

Hidden Inflation in Supply Chains: Theory and Evidence*

Robert Minton[†] (Job Market Paper)

Brian Wheaton[‡]

November 14, 2022

Abstract

How much and how quickly do supply chains transmit commodity price movements into inflation? We study this question theoretically and empirically, demonstrating that commodity price movements propagate gradually to prices throughout the economy and generate lasting effects on inflation. First, we set up a dynamic supply-side pricing model featuring a production network and Calvo pricing. We show theoretically that commodity price movements pass through fully but much more gradually to firms using the commodity indirectly through their supply chains, even if they are forward-looking about how their suppliers' prices will react in the future. Second, we confirm these theoretical results empirically and show robustness to a variety of identification strategies. Third, we structurally estimate the model, explaining these findings as a combination of heterogeneity in firms' frequencies of price adjustment and slow propagation to indirect users resulting from the presence of a production network. We find that indirect users are less forward-looking about input cost changes resulting from non-oil commodity price movements, slowing down pass-through relative to oil price movements. Finally, we demonstrate how practitioners can use our model to measure the component of aggregate inflation induced by supply chain effects of oil price movements. These effects take years to manifest fully (1 year to reach 75%), are twice as large as the direct effect on consumer gasoline prices, and account for 15% of the monthly variation in core PCE inflation. Removing all of oil's effects in the current inflationary episode reduces measured year-over-year PCE inflation to a peak of 5%.

*We are especially grateful for invaluable feedback and advice from Robbie's advisors: Xavier Gabaix, Ed Glaeser, Andrei Shleifer, and Ludwig Straub. We are also grateful to Maxim Alekseev, Isaiah Andrews, Pol Antras, Robert Barro, Adrien Bilal, Raj Chetty, Gabe Chodorow-Reich, Jason Furman, Jordi Galí, Elhanan Helpman, Larry Katz, Greg Mankiw, Marc Melitz, Jeff Miron, Ariel Pakes, Giacomo Ponzetto, Ken Rogoff, Emanuel Schertz, Roman Sigalov, Jim Stock, Elie Tamer, 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

How much and how quickly do supply chains transmit commodity price movements into inflation? This question lies at the heart of many debates about the causes of inflation, both now and in the 1970s/80s, because commodity price increases may have lasting effects on aggregate inflation if they take time to propagate through supply chains. It is relatively straightforward to study how oil price movements, for example, affect prices charged for petroleum refinery output, but oil is used throughout the economy, both as a raw material and as an energy input. We build a theoretical and empirical framework to assess how quickly and fully commodity price movements propagate through supply chains, demonstrating their dynamic and predictable effects on aggregate inflation.

First, we study supply chain propagation theoretically, using a model that we can take directly to the data. In a dynamic supply-side pricing model with Calvo pricing and a production network, we show that the extent and pace of pass-through is influenced by a number of factors. How much a price movement in sector j affects prices in sector i depends on the overall share of cost in sector i that ultimately lies in sector j , inclusive of supply chain connections. The speed of pass-through depends on sector i 's "downstreamness" from sector j , a notion that captures how many links in the supply chain sector i is from sector j (on average, as many supply chain links between sectors may exist). Finally, forward-lookingness plays a key role in speeding up pass-through. With rational expectations (and full attention), downstream sectors adjust their prices to account for anticipated future price changes by their suppliers – that is, before the actual cost changes have worked their way down the supply chain to reach the downstream sectors. If firms are myopic – not forward-looking – they wait until the cost shocks have filtered down to them step-by-step before adjusting their prices, leading to a much higher degree of effective price rigidity.

We then test our propositions in the data. In 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 unrefined oil – i.e., first-order exposure to unrefined oil. The Synthetic Rubber Manufacturing industry has approximately 20% of its total costs in oil, but nearly none of it constitutes direct purchase of unrefined oil. Most of this originates from purchases of refined oil from petroleum refineries (i.e., second-order exposure to unrefined oil), purchases of petrochemical products from the Petrochemical Manufacturing industry, which itself purchased most inputs from refineries (i.e., third-order exposure to unrefined oil), and beyond.

Using these measures of exposure, we first pool all variation in oil prices since 1997, and we

show that long-run pass-through from both direct and indirect exposure to oil price movements is strongly significant and statistically indistinguishable from 100%. Pass-through from direct exposure occurs mostly on impact, whereas pass-through from indirect exposure occurs over a slower period of 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 surprise effects of OPEC announcements. Results are also analogous when we investigate the effects of other commodity shocks beyond oil, though pass-through is even more gradual than for 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 our model’s 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.

We then turn to a structural estimation using our dynamic pricing model. We implement a generalized method of moments (GMM) procedure to determine what features of the model are required to fit the data. We find that (1) substantial heterogeneity in industries’ frequencies of price adjustment, (2) a markup over marginal cost, and (3) context-dependent forward-lookingness are required to fit the pass-through patterns observed in the data. In particular, industries tend to behave in line with rational expectations during oil price shocks while behaving myopically during shocks to other commodities. We also find clear evidence of rational expectations in response to oil price movements that is 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 inflation

purged of oil's influence.

