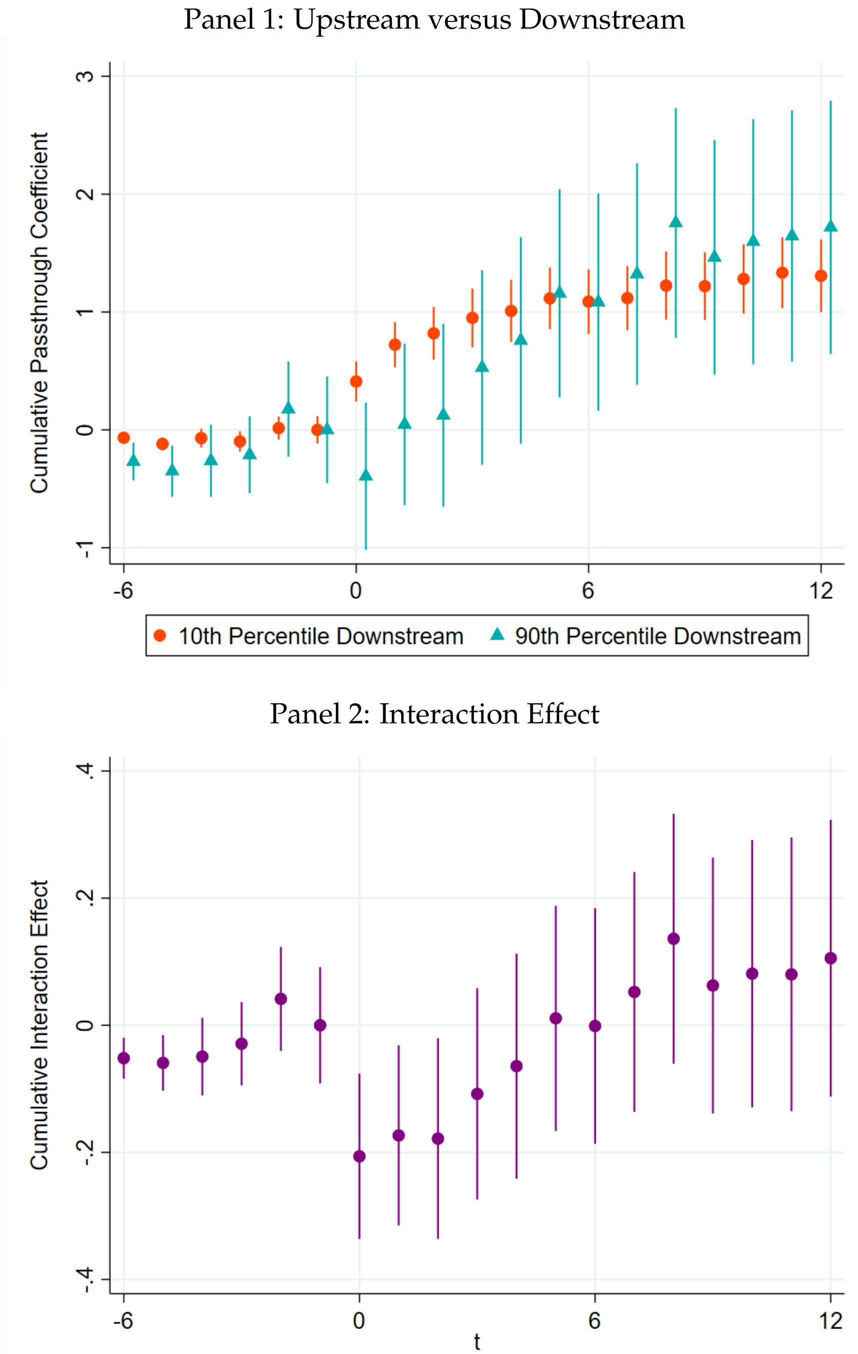

Note: Regression specifications correspond to column (4) of Table 1, with shock terms interacted with our measure of downstreamness. Consequently, the top panel plots monthly cumulative price pass-through of crude oil price shocks for both upstream (10th percentile downstreamness) and downstream (90th percentile downstreamness) firms. The bottom panel plots the cumulated interaction coefficient itself. Downstream firms have delayed pass-through by a few months relative to upstream firms.

4.3.1 Specifications for Semi-Structural Tests and Structural Estimation

Recall that Figure 2 characterized New Keynesian impulse response functions (IRFs) measuring how much each industry changes prices in response to a commodity price shock under a rational expectations assumption. We develop regression specifications to test how close these IRFs are to the industry pass-through patterns we find in the data. Denote by

$$\left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}}\right)_{\text{RE}}(j) \tag{11}$$

the model-derived effect under rational expectations (RE) of a change in the log oil price j periods ago on industry i's current price. Note that this measure does not depend on t, only j.[7] The impulse response function accumulates these effects. Formally, the cumulative effect on industry i's price by time h of a log point oil shock in period 0 is

$$IRF^{RE}_{i,h} = \sum_{j=0}^{h} \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}}\right)_{\text{RE}}(j).$$

Now, the data is not nearly as beautiful as an IRF, since oil prices move essentially every period. At any moment in time, an industry is affected by the current oil price movement, last month's oil price movement, etc. The New-Keynesian model tells us exactly how log industry prices respond to oil price movements in previous periods, suggesting that we run the regression

$$\Delta P_{i,t} = \lambda_t + \sum_{j=0}^{K} \beta_j \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}}\right)_{\text{RE}}(j)\Delta P_{Oil,t-j} + \epsilon_{i,t}. \tag{12}$$

We should have $\beta_j = 1$ for all j if the model is correct and the treatment is uncorrelated with the residual, conditional on time fixed effects. Now, as in our regressions based on the flexible pricing model, $\epsilon_{i,t}$ contains a network object comprised of industry wages, TFP, and desired markups. Unlike before, however, $\epsilon_{i,t}$ contains not just the contemporaneous values of these variables; it also contains agents' expected future values of the variables. A specific worry arising from this new channel, if we did not include time fixed effects in our regressions, is that an oil shock might lead to action by the monetary authority, which would in turn affect expected future values of aggregate nominal variables.

We also note that the IRFs in the New Keynesian model depend on whether the shock is persistent and unexpected. Fortunately, real oil price movements are approximately a random walk over a horizon of one year, and the real oil price movement is relevant for us (rather than the nominal price movement) because of our use of time fixed effects. We have verified this persistence in

[7] This argument holds so long as we use a single input-output table to compute the model IRF. Because the input-output structure of the economy changes over time, the model-derived IRFs change with the input-output table used to compute them.

a time series analysis for our sample, and this degree of persistence has been found elsewhere in the literature (see, e.g., Alquist et al. (2013)).

We deliver our results in impulse response form. Formally, the impulse response in the model for industry i at horizon h was already defined as $IRF_{i,h}$, and the impulse response function in the data for industry i at horizon h is

$$IRF_{i,h}^{\text{RE, Data}} = \sum_{j=0}^{h} \hat{\beta}_j \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}} \right)_{\text{RE}} (j),$$

which is a function of our estimates of the coefficients in equation 12. Note that we have not estimated IRFs for each industry explicitly in the data. We are utilizing that regression 12 is ultimately a difference-in-differences specification estimating a relative effect: how much *more* industry prices increase in the data, relative to the model, at certain horizons. Now, it is infeasible to show IRFs for every industry in our regressions, so we summarize our results in IRFs for synthetic industries. We will show our results by plotting an IRF for a synthetic industry comprised of the top 10th percentile of upstream industries affected by the shock, a synthetic industry comprised of the top 10th percentile of downstream industries affected by the shock, and a synthetic industry comprised of how the average industry is affected by the shock. The downstreamness split is simply made on the median of downstreamness defined under the rational expectations model. Because speed of pass-through is slower for downstream industries than upstream industries, the downstream IRF gives us a better sense of how well the model fits the data when substantial variation in pass-through is generated from distant rather than recent oil shocks. The upstream IRF gives us a sense of how well the model fits the data when pass-through is generated primarily from recent oil shocks.

Finally, we note that all estimation exercises for oil in this section will exclude petroleum refineries from the regression, as they are a notable outlier with substantial power to influence estimation of β_0 and β_1 in particular. We do not think there is substantial loss with this exclusion, as the direct effects plotted in our reduced-form analysis already showed how quickly and fully petroleum refineries pass-through oil price movements.

Now, in Figure 2 we displayed the solution to our New Keynesian model under a rational expectations assumption. But we could solve the model under the assumption that agents are fully myopic, and this would lead to different impulse response functions for each industry. It is clear that we can decompose the model-derived dynamic effects of oil price changes given in expression 11 into (1) a myopic component and (2) a component due to rational expectations:

$$\left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}} \right)_{\text{RE}} (j) = \underbrace{\left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}} \right)_{\text{Myopic}} (j)}_{(1)} + \underbrace{\left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}} \right)_{\text{RE}} (j) - \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}} \right)_{\text{Myopic}} (j)}_{(2)}$$

If the rational expectations component of the model solution has any predictive power for prices, we can also run the regression

$$\Delta P_{i,t} = \lambda_t + \sum_{j=0}^{K} \beta_{j,\text{Myopic}} \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}} \right)_{\text{Myopic}} (j) \Delta P_{Oil,t-j}$$
$$+ \sum_{j=0}^{K} \beta_{j,\text{REGap}} \left\{ \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}} \right)_{\text{RE}} (j) - \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}} \right)_{\text{Myopic}} (j) \right\} + \epsilon_{i,t},$$

to test whether $\beta_{j,\text{Myopic}} = \beta_{j,\text{REGap}} = 1$ for all j. In particular, $\beta_{j,\text{REGap}} \neq 0$ tells us that the rational expectations component of the New Keynesian model has predictive power for pass-through of commodity shocks. It is worth considering how this test of forward-lookingness works more carefully. This regression captures the experiment in which two hypothetical industries with the same path of pass-through as predicted by the myopic model differ in the degree to which rational expectations affects the timing of their pass-through; formally, it controls for the path of myopic pass-through, meaning $\beta_{j,\text{REGap}}$ identifies the effect of increasing the degree of pass-through due to forward-lookingness from a commodity shock j periods earlier, holding constant the path of pass-through under the myopic model and the degree of forward-lookingness about pass-through of oil shocks occurring outside of j periods ago.

Now, we set up our basic framework for generalized method of moments (GMM) estimation of parameters either already in or that will be introduced to our baseline New Keynesian model. These parameters are generically denoted $\boldsymbol{\alpha}$. For example, we might have $\alpha = m_f$. Given a vector of parameters $\boldsymbol{\alpha}$, the New Keynesian can be solved to yield

$$\left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}} \right) (j; \boldsymbol{\alpha}).$$

In this expanded notation, for example, if $\alpha = m_f$, we have

$$\left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}} \right)_{\text{RE}} (j) = \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}} \right) (j; 1).$$

GMM requires a set of moments to estimate α. A sensible moment condition is

$$m(\boldsymbol{\alpha_0}) = \text{E}[\boldsymbol{IRF}_i^{Data}(\boldsymbol{\alpha_0}) - \boldsymbol{IRF}_i(\boldsymbol{\alpha_0})] = 0,$$

where $\boldsymbol{IRF}_i = (IRF_{i,0}, IRF_{i,1}, ..., IRF_{i,K})'$, K represents our choice of how many lags to estimate in regression specification 12,

$$IRF_{i,h}(\boldsymbol{\alpha}) = \sum_{j=0}^{h} \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}} \right) (j; \boldsymbol{\alpha}), \quad IRF_{i,h}^{\text{Data}}(\boldsymbol{\alpha}) = \sum_{j=0}^{h} \hat{\beta}_j(\boldsymbol{\alpha}) \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}} \right) (j; \boldsymbol{\alpha}),$$

and the $\hat{\beta}_j(\boldsymbol{\alpha})$ come from estimation of the $\boldsymbol{\alpha}$-generalized form of regression 12:

$$\Delta P_{i,t} = \lambda_t + \sum_{j=0}^{K} \beta_j(\boldsymbol{\alpha}) \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}} \right) (j; \boldsymbol{\alpha}) \Delta P_{Oil,t-j} + \epsilon_{i,t}. \tag{13}$$

Intuitively, our GMM procedure is as follows: given $\boldsymbol{\alpha}$, we will estimate the model, run regression 13 with K lags, construct IRFs for each industry, and then choose the $\boldsymbol{\alpha}$ that gets the estimated IRFs in the data as close as possible to the model-generated IRFs. The sample average of our moments is

$$\hat{m}(\boldsymbol{\alpha}) = \frac{1}{I} \sum_{i=1}^{I} \boldsymbol{IRF}_i^{Data}(\boldsymbol{\alpha}) - \boldsymbol{IRF}_i(\boldsymbol{\alpha}).$$

We can think about this sample average as the difference in the estimated and model IRFs for a synthetic industry whose IRF is the average of all industry IRFs in the data. A continuously updating GMM estimator is

$$\hat{\boldsymbol{\alpha}} = \arg\min_{\boldsymbol{\alpha}} \hat{m}(\boldsymbol{\alpha})' \widehat{\mathrm{Var}}(\hat{m}(\boldsymbol{\alpha}))^{-1} \hat{m}(\boldsymbol{\alpha}). \tag{14}$$

The following proposition is useful for intuition:

Proposition 5. *The CUGMM estimator given in 14 can be written as*

$$\hat{\alpha} = \arg\min_{\boldsymbol{\alpha}} \left(\hat{\boldsymbol{\beta}}(\boldsymbol{\alpha}) - \mathbf{1} \right)' \widehat{\mathrm{Var}}(\hat{\boldsymbol{\beta}}(\boldsymbol{\alpha}))^{-1} (\hat{\boldsymbol{\beta}}(\boldsymbol{\alpha}) - \mathbf{1}),$$

with $\hat{\boldsymbol{\beta}}(\boldsymbol{\alpha})$ and $\widehat{\mathrm{Var}}(\hat{\boldsymbol{\beta}}(\boldsymbol{\alpha}))$ the estimated coefficients and covariance matrix, respectively, from regression 13. The optimal weight matrix is $\mathrm{Var}(\boldsymbol{\beta}(\boldsymbol{\alpha}_0))^{-1}$.

Therefore, the CUGMM estimator is equivalent to attempting to match the β's from regression 13 to a vector of 1s, which should be true under the null that the New Keynesian model and our statistical assumptions are correct. As a result of the proposition, clustering our GMM standard standard errors by industry is as simple as clustering standard errors by industry when we estimate regression 13 for each $\boldsymbol{\alpha}$ and using the resulting estimate of the covariance matrix of $\hat{\beta}$'s in forming the weight matrix.

In practice, we use the more standard 2-step GMM setup. In step 1, we estimate

$$\hat{\boldsymbol{\alpha}}_{\text{Step 1}} = \arg\min_{\boldsymbol{\alpha}} \hat{m}(\boldsymbol{\alpha})' \hat{m}(\boldsymbol{\alpha}).$$

For step 2, we use the estimated weight matrix, which using similar arguments to those underlying

Proposition 5, entails estimating

$$\begin{aligned}\hat{\boldsymbol{\alpha}}_{\text{Step 2}} &= \arg\min_{\boldsymbol{\alpha}} \left(\hat{\boldsymbol{\beta}}(\boldsymbol{\alpha}) - \mathbf{1}\right)' \widehat{\mathrm{Var}}(\hat{\boldsymbol{\beta}}(\hat{\boldsymbol{\alpha}}_{\text{Step 1}}))^{-1}(\hat{\boldsymbol{\beta}}(\boldsymbol{\alpha}) - \mathbf{1}) \\ &= \arg\min_{\boldsymbol{\alpha}} \hat{m}(\boldsymbol{\alpha})' \widehat{\mathrm{Var}}(\hat{m}(\hat{\boldsymbol{\alpha}}_{\text{Step 1}}))^{-1} \hat{m}(\boldsymbol{\alpha}).\end{aligned}$$

This procedure differs from a step 1 in which we instead estimate

$$\hat{\boldsymbol{\alpha}}_{\text{Step 1}} = \arg\min_{\boldsymbol{\alpha}} (\hat{\boldsymbol{\beta}}(\boldsymbol{\alpha}) - \mathbf{1})'(\hat{\boldsymbol{\beta}}(\boldsymbol{\alpha}) - \mathbf{1})$$

because the weight matrix will differ in step 2. Intuitively, we are prioritizing matching IRFs rather than matching β's to 1, while in the CUGMM setup these goals would be the same.

4.3.2 Empirics Assessing the Profit Share of Payments to Capital

We first determine what share of payments to capital are profits. If payments to capital are primarily profits, they should not affect an industry's cost shares. While assumptions on this share led to only small changes in our reduced-form results, and no meaningful changes in our reduced-form instrumental variable results, as we discuss in Appendix C.4, they will meaningfully affect the results in this subsection. We discuss the cause of this discrepancy after defining the relevant concepts. Recall that industry i's intermediate input share in industry j is

$$s_i^j = \frac{P_j X_i^j}{C_i},$$

where C_i were the total input costs in industry i. Total input costs can be decomposed into intermediate input expenditures, labor compensation, and payments to capital. In our fully reduced-form analysis, we used $C_i = \sum_j P_j X_i^j + W_i L_i + R_i K_i$, or input costs inclusive of payments to capital. In Appendix C.4, we showed that our results were not meaningfully altered if we instead considered $C_i = \sum_j P_j X_i^j + W_i L_i$, input costs excluding payments to capital. Now, we set up the notation for our first structural estimation in the next subsection, noting that we can generally consider the intermediate input cost shares

$$s_i^j(\alpha) = \frac{P_j X_i^j}{\sum_j P_j X_i^j + W_i L_i + (1-\alpha) R_i K_i}.$$

The parameter α therefore captures the profit share of payments to capital. Arranging these shares across inputs within industry and then stacking across industries, we get the $\alpha-$dependent input-output matrix $\boldsymbol{\Phi}(\alpha)$ that implies α-dependent New Keynesian solutions

$$\left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}}\right)_{\text{RE}} (j; \alpha).$$

Before conducting any GMM estimation, we first run regression specification 13 for $\alpha = 0$ and $\alpha = 1$ under rational expectations. The results motivate why the choice of α is important. We show our results in Figure 9. First, notice that setting $\alpha = 1$ yields IRFs that already match the data remarkably well. Setting $\alpha = 0$ still leads to measures that imply strongly statistically significant pass-through in the data, but the scale of the IRFs is wrong; pass-through in the data is higher than what the model suggests. Part of the finding that increasing α increases pass-through in the model is mechanical, since the vast majority of payments to capital are positive, and so increasing α increases the industries' cost shares (both direct and network) in the commodity. Part of the finding, however, is not mechanical, because there is substantial industry heterogeneity in payments to capital, and so the cross-sectional variation defined by our model-defined treatments differs meaningfully with α.

Now, we return to the question of why there were such small effects on our reduced-form empirics from the change of $\alpha = 0$ to $\alpha = 1$. With model-defined dynamics, we need to think about which industries identify which coefficients in regression 12. Upstream industries are the most affected by oil shocks—they experience the largest price changes in response—and they have the fastest frequencies of price adjustment on average. Therefore, they primarily identify the first several lags in regression 12. The bulk of model-implied variation in later lags comes from more downstream industries, which on average increase prices less and more gradually in response to oil price movements. These industries are not as responsible for identifying any reduced-form coefficients, since the reduced-form regressions do not have dynamic, model-defined treatments and therefore maximize fit (minimizing the sum of squared error) by prioritizing fit of upstream industries whose prices are more quickly and substantially moved by oil price changes. We confirm this hypothesis in appendix D.1, finding pass-through for sufficiently downstream industries that is larger than 1 when long-run pass-through for industries' is defined using intermediate input shares inclusive of payments to capital in cost.