Using our GMM-calibrated model, we show that the supply chain effects of oil price movements cause about twice as much inflation as the direct effects of oil on the gasoline prices consumers pay – but these supply chain effects take years to manifest fully in inflation. Purging all of oil's effects from aggregate inflation reduces measured year-over-year inflation during the current episode to 5%. More generally, inclusive of supply chain effects, oil price movements explain 34% of the monthly variation in personal consumption expenditures (PCE) inflation and 16% of the variation in Core PCE inflation. We produce a historical inflation series beginning in 1960 that removes the direct and indirect effects of oil price movements.

The remainder of this paper is organized as follows. Section 2 reviews the relevant literature on production networks. Section 3 outlines the setup and results of our dynamic pricing model. In Section 4, we discuss data used in our estimation. In Section 5, we present the results of our empirical analysis evaluating the extent and pace of commodity price pass-through, including a variety of robustness checks and heterogeneity analysis. In Section 6, we reconcile our empirical results with the model by structurally estimating the degree of forward-lookingness, the aggregate markup, and the importance of heterogeneous frequencies of price adjustment. In Section 7, we focus on our application: stripping the full network effects of oil price movements from core inflation. We also discuss a more simple “strawman” model without a network and how it fails to match the data. In Section 8, we conclude by summarizing our findings and discussing some of their implications.

2 Literature Review

Our paper contributes to four strands of existing literature. First, we contribute to the literature finding that sectoral shocks can generate meaningful aggregate fluctuations. We show this in the context of oil price movements, which generate indirect inflationary effects (through the supply chain) twice as large as their direct effect. Second, 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. Third, 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, as we develop a way to estimate firm forward-lookingness from observational macroeconomic data.

We build on a large literature finding that sectoral shocks can generate meaningful aggregate fluctuations. Horvath (1998) presents a model in which 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), Gabaix (2011), Acemoglu et al. (2012), and Baqaee and Farhi (2019) present additional modeling evidence of the importance of accounting explicitly for the network structure of the economy. 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 can be responsible for amplification. Acemoglu, Ozdaglar, and Tahbaz-Salehi (2017) further document the importance of industry heterogeneity.

We also build on the recent but growing literature performing empirical tests of production network models. 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 causal tests of production network models with techniques standard in the modern applied-econometrics literature.

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.¹

Two papers with findings related to some of our empirical results 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

¹Their 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.

of PPI inflation across countries. Smets et al. empirically analyze the network price pass-through patterns of estimated micro-level shocks in a structural Bayesian framework, finding in a horse-race that the data prefer a model with network propagation to a model without such propagation.

We also contribute to the literature 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.² Theoretical results from Afrouzi and Bhattarai (2022) characterize pass-through speed in production networks more generally and are closely related to our findings in Proposition 2. In addition to our theoretical results linking pass-through speed of a shock to an industry’s downstreamness from the shock (in Proposition 3), we also empirically test and confirm the hypothesis that nominal rigidities compound in supply chains.

It is important to verify whether nominal rigidities compound in supply chains because this conjecture increasingly underlies policy advice. Rubbo (2020) and La’O and Tahbaz-Salehi (2022), for instance, 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). Our contribution is to show that the network compounding of nominal rigidities actually does occur in the data.

Finally, we contribute to the literature on how deviations from rational expectations can affect macroeconomic outcomes. As illustrated by Carlstrom, 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

²Basu 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.

data on the macroeconomy (as opposed to microeconomic experiments) and, in so doing, provide empirical evidence of how deviations from rational expectations are present in the data.

Other deviations from a fully rational, attentive, and informed benchmark model have also been studied in the literature. Sims (2003), Woodford (2003), Mankiw and Reis (2002), and Gabaix (2019) discuss inattention and information rigidities. Bordalo, Gennaioli, and Shleifer (2018, 2022) explore an alternative model of expectations, “diagnostic expectations.” Angeletos and Huo (2021) develop an equivalence between informational frictions and backward-looking, myopic expectations. We refer the interested reader to these papers for a more in-depth review of this extensive literature.

3 Model

Our modeling seeks to guide all the empirical analysis we perform in our study of price propagation. We first develop a proposition characterizing the long-run pass-through of sectoral price changes to all industry prices. These long-run pass-through measures capture the fact that, even if industries do not directly purchase inputs from a sector for use in production – if, that is, they are instead exposed to the sector’s price movements only indirectly through their suppliers’ use (or their suppliers’ suppliers’ use, etc.) – they should still change prices in response to the input price change. We further develop propositions about 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. Our model most closely follows the supply-side 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.

Finally, we use a simple, linear network to develop intuition about the role of myopia and compounding price rigidities 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. Our simple calibration of the linear network model allows for convenient graphical illustrations of propositions for the general network.