Before turning to estimating whether rational expectations yields meaningful predictive power for industry price responses to oil shocks, we structurally estimate α. Our resulting estimates using all variation are $\hat{\alpha} = .99$, with a standard error of 0.08. Using our Kanzig IV, the estimate is $\hat{\alpha} = 1$, with a standard error of 0.07. This estimates can alternatively be construed as measurements of the average steady state markup of prices over marginal cost, $\frac{1}{I}\sum_{i=1}^{I}\mu_i^{SS}\frac{\sigma_i}{\sigma_i-1}$. Construed this way, we estimate a steady state markup of 1.303 (SE = 0.066) using all oil variation and a steady state markup of 1.312 (SE = 0.065) using Kanzig variation. We show the fit of the model IRFs in Figure 10. We broadly cannot reject the hypothesis that the IRFs predicted by the model lie within the IRFs generated in the data. For our synthetic average and downstream industries, we see that pass-through under the Kanzig IV variation seems to differ from the model at higher lags. We show later that this is at least in part because our model-defined dynamic treatments are mismeasured at higher lags for the Kanzig shocks; Kanzig shocks generate smaller oil price movements

Figure 9 – New Keynesian IRFs in the Data Relative to the Model

Panel A: Cost Excluding Payments to Capital

Panel B: Cost Including Payments to Capital

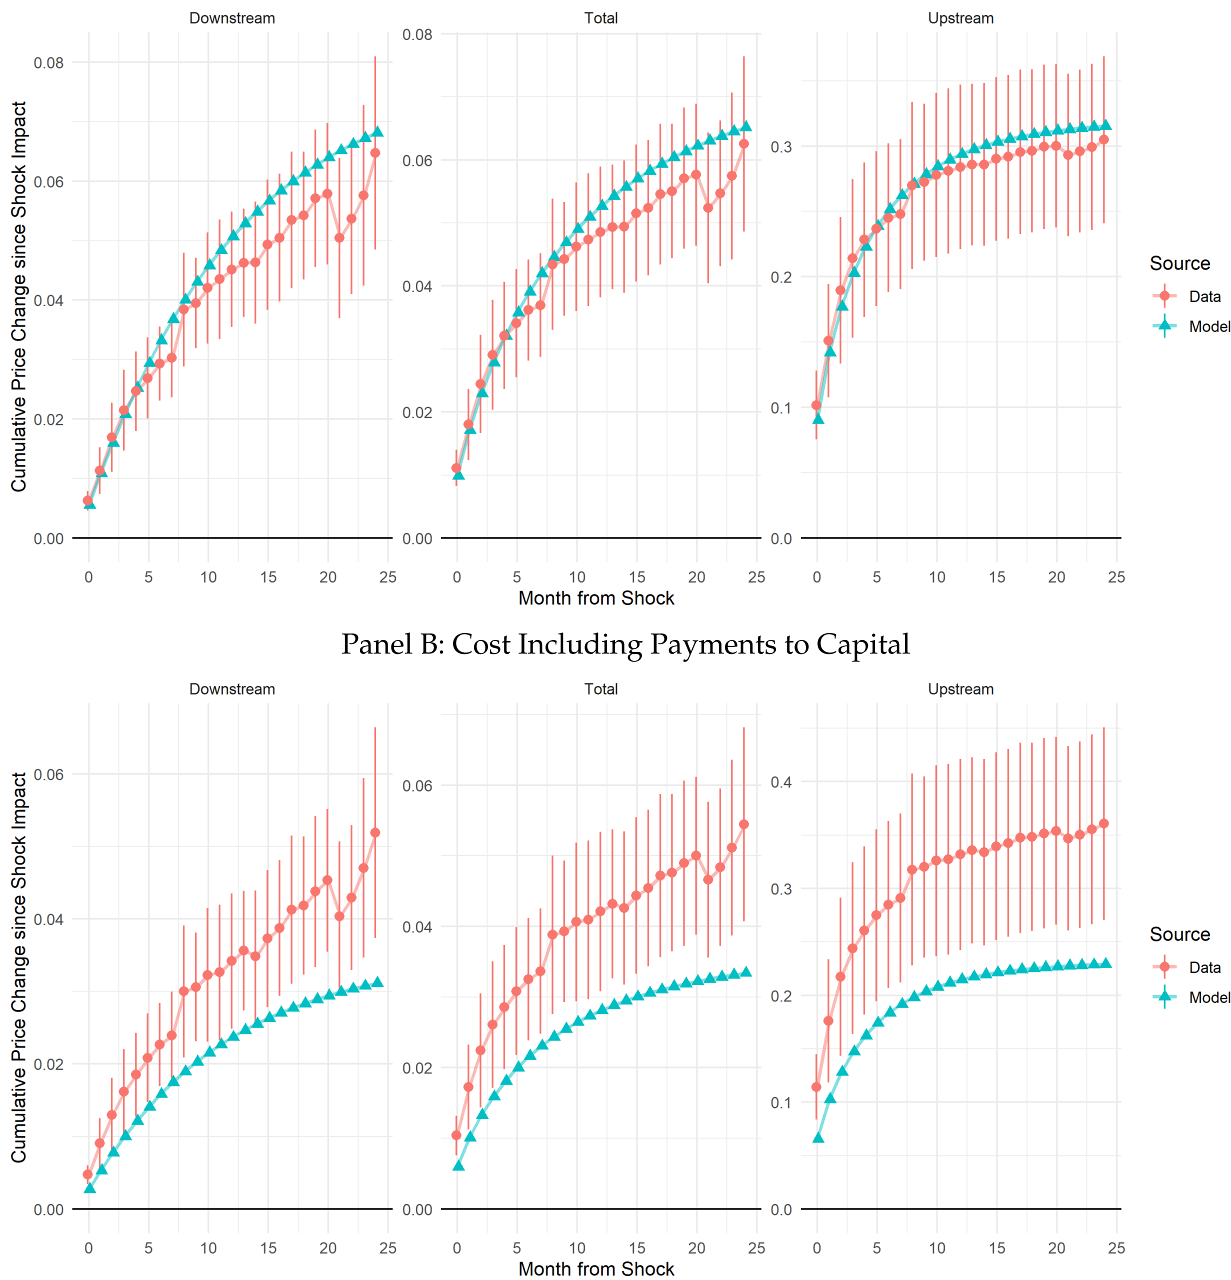

We plot the estimated fit of IRFs generated under the Calvo New Keynesian model for a selection of synthetic industries: the 10% most exposed upstream and downstream industries, and the equal-weighted industry. Long-run pass-through is more accurate when marginal cost excludes payments to capital.

than oil price movements in general, and we will find that firms are more myopic with respect to these shocks.

4.3.3 Regression Specifications Testing Rational Expectations

As before, we present our results in IRF form. This time, rather than plotting separate results for upstream and downstream industries, we will merely focus on a synthetic industry comprised of the top 10% industries affected by rational expectations in the model. We plot separate IRFs for the myopic component and the rational expectations gap. Note that

$$IRF_{i,h}^{\text{REGap}} = \sum_{j=0}^{h} \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}}\right)_{\text{RE}} (j) - \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}}\right)_{\text{Myopic}} (j) \to 0 \quad \text{as } h \to \infty,$$

which follows from long-run pass-through of a persistent shock being the same under rational expectations and myopia. We define

$$IRF_{i,h}^{\text{Myopia}} = \sum_{j=0}^{h} \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}}\right)_{\text{Myopic}} (j),$$

and it follows from our decomposition that

$$IRF_{i,h}^{RE} = IRF_{i,h}^{\text{Myopia}} + IRF_{i,h}^{\text{REGap}}.$$

Formally, we will show how our IRFs estimated in the data,

$$\begin{aligned} IRF_{i,h}^{\text{REGap, Data}} &\equiv \sum_{j=0}^{h} \hat{\beta}_{j,\text{REGap}} \left[\left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}}\right)_{\text{RE}} (j) - \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}}\right)_{\text{Myopic}} (j)\right] \\ IRF_{i,h}^{\text{Myopia, Data}} &\equiv \sum_{j=0}^{h} \hat{\beta}_{j,\text{Myopic}} \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}}\right)_{\text{Myopic}} (j), \end{aligned}$$

compare to $IRF_{i,h}^{\text{Myopia}}$ and $IRF_{i,h}^{\text{REGap}}$.

We show our results in Figure 11. Using both all variation in oil prices and Kanzig-induced variation, we see that there is a strongly statistically significant effect of forward-lookingness on pass-through that tends towards zero as predicted by the model. There is some evidence, although not enough statistically to interpret with confidence, that forward-lookingness reverts to zero too quickly when oil price movements are induced by OPEC announcements. This may suggest that while more upstream industries are forward-looking about this variation, it may not generate strong enough oil price movements for more downstream industries to pay attention. We test whether firms are more myopic with respect to Kanzig variation now.

Figure 10 – New Keynesian IRFs in the Data Relative to the Model

Panel A: All Variation, $\hat{\alpha} = .99$, (SE = 0.08)

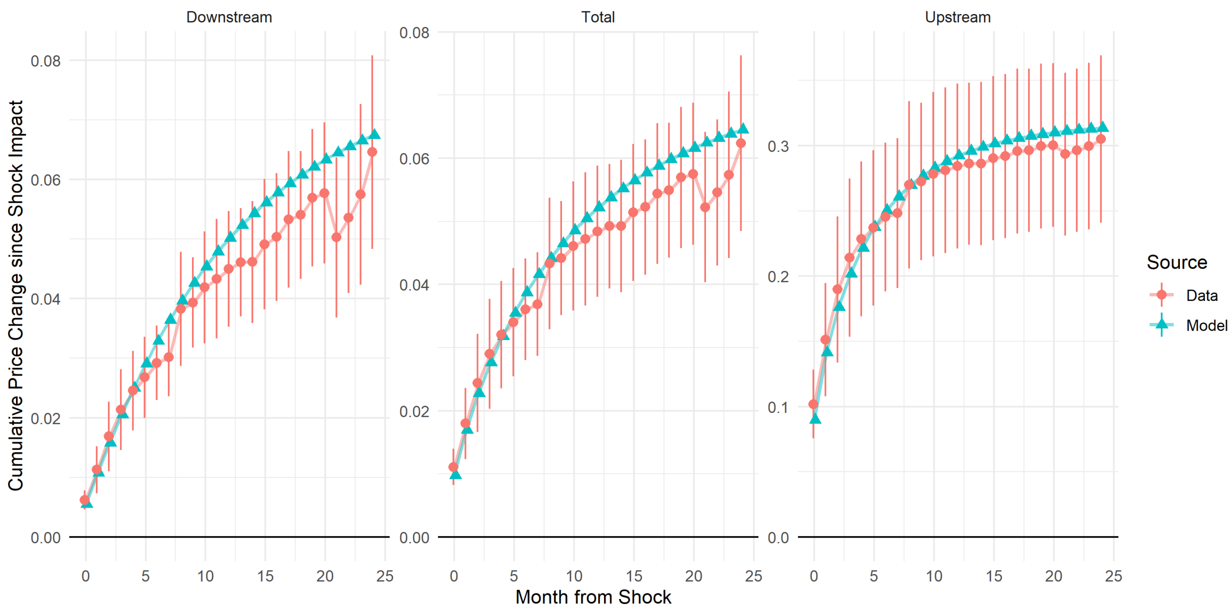

Panel B: Kanzig IV Variation, $\hat{\alpha} = 1$, (SE = 0.07)

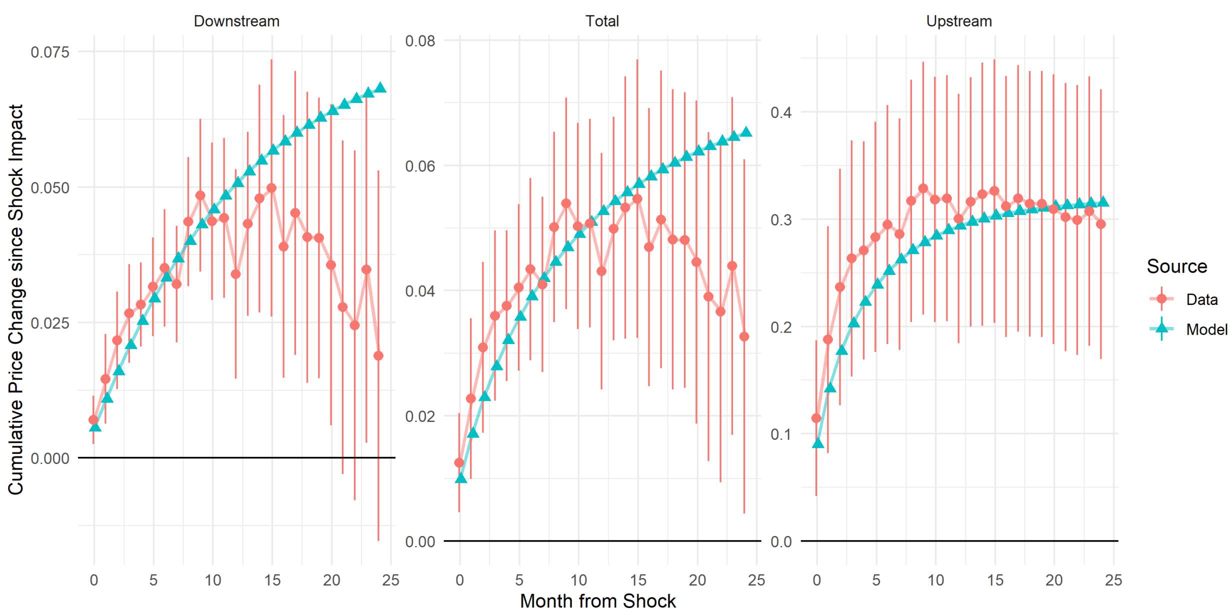

We plot the estimated fit of IRFs generated under the GMM-estimated Calvo New Keynesian model for a selection of synthetic industries: the 10% most exposed upstream and downstream industries, and the equal-weighted industry. Outside of higher lags of the IRFs estimated using Kanzig-induced variation for the synthetic downstream and equal-weighted industry, we largely cannot reject that model-predicted IRFs are different from those in the data.

Figure 11 – Tests of Baseline New Keynesian Model

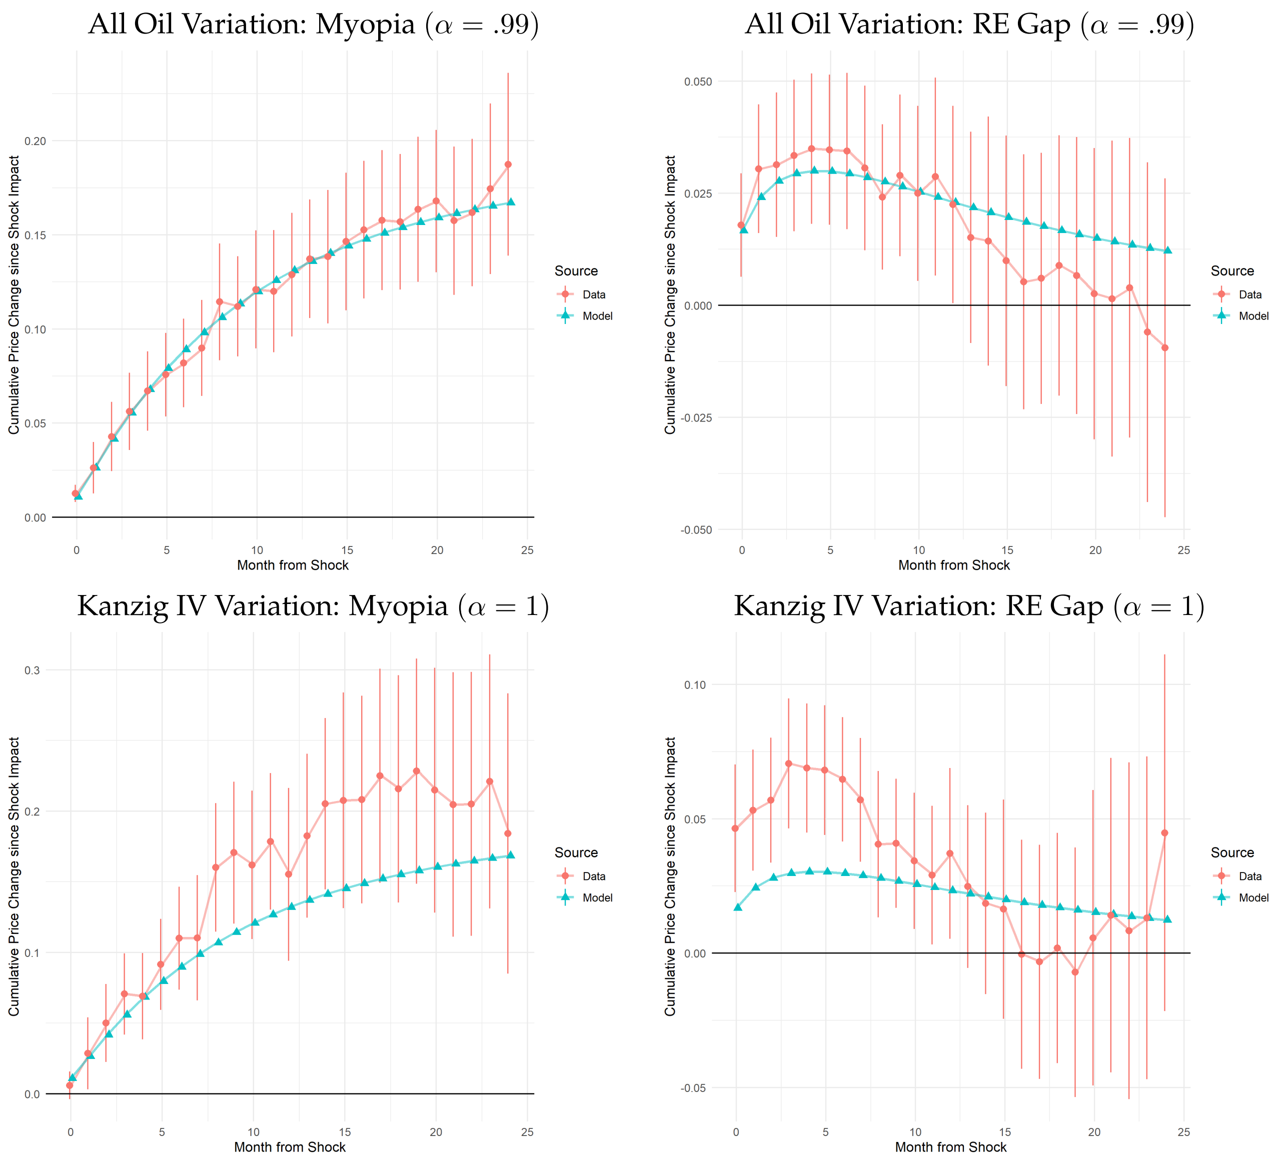

We plot the results of our test that pass-through due to forward-lookingness under rational expectations is present in the data. Statistical significance in both right panels implies we cannot reject that firms are forward-looking about the gradual pass-through of upstream shocks to their marginal costs.

We use our GMM framework to estimate myopia and the profit share of capital payments. This exercise helps us answer two key questions: (1) while we have found evidence for forward-lookingness in the data, we do not know if this evidence is consistent with full rational expectations, $m_f = 1$, or some degree of myopia, $0 < m_f < 1$, and (2) even if we have found evidence of forward-lookingness conditional on our previous estimates of α, the measurement of the rational expectations gap varies with α—it may be that a different joint estimate of α and myopia m_f provide a better fit to the data.

Using all oil variation, we estimate $m_f = .91$, with a standard error of 0.113, and $\alpha = 0.97$, with a standard error of 0.089. Using Kanzig IV variation, we estimate $m_f = .32$, with a standard error of 0.22, and $\alpha = 0.83$, with a standard error of 0.11. The estimates of the profit share of capital payments are broadly consistent with those we estimated before. Now, while the estimate of the degree of myopia is noisy using the instrument variation, we can reject that $m_f > 0.75$ at the 5% level. Compared to the relatively well-estimated myopia of 0.91 in the case using all variation, this suggests firms in the network are less forward-looking about OPEC-induced oil price movements. This may not be surprising, given that OPEC-induced oil price movements are typically smaller than the large swings frequently observed in oil prices. We show Kanzig-shock induced variation in oil prices in Appendix D.2, confirming that they indeed generate oil price movements almost entirely during small oil price change episodes. To assess the validity of this hypothesis more formally, we analyze large and small oil price variation separately, no longer using instrumental variation, and we also assess the degree of forward-lookingness when the prices of other commodities move.

4.3.4 Large versus Small Oil Shocks, and Non-oil Commodities

We can replicate our analysis for large and small oil shocks. Immediately, there arises the question of how to define large and small shocks. In particular, do we mean the oil price change occurred during a time period with large oil price movements, or do we mean the specific monthly oil price change was large? This definition turns out to be important for our result. For myopia, our hypothesis is that firms will be paying more attention to *all* oil price movements during an episode in which oil price movements have been large. Our baseline definition is that the economy is in a large oil shock period in period t if the absolute value of the six month log change in the oil price between t and $t-6$ is larger than the median six month absolute change during our sample period. In this case, we estimate $m_f = 0.77$ with a standard error of 0.11 for large oil shock periods and $m_f = 0.01$ with a standard error of 0.05 for small oil shock periods. Our results are not meaningfully altered if we widen the definition of a shock period to be about absolute year-over-year changes in the oil price.

Our results are altered, however, if we define a large oil shock period as one where the month-over-month absolute change in the oil price is larger than the median during our sample period.

In this case, we cannot reject high levels of forward-lookingness (high m_f) for both large and small shocks. We think the reason for this result, given the results for larger episode definitions, is that firms still pay close attention to small oil shocks during months surrounding large oil shocks. Intuitively, firms may use a binary threshold: either they are paying attention to upstream oil shocks because they are salient and have been large, or they are not paying attention because they have largely been small, or, if they have been large, they have moved in opposite directions, canceling out over a six month or one year horizon.

We now repeat our analysis for a shock to non-oil commodity prices, pooling commodities as we did in our reduced-form analysis. Our GMM estimates for matching non-oil commodity shock IRFs are $\hat{\alpha} = .6$, with a standard error of 0.31, and $\hat{m_f} = 0$, with a standard error of 0.17. Now, the profit share of payments to capital is noisily estimated for non-oil commodities, but we do have some precision on our estimate of the degree of myopia. It appears that firms are much more myopic when passing through shocks to non-oil commodities than they are when passing through shocks to oil.

We show the model fit for non-oil commodity shocks in Figure 12. Panel 1 shows model fit for the model estimated using GMM. Panels 2 and 3 show model fit for the rational expectations gap, which entails estimating the model under full myopia and rational expectations using our GMM-estimated α. Panels 2 and 3 are a different way of assessing whether there is evidence any forward-lookingness in industry pass-through of non-oil commodity price shocks. Here, we find statistically significant evidence that industries predicted to pass through shocks due to forward-lookingness actually do increase prices faster in the data. This evidence is substantially weaker than the evidence we found for oil, even in the case where we used Kanzig instrumental variation. Taken together, we think our GMM estimates, combined with the estimates in Panel 3, suggest that firms are more myopic when passing through shocks to non-oil commodities.

We summarize all of our GMM estimates in the paper in Table 2. Taken together, we conclude that our result that firms are forward-looking about oil price movements is driven by oil price movements during large shock episodes. For small oil price episodes, oil price movements generated by the Kanzig IV variation, or non-oil commodity price movements, firms appear to be less-forward looking, leading to slower pass-through. Myopia therefore seems to amplify price rigidity in the data.

5 Application: Network Oil Inflation

As we have shown, oil and other commodity shocks generate inflation beyond changing prices for products in which they are directly used. For instance, the main personal consumption expenditure (PCE) category for the petroleum refineries industry is "gasoline and other motor fuel." An inflation measure that subtracts only this component of oil inflation from aggregate PCE infla-

Figure 12 – New Keynesian IRFs in the Data Relative to the Model: Non-Oil Commodities

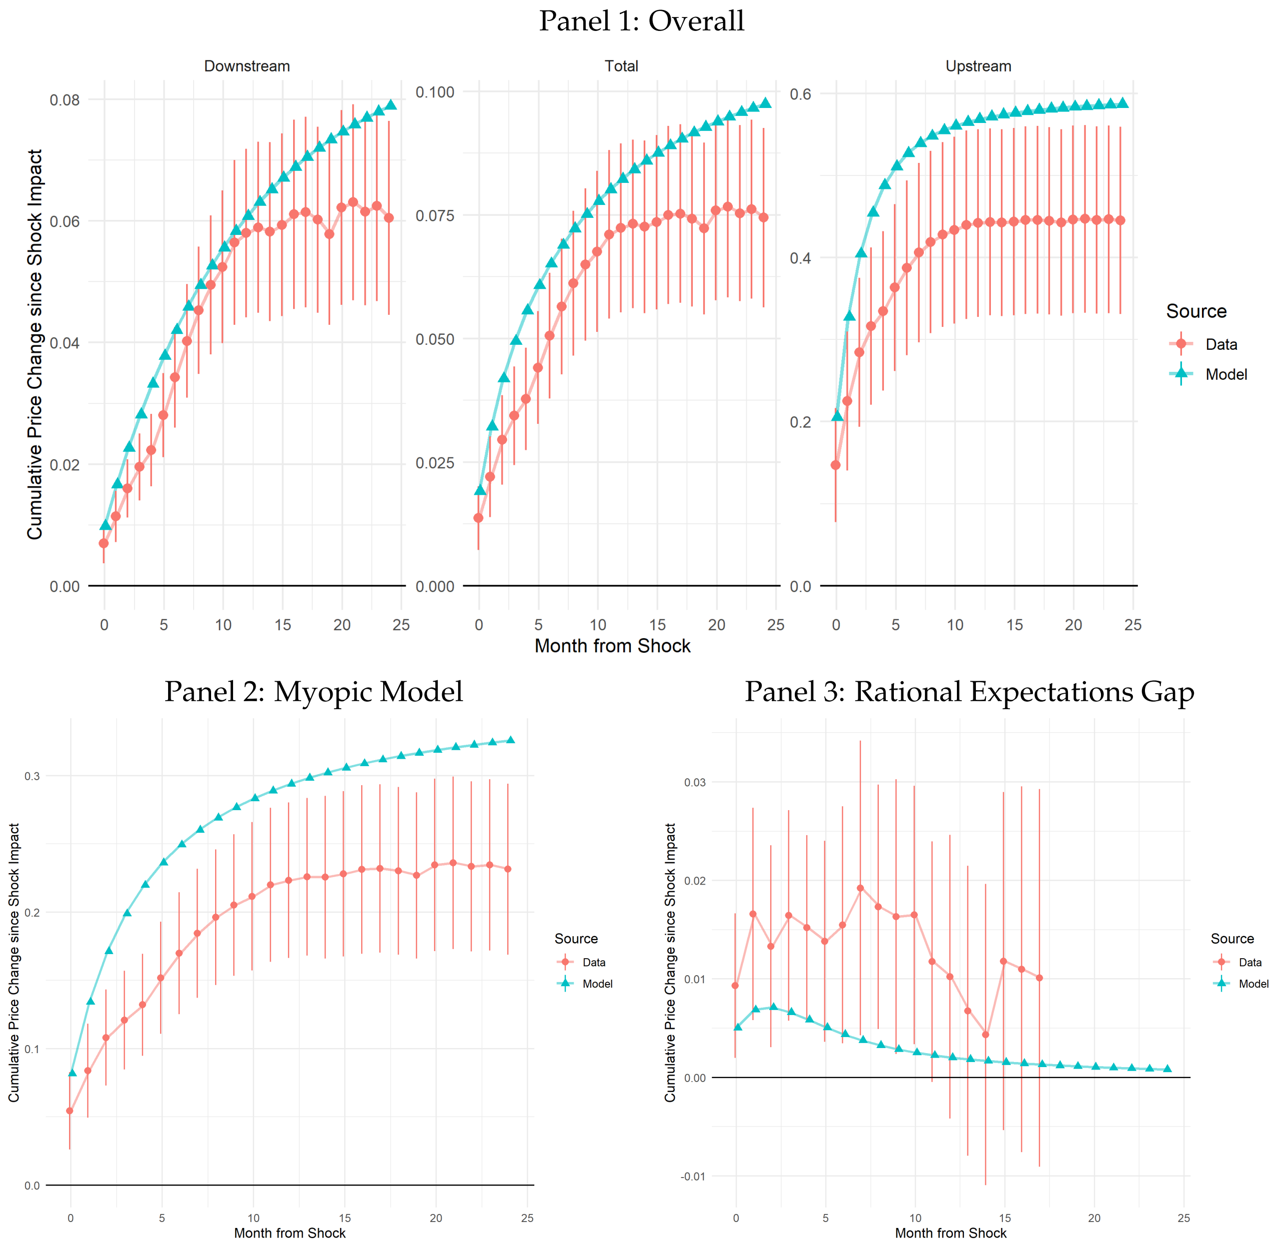

Note: Our regression tests of the New Keynesian model repeated for non-oil commodities. There is not enough variation in the rational expectations gap to estimate coefficients beyond the 17th lag, as iterations to industry IRFs are all negligible by that point. This can be observed in the very small slope of the model-defined IRF for the rational expectations gap.

Table 2 – GMM Estimates

	(1) Oil	(2) Oil	(3) Kanzig	(4) Kanzig	(5) Oil - Large	(6) Oil - Small	(7) Non-Oil
Profit Share	0.99	0.97	1.00	0.83			0.6
	(0.080)	(0.089)	(0.065)	(0.107)			(0.312)
Myopia		0.91		0.32	0.77	0.01	0.00
		(0.113)		(0.218)	(0.114)	(0.050)	(0.174)

Note: All GMM estimates in the main body of the paper. We find relative stability of the profit share of payments to capital across all oil and Kanzig-induced oil price variation. Myopia varies with the shock.

tion does not fully purge oil inflation from aggregate inflation, as all of oil's network uses remain embedded in the aggregate inflation measure. Moreover, because network propagation of shocks takes time to occur, aggregate PCE inflation purging "gas and other motor fuel" will still be predictable using the network component of oil inflation. This section therefore undertakes the study of whether Core PCE inflation is predictable using the network component of oil inflation. Further, we assess how much of aggregate PCE inflation's variation can be explained using network oil inflation.

Earlier in the paper, we noted that the Calvo New Keynesian model, with a given parameterization $\boldsymbol{\alpha}$, tells us how much an industry's price should respond, and over what horizon it should respond, to an oil price movement. We tested how close these predictions were to those in the data using the regression

$$\Delta P_{i,t} = \lambda_t + \sum_{j=0}^{K} \beta_j(\boldsymbol{\alpha}) \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}} \right) (j; \boldsymbol{\alpha}) \Delta P_{Oil,t-j} + \epsilon_{i,t},$$

which led to estimates $\hat{\beta}_j(\boldsymbol{\alpha})$ informing us how much more or less industries' prices moved relative to the model's predictions, where the model and our statistical assumptions being correct would lead to $\hat{\beta}_j(\boldsymbol{\alpha})$ statistically indistinguishable from 1. The inclusion of a time fixed effect gave our coefficients this relative interpretation, and we will need to apply a correction when aggregating, which we will discuss shortly.

Now, in the BEA's benchmark input-output tables, we observe the personal consumption expenditure share of each industry. Our basic aggregation exercise then uses that the response of aggregate inflation to a sequence of oil price shocks as implied by the model is

$$\hat{\Pi}_t = \text{Intercept} + \sum_i PCEShare_i \sum_{j=0}^{K} \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}} \right) (j; \boldsymbol{\alpha}) \Delta P_{Oil,t-j}.$$

The IRF of aggregate inflation to a unit shock to the log oil price is then just

$$IRF_{t,h} = \text{Intercept} + \sum_i PCEShare_i \sum_{j=0}^{h} \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}}\right)(j;\boldsymbol{\alpha}).$$

These predictions can be compared to the data by plotting them against

$$IRF_{t,h}^{Data} = \text{Intercept} + \sum_i PCEShare_i \sum_{j=0}^{h} \hat{\beta}_j(\boldsymbol{\alpha}) \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}}\right)(j;\boldsymbol{\alpha}).$$

which is analogous to our previous figures. We previously considered IRFs in synthetic industries comprised of the most affected upstream and downstream industries, along with an equal-weighted industry; instead, our synthetic industry now is the PCE-weighted average industry. We will use the parameterization $\boldsymbol{\alpha}$ associated with the best model fit estimated in our earlier generalized method of moments analysis, $m_f = 0.91$ and $\alpha = .97$.

To find the level effect from the relative effect in our difference-in-differences regressions, we require an intercept. We can perform this exercise simply using our IV specification with the Kanzig (2021) shocks, given that the time-series variation in these shocks is valid for identification. What we find is that, in response to oil shocks, a hypothetical industry with a zero network cost share in oil does not experience any inflation. We might have expected otherwise if we thought that Federal Reserve policy responded strongly to oil price shocks, since even a hypothetical industry that does not use oil anywhere in its supply chain would be affected by such a policy. This does not appear to be the case in the data over our main sample period, which starts in 1997. Therefore, we will use the prediction

$$\hat{\Pi}_t = \sum_i PCEShare_i \sum_{j=0}^{K} \left(\frac{\partial \Delta P_i}{\partial \Delta P_{Oil}}\right)(j;\boldsymbol{\alpha}) \Delta P_{Oil,t-j}.$$

Now, our structural estimation excluded petroleum refineries, a major outlier industry that is heavily exposed to oil price movements. Correct pass-through predictions in this industry, however, are important for aggregate inflation because the PCE share of petroleum refineries is large—as previously noted, this is where consumer purchases of gasoline appear.

We plot our model-defined response of aggregate PCE inflation to a unit oil shock against the variants in the data, both including and excluding petroleum refineries, in Figure 13. Note that the response of aggregate inflation is large: 0.06 log points over two years for a one log point shock to oil prices. A one log point shock to oil prices is very large but not unseen in the data; in such cases, aggregate inflation moves by about 6 percentage points. Now, because petroleum refinery pass-through occurs in just a few months, we see that the first few coefficients are heavily influenced, generating high precision in estimates. Estimates at higher lags remain close to the model-implied

Figure 13 – Model Tests for Aggregate PCE Inflation

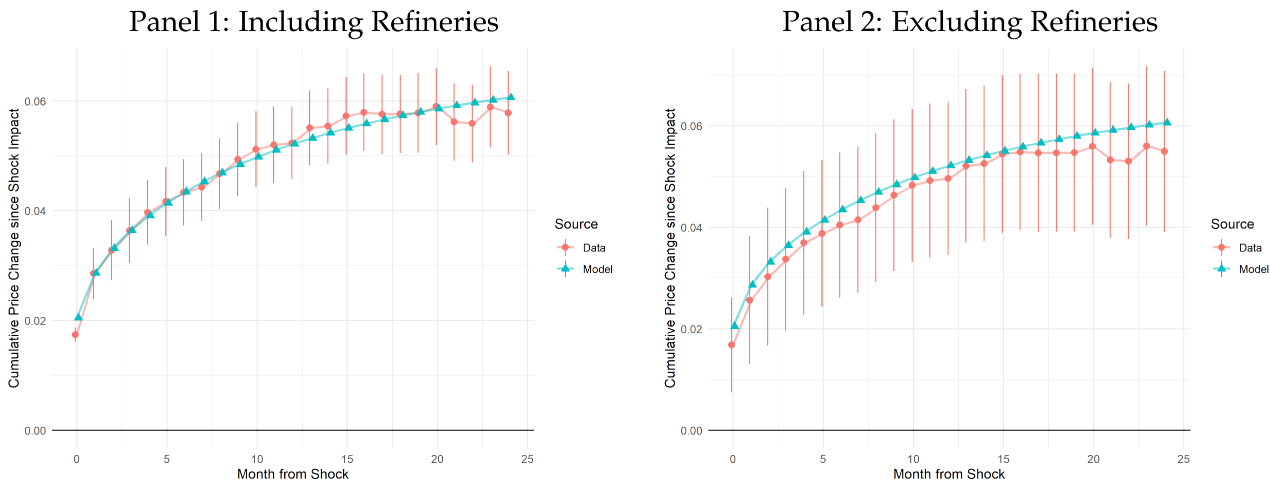

predictions, reflecting how the model fits the data closely, as previously discussed. The model fit remains good regardless of whether refineries are included or excluded in the empirical test.

Now, consumer use of gasoline appears in the data as personal consumption expenditures from the petroleum refining sector. So we can further distinguish the effect of oil prices on inflation that stem from consumer purchases of gasoline from all other effects of oil on inflation. The effect on gas prices is typically considered to be direct oil inflation, which is thought to be purged from Core PCE inflation relative to overall PCE inflation. The effect stemming from all other industries' price responses is indirect oil inflation, which is not explicitly removed from Core PCE inflation, and our results imply that it will result in further inflation from oil price movements.

We can visualize our model's predictions for the contribution of oil prices to the petroleum refinery and non-refinery components of aggregate inflation. Figure 14 reports our results. In Panel 1, we see that, of the about 6.75 percentage points of aggregate inflation resulting from a unit shock to log oil prices, just over 2 percentage points comes from consumer gas purchases (PCE share in petroleum refineries), while just over 4.5 percentage points comes from consumer purchases from all other industries. In panel 2, we assume that oil prices remain constant starting in May 2022. We see that there is an additional two percentage points of inflation that will be realized over the following four years if oil prices do not decline.

Now, there is an important difference between our predictions for aggregate PCE inflation and the actual effects on official PCE inflation. Official PCE inflation is constructed primarily using price measurements from the consumer price indices, while our empirical tests were conducted on a measure of aggregate PCE inflation using price measurements from the producer price indices. We can test whether network oil inflation using our measures passes through to official PCE inflation by regressing official PCE inflation on our network oil inflation measures.

Figure 14 – Inflationary Effects of Oil

Panel 1: Components of Aggregate Oil Inflation Predicted by the Model

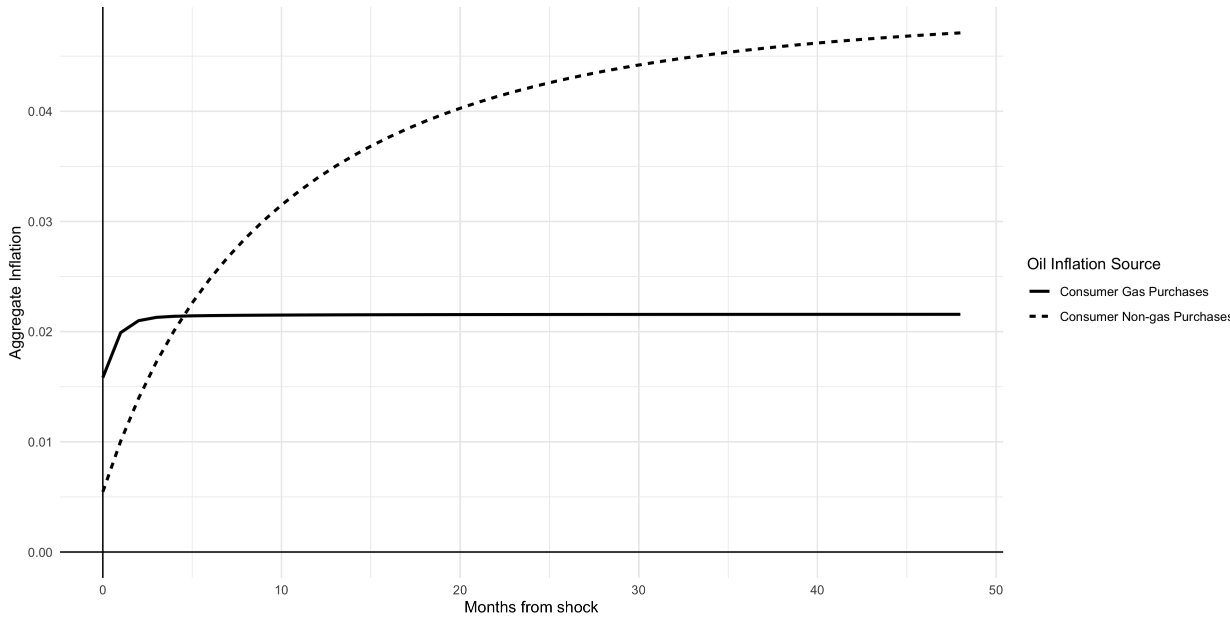

Panel 2: Accumulated and Unrealized Inflation since January 2020

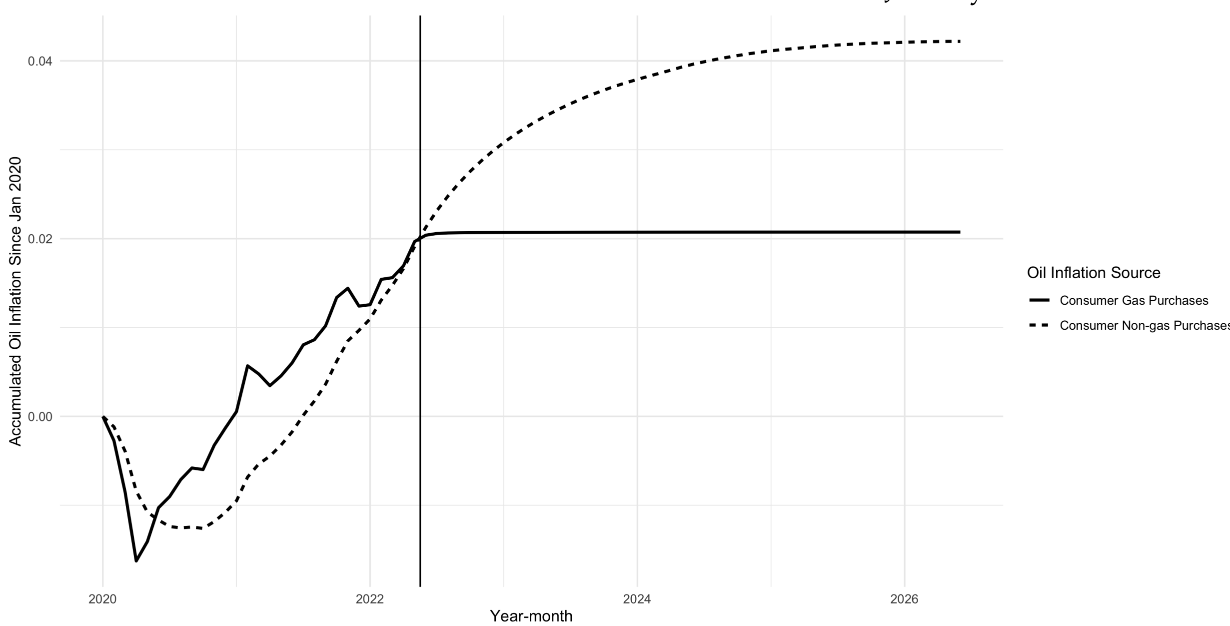

Note: Model predicted IRFs of aggregate inflation to a log point shock to oil prices. Because inflation result from indirect network effects of oil price increases takes time to occur, we see that more than 2 percentage points of aggregate inflation have not yet been realized as a result of the run-up in oil prices prior to and including May 2022. The actual inflation realized through indirect effects of oil price increases after May 2022 may be lower if oil prices begin to decline.

Table 3 – Predicting Inflation with Network Oil Inflation

Panel 1: Total Inflation

Dependent Variable: Total PCE Inflation	(1) All Variation	(2) Kanzig Variation	(3) All Variation	(4) Kanzig Variation
Direct Oil Inflation			0.222***	0.362***
			(0.074)	(0.133)
Indirect Oil Inflation	1.578***	1.246***	1.341***	0.856***
	(0.108)	(0.190)	(0.129)	(0.258)
R-Squared	0.2906	0.2777	0.3006	0.2836
Observations	758	758	758	758

Panel 2: Core Inflation

Dependent Variable: Core PCE Inflation	(1) All Variation	(2) Kanzig Variation	(3) All Variation	(4) Kanzig Variation
Direct Oil Inflation			-0.101*	-0.053
			(0.060)	(0.105)
Indirect Oil Inflation	0.782***	0.387**	0.890***	0.444*
	(0.091)	(0.172)	(0.120)	(0.229)
R-Squared	0.1134	0.0845	0.1167	0.0871
Observations	758	758	758	758

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. This table shows the results of simple regressions of official PCE inflation on our measures of network oil inflation and a constant. Robust standard errors in parentheses. Direct and indirect network oil inflation are both predictive of total PCE inflation. Direct oil inflation is not predictive of core PCE inflation, but indirect network oil inflation still is.

Our results are shown in Table 3. We see in Panel 1 that indirect oil inflation explains a substantial fraction of official PCE inflation: the R^2 is 29%. The predictive power of indirect oil inflation remains substantial if we include direct oil inflation, the component of oil inflation due to consumer gas purchases; together, direct and indirect oil inflation — i.e., network oil inflation — explains 30% of the variation in official PCE inflation. We noted earlier that official core PCE inflation tries to purge the effects of oil inflation from official PCE inflation. Panel 2 of Table 3 reveals that it is only partially successful in doing this. While direct oil inflation indeed has little to no predictive power for core PCE inflation, indirect oil inflation is still strongly predictive. We retain an R^2 of 12% when explaining official core PCE inflation with network oil inflation.