3.1 Model Setup

The setup is standard Dixit-Stiglitz at the industry level. There is a continuum of firms $j \in [0, 1]$ in each industry $i \in \{1, \dots, I\}$. Each firm produces output denoted $Y_{i,j,t}$ and these varieties are

transformed³ into an industry aggregate according to

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

The industry's price index is then

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

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 bundled varieties from all modeled industries, denoted $\mathbf{X}_{i,j,t} = (X_{i,j,t}^1, \dots, X_{i,j,t}^I)$ and labor $L_{i,j,t}$.⁴ 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. Formally, then, $Y_{i,j,t} = A_{i,t}F_i(L_{i,j,t}, \mathbf{X}_{i,j,t}, Z_{i,j,t})$. We do not allow the wage to vary by firm within sector, which can be microfounded with perfect substitutability in labor supply across firms within sector.

Firms minimize input costs subject to producing a given level of output, yielding the cost function $C_i(W_{i,t}, \mathbf{P}_t, P_{Z,t}, Y_{i,j,t}/A_{i,t})$. Therefore, marginal cost is

$$MC_{i,t} = \frac{1}{A_{i,t}} C_i(W_{i,t}, \mathbf{P}_t, P_{Z,t}, 1). \quad (1)$$

Marginal cost does not vary across firms within industry and depends on the industry-specific wage, $W_{i,t}$, the vector of modeled industry prices $\mathbf{P}_t = (P_{1,t}, \dots, P_{I,t})$, the commodity price $P_{Z,t}$, and the level of TFP, $A_{i,t}$.

Knowing marginal cost, we can consider the firm's optimal pricing problem. The setup is standard, following the textbook treatment in Galí (2010) at the industry level, but allows for deviations from rational expectations as in Gabaix (2020). Each firm j in industry 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. The elasticity of substitution between varieties in each industry, σ_i , is constrained to be greater than 1 for a well-defined monopoly profit maximization problem. Denote the stochastic discount factor, the relevant discount rate for firms, between periods t and $t + k$ by $SDF_{t,t+k}$.⁵ The optimal reset

³They can be transformed by a competitive industry dedicated to this task, or by each firm separately whenever it produces using the industry bundle.

⁴Labor can be thought of as a value-added commodity that jointly includes labor and capital.

⁵One example SDF is $SDF_{t,t+k} = \delta^k \frac{U'(C_{t+k})}{U'(C_t)}$, where C is aggregate consumption, U is the utility function, and δ is the consumer's discount factor.

price will not vary across firms within an industry, and it solves

$$\max_{P_{i,j,t}^*} \sum_{k=0}^{\infty} \theta_i^k \tilde{E}_t [SDF_{t,t+k} Y_{i,j,t+k} (P_{i,j,t}^* - MC_{i,t+k})] \quad (2)$$

$$\text{s.t. } Y_{i,j,t+k} = Y_{i,t+k} \left(\frac{P_{i,t+k}}{P_{i,j,t}^*} \right)^{\sigma_i} . \quad (3)$$

Now, \tilde{E}_t is the potentially myopic expectations operator given the period t information set and will be defined shortly in terms of log-linearized variables. 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

$$P_{i,t} = \left(\theta_i P_{i,t-1}^{1-\sigma_i} + (1 - \theta_i) (P_{i,t}^*)^{1-\sigma_i} \right)^{\frac{1}{1-\sigma_i}} . \quad (4)$$

Our supply-side equilibrium model can be summarized in three equations:

Definition 1 (Industry Equilibrium). *The law of motion for industry prices follows (4), firms with the opportunity to change prices solve the maximization problem (2) subject to (3), and marginal cost is defined by (1).⁶*

Note that, at this point, we have not specified the stochastic discount factor, the labor supply conditions that determine industry wages in equilibrium, the monetary rule, etc. Our results that follow must hold regardless of these specifications.

3.2 Log-linearized Model

We log-linearize our industry equilibrium system around a zero-growth and zero-inflation steady state. We will denote the log deviation of a variable from its log steady state value 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 $\hat{p}_{i,t}$.

We now define the (potentially) myopic expectation of the deviation of a random variable from steady state. When taking the 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

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

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. When there is myopia, firms discount expected future

⁶The law of motion is required in the definition of industry equilibrium because, even though all firms are optimizing in every period, only some share of firms has the opportunity to change prices in each period. The industry price index evolves according to the optimal reset price, weighted by the share of firms that can change prices, and the previous period's price, weighted by the share of firms that cannot change prices.

disturbances more relative to the rational agent. In the case of our model, this will mean that firms may neglect drift in their marginal cost induced by gradual pass-through of shocks through the production network. Beyond being of independent interest, the parameter m_f will enable us to test whether rational expectations is operative in the data.

Now, we develop our log-linearized three equation model. To do so, we require some additional notation. Define industry cost shares in each input (in steady state) as

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

The matrix Φ is commonly called the input-output matrix; it captures how much each industry spends on inputs from every other industry as a fraction of cost, conveniently summarizing the complex input-output linkages of the economy. Now, it is a straightforward application of the envelope theorem on the cost function that

$$\widehat{mc}_{i,t} = -\hat{a}_{i,t} + s_i^L \hat{w}_{i,t} + s_i^Z \hat{p}_{Z,t} + \sum_{k=1}^I \Phi_{i,k} \hat{p}_{k,t}. \quad (5)$$

The log-linearizations of the law of motion for industry prices and the first-order condition determining the optimal reset price are well-known in the New Keynesian theory. With the addition of myopia, these equations log-linearize to

$$\hat{p}_{i,t} = \theta_i \hat{p}_{i,t-1} + (1 - \theta_i) \hat{p}_{i,t}^* \quad (6)$$

$$\hat{p}_{i,t}^* = (1 - \theta_i \delta m_f) \sum_{k=0}^{\infty} (\theta_i \delta m_f)^k \mathbb{E}_t \widehat{mc}_{i,t+k}. \quad (7)$$

In the case where $m_f = 1$, equations (6) and (7) yield the standard Phillips curve (in nominal marginal cost), just at the industry level. The network enters through equation (5). An industry's price depends on its marginal cost, which in turn depends (potentially) on the vector of all industries' prices.

Note that a simplification of equation (7) is possible under complete myopia, $m_f = 0$. In this case, $\hat{p}_{i,t}^* = \widehat{mc}_{i,t}$. As anticipated earlier, under complete myopia, firms will wait until their marginal costs have adjusted to alter their optimal reset prices. With forward-lookingness, firms will anticipate future changes in marginal cost, passing them through at least partially to their optimal reset prices in the present.

3.3 Long-run Impact of a Commodity Price Shock

We now characterize the extent and pace of pass-through of the commodity shock to all industry prices. We begin by characterizing the extent of pass-through. We will use the notation $diag(\cdot)$ to

denote the diagonal matrix with its argument ordered across industries on the diagonal.

Proposition 1 (Long-run Impact). *Suppose there is a one-time, persistent, and unexpected shock to the commodity price at time $t = 0$, $\hat{p}_{Z,0}$. The economy was in steady state, and there are no shocks to TFP. To a first-order approximation, if all industries can change prices at some point, $\theta_i < 1$ for all i , the long-run equilibrium response of prices (holding wages constant) is*

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

If wages adjust in general equilibrium,

$$\hat{p}_\infty = (\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z \hat{p}_{Z,0} + (\mathbf{I} - \Phi)^{-1} \text{diag}(s_i^L) \hat{w}_\infty.$$

It is of course also true that, under flexible pricing ($\theta_i = 0$ for all i), the immediate partial equilibrium response of all prices to a commodity price movement is the long-run partial equilibrium response. As far as we know, Proposition 1 has not been stated this way elsewhere in the literature, but it is directly related to many existing results, particularly those from Baqaee and Farhi (2017b) Propositions 2 and 3, Baqaee and Farhi (2019) Theorem 3, and, if we were to aggregate appropriately across industries, Hulten's theorem.

We can test Proposition 1 directly in the data. In particular, we can compute $(\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z$ in the data, and so we know, to a first-order, how each industry's price will respond to a persistent and unexpected commodity price movement in the long-run. We develop some intuition using the geometric sum formula for matrices. We can decompose industries' exposures to the commodity price movement:

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

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

$$\text{NetworkExposure}_{i,k} = \left[\Phi^{k-1} \mathbf{s}^Z \right]_i.$$

It is clear that $\text{NetworkExposure}_{i,1}$ 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 input costs are comprised of purchases of the commodity, then $\text{NetworkExposure}_{i,1} = .1$. Through simple matrix multiplication, it also follows that $\text{NetworkExposure}_{i,2}$ is an industry's *second-order* 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, $\text{NetworkExposure}_{i,2} = .2 \times .8 = .16$. Analogously, $\text{NetworkExposure}_{i,3}$ is an industry's *third-order* exposure to the commodity through its suppliers' suppliers' costs, etc.

A key prediction of the model is that, to a first order in partial equilibrium, the full network exposure is a sufficient statistic for long-run pass-through. Put differently, the distribution of network exposures at various orders is irrelevant for long-run pass-through, conditional on the full network exposure.

3.4 Speed of Adjustment and “Downstreamness” Measures

Next, we seek to develop empirical measures of the speed of pass-through in an industry. Remarkably, we will find that the speed of pass-through is quite intuitive for a special case; in this special case, we can formalize how pass-through depends on the distribution of an industry’s orders of network exposure to the shock. It is helpful to begin first with a characterization of how industry prices evolve over time in response to a commodity price movement. We assume that the stochastic discount factor between periods t and $t + k$ in steady state is δ^k , which follows from a large class of household setups; δ can equivalently be given in terms of the steady state interest rate. Define the continuous time parameters $\phi_i = -\ln \theta_i$ and $\psi_i = \phi_i - \ln \delta - \ln m_f$. We call ϕ_i the rate of price adjustment and ψ_i the discount rate. Note that, under complete myopia, now we have $\psi \rightarrow \infty$, while under rational expectations and a steady state interest rate of 0, we have $\psi_i = \phi_i$.