Having confirmed that our measures of network oil inflation pass through to official PCE inflation, we show how much aggregate inflation is changed if the entire contribution of oil prices to aggregate inflation is removed. We subtract the components of oil-induced inflation in gas prices and non-gas prices from official PCE inflation, showing our results in Figure 15. In Panel 1, we see that the inflation series purged of the network contribution of oil does not change the interpretation that 1970s inflation was not driven mechanically by oil. Note that we cannot say whether oil was at least partially responsible for causing runaway inflation expectations, which may generate movements in aggregate inflation beyond those mechanically caused by network oil inflation. In Panel 2, we focus on the COVID period. Early in COVID, we see that there was a spike in inflation if network oil inflation is removed, which reflects that there should have been much less inflation if the large oil price decline early in COVID fully passed through to aggregate inflation and no other prices changed. After the large oil price decline early in COVID, there was a large run-up of oil prices. We see that removing this component from official PCE inflation does not change the finding that inflation increased substantially over 2021, but it does reduce the overall amount of year-over-year inflation by 2022 to 4.5 percentage points from above 6 percentage points. We emphasize that our analysis is silent about whether the oil price movements during the COVID crisis are demand or supply driven.

6 Conclusion

We study the extent and pace of price propagation through supply chains, finding statistically-significant evidence of full pass-through of cost shocks into prices for both industries directly exposed to the shocks and those exposed only indirectly through a complex network of industry linkages. The pass-through to directly-exposed industries, however, occurs primarily in the month of the cost shock's impact, whereas pass-through to indirectly-exposed industries takes six to eight months or more to be realized. This remains true regardless of whether we study all variation in oil prices, specific cases of major oil price movements, or instrumental-variables variation driven by OPEC announcements, as distilled by the oil price shock series from Kanzig (2021). Full pass-through remains evident if we widen our purview beyond oil and examine all other non-oil commodities.

In order to understand the difference in the pace of pass-through between directly-exposed (i.e., upstream) and indirectly-exposed (i.e., downstream) industries, we turn to the framework of a network-enriched Calvo New Keynesian model. The model reveals that nominal rigidities are compounded throughout the supply chain, converting even minor price rigidities at the industry level into substantial aggregate price rigidity. The predictions of the model match our reduced-form findings well.

We further show that the model reveals this compounding of rigidities is intensified in the context of myopia – as opposed to fully rational expectations – on the part of firms. A fully rational (and attentive) firm will observe its far-upstream suppliers' suppliers being hit by a cost shock and adjust their prices when they next have the opportunity to do so. A myopic firm will wait for the shock to trickle through the supply chain and reach the firm itself before choosing to make such a

Figure 15 – Contribution of Oil to Year-over-year Official PCE Inflation

Panel 1: Since January 1960

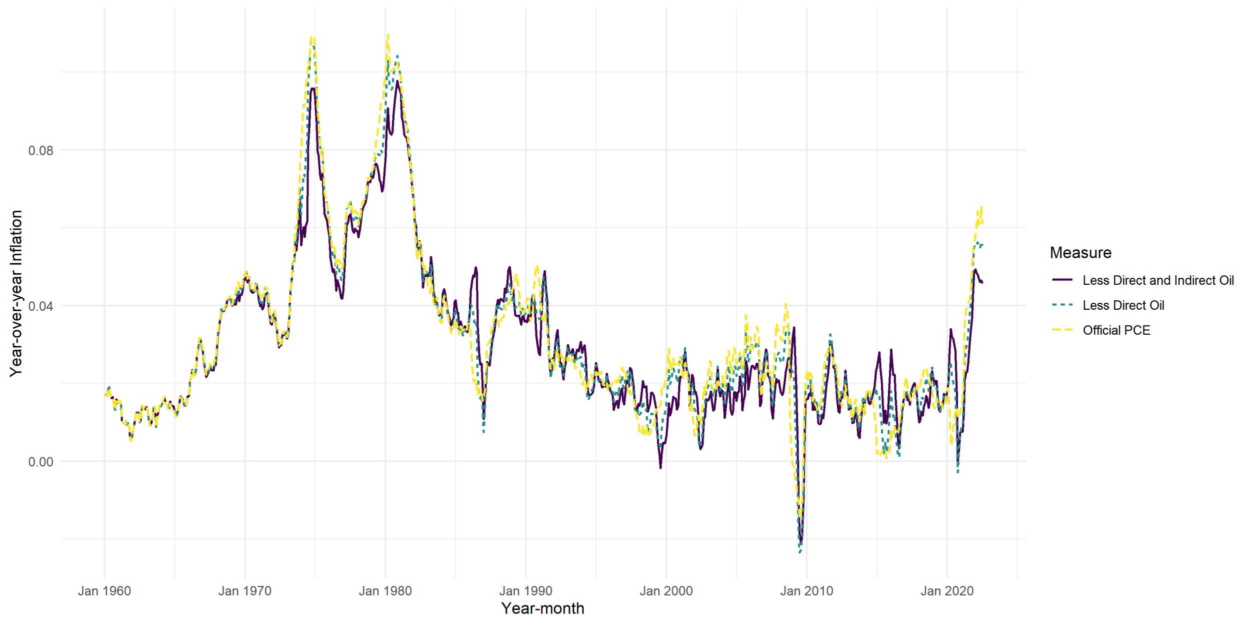

Panel 2: Since January 2019

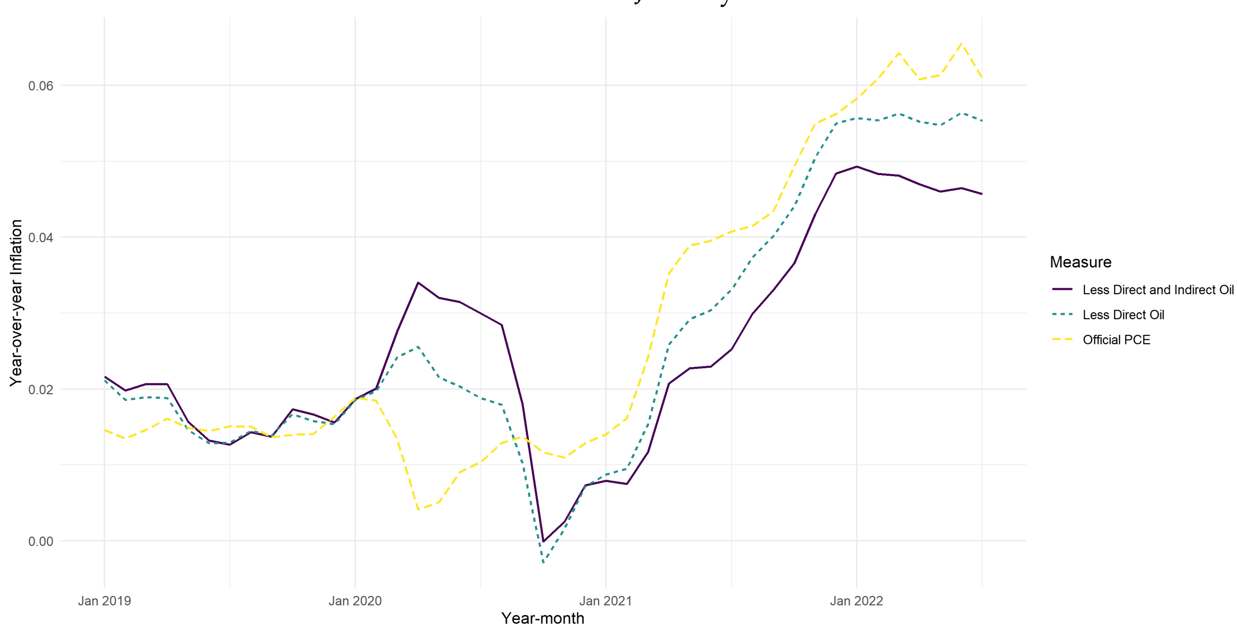

Note: Official PCE inflation plotted against the same inflation series removing direct (through gas purchases) and indirect (through all other industry price movements) oil inflation.

price adjustment. Empirically, we show that price pass-through responses to large oil shocks are consistent with fully rational expectations, whereas responses to small oil shocks and shocks to non-oil commodities are consistent with more myopic behavior.

Finally, in an application of our findings and our framework, we observe that measures of core inflation designed to exclude price variation induced by oil shocks will still contain plenty of such variation indirectly through the production network. We show that indirect oil inflation has statistically-significant and non-trivial (R^2 = 12%) predictive power for official Core PCE inflation, as well as even greater predictive power (R^2 = 30%) for total PCE inflation. As such, we create a revised measure which strips this variation out of core inflation. This results in a noticeably lower level of core inflation in 2021 and 2022 (a peak of 5% instead of 6.5%) but does not overturn the result that inflation in those years was well above the normal range of recent decades.

References

Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi, "Microeconomic Origins of Macroeconomic Tail Risks," *Microeconomic Origins of Macroeconomic Tail Risks*, 2017, *107* (1), 54–108.

_ , U. Akcigit, and W Kerr, "Networks and the Macroeconomy: An Empirical Exploration," *NBER Macroeconomics Annual*, 2021, *30* (1), 275–335.

_ , V. Carvalho, and A. Tahbaz-Salehi, "The Network Origins of Aggregate Fluctuations," *Econometrica*, 2012, *80* (5), 1977–2016.

Alquist, R., L. Kilian, and R.J. Vigfusson, "Forecasting the Price of Oil," *Handbook of Economic Forecasting*, 2013, *2*, 427–507.

Atalay, E., "How Important are Sectoral Shocks?," *American Economic Journal: Macroeconomics*, 2017, *9* (4), 254–280.

Auer, R.A., A.A. Levchenko, and P. Saure, "International Inflation Spillovers through Input Linkages," *Review of Economics and Statistics*, 2019, *101* (3), 507–521.

Baqaee, D.R., "Cascading Failures in Production Networks," *Econometrica*, 2018, *86* (5), 1819–1838.

_ and E. Farhi, "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem," *Econometrica*, 2019, *87* (4), 1155–1203.

Barrot, J.N. and J. Sauvagnat, "Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks," *Quarterly Journal of Economics*, 2016, *131* (3), 1543–1592.

Bartelme, D. and Y. Gorodnichenko, "Linkages and Economic Development," Working Paper 21251, National Bureau of Economic Research 2015.

Basu, S., "Intermediate Goods and Business Cycles: Implications for Productivity and Welfare," *American Economic Review*, 1995, *85* (3), 512–531.

Blanchard, O.J. and C.M. Kahn, "The Solution of Linear Difference Models under Rational Expectations," *Econometrica*, 1980, *48* (5), 1305–1311.

Blinder, Alan, Elie RD Canetti, David E Lebow, and Jeremy B Rudd, *Asking about prices: a new approach to understanding price stickiness*, Russell Sage Foundation, 1998.

Boehm, C.E., A. Flaaen, and N. Pandalai-Nayar, "Input Linkages and the Transmission of Shocks: Firm Level Evidence from the 2011 Tohoku Earthquake," *Review of Economics and Statistics*, 2021, *101* (1), 60–75.

Caliendo, L., F. Parro, and A. Tsyvinski, "Distortions and the Structure of the World Economy," Working Paper 23332, National Bureau of Economic Research 2021.

Carlstrom, C.T., T.S. Fuerst, and M. Paustian, "How Inflationary is an Extended Period of Low Interest Rates?," Working Paper 12-02, Federal Reserve Bank of Cleveland 2012.

Carvalho, Carlos, "Heterogeneity in price stickiness and the real effects of monetary shocks," *Frontiers in Macroeconomics*, 2006, *2* (1).

Carvalho, V.M. and X. Gabaix, "The Great Diversification and its Undoing," *American Economic Review*, 2013, *103* (5), 1697–1727.

_ , M. Nirei, Y.U. Saito, and A. Tahbaz-Salehi, "Supply Chain Disruptions: Evidence from the Great East Japan Earthquake," *Quarterly Journal of Economics*, 2021, *136* (2), 1255–1321.

DelNegro, M., M.P. Giannoni, and C. Patterson, "The Forward Guidance Puzzle," FRB of New York Staff Report 574 2012.

Foerster, A.T., P.D.G. Sarte, and M.W. Watson, "Sectoral versus Aggregate Shocks: A Structural Factor Analysis of Industrial Production," *Journal of Political Economy*, 2011, *119* (1), 1–38.

Gabaix, Xavier, "A Behavioral New Keynesian Model," *American Economic Review*, 2020, *110* (8), 2271–2327.

Gordon, Robert J., "Output Fluctuations and Gradual Price Adjustment," *Journal of Economic Literature*, 1981, *19* (2), 493–530.

Horvath, M., "Cyclicality and Sectoral Linkages: Aggregate Fluctuatiohns from Independent Sectoral Shocks," *Review of Economic Dynamics*, 1998, *1* (4), 781–808.

_ , "Sectoral Shocks and Aggregate Fluctuations," *Journal of Monetary Economics*, 2000, *45* (1), 69–106.

Hulten, C.R., "Growth Accounting with Intermediate Inputs," *Review of Economic Studies*, 1978, *45* (3), 511–518.

Kanzig, D.R., "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements," *American Economic Review*, 2021, *111* (4), 1092–1125.

Kohli, M., "The Leontief-BLS Partnership: A New Framework for Measurement," *Monthly Labor Review*, 2001, *124* (6), 29–37.

La'O, J. and A. Tahbaz-Salehi, "Optimal Monetary Policy in Production Networks," *Econometrica*, 2022, *90* (3), 1295–1336.

Leontief, W.W., "Quantitative Input and Output Relations in the Economic Systems of the United States," *The Review of Economic Studies*, 1936, *18* (3), 105–125.

Liu, E., "Industrial Policies in Production Networks," *Quarterly Journal of Economics*, 2021, *134* (4), 1883–1948.

Long, J.B. Jr. and C.I. Plosser, "Real Business Cycles," *Journal of Political Economy*, 1983, *91* (1), 39–69.

Nakamura, Emi and Jon Steinsson, "Monetary non-neutrality in a multisector menu cost model," *The Quarterly journal of economics*, 2010, *125* (3), 961–1013.

Pasten, E., R. Schoenle, and M. Weber, "Price Rigidity and the Origins of Aggregate Fluctuations," Working Paper 23750, National Bureau of Economic Research 2017.

Quesnay, F., *Tableau Economique* 1758.

Rubbo, E., "Networks, Phillips Curves and Monetary Policy," Working Paper 2020.

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

_ **, J. Tielens, and J. Van Hove**, "Pipeline Pressures and Sectoral Inflation Dynamics," Working Paper 351, National Bank of Belgium 2018.

A Appendix: Proofs

A.1 Flexible Pricing Model

The cost function is defined by the cost minimization problem:

$$C_i(\boldsymbol{P}_t, W_{i,t}, P_{Z,t}, Y_{i,j,t}/A_{i,t}) = \min_{\boldsymbol{X}_{i,j,t}, L_{i,j,t}, Z_{i,j,t}} \boldsymbol{P}_t' \boldsymbol{X}_{i,j,t} + W_{i,t} L_{i,j,t} + P_{Z,t} Z_{i,j,t}$$
$$\text{s.t. } A_{i,t} F_i(\boldsymbol{X}_{i,j,t}, L_{i,j,t}, Z_{i,j,t}) = Y_{i,t}.$$

It follows from constant returns to scale that

$$MC_{i,t} = C_i(\boldsymbol{P}_t, W_{i,t}, P_{Z,t}, 1/A_{i,t}),$$

which does not vary by firm within an industry. Total cost does, however, because there will be variation among firms' prices in the industry. We have

$$Y_{i,j,t} MC_{i,t} = C_i(\boldsymbol{P}_t, W_{i,t}, P_{Z,t}, Y_{i,j,t}/A_{i,t}).$$

Now it follows immediately from the profit maximization problem that

$$P_{i,t} = \mu_{i,t} \frac{\sigma_i}{\sigma_i - 1} MC_{i,t}.$$

Therefore,

$$d \ln P_{i,t} = d \ln \mu_{i,t} + d \ln MC_{i,t}.$$

It is a straightforward application of the envelope theorem on the cost function that

$$d \ln MC_{i,t} = -d \ln A_{i,t} + \boldsymbol{s}_i' d \ln \boldsymbol{P}_t + s_i^L d \ln W_{i,t} + s_i^Z d \ln P_{Z,t}.$$

So

$$d \ln P_{i,t} = d \ln \mu_{i,t} - d \ln A_{i,t} + \boldsymbol{s}_i' d \ln \boldsymbol{P}_t + s_i^L d \ln W_{i,t} + s_i^Z d \ln P_{Z,t}.$$

Stacking across industries,

$$d \ln \boldsymbol{P}_t = d \ln \boldsymbol{\mu}_t - d \ln \boldsymbol{A}_t + \boldsymbol{\Phi} d \ln \boldsymbol{P}_t + \boldsymbol{s}^{L\prime} \boldsymbol{I} d \ln \boldsymbol{W}_t + \boldsymbol{s}^Z d \ln P_{Z,t}.$$

First note that the unobserved TFP and markup changes can therefore be measured as

$$d \ln \boldsymbol{\mu}_t - d \ln \boldsymbol{A}_t = d \ln \boldsymbol{P}_t - \left(\boldsymbol{\Phi} d \ln \boldsymbol{P}_t + (\boldsymbol{s}^L)' \boldsymbol{I} d \ln \boldsymbol{W}_t + \boldsymbol{s}^Z d \ln P_{Z,t} \right),$$

as claimed. Gathering terms and inverting, we also have

$$d\ln \boldsymbol{P}_t = (\boldsymbol{I} - \boldsymbol{\Phi})^{-1}\left(d\ln \boldsymbol{\mu}_t - d\ln \boldsymbol{A}_t + (\boldsymbol{s}^L)'\boldsymbol{I} d\ln \boldsymbol{W}_t + \boldsymbol{s}^Z d\ln P_{Z,t}\right).$$

A.2 Calvo New Keynesian Model

We find the first-order condition of the equation determining the optimal reset price in an industry and log-linearize it (see Gali (2010)), yielding

$$\hat{p}^*_{i,t} = (1-\theta_i\beta m_f)\sum_{k=0}^{\infty}(\theta_i\beta m_f)^k \mathrm{E}_t[\widehat{mc}_{i,t+k} + \hat{\mu}_{i,t+k}].$$

At this point, it is fine if $m_f = 0$, in which case

$$\hat{p}^*_{i,t} = \widehat{mc}_{i,t} + \hat{\mu}_{i,t},$$

meaning future deviations of marginal cost and desired markups from steady state do not enter the equation determining the deviation of the optimal reset price from steady state. Clearly, $m_f = 1$ is associated with rational expectations. Compactly,

$$\hat{p}^*_{i,t} = \theta_i\beta m_f \mathrm{E}_t[\hat{p}^*_{i,t+1}] + (1-\theta_i\beta m_f)(\widehat{mc}_{i,t} + \hat{\mu}_{i,t}).$$

The industry price index log-linearizes to

$$\hat{p}_{i,t} = \theta_i\hat{p}_{i,t-1} + (1-\theta_i)\hat{p}^*_{i,t}.$$

Combining,

$$\mathrm{E}_t[\hat{p}_{i,t+1}] = \frac{1+\theta_i^2\beta m_f}{\theta_i\beta m_f}\hat{p}_{i,t} - \frac{1}{\beta m_f}\hat{p}_{i,t-1} - \frac{(1-\theta_i)(1-\theta_i\beta m_f)}{\theta_i\beta m_f}(\widehat{mc}_{i,t} + \hat{\mu}_{i,t}).$$

Stacking across industries,

$$\mathrm{E}_t[\hat{\boldsymbol{p}}_{t+1}] = \mathrm{diag}\left(\frac{1+\theta_i^2\beta}{\theta_i\beta m_f}\right)\hat{\boldsymbol{p}}_t - \frac{1}{\beta m_f}\hat{\boldsymbol{p}}_{t-1} - \mathrm{diag}\left(\frac{(1-\theta_i)(1-\theta_i\beta m_f)}{\theta_i\beta m_f}\right)(\widehat{\boldsymbol{mc}}_t + \hat{\boldsymbol{\mu}}_t).$$

Recall that

$$\widehat{\boldsymbol{mc}}_t = \boldsymbol{\Phi}\hat{\boldsymbol{p}}_t + \boldsymbol{\Phi}_{\boldsymbol{Z}}\hat{\boldsymbol{p}}_{Z,t} + \mathrm{diag}(s_i^L)\hat{\boldsymbol{w}}_t - \hat{\boldsymbol{a}}_t.$$

Therefore,

$$\begin{aligned}\mathrm{E}_t[\hat{\boldsymbol{p}}_{t+1}] = &\left(\mathrm{diag}\left(\frac{1+\theta_i^2\beta m_f}{\theta_i\beta m_f}\right) - \mathrm{diag}\left(\frac{(1-\theta_i)(1-\theta_i\beta m_f)}{\theta_i\beta m_f}\right)\boldsymbol{\Phi}\right)\hat{\boldsymbol{p}}_t - \frac{1}{\beta m_f}\hat{\boldsymbol{p}}_{t-1} \\ &- \mathrm{diag}\left(\frac{(1-\theta_i)(1-\theta_i\beta m_f)}{\theta_i\beta m_f}\right)(\boldsymbol{\Phi}_{\boldsymbol{Z}}\hat{\boldsymbol{p}}_{Z,t} + \mathrm{diag}(s_i^L)\hat{\boldsymbol{w}}_t - \hat{\boldsymbol{a}}_t + \hat{\boldsymbol{\mu}}_t).\end{aligned}$$

Note for intuition that if we have one-time, persistent shocks to commodities at time 0, long-run prices converge to

$$\hat{\boldsymbol{p}}_\infty = (\boldsymbol{I} - \boldsymbol{\Phi})^{-1}\boldsymbol{\Phi}_{\boldsymbol{Z}}\hat{\boldsymbol{p}}_{Z,0},$$

which was the solution for the on-impact effect of commodity shocks in the flexible price model. Now, define

$$\hat{\boldsymbol{x}}_{t+1} = \begin{bmatrix}\hat{\boldsymbol{p}}_t \\ \hat{\boldsymbol{p}}_{t+1}\end{bmatrix}, \quad \hat{\boldsymbol{e}}_t = \begin{bmatrix}\hat{\boldsymbol{\mu}}_t \\ \hat{\boldsymbol{a}}_t \\ \hat{\boldsymbol{w}}_t \\ \hat{\boldsymbol{p}}_{Z,t}\end{bmatrix}.$$

Then

$$\mathrm{E}_t\hat{\boldsymbol{x}}_{t+1} = \boldsymbol{B}_x\hat{\boldsymbol{x}}_t + \boldsymbol{B}_e\hat{\boldsymbol{e}}_t,$$

The solution proceeds as follows. Perform an eigendecomposition of $\boldsymbol{B}_x$:

$$\boldsymbol{B}_x = V\Lambda V^{-1}.$$

In R, the authors' preferred programming language for solving this model, the eigenvalues in V with magnitude greater than 1 are stacked first in the resulting decomposition. Define $\tilde{\boldsymbol{x}}_t = V^{-1}\hat{\boldsymbol{x}}_t$ and $\tilde{\boldsymbol{B}}_e = V^{-1}\boldsymbol{B}_e$. Then

$$\mathrm{E}_t\begin{bmatrix}\tilde{\boldsymbol{x}}_{1,t+1} \\ \tilde{\boldsymbol{x}}_{2,t+1}\end{bmatrix} = \begin{bmatrix}\Lambda_1 & \boldsymbol{0} \\ \boldsymbol{0} & \Lambda_2\end{bmatrix}\begin{bmatrix}\tilde{\boldsymbol{x}}_{1,t} \\ \tilde{\boldsymbol{x}}_{2,t}\end{bmatrix} + \begin{bmatrix}\tilde{\boldsymbol{B}}_1 \\ \tilde{\boldsymbol{B}}_2\end{bmatrix}\hat{\boldsymbol{e}}_t,$$

where the diagonal elements of Λ_1 are all greater than 1 and the diagonal elements of Λ_2 are all less than 1. The model has a unique solution if the diagonal of Λ_1 is the same size as $\hat{\boldsymbol{p}}$ (Blanchard and Kahn 1980), which turns out to be the case for all the input output tables published by the BEA. It is not necessary to give the general conditions for solvability of this model for our purposes, and so we do not undertake such a proof here.