Proposition 2 (Transitory Dynamics and Speed of Adjustment). *Suppose there is a one-time, persistent, and unexpected shock to the commodity price at time $t = 0$, $\hat{p}_{Z,0}$. The economy was in steady state, and there are no shocks to TFP. Then the response of prices (holding wages constant) is*

$$\hat{p}_t \approx (\mathbf{I} - e^{-\mathbf{B}t}) \hat{p}_\infty,$$

(with equality in continuous time) where $\hat{p}_\infty = (\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z \hat{p}_{Z,0}$. Defining $\bar{\phi} = \text{diag}(\phi_i)$ and $\bar{\psi} = \text{diag}(\psi_i)$, \mathbf{B} satisfies

$$(\bar{\phi}^{-1} \mathbf{B} - (\mathbf{I} - \bar{\psi} \bar{\phi}^{-1})) \mathbf{B} = \bar{\psi} (\mathbf{I} - \Phi),$$

and we take the roots with all positive eigenvalues.⁷ In the myopic case, $\mathbf{B} = \bar{\phi} (\mathbf{I} - \Phi)$. Under rational expectations, when the interest rate r is 0, we have $\mathbf{B} = (\bar{\phi}^2 (\mathbf{I} - \Phi))^{1/2}$, and more generally, $\mathbf{B} = (\bar{\phi}^2 (\mathbf{I} - \Phi))^{1/2} + O(r)$.

Proposition 2 yields a general method for determining the time path of pass-through of price movements to all industry prices, given ϕ , ψ , Φ , and \mathbf{s}^Z . In two special cases – complete myopia

⁷If $\phi_i = \phi$ for all i , then $\psi_i = \psi$ for all i , \mathbf{B} satisfies the equation

$$(\phi \mathbf{I} - \mathbf{B})(\mathbf{B} + \psi \mathbf{I}) = \phi \psi \Phi,$$

and we take the roots with all positive eigenvalues:

$$\mathbf{B} = \frac{\psi - \phi}{2} \mathbf{I} + \left(\left(\frac{\psi + \phi}{2} \right)^2 \mathbf{I} - \phi \psi \Phi \right)^{1/2}.$$

and rational expectations with a steady state interest rate of 0 – we actually know the functional form of the solution. Proposition 2, under rational expectations, is closely related to theoretical results in parallel work by Afrouzi and Bhattacharai (2022).

Now, we turn to further characterizations of the speed of pass-through in a sector. Denote by e^i the standard basis vector in \mathbb{R}^I (a column vector of 0's with a 1 in element i). Define the fraction of long-run pass-through in sector i at time t by $b_{i,t}$, which satisfies

$$\hat{p}_{i,t} = b_{i,t} \hat{p}_{i,\infty},$$

where $\hat{p}_{i,\infty} = [(\mathbf{I} - \mathbf{\Phi})^{-1} \mathbf{s}^Z]_i \hat{p}_{Z,0}$. It follows from Proposition 2 that

$$b_{i,t} \approx 1 - \frac{(e^i)' e^{-\mathbf{B}t} \hat{p}_\infty}{(e^i)' \hat{p}_\infty},$$

with equality in continuous time. In industry i at time t , the fraction of long-run pass-through not yet achieved is $1 - b_{i,t}$. We can measure the average time it takes to pass-through a commodity price increase in industry i as follows:

Definition 2 (Duration of Pass-through). *The duration of pass-through in sector i is*

$$D_i = \int_0^\infty (1 - b_{i,s}) ds.$$

Clearly, if $b_{i,t} = 1 - e^{-\phi_i t}$, the pass-through rate without a network, then $D_i = 1/\phi_i$, which is declining in the rate of price adjustment ϕ_i . We can characterize $D(i)$ generally and produce a specific functional form in special cases, applying Proposition 2:

Proposition 3 (Duration of Pass-through is “Downstreamness”). *The duration of pass-through in industry i is*

$$D_i \approx \frac{(e^i)' \mathbf{B}^{-1} \hat{p}_\infty}{(e^i)' \hat{p}_\infty},$$

with equality in continuous time. Suppose the frequency of price adjustment does not vary across industries. Then, under complete myopia, where $\mathbf{B} = \phi(\mathbf{I} - \mathbf{\Phi})$,

$$D_i \approx \frac{1}{\phi} \frac{(e^i)' (\mathbf{I} - \mathbf{\Phi})^{-2} \mathbf{s}^Z}{(e^i)' (\mathbf{I} - \mathbf{\Phi})^{-1} \mathbf{s}^Z} = \frac{1}{\phi} \frac{(e^i)' \sum_{n=1}^\infty n \mathbf{\Phi}^{n-1} \mathbf{s}^Z}{(e^i)' \sum_{n=1}^\infty \mathbf{\Phi}^{n-1} \mathbf{s}^Z},$$

and, under rational expectations and a steady state interest rate of 0, where $\mathbf{B} = \phi(\mathbf{I} - \mathbf{\Phi})^{1/2}$,

$$D_i \approx \frac{1}{\phi} \frac{(e^i)' (\mathbf{I} - \mathbf{\Phi})^{-3/2} \mathbf{s}^Z}{(e^i)' (\mathbf{I} - \mathbf{\Phi})^{-1} \mathbf{s}^Z} = \frac{1}{\phi} \frac{(e^i)' \sum_{n=1}^\infty \binom{-3/2}{n-1} \mathbf{\Phi}^{n-1} \mathbf{s}^Z}{(e^i)' \sum_{n=1}^\infty \mathbf{\Phi}^{n-1} \mathbf{s}^Z}.^8$$

⁸Unfortunately, the coefficient $\binom{-3/2}{n-1}$ does not admit a clean functional form. For values $n \in \{1, 2, \dots, 5\}$, it is

Both statements hold with equality in continuous time.

Under complete myopia, ϕD_i is an intuitive measure of downstreamness; it is the weighted average of an industry's orders of exposure to the commodity. It is 1 if an industry only uses the commodity directly, 2 if an industry only uses the commodity through its suppliers, 3 if an industry only uses the commodity through its suppliers' suppliers, etc. But the measure allows for complex linkages between the industry and the commodity. Interestingly, this is the measure of downstreamness defined in Antras and Chor (2021). Our result directly links this measure to pass-through duration in a network when firms are completely myopic.

Under rational expectations, downstreamness is similar, but more downstream industries pass through the shock faster. This manifests in orders of exposure larger than 1 being weighted by smaller terms than they are under complete myopia. Formally, under rational expectations, $D_i \phi$ is 1 if an industry only uses the commodity directly, $3/2$ if the industry only uses the commodity through its suppliers, $15/8$ if an industry only uses the commodity through its suppliers' suppliers, etc., with coefficients generally given by a binomial coefficient $\binom{-3/2}{n-1}$ (where n is the order of exposure). Importantly, the network still slows pass-through, even under rational expectations.

One might ask whether our intuition for downstreamness remains meaningful if the frequency of price adjustment varies across sectors. Plugging in our solution from Proposition 2 under myopia (now for a vector of price adjustment rates ϕ), duration becomes

$$D_i \approx \frac{(e^i)'(\mathbf{I} - \Phi)^{-1} \bar{\phi}^{-1} (\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z}{(e^i)'(\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z},$$

with equality in continuous time. In general, $D_i \neq \frac{1}{\phi_i} \frac{(e^i)'(\mathbf{I} - \Phi)^{-2} \mathbf{s}^Z}{(e^i)'(\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z}$, though this was the case under homogeneous θ . Intuitively, two sectors with the same frequency of price adjustment may vary still in pass-through speed when one has more price-flexible suppliers than the other.

When moving to our empirical work, an important question will be whether more downstream industries pass through the commodity price increase more slowly simply because they have lower rates of price adjustment. To distinguish these concepts, we define a second notion of downstreamness. We must first develop a notion of how quickly a firm would pass through an input price increase in isolation, i.e. in an economy without a network. Intuitively, the pass-through rate should be $b_{i,t} = 1 - e^{-\phi_i t}$, where ϕ_i is the rate of price adjustment in the industry.⁹

Connecting to our previous result, this simplified model would have $\phi_i D_i = 1$, so that pass-

approximately 1, 1.5, 1.875, 2.188, and 2.461, respectively. For our purposes, it is noteworthy that these values are less than n for $n > 1$, so that duration is lower under rational expectations and a 0 interest rate than it is under myopia.

⁹This turns out to be the case in the hyper forward-looking limit of the network model. Denote by $b_{i,t}(m_f)$ the fraction of long-run pass-through in industry i at time t in the model where myopia is set to m_f . Recall that $\theta_i \delta m_f$ was the discount rate used by the firm, and so there is no discounting as $m_f \rightarrow 1/(\theta_i \delta)$ (or as $\psi \rightarrow 0$ in continuous time).

through speed is as if all industries only use the commodity directly (i.e., not through their suppliers at all). Define the time it takes industry i to reach a fraction X of long-run pass-through under myopia of degree m_f , denoted $t_i(X, m_f)$, implicitly by

$$b_{i,t_i(X,m_f)} = X.$$

It is clear that, in the simplified economy where $b_{i,t} = 1 - e^{-\phi_i t}$, we have $t_i(X; m_f) = -\ln(1-X)/\phi_i$, which does not depend on the degree of myopia.

Definition 3 (Excess Pass-through Time due to Downstreamness). *The excess time required to reach a fraction X of long-run pass-through due to the presence of the network is*

$$T_i(X, m_f) = t_i(X, m_f) + \ln(1 - X)/\phi_i.$$

This measure will allow us to perform joint regression analyses in the data where we test whether industries with the same rate of price adjustment exhibit slower pass-through if they are more downstream. Empirically, we will usually set $X = 0.5$ and $m_f = 1$, so that the interpretation of T_i is the half-life of pass-through due to downstreamness under rational expectations.