Now the explosive eigenvalues can be solved under a transversality condition and a growth restriction on exogenous shocks. We have

$$\mathrm{E}_t[\tilde{\boldsymbol{x}}_{1,t+1}] = \Lambda_1\tilde{\boldsymbol{x}}_{1,t} + \tilde{\boldsymbol{B}}_1\hat{\boldsymbol{e}}_t,$$

which can be forward solved to get

$$\tilde{x}_{1,t} = -\sum_{j=0}^{\infty} \left(\Lambda_1^{-1}\right)^{j+1} \tilde{\boldsymbol{B}}_1 \mathrm{E}_t[\hat{\boldsymbol{e}}_{t+j}] + \lim_{j\to\infty} \left(\Lambda_1^{-1}\right)^{j} \mathrm{E}_t[\tilde{\boldsymbol{x}}_{1,t+j}].$$

The required transverality condition is

$$\lim_{j\to\infty} \left(\Lambda_1^{-1}\right)^{j} \mathrm{E}_t[\tilde{\boldsymbol{x}}_{1,t+j}] = 0.$$

We also require that assume that shocks do not grow at an exponential rate, so that

$$\sum_{j=0}^{\infty} \left(\Lambda_1^{-1}\right)^{j+1} \tilde{\boldsymbol{B}}_1 \mathrm{E}_t[\hat{\boldsymbol{e}}_{t+j}]$$

is finite. Then

$$\tilde{x}_{1,t} = -\sum_{j=0}^{\infty} \left(\Lambda_1^{-1}\right)^{j+1} \tilde{\boldsymbol{B}}_1 \mathrm{E}_t[\hat{\boldsymbol{e}}_{t+j}].$$

If the shocks satisfy $\mathrm{E}_t[\hat{\boldsymbol{e}}_{t+1}] = \boldsymbol{\rho}\hat{\boldsymbol{e}}_t$, with the eigenvalues of $\boldsymbol{\rho}$ all less than 1, we have

$$\tilde{x}_{1,t} = -\sum_{j=0}^{\infty} \left(\Lambda_1^{-1}\right)^{j+1} \tilde{\boldsymbol{B}}_1 \boldsymbol{\rho}^j \hat{\boldsymbol{e}}_t.$$

A special case is $\boldsymbol{\rho} = \rho\boldsymbol{I}$, a useful assumption when we shock only one dimension of $\hat{e}_t$, in which case

$$\tilde{x}_{1,t} = -\Lambda_1^{-1} \sum_{j=0}^{\infty} \left(\Lambda_1^{-1}\rho\right)^{j} \tilde{\boldsymbol{B}}_1 \hat{\boldsymbol{e}}_t = -\Lambda_1^{-1}(\boldsymbol{I} - \Lambda_1^{-1}\rho)^{-1} \tilde{\boldsymbol{B}}_1 \hat{\boldsymbol{e}}_t.$$

Now we can turn to the eigenvalues that are less than 1. Our solution will come from the initial conditions on prices, as lagged prices are a state variable. Rewrite V as

$$V = \begin{bmatrix} V_{11} & V_{12} \\ V_{21} & V_{22} \end{bmatrix}.$$

Then

$$\hat{\boldsymbol{p}}_t = V_{22}V_{12}^{-1}\hat{\boldsymbol{p}}_{t-1} + (V_{21} - V_{22}V_{12}^{-1}V_{11})\tilde{\boldsymbol{x}}_{1,t}.$$

Under a fully persistent shock normalized to occur in period 0, $\hat{\boldsymbol{e}}_t = \hat{\boldsymbol{e}}_0$ for all $t \geq 0$. So

$$\hat{\boldsymbol{p}}_t = V_{22}V_{12}^{-1}\hat{\boldsymbol{p}}_{t-1} + (V_{21} - V_{22}V_{12}^{-1}V_{11})\tilde{\boldsymbol{x}}_0.$$

The long-run pass-through is

$$\hat{\boldsymbol{p}}_\infty = (\boldsymbol{I} - V_{22}V_{12}^{-1})^{-1}(V_{21} - V_{22}V_{12}^{-1}V_{11})\tilde{\boldsymbol{x}}_0,$$

which, as we have already shown, is

$$\hat{\boldsymbol{p}}_\infty = (\boldsymbol{I} - \boldsymbol{\Phi})^{-1}\boldsymbol{\Phi}_Z\boldsymbol{1}$$

when all commodity prices are shocked by 1 log point. This is long-run pass-through under the flexible model. We will focus on the case of having just one exogenous commodity, oil, and shocking that commodity while leaving desired markups, TFP, and wages constant. In this case,

$$\hat{\boldsymbol{p}}_\infty = (\boldsymbol{I} - \boldsymbol{\Phi})^{-1}\boldsymbol{\Phi}_{Oil}\hat{p}_{Oil,0}.$$

Then we have

$$\hat{\boldsymbol{p}}_t - \hat{\boldsymbol{p}}_\infty = (V_{22}V_{12}^{-1})^t(\hat{\boldsymbol{p}}_0 - \hat{\boldsymbol{p}}_\infty),$$

with the initial IRF condition $\hat{\boldsymbol{p}}_{-1} = 0$ pinning down

$$\hat{\boldsymbol{p}}_0 = (V_{21} - V_{22}V_{12}^{-1}V_{11})\tilde{\boldsymbol{x}}_0.$$

Therefore,

$$\hat{\boldsymbol{p}}_t = (V_{22}V_{12}^{-1})^t\hat{\boldsymbol{p}}_0 + (\boldsymbol{I} - (V_{22}V_{12}^{-1})^t)\hat{\boldsymbol{p}}_\infty.$$

Now, in the continuous time approximation, $\hat{\boldsymbol{p}}_0 \approx 0$, and in this case we can separate the timing of pass-through due to direct and indirect exposure to oil as we did in our reduced form regressions. Formally, recall we can write

$$\hat{\boldsymbol{p}}_\infty = \underbrace{\boldsymbol{\Phi}_{Oil}}_{\text{Direct}}\hat{p}_{Oil,0} + \underbrace{((\boldsymbol{I} - \boldsymbol{\Phi})^{-1} - \boldsymbol{I})\boldsymbol{\Phi}_{Oil}}_{\text{Indirect}}\hat{p}_{Oil,0}.$$

Therefore, we have

$$\hat{\boldsymbol{p}}_t = (V_{22}V_{12}^{-1})^t\hat{\boldsymbol{p}}_0 + (\boldsymbol{I} - (V_{22}V_{12}^{-1})^t)\left(\underbrace{\boldsymbol{\Phi}_{Oil}}_{\text{Direct}}\hat{p}_{Oil,0} + \underbrace{((\boldsymbol{I} - \boldsymbol{\Phi})^{-1} - \boldsymbol{I})\boldsymbol{\Phi}_{Oil}}_{\text{Indirect}}\hat{p}_{Oil,0}\right).$$

The value of $(V_{22}V_{12}^{-1})^t$ can be efficiently computed using the eigendecomposition

$$(V_{22}V_{12}^{-1}) = \tilde{V}\tilde{\Lambda}\tilde{V}^{-1},$$

so that

$$(V_{22}V_{12}^{-1})^t = \tilde{V}\tilde{\Lambda}^t\tilde{V}^{-1}.$$

Putting everything together, we have

$$\hat{\boldsymbol{p}}_t = \tilde{V}\tilde{\Lambda}^t\tilde{V}^{-1}(V_{21} - V_{22}V_{12}^{-1}V_{11})\tilde{\boldsymbol{x}}_0 + (\boldsymbol{I} - \tilde{V}\tilde{\Lambda}^t\tilde{V}^{-1})\left(\underbrace{\boldsymbol{\Phi}_{Oil}}_{\text{Direct}}\hat{p}_{Oil,0} + \underbrace{((\boldsymbol{I} - \boldsymbol{\Phi})^{-1} - \boldsymbol{I})\boldsymbol{\Phi}_{Oil}}_{\text{Indirect}}\hat{p}_{Oil,0}\right),$$

with

$$\tilde{\boldsymbol{x}}_0 = -\Lambda_1^{-1}(\boldsymbol{I} - \Lambda_1^{-1})^{-1}\tilde{\boldsymbol{B}}_1(0, 0, ..., 0, \hat{p}_{Oil,0})',$$

where the last vector represents that we are not shocking desired markups, TFP, or wages but are shocking oil prices by $\hat{p}_{Oil,0}$.

A.3 Model in Continuous Time

Define ϕ_i as the instantaneous probability that a firm in industry i can update prices, and $\tilde{E}_t$ as the (potentially myopic) expectations operator, to be defined in more detail later following Gabaix (2020).

The optimal reset price for any firm j in industry i is the argmax of

$$\max_{P_{i,t}^*} \int_{\tau=t}^{\infty} e^{-\phi_i(\tau-t)}\tilde{\mathrm{E}}_t\left[SDF_{t,\tau}X_{i,j,\tau}\left(\frac{P_{i,t}^*}{\mu_{i,\tau}} - MC_{i,\tau}\right)\right]d\tau$$

subject to the demand conditions

$$X_{i,j,\tau} = Y_{i,\tau}\left(\frac{P_{i,\tau}}{P_{i,t}^*}\right).$$

The first-order condition log-linearizes to

$$\int_{\tau=t}^{\infty} e^{-\phi_i(\tau-t)}\tilde{\mathrm{E}}_t\left[SDF_{t,\tau}^{SS}(\hat{p}_{i,t}^* - \hat{\mu}_{i,\tau} - \widehat{mc}_{i,\tau})\right]d\tau = 0,$$

where a superscript SS denotes a variable's steady state value. Now, myopia in Gabaix (2020) is defined in our setting as

$$\tilde{\mathrm{E}}_t[\hat{f}_\tau] = e^{-\tilde{m}_f(\tau-t)}\mathrm{E}_t[\hat{f}_\tau],$$

where E_t is the rational expectations operator, $\hat{f}$ is the deviation of any of our variables above from steady state, and $\tau \geq t$. Therefore, for $\tilde{m}_f > 0$, firms neglect future deviations of variables of interest from steady state in making their optimization decisions, and for $\tilde{m}_f = 0$ we recover rational expectations.

Now, for a standard household problem with no myopia, we have $SDF_{t,\tau}^{SS} = e^{-\rho(\tau-t)}$, where ρ is the discount rate. We could allow for separate household myopia here but do not for now.

Therefore, the above equation simplifies to

$$\mathrm{E}_t \int_{\tau=t}^{\infty} e^{-(\phi_i+\rho+\tilde{m}_f)(\tau-t)} \left[(\hat{p}^*_{i,t} - \hat{\mu}_{i,\tau} - \widehat{mc}_{i,\tau})\right] = 0,$$

or

$$\hat{p}^*_{i,t} = (\phi_i+\rho+\tilde{m}_f)\mathrm{E}_t \int_{\tau=t}^{\infty} e^{-(\phi_i+\rho+\tilde{m}_f)(\tau-t)} \left[\hat{\mu}_{i,\tau} + \widehat{mc}_{i,\tau})\right] d\tau,$$

The industry price index satisfies

$$\dot{\hat{p}}_{i,t} = \phi_i(\hat{p}^*_{i,t} - \hat{p}_{i,t}).$$

A.4 Linear Model

Eliminating hats to simplify notation, it is clear from our continuous time model (eliminating changes in wages and exogenous markups and setting $\phi_i = \phi$ for all i) that

$$\begin{aligned} p^*_{n,t} &= \psi \mathrm{E}_t \int_{\tau=t}^{\infty} e^{-\psi(\tau-t)} \left[s p_{n-1,t})\right] d\tau \\ \dot{p}_{n,t} &= \phi(p^*_{n,t} - p_{n,t}), \end{aligned}$$

with $\psi = \phi_i + \rho + \tilde{m}_f$. We focus on the case of a one-time, permanent increase in the commodity price and no additional future shocks, so that rational expectations drop out (though myopia will still play a role). In the case of $\psi \to \infty$, we get

$$\begin{aligned} p^*_{n,t} &= s p_{n-1,t} \\ \dot{p}_{n,t} &= \phi(p^*_{n,t} - p_{n,t}), \end{aligned}$$

or just

$$\dot{p}_{n,t} = \phi(s p_{n-1,t} - p_{n,t}).$$

We conjecture that the solution will take the form

$$p_{n,t} = b_{n,t} p_{n,\infty}, \quad \text{with } b_{n,t} = 1 - e^{-\phi t} q_n(\phi t),$$

It follows that

$$\dot{b}_{n,t} = \phi(b_{n-1,t} - b_{n,t}),$$

and plugging in our guess, this implies

$$\dot{q}_n(\phi t) = q_{n-1}(\phi t).$$

Because the commodity shock is persistent, we have $q_0(t) = 0$ for all t, and in continuous time we have the boundary condition that $q_n(0) = 1$ for all n, assuming WLOG that the shock occurs at time $t = 0$. Successively integrating, we have

$$\begin{aligned} q_1(\phi t) &= 1 \\ q_2(\phi t) &= 1 + \phi t \\ q_3(\phi t) &= 1 + \phi t + \frac{(\phi t)^2}{2} \\ q_4(\phi t) &= 1 + \phi t + \frac{(\phi t)^2}{2} + \frac{(\phi t)^3}{6}, \end{aligned}$$

and therefore the result that

$$q_n(\phi t) = \sum_{i=0}^{n-1} \frac{(\phi t)^i}{i!}.$$

Now, we define $t_n(X)$ implicitly as satisfying

$$b_{n,t_n(X)} = 1 - e^{-\phi t_n(X)} \sum_{i=0}^{n-1} \frac{(\phi t_n(X))^i}{i!} = 1 - e^{-X}$$

In the limit taking $X \to \infty$,

$$e^{-\phi t_n(X)} \phi t_n(X)^{n-1} \sim (n-1)! e^{-X}.$$

So

$$\phi t - (n-1)\ln(\phi t) \sim X - \ln((n-1)!).$$

Now we certainly know that

$$\phi t \sim \phi t - (n-1)\ln(\phi t),$$

therefore, by transitivity,

$$\phi t = X + o(X)$$

so

$$\phi t = X + (n-1)\ln X + o(n \ln X),$$

and dividing by ϕ gives our result characterizing $t_n(X)$.

Now, we discuss the case where we include some degree of forward-lookingness, i.e. we do not let $\psi \to \infty$. We work with the system

$$\begin{aligned} p^*_{n,t} &= \psi \mathrm{E}_t \int_{\tau=t}^{\infty} e^{-\psi(\tau-t)} \left[s p_{n-1,t}) \right] d\tau \\ \dot{p}_{n,t} &= \phi(p^*_{n,t} - p_{n,t}). \end{aligned}$$

We conjecture that

$$p_{n,t} = b_{n,t}p_{n,\infty}, \quad p^*_{n,t} = b^*_{n,t}p_{n,\infty},$$

with

$$b_{n,t} = 1 - e^{-\phi t}q_n(t), \quad b^*_{n,t} = 1 - e^{-\phi t}q^*_n(t).$$

Plugging in the first step of our guess, we can rewrite our system as

$$\begin{aligned} b^*_{n,t} &= \psi \mathrm{E}_t \int_{\tau=t}^{\infty} e^{-\psi(\tau-t)} b_{n-1,s} d\tau \\ \dot{b}_{n,t} &= \phi(b^*_{n,t} - b_{n,t}). \end{aligned}$$

Plugging in the second step of our guess, we can simplify the system further to

$$q^*_{n,t} = \psi \int_{u=0}^{\infty} e^{-(\psi+\phi)u} q_{n-1}(t+u)du \tag{15}$$

$$\dot{q}_{n,t} = \phi q^*_n(t). \tag{16}$$

The first equality follows from

$$\begin{aligned} e^{-\phi t}q^*_n(t) = 1 - b^*_{n,t} &= 1 - \int_t^{\infty} e^{-\psi(s-t)}\psi b_{n-1,s} ds \\ &= \int_t^{\infty} e^{-\psi(s-t)}\psi(1 - b_{n-1,s})ds \\ &= \int_t^{\infty} e^{-\psi(s-t)}\psi e^{-\phi s} q_{n-1}(s)ds \\ &= e^{-\phi t}\int_0^{\infty} e^{-(\psi+\phi)u}\psi q_{n-1}(t+u)du. \end{aligned}$$

By successive integration and induction, we can show that $q_n(t)$ and $q^*_n(t)$ are polynomials of degree $n-1$ and $n-2$, respectively. For example,

$$\begin{aligned} q^*_2(t) &= \frac{\psi}{\psi+\phi}, \quad q_2(t) = \frac{t\psi\phi + \psi + \phi}{\psi+\phi} \\ q^*_3(t) &= \frac{\psi\left(\phi^2(t\psi+1) + \psi\phi(t\psi+3) + \psi^2\right)}{(\psi+\phi)^3} \\ q_3(t) &= \frac{\phi^3\left(t^2\psi^2 + 2t\psi + 2\right) + \psi\phi^2\left(t^2\psi^2 + 6t\psi + 6\right) + 2\psi^2\phi(t\psi+3) + 2\psi^3}{2(\psi+\phi)^3}, \end{aligned}$$

etc. Again, we define $t^F_n(X)$ implicitly as satisfying

$$b_{n,t^F_n(X)} = 1 - e^{-X}.$$

Because $q_n(t)$ and $q^*_n(t)$ are polynomials of the aforementioned degree, $q^*_n(t) \sim C^*_n t^{n-2}$ and $q_n(t) \sim$

$C_n t^{n-1}$ for large t (implicitly, large X, as before). Then equation 15 yields $C_n^* = C_{n-1}\frac{\psi}{\psi+\phi}$, and equation 16 yields $C_n = \frac{\phi}{n-1}C_n^*$. Combining, $C_n = \frac{1}{n-1}\frac{\psi\phi}{\psi+\phi}C_{n-1}$. Our initial condition $q_0 = 1$ yields $C_1 = 1$. Therefore,

$$q_n(t) \sim \frac{1}{(n-1)!}\left(\frac{\phi\psi}{\psi+\phi}\right)^{n-1} t^{n-1}.$$

Using similar arguments to before, we therefore have

$$t_n^F(X) = \frac{X}{\phi} + (n-1)\frac{\ln\left(X\frac{\psi}{\psi+\phi}\right)}{\phi} + o\left(n\frac{\ln\left(X\frac{\psi}{\psi+\phi}\right)}{\phi}\right).$$

A.5 Proof of Proposition 5

Define

$$\frac{\partial\Delta\overline{P}}{\partial\Delta\boldsymbol{P}_{Oil}}(\boldsymbol{\alpha}) = \frac{1}{N}\sum_{i=1}^{N}\frac{\partial\Delta P_i}{\partial\Delta\boldsymbol{P}_{Oil}}(\boldsymbol{\alpha}),$$

where

$$\frac{\partial\Delta P_i}{\partial\Delta\boldsymbol{P}_{Oil}}(\boldsymbol{\alpha}) = \left(\frac{\partial\Delta P_i}{\partial\Delta P_{Oil}}(0,\boldsymbol{\alpha}), \frac{\partial\Delta P_i}{\partial\Delta P_{Oil}}(1,\boldsymbol{\alpha}), ..., \frac{\partial\Delta P_i}{\partial\Delta P_{Oil}}(K,\boldsymbol{\alpha})\right)'.$$

Let S be a lower triangular matrix of 1's (including the diagonal), and define

$$G = S\mathrm{diag}\left(\frac{\partial\Delta\overline{P}}{\partial\Delta\boldsymbol{P}_{Oil}}(\boldsymbol{\alpha})\right).$$

Then

$$G\mathbf{1} = \frac{1}{N}\sum_{i=1}^{N}\boldsymbol{IRF}_i(\boldsymbol{\alpha}), \quad G\hat{\boldsymbol{\beta}}(\boldsymbol{\alpha}) = \frac{1}{N}\sum_{i=1}^{N}\boldsymbol{IRF}_i^{Data}(\boldsymbol{\alpha}),$$

and

$$\hat{m}(\boldsymbol{\alpha}) = G(\hat{\boldsymbol{\beta}}(\boldsymbol{\alpha}) - \mathbf{1}).$$

Then

$$\widehat{\mathrm{Var}}(\hat{m}(\boldsymbol{\alpha})) = G\widehat{\mathrm{Var}}(\hat{\boldsymbol{\beta}}(\boldsymbol{\alpha}))G'$$

and

$$\hat{m}(\boldsymbol{\alpha})'\widehat{\mathrm{Var}}(\hat{m}(\boldsymbol{\alpha}))^{-1}\widehat{\mathrm{Var}}(\hat{m}(\boldsymbol{\alpha})) = (\hat{\boldsymbol{\beta}}(\boldsymbol{\alpha}) - \mathbf{1})'G'(G')^{-1}\widehat{\mathrm{Var}}(\hat{\boldsymbol{\beta}}(\boldsymbol{\alpha}))^{-1}G^{-1}G(\hat{\boldsymbol{\beta}}(\boldsymbol{\alpha}) - \mathbf{1}).$$

The result follows.