3.5 Illustration: Linear Network

In this section, we develop intuition on how expectations influence the dynamics of commodity shock pass-through in a linear network. A linear network of length N is comprised of industries $n \in \{1, \dots, N\}$, 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. An example linear network is shown in Figure 1.

We make the assumption that industry wages 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 . It follows that long-run pass-through is $\hat{p}_{n,\infty} = s^n$. Therefore, if $s = .5$, so that 50% of costs are in intermediate inputs, $\hat{p}_{1,\infty} = .5$, $\hat{p}_{2,\infty} = .25$, 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).

The rate of pass-through in the hyper forward-looking limit is

$$\lim_{m_f \rightarrow 1/(\theta_i \delta)} b_{i,t}(m_f) \approx 1 - e^{-\phi_i t},$$

with equality in continuous time. The intuition here is that, when a firm focuses on the long-run increase in its marginal costs immediately upon shock impact, there is no role for gradual pass-through of the shock by suppliers, as the drift in actual marginal cost is not relevant for the firm's pricing decision.

Now, we visualize how much longer pass-through takes to occur in each sector in the supply chain. We set $\delta = .96$ (a standard annual discount factor), and we set $\theta = .5$ (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. For each industry, we plot the fraction of long-run pass-through over time: $b_{n,t}$ such that $\hat{p}_{n,t} = b_{n,t}\hat{p}_{n,\infty}$.

Before discussing our results, we can apply Proposition 3 to make predictions about how the speed of pass-through varies with n and myopia, m_f :

Corollary 1. *The duration of pass-through in sector n simplifies to n/ϕ under complete myopia. This is increasing in n , meaning more downstream sectors pass through the shock more slowly. Under rational expectations and an interest rate of 0, the duration in sector n is $\frac{1}{\phi} \left| \binom{-3/2}{n-1} \right|$, which is still increasing in n . We have that the duration under complete myopia is longer than the duration under rational expectations for all $n > 1$ and is equal when $n = 1$.*

Figure 2 shows our results – we use the continuous time model directly in this case so that our propositions can be applied exactly. 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 3. 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%. We note that, even under rational expectations, pass-through is still more gradual as n increases.

4 Data

In order to test whether full pass-through of commodity shocks deep into the production network actually exists in practice, and whether it is gradual, 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). We initially focus on the effects of oil price changes in a setting with an arbitrary lag structure, motivated by Proposition 1, but we branch out to perform a variety of other analyses using other sources of variation and regression setups. We then perform empirical analyses with the lag structure of pass-through determined by the model.

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. 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 core analysis, we make use of input-output tables from 1997 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 3 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.¹⁰ 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.

¹⁰Some 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.

Additionally, we obtain data on industry-level wages from the Quarterly Census of Employment and Wages (QCEW). These, too, can be merged with the BEA input-output tables in the same way as the industry-level PPIs. And we obtain data compiled by Pasten, Schoenle, and Weber (2017) on the frequency of price adjustment.¹¹ The authors gained access to the BLS price micro data and computed the average frequency of price changes by industry.

5 Empirics Assessing Long-run Pass-through

Our analysis seeks to test our theoretical propositions in a more “model-free” way than our later structural analysis. We first utilize a shift-share design inspired by Proposition 1 to determine whether commodity price increases pass through to industry prices, regardless of whether the industry uses the commodity directly or indirectly through its supplier network. We also assess whether pass-through is gradual. Finally, we test whether industries more downstream from the commodity price increase experience slower pass-through.

5.1 Regression Specifications

We begin by testing the long-run pass-through implications of Proposition 1 with an arbitrary lag structure for short-run pass-through. 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 using pooled variation, we regress the price change in an industry on the input cost change due to movements in the price of a commodity or commodities of interest, the input cost change due to movements in wages, and a time fixed effect. That is,

$$\Delta P_{i,t} = \lambda_t + \sum_{h=-6}^{24} \beta_h [(\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z]_i \Delta P_{Z,t-h} + \sum_{h=-6}^{24} \gamma_h [(\mathbf{I} - \Phi)^{-1} \text{diag}(s_i^L) \Delta \mathbf{w}_{t-h}]_i + \epsilon_{i,t}, \quad (8)$$

where $P_{i,t}$ denotes the log price of industry i at time t , λ_t is a time fixed-effect, $[(\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z]_i$ represents the network cost shares of industry i in commodity Z , $\Delta P_{Z,t}$ is the change in the price of commodity Z over period t , $[(\mathbf{I} - \Phi)^{-1} \text{diag}(s_i^L) \Delta \mathbf{w}_t]_i$ represents input cost changes due to wage movements in various sectors whose output industry i utilizes (a necessary control suggested by Proposition 1), and $\epsilon_{i,t}$ is an error term. Note that $\sum_h \beta_h = 1$ corresponds to full pass-through in the long run (here defined as 24 months), as the right-hand-side variable of interest corresponds to the size of the cost shock experienced by industry i ; $\beta_0 = 1$ corresponds to full pass-through on impact, consistent with completely flexible pricing. Note that this specification has a shift-share interpretation, where the shares here are the cost shares originating from the input-output table and the shifts are commodity price changes. Consequently, the identification assumption for plim

¹¹We kindly thank the authors for sharing their data with us.