B Appendix: Processing of the BEA's Input-output Tables

This section outlines our processing of the BEA's input-output tables, which follows the BEA's guidance as provided in the "Concepts and Methods of the U.S. Input-Output Accounts," originally published in 2006 and updated in April 2009 (the most recent documentation available on the BEA's website as of the writing of this paper).

The BEA publishes several datasets in its input-output accounts that are worth distinguishing. It publishes benchmark Make and Use tables every 5 years, constructed primarily using the microdata underlying the Economic Census, which is run every 5 years (1977, 1982, ...). The BEA publishes before- and after-redefinitions versions of each of these files. The Make tables measure how much each industry i produces each commodity j. Commodities are distinct from the commodities we describe in the paper, which comprise primarily upstream goods such as those from the Oil and Gas Extraction sector. In particular, the Make tables measure how much each industry produces their primary output, but they also tell us how much each industry produces outputs primarily sold by other industries. The Use tables measure how much each industry i purchases each commodity j, but commodity j might be produced by the main industry producer of that commodity or by another industry that produces that commodity. Creating an industry-by-industry input-output table, which is required by our model, therefore requires combining information in both the Make and Use tables.

Before-redefinitions Make and Use tables represent the BEA's best efforts to create input-output measures based on the raw data. Redefinitions are made in the After-redefinition tables "when the input structure for a secondary product of an industry differs significantly from the input structure for the primary product of that industry" (p. 4-6 of the BEA's documentation). For example, the hotel industry often runs restaurants, and the input mix for restaurants differs substantially from that of hotels. The BEA tries to reallocate restaurant output from the hotel sector to the restaurant sector in the after-redefinitions tables. We elect to use before-redefinitions tables in our analysis because they accord better with the standard PPI data published by the BEA. Specifically, the producer price index for an industry in principle represents a weighted average of prices of all products and services an industry supplies. If an industry produces outputs that are primarily produced by other sectors, the prices of these outputs are contained in the industry's PPI. After-redefinitions tables could be used in combination with the BLS's publications on PPIs by major industry products; for instance, the BLS publishes a primary PPI dataset for the primary outputs sold by an industry, and these primary PPIs may be a good match to the after-redefinitions input-output tables. Because the BEA's formal methodology for redefinitions is obscure, however, it is difficult to know how good a match the primary PPIs are with the after-redefinitions input-output tables.

Our harmonization of the BEA's Make and Use tables to produce Before-redefinitions industry-by-industry input-output tables follows exactly the BEA documentation starting on page 12-21, and so we refer the reader there for our methodology. In 2007 and 2012, the BEA publishes industry-by-industry input-output tables before-redefinitions in the Total Requirements format, which represents the BEA's measure of our Leontief inverse object, $(\boldsymbol{I} - \boldsymbol{\Phi})^{-1}$. We are able to replicate the BEA's Total Requirements tables for 2007 and 2012. For other NAICS years, 1997 and 2002, we use the same methodology that replicated the BEA's published industry-by-industry total requirements tables before redefinitions, but we cannot verify that they are the same as what the BEA would have published. For our case study of the 1979 oil shock, we replicate this same procedure on the Make and Use tables before-redefinitions published in 1977 to create an industry-by-industry input-output table before redefinitions.

C Appendix: Additional Robustness of Reduced-form Empirics

C.1 Binscatter Evidence that Visualizes Identifying Variation

To illustrate that our finding of full pass-through is not merely an artefact of complex regression specifications, we plot some simple binscatters of industry price changes on industry oil cost changes. We demean these variables by their average value in each time period. We then split the data into 100 quantiles of industry oil cost changes and show the results for a variety of time horizons. Figure 16 shows that - regardless of whether one examines a one-month horizon, a three-month horizon, a six-month horizon, or a one-year horizon - there is robust evidence of a high degree of pass-through. In particular, the slope increases with the time horizon, and by the one-year horizon, the slope of the line of best fit through the binscatter is approximately one, revealing evidence of full pass-through.

These binscatters also reveal little to no evidence for heterogeneity on the size or sign of cost shocks industries are exposed to. The slope does not appear to vary on either side of the origin, nor does it appear to be steeper for larger shocks than smaller ones - at least in the overall data. In Figure 17, I split the sample of industries into upstream and downstream industries (i.e., industries with below- and above-median measures of downstreamness). It is evident that the upstream industries have no heterogeneity in pass-through on the sign of the cost shock they experience, whereas the downstream industries to exhibit such a heterogeneity. pass-through is lower for positive cost shocks than negative ones. This is consistent with either a higher ability of downstream industries to substitute across inputs in the face of price increases or a reluctance on the part of downstream, consumer-facing industries to raise the ire of consumers through large or frequent price increases.

Figure 16 – Binscatters at Various Time Horizons

Panel 1: One Month

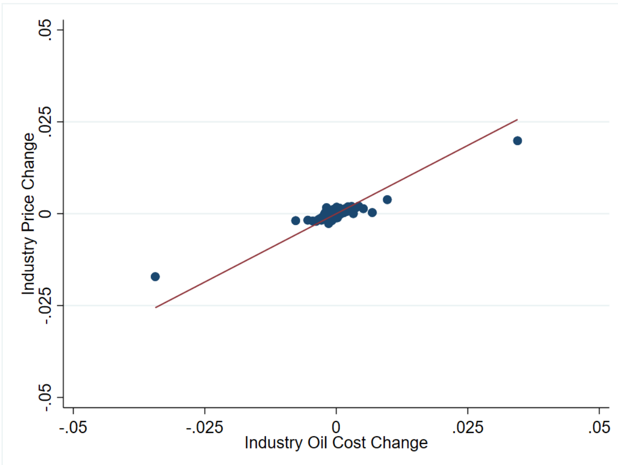

Panel 2: Three Months

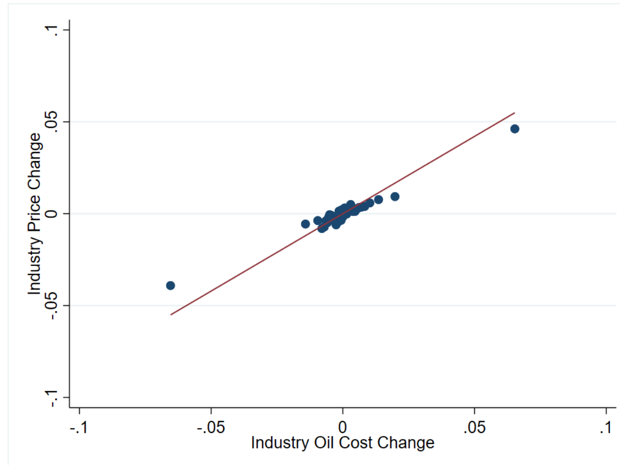

Panel 3: Six Months

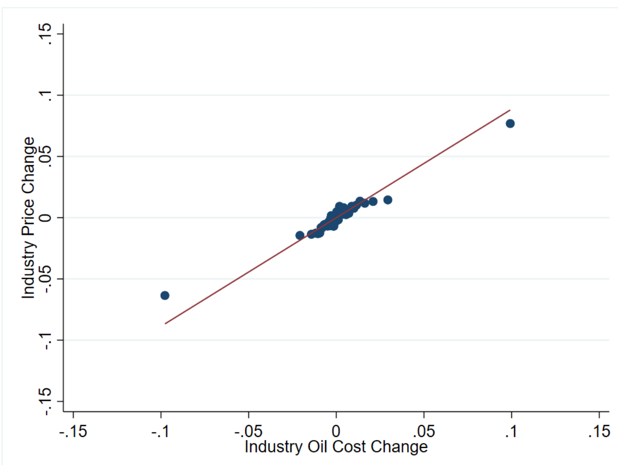

Panel 4: One Year

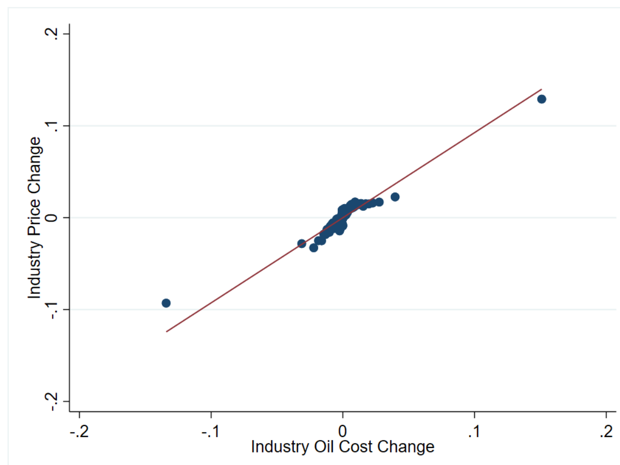

Note: These plots are 100-quantile binscatters displaying how (de-meaned) industry price changes vary with (de-meaned) industry oil cost changes. Lines of best fit are included. The slope of these lines can be interpreted as the fraction of cost increases that are passed through into prices over the corresponding time horizon.

Figure 17 – Binscatters: Upstream vs. Downstream Industries

Panel 1: Upstream Industries

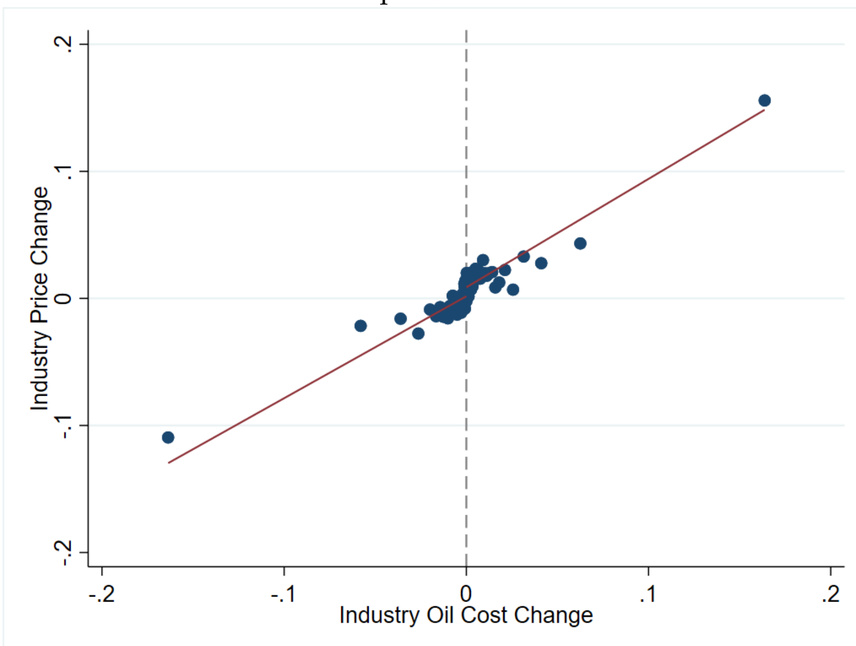

Panel 2: Downstream Industries

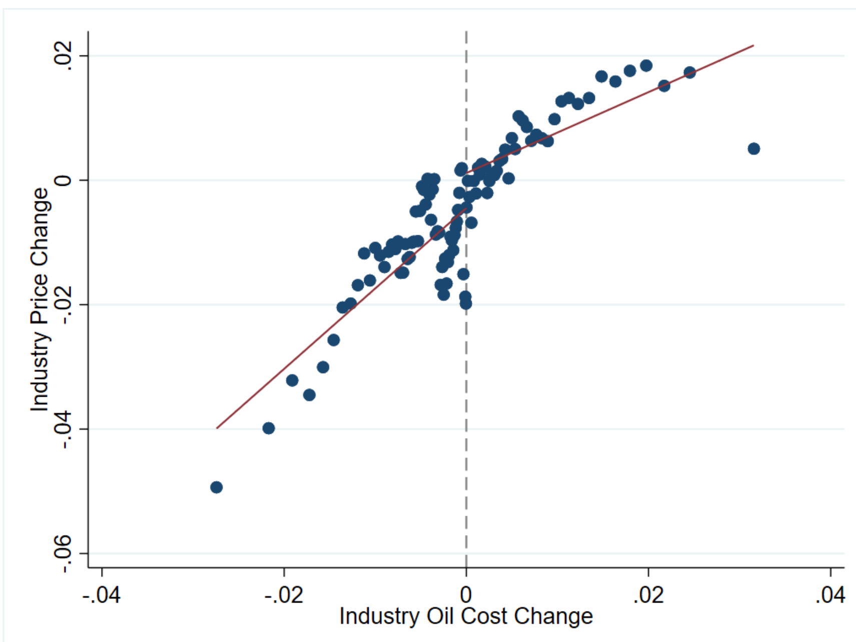

Note: These plots are 100-quantile binscatters displayihng how (de-meaned) industry price changes vary with (de-meaned) industry oil cost changes. Lines of best fit with a regression discontinuity at zero are included. For upstream industries (industries with below-median downstreamness), the slope is no different for negative and positive cost shocks. For downstream consumer-facing industries (industries with above-median downstreamness), there is evidence of lesser pass-through of cost increases.

C.2 Wage-inclusive Shocks

It is possible to replicate our main results after recomputing measures of our right-hand-side shock variable that make use of data on seasonally-adjusted sectoral wages instead of assuming exogeneity of the wage variation. Specifically, as revealed in Proposition 1, industry price shocks are a function of both price changes in the underlying commodities and wage changes. In the main specifications in the body of the paper, we assume that these wage changes are uncorrelated with price changes. To assess the robustness of our main results to violation of this assumption, we directly control for these industry wage changes using data from the Quarterly Census of Employment and Wages (QCEW). Figure 18 shows that this scarcely changes the findings.

Figure 18 – Month-by-Month pass-through of Oil Price Changes

Panel 1: All Variation

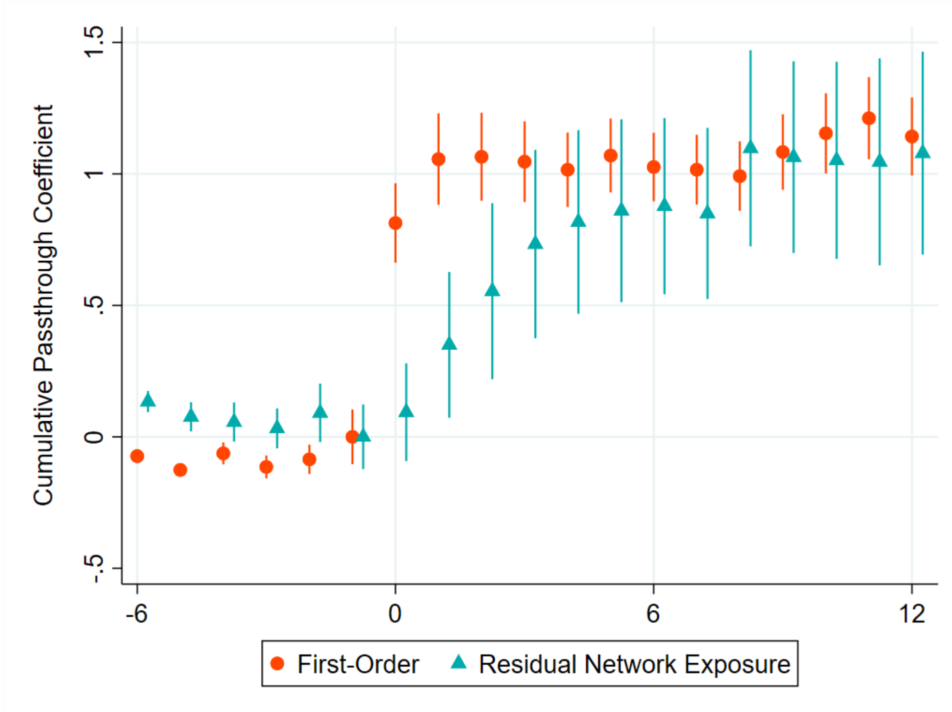

Panel 2: Kanzig IV Variation

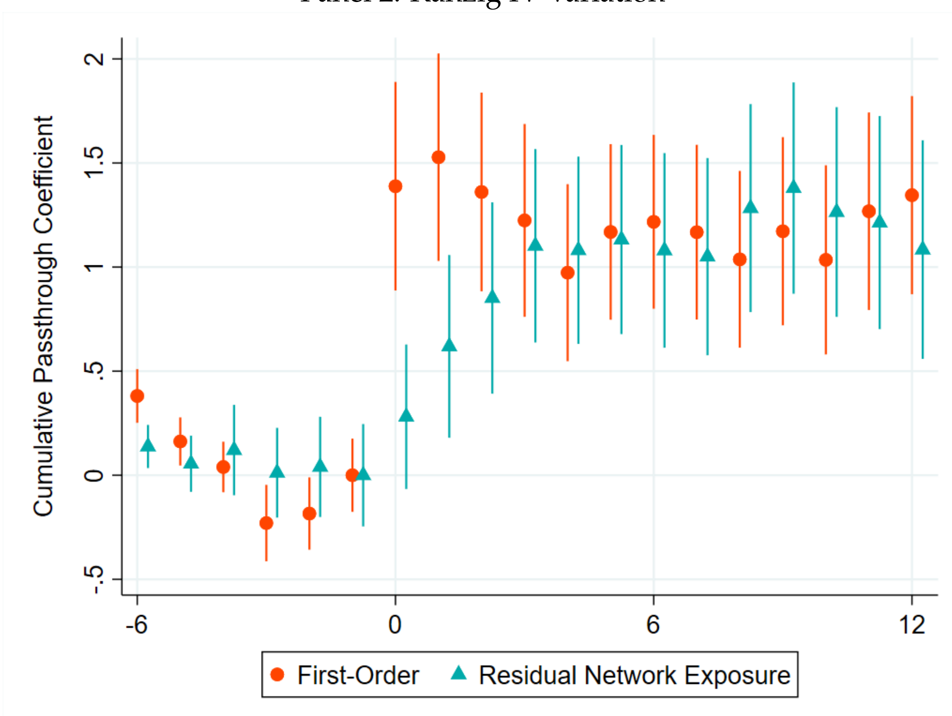

Note: Regression specifications correspond to Equation (7). Red coefficients plot monthly price pass-through of first-order exposure to wage-inclusive crude oil shocks; blue coefficients plot monthly price pass-through of second- and higher-order network exposure to wage-inclusive crude oil shocks.

C.3 Controlling for Gas and Electricity Exposure

Industries with high exposure to oil may have the ability to substitute away to gas and/or electricity when the price of oil rises relative to the prices of those commodities. To the extent this occurs, it might bias our pass-through coefficients if we do not control for the latter. We replicate our main results controlling for gas and electricity exposure to show that the main findings on oil are not driven by these related sectors. These findings are displayed in Figure 19; again, the main results are scarcely changed.

Figure 19 – Month-by-Month pass-through of Oil Price Changes

Panel 1: All Variation

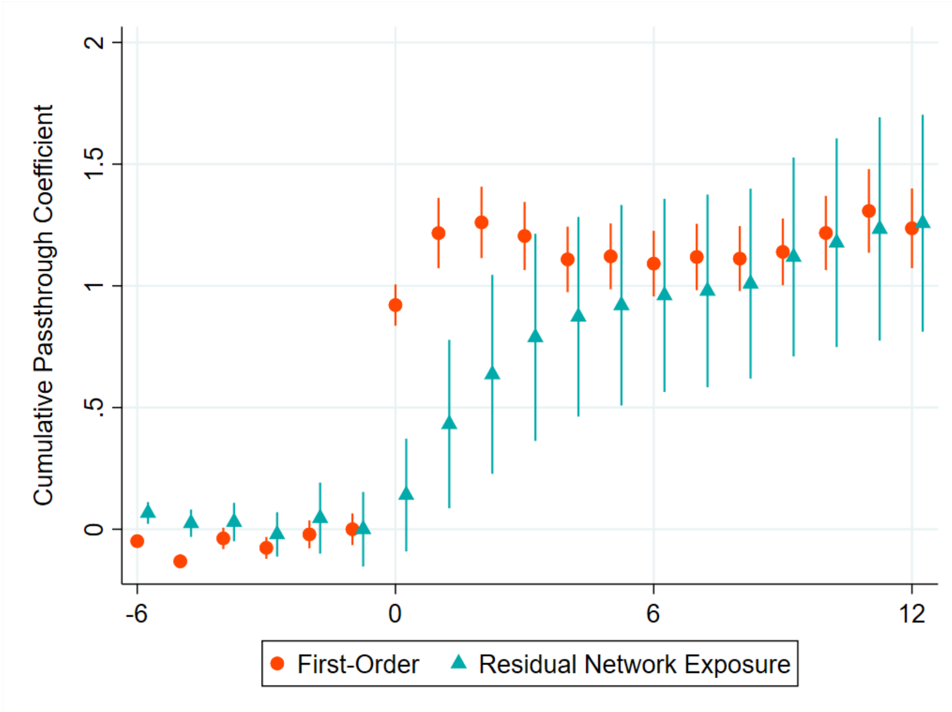

Panel 2: Kanzig IV Variation

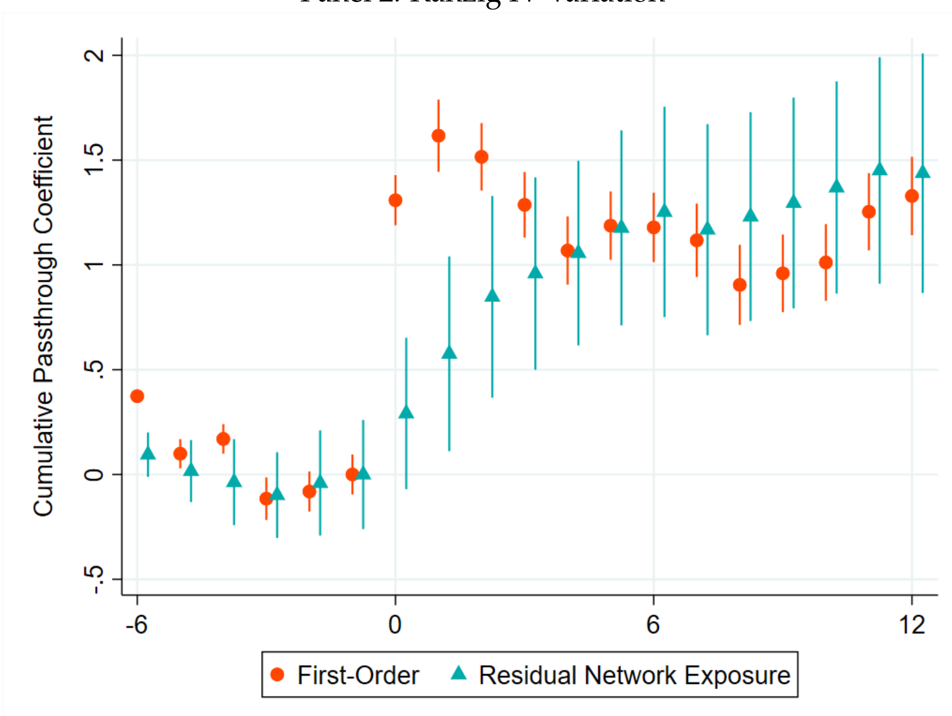

Note: Regression specifications correspond to Equation (7) with controls for gas and electricity shock exposure. Red coefficients plot monthly price pass-through of first-order exposure to crude oil price shocks; blue coefficients plot monthly price pass-through of second- and higher-order network exposure to crude oil price shocks.

C.4 Cost Shares Excluding Payments to Capital

The cost shares utilized in our main regression specifications divide the network cost of oil by the summed network cost of all commodities plus the cost of labor and capital. We compute alternative shares excluding the cost of capital from the denominator and repeat our main analysis, yielding Figure 20. Once again, the extent and pace of pass-through is little changed from baseline.

Figure 20 – Month-by-Month pass-through of Oil Price Changes

Panel 1: All Variation

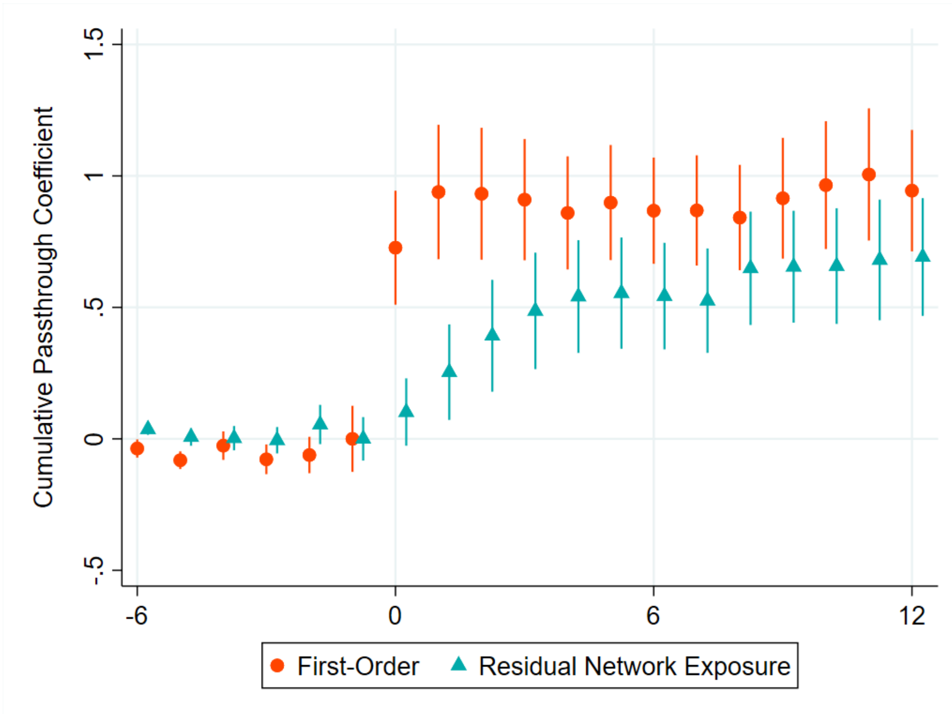

Panel 2: Kanzig IV Variation

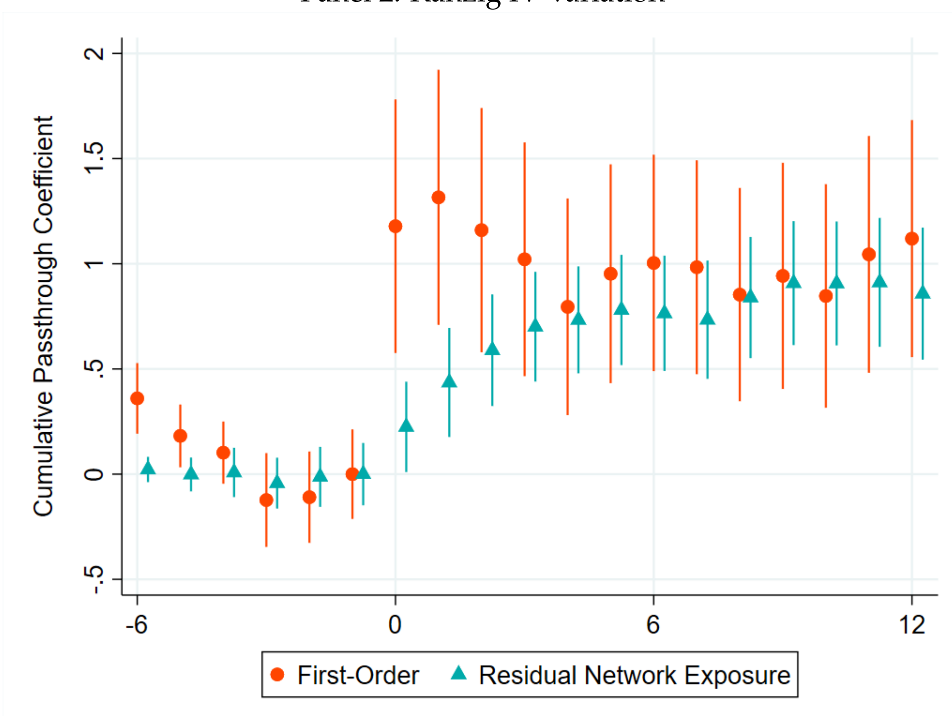

Note: Regression specifications correspond to Equation (7), albeit with alternative oil cost shares that do not include capital in the denominator. Red coefficients plot monthly price pass-through of first-order exposure to crude oil price shocks; blue coefficients plot monthly price pass-through of second- and higher-order network exposure to crude oil price shocks.

C.5 Non-commodity Shocks

Having previously studied oil shocks and non-oil commodity shocks, we replicate our main results using shocks driven by non-commodities (i.e., sectors further downstream than commodities with NAICS codes beginning in 23 or higher). Once again, the results are similar to our other findings, as seen in Figure 21.

Figure 21 – Month-by-Month pass-through of Downstream Price Changes

Panel 1: Full Network

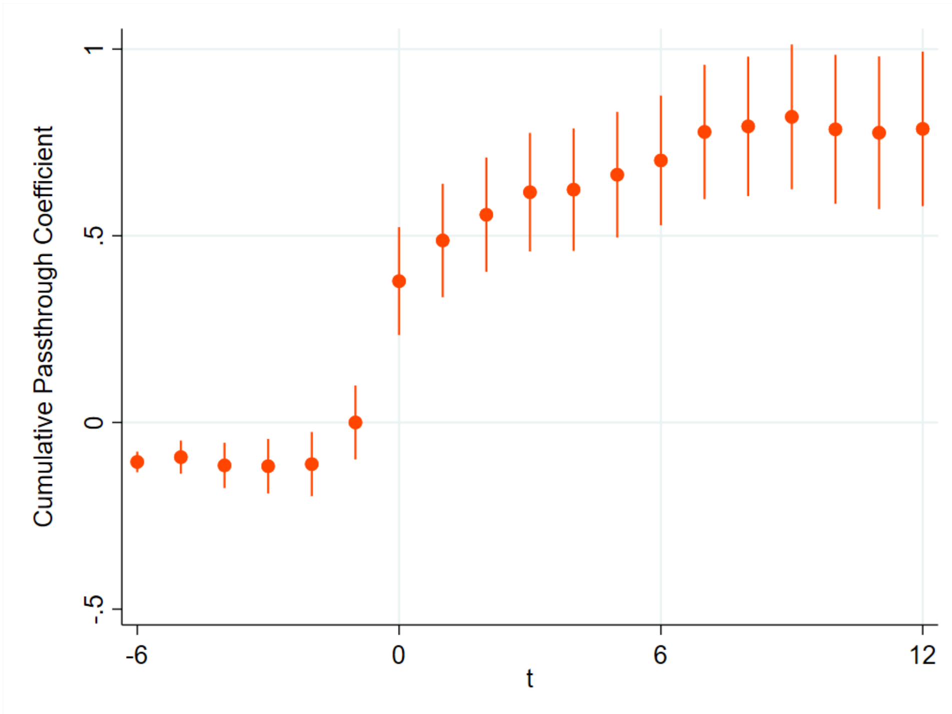

Panel 2: Direct vs. Indirect

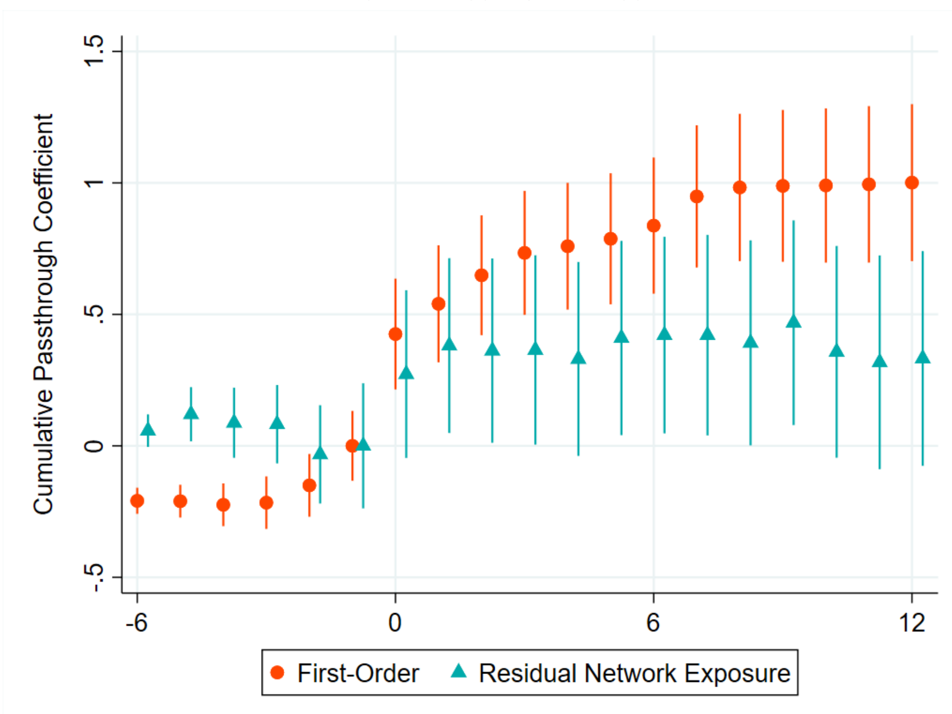

Note: Regression specifications correspond to Equations (6) and (7). Red coefficients plot monthly price pass-through of first-order exposure to non-commodity (NAICS greater than 22) price shocks; blue coefficients plot monthly price pass-through of second- and higher-order network exposure to non-commodity price shocks.

C.6 Permutation Tests

In order to confirm the validity of our standard errors, as an alternative method of generating p-values, we apply permutation tests to our main specifications. In particular, we randomly permutate treatment across industries 1000 times. We run our main specifications on these placebo variations and compare the magnitude of the coefficients resulting from these regressions with the magnitude of the actual coefficients, yielding information on the likelihood with which the actual coefficients resulted from pure chance. In particular, the permutation tests are performed on the cumulative one-year pass-through coefficient. Panels 1a and 1b of Figure 22 correspond to the regression specification given by Equation (6) measuring pass-through of total network exposure. Panel 1a uses all oil price variation, whereas Panel 1b uses the Kanzig IV variation. In both cases, the p-value of the actual regression coefficient is $p < 0.001$. Panels 2a and 2b correspond to the regression specification given by Equation (7) measuring pass-through of direct and indirect network exposure separately. Again, both direct and indirect network exposure are strongly robust to the permutation test, yielding p-values below 1% in all cases.

Figure 22 – Results of Permutation Tests

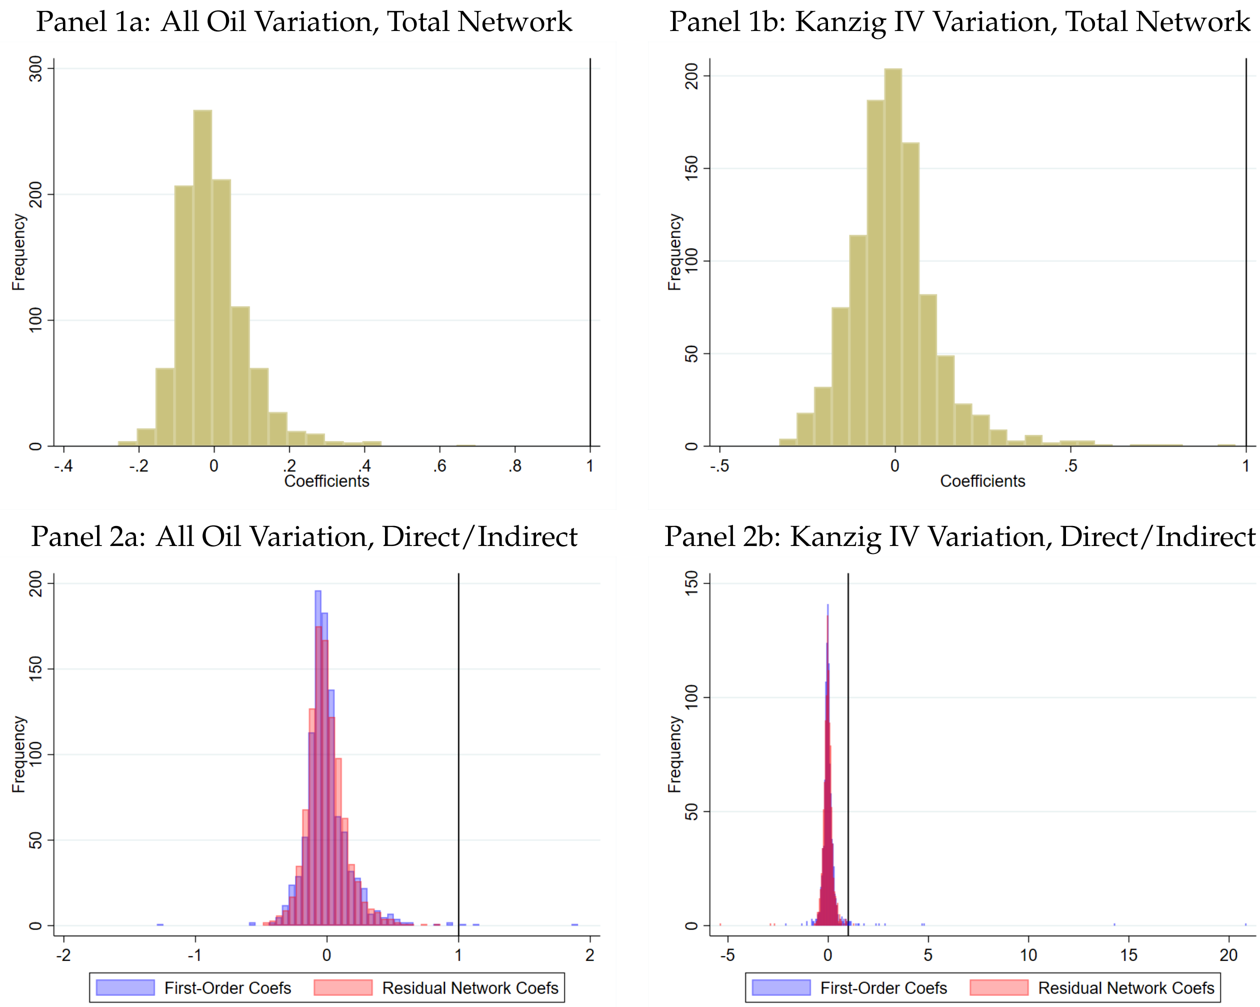

Note: These plots display the results of 1000-repetition permutation tests on the 12-month cumulative pass-through coefficient from the main specifications, as given by Equation (6) in the top panel and (7) in the bottom panel.

D Appendix: Robustness of Structural Estimation

D.1 The Profit Share of Payments to Capital

Step 1 of the GMM estimation yields the loss function displayed in Panel 1a of Figure 23. Step 2, using the estimate of the optimal weight matrix, changes the loss function to that displayed in Panel 2a. The analogous figures using the Kanzig IV variation are shown in Panels b. We see that the estimates of α are broadly stable across steps of the estimation process, suggesting reweighting of noisily estimated components of the IRFs in the data is not substantially affecting estimation.

Figure 23 – Loss Functions for Estimation of α

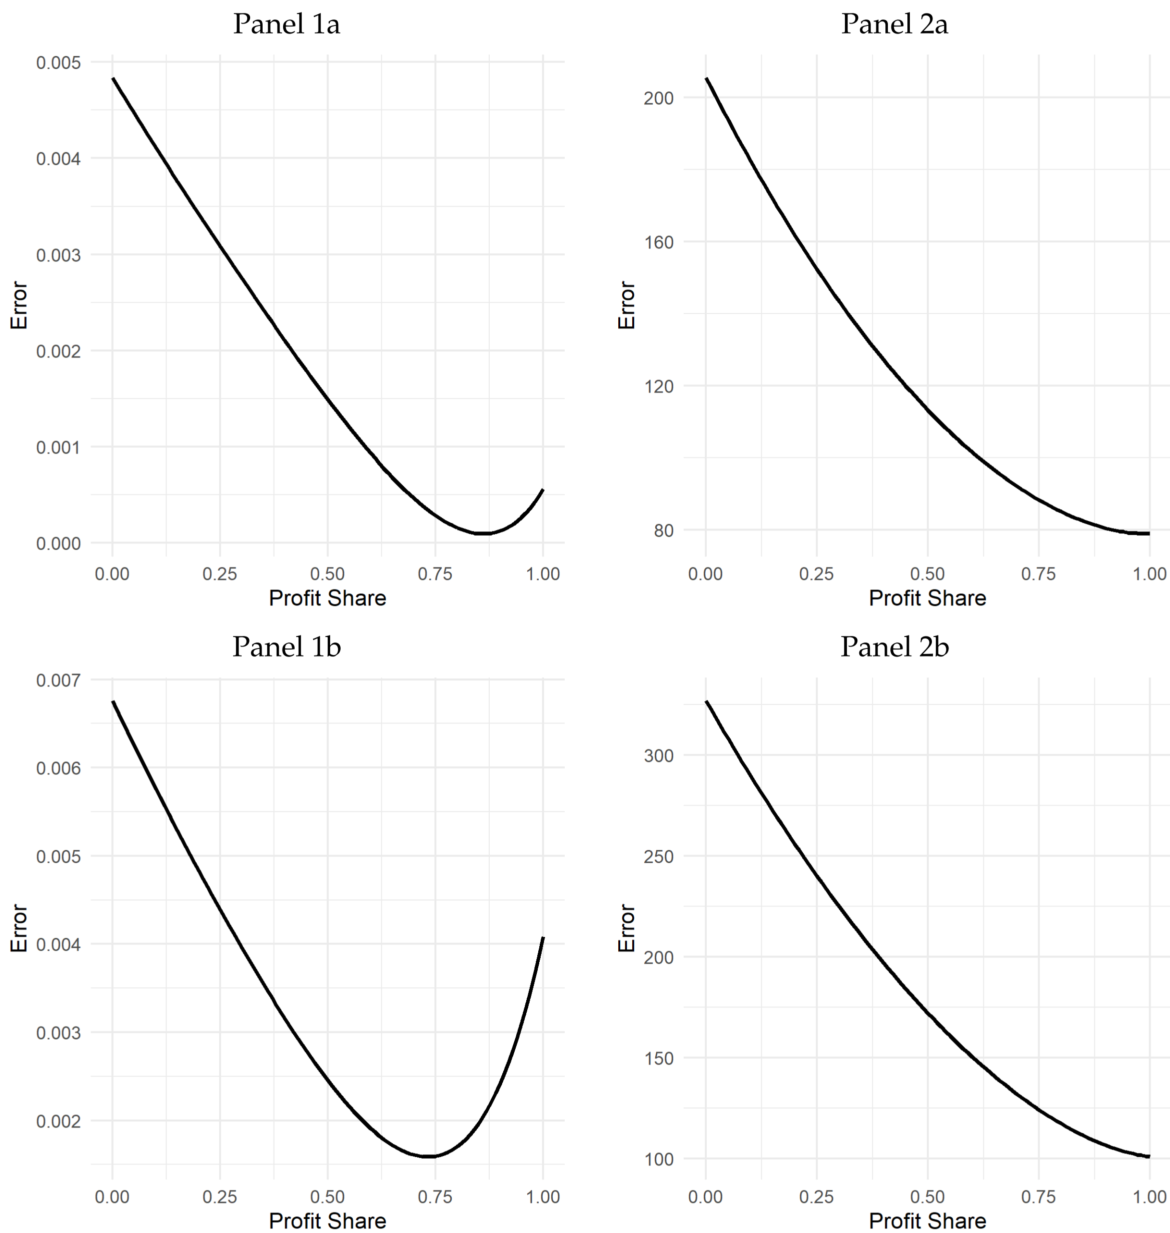

Note: Loss functions for GMM estimation of α. Panels a display the step 1 and 2 loss functions using all oil variation, while panels b display the step 1 and 2 loss functions using Kanzig variation.

In our structural work, we noted that different assumptions about the profit share of payments to capital meaningfully affected overall pass-through of oil price movements as predicted by the Calvo New Keynesian model. We also noted that our reduced-form results were much less influence by assumptions about this parameter. We hypothesized that, had we considered higher-order pass-through of oil price shocks in reduced form, including payments to capital in industry costs would have yielded long-run pass-through estimates larger than one. In Figure 24, we confirm this hypothesis. When we include payments to capital in industries' costs, third- and higher-order pass-through of oil price movements is around 1.6, and fourth- and higher-order pass-through of oil price movements is around 2. This means that industries 3 and higher or 4 and higher links away from oil in the supply chain increase prices 1.6 and 2 times more, respectively, than predicted by the flexible pricing model. Further, there is evidence of a continued upward trajectory of these industries' prices, as predicted by the Calvo New Keynesian model.

Figure 24 – Higher-order Pass-through in Reduced-form

Panel 1: Third-order Pass-through

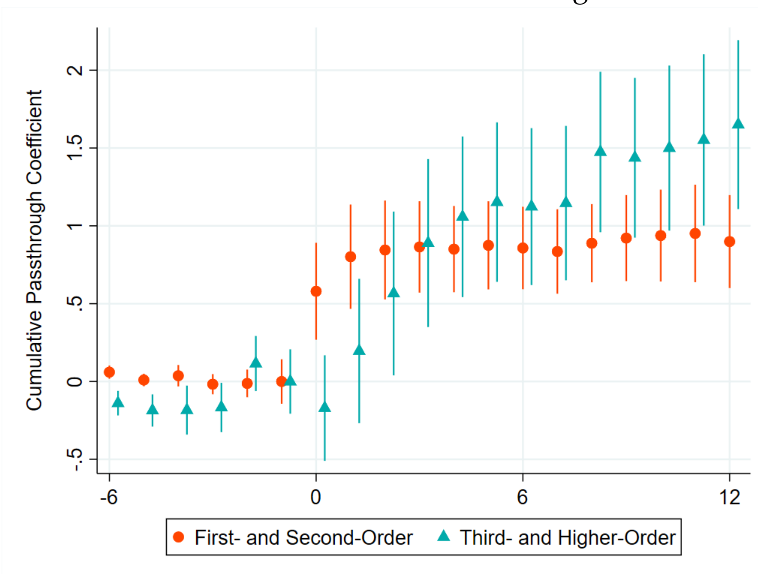

Panel 2: Fourth-order Pass-through

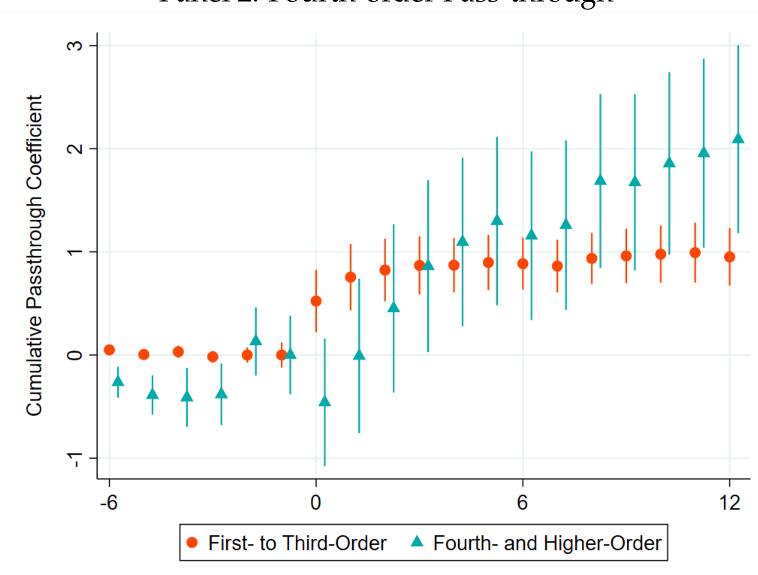

D.2 Joint Estimation of the Profit Share and Myopia for Oil

Step 1 of the GMM estimation yields the loss function displayed in Panel 1a of Figure 25. Step 2, using the estimate of the optimal weight matrix, changes the loss function to that displayed in Panel 2a. The analogous figures using the Kanzig IV variation are shown in Panels b. We see that the estimates of α and m_f are broadly stable across steps of the estimation process, suggesting re-weighting of noisily estimated components of the IRFs in the data is not substantially affecting estimation.

Figure 25 – Loss Functions for Joint Estimation of α and m_f (Oil)

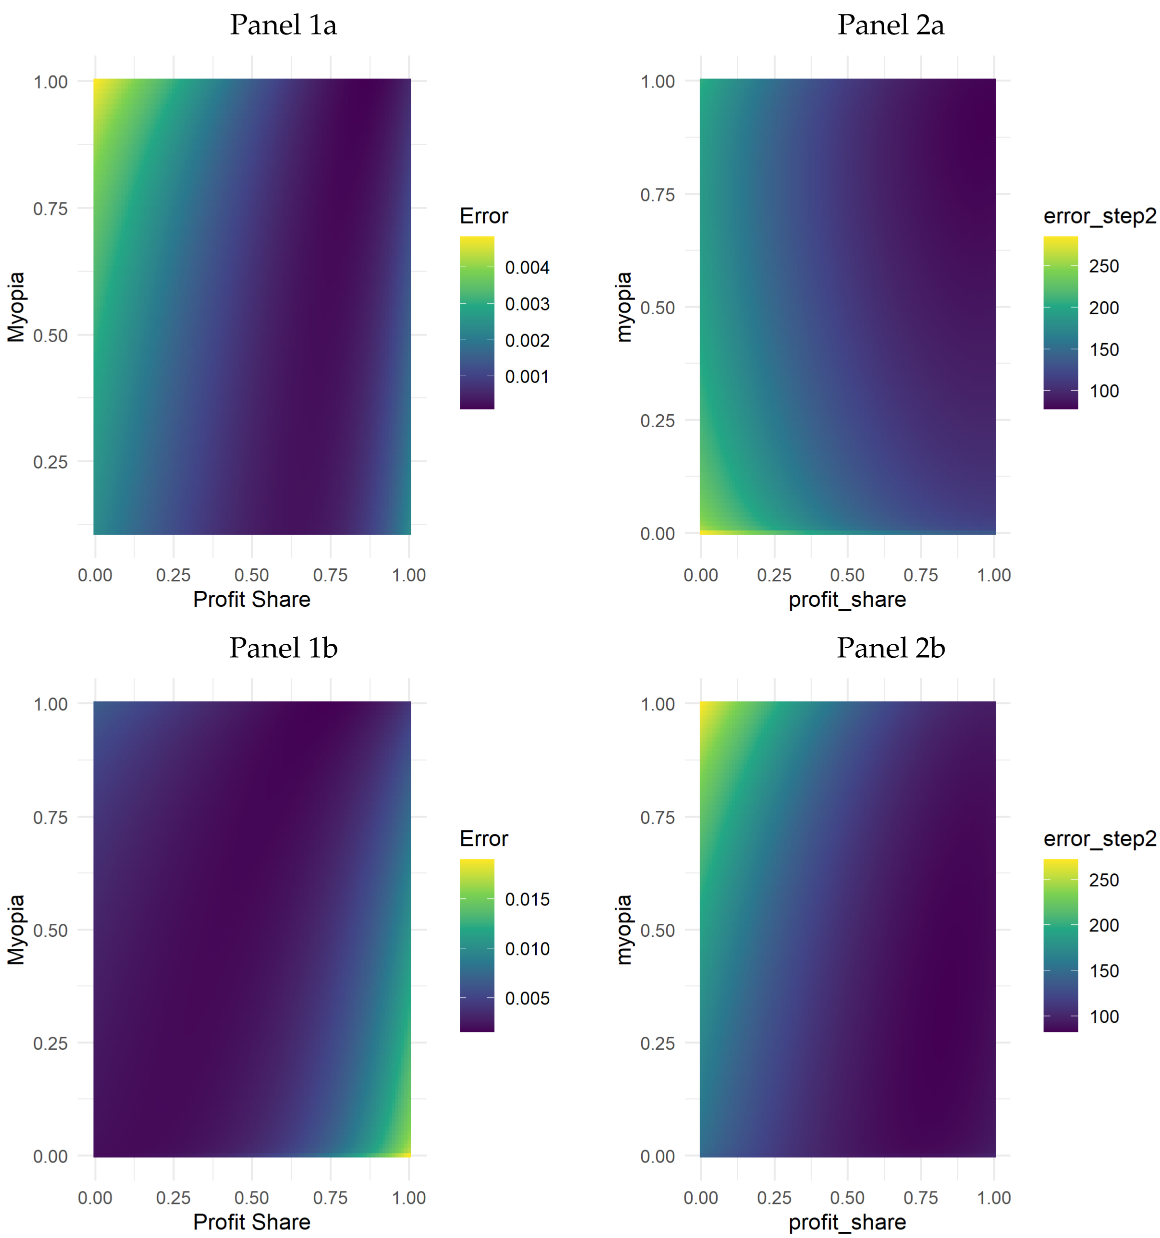

Note: Loss functions for joint GMM estimation of α and m_f. Panels a display the step 1 and 2 loss functions using all oil variation, while panels b display the step 1 and 2 loss functions using Kanzig variation.

Figure 26 – Kanzig-induced Variation in the Oil Price

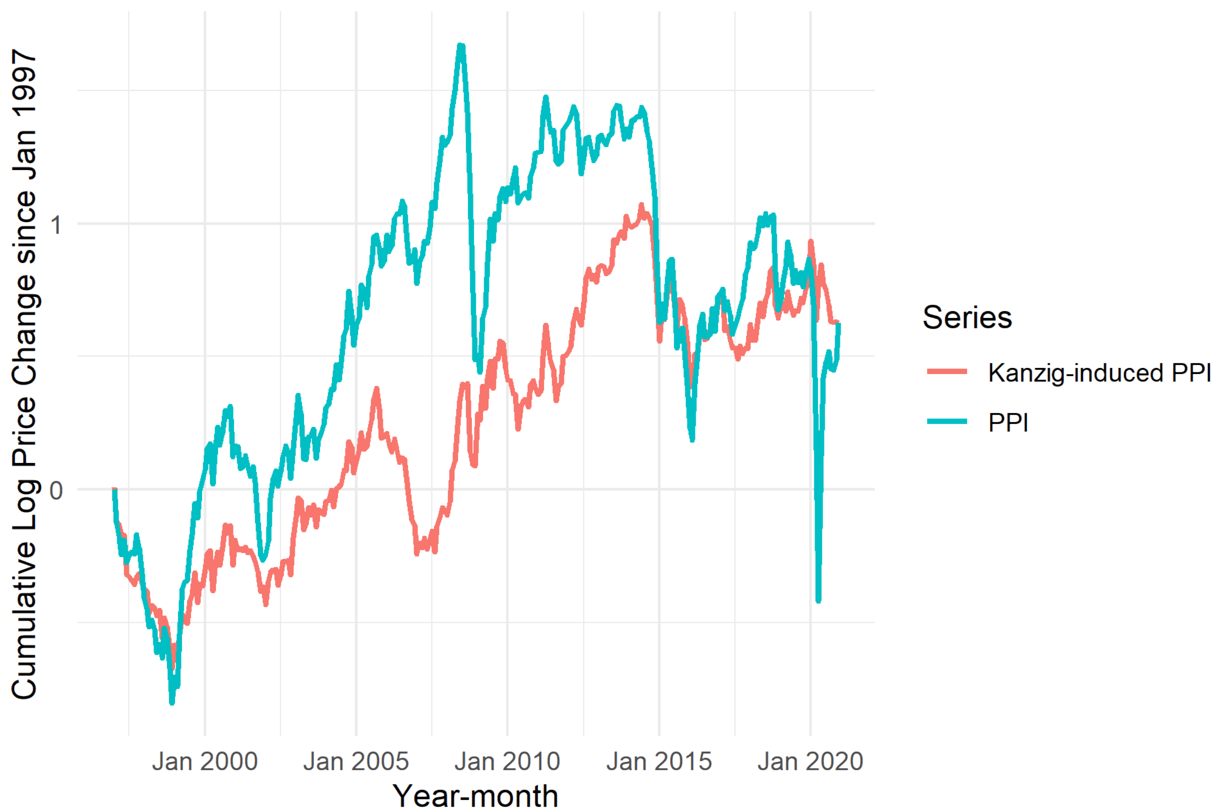

Note: The cumulative change in the log WTI oil spot price is plotted against the fitted values of a regression of the monthly log change in the WTI spot price on the Kanzig shocks.

We also assess whether Kanzig-induced changes in the oil price are smaller than changes in the oil price generally. We confirm that this is the case in Figure 26. Outside of the shale boom, large movements in oil prices induced using the Kanzig variation occur during small oil price change episodes, confirming that myopia estimated using Kanzig variation should be more in line with myopia estimated using oil price movements during small oil price change episodes.

D.3 Joint Estimation of the Profit Share and Myopia for Commodities

Step 1 of the GMM estimation yields the loss function displayed in Panel 1 of Figure 27. Step 2, using the estimate of the optimal weight matrix, changes the loss function to that displayed in Panel 2. We see that the estimates of α and m_f are not broadly stable across steps of the estimation process, suggesting re-weighting of noisily estimated components of the IRFs in the data is substantially affecting estimation.

Figure 27 – Loss Functions for Joint Estimation of α and m_f (Non-oil Commodities)

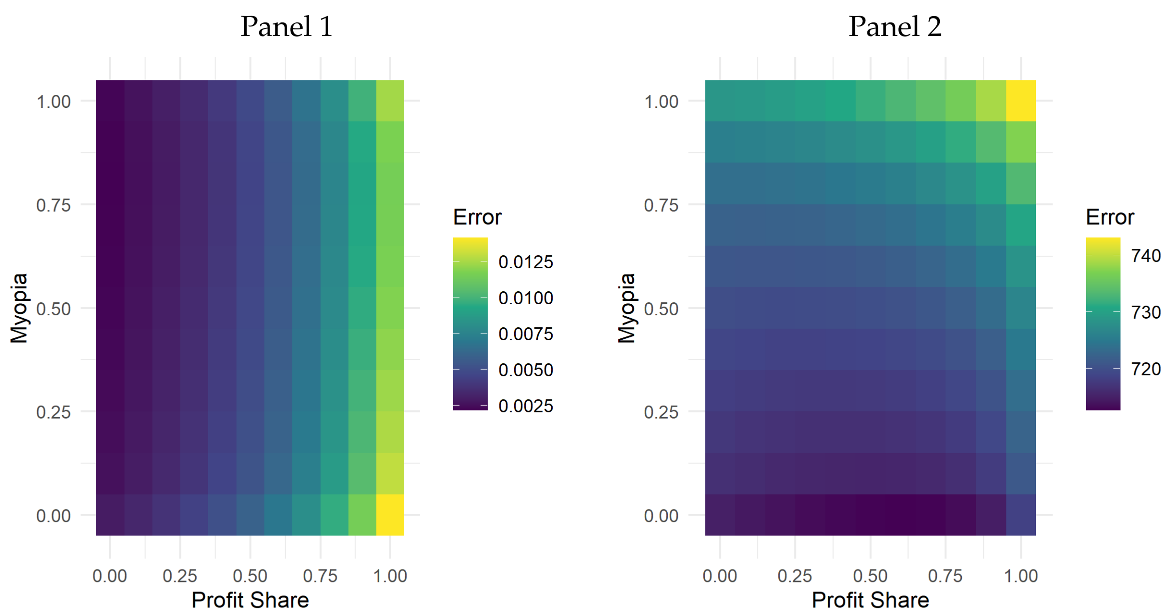

Note: Loss functions for joint GMM estimation of α and m_f for non-oil commodities, with step 1 and step 2 loss functions displayed in their respective panels.

D.4 Estimation of Myopia for Large and Small Oil Shocks

Step 1 of the GMM estimation yields the loss function displayed in Panel 1 of Figure 28. Step 2, using the estimate of the optimal weight matrix, changes the loss function to that displayed in Panel 2. We see that the estimates of m_f are broadly stable across steps of the estimation process, suggesting re-weighting of noisily estimated components of the IRFs in the data is not substantially affecting estimation.

Figure 28 – Loss Functions for Estimation of m_f (Large Versus Small Oil Shocks)

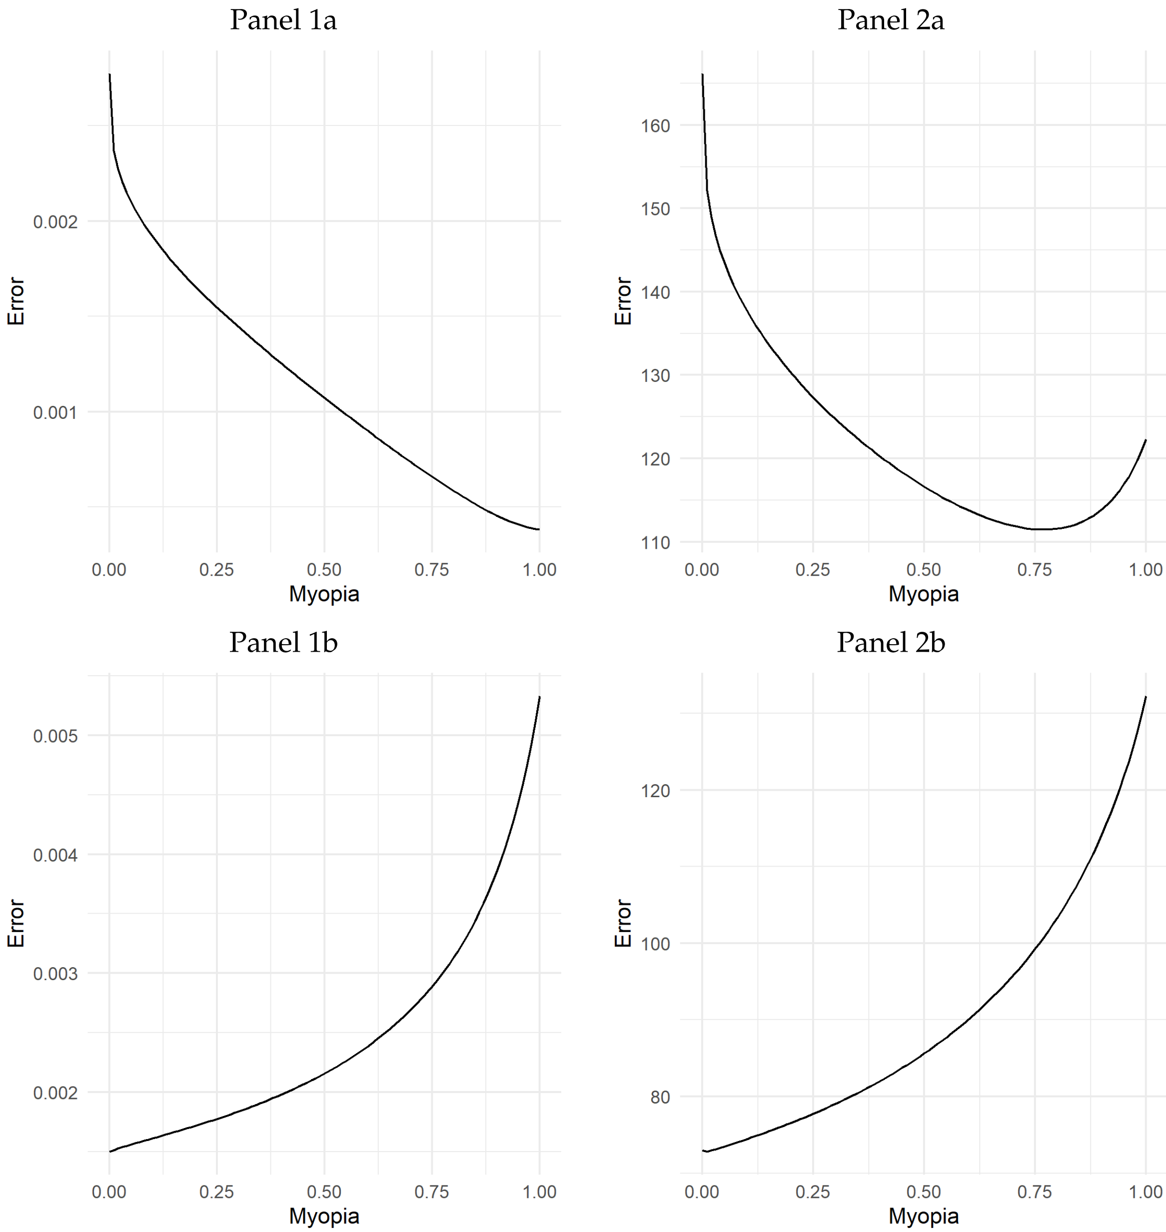

Note: Loss functions for GMM estimation of m_f for large relative to small oil shocks. Panels a display 1 and 2 step loss functions for large oil shocks, while panels b display loss functions for small oil shocks.