$\hat{\beta}_h = \beta_h$ is given by

$$E[\Delta P_{Z,t-h} \mu_t] = 0,$$

where $\mu_t \equiv E[(\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z]_i \epsilon_{i,t}$ is a cross-industry average of the product of the commodity- Z network cost share and the unobserved component $\epsilon_{i,t}$. In intuitive terms, if industry prices $p_{i,t}$ in high commodity- Z share industries grow faster for omitted reasons ($\Rightarrow \mu_t > 0$) in periods when positive shocks to the price of commodity Z also tend to occur ($\Delta P_{Z,t-h} > 0$) – or vice versa for negative shocks – the condition will not hold.

Recent work by Goldsmith-Pinkham, Sorkin, and Swift (2020) and Borusyak, Hull, and Jaravel (2022) has focused on the econometrics of shift-share designs. Our preceding identification assumption corresponds directly to the exogeneity condition in Borusyak, Hull, and Jaravel (2022) in the case of a macro-level shock and arbitrarily many cross-sectional units and time periods. The condition re-frames the identifying assumption from a difference-in-differences panel regression into a time series moment by defining the average over the cross-section μ_t . Consequently, they argue that such shift-share designs are valid provided the shift term is exogenous. In our case, the shift term is the price movement in commodity Z . We argue that this exogeneity assumption is reasonable, particularly in the case of commodities such as oil, for which the price is set in large global markets with the bulk of supply originating from foreign sources. Furthermore, we probe the validity of the assumption by studying specific contexts wherein it is even more likely to be true. That is, in robustness checks, we focus on specific subsets of variation most likely to be exogenous such as the 1979 oil shock or the 2014 shale boom, and we run an instrumental variables (IV) version of the preceding regression, instrumenting oil price changes with Kanzig’s (2021) series of exogenous oil price shocks induced by OPEC announcements.

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:

$$\begin{aligned} \Delta P_{i,t} = & \lambda_t + \sum_{h=-6}^{24} \eta_h [\mathbf{s}^Z]_i \Delta P_{Z,t-h} + \sum_{h=-6}^{24} \rho_h [(\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z - \mathbf{s}^Z]_i \Delta P_{Z,t-h} \\ & + \sum_{h=-6}^{24} \gamma_h [(\mathbf{I} - \Phi)^{-1} \text{diag}(\mathbf{s}_i^L) \Delta \mathbf{w}_{t-h}]_i + \epsilon_{i,t}. \end{aligned} \quad (9)$$

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. The coefficients η_h correspond to the extent of pass-through from direct exposure to cost shocks; the coefficients ρ_h correspond to the extent of pass-through from indirect exposure to cost shocks.

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,

$$\sum_{j=0}^t \Delta P_{i,j} = \lambda_t + \beta_t [(\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z]_i + \epsilon_{i,t} \quad (10)$$

where $P_{i,t}$ denotes the price of industry i at time t , $\sum_{j=0}^t \Delta P_{i,j}$ 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, $[(\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z]_i$ represents the network cost share of commodity or commodities \mathbf{Z} in industry i , and $\epsilon_{i,t}$ is an error term.¹² Here, the estimates β_t should eventually align with the change in the commodity price under full pass-through in the long run (rather than being equal to 1 in the long run). Concretely, if there is a cumulative $\sum_{j=0}^t \Delta P_{Z,j}$ log-point increase in the price of our commodity of interest, full pass-through and flexible pricing would imply a coefficient value of $\beta_t = \sum_{j=0}^t \Delta P_{Z,j}$. 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 $\sum_{j=0}^t \Delta P_{Z,j}$ 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:

$$\sum_{j=0}^t \Delta p_{i,j} = \lambda_t + \eta_t [\mathbf{s}^Z]_i + \rho_t [(\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z - \mathbf{s}^Z]_i + \epsilon_{i,t} \quad (11)$$

We note that in all of our regressions, the shocked industries will be excluded from the regression analysis.¹³ Finally, we provide some intuition for interpreting our regression coefficients in light of our inclusion of time fixed effects. Consider regression specification (11), our case study specification separating first-order from higher-order exposure. The coefficient ρ_h 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

¹²To ease the additional notational burden, we exclude general equilibrium wage controls in the case study specifications. As we discuss in Appendix C.2, there are no meaningful differences in our results if we exclude the wage controls from the regression.

¹³For 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, 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.

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.

5.2 Main Results

We begin by using all variation in oil prices from 1997 onward, running the regression specification given by Equation (8). 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.¹⁴

The red coefficients in Panel 1 of Figure 4 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 several months, pass-through increases gradually 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. The blue coefficients in Panel 1 of Figure 4 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, with some evidence of slower pass-through for non-oil commodity price movements.

Panel 2 of Figure 4 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 that correlated movements in non-oil commodity prices are not driving our findings for oil price movements (or vice versa).

One might worry the results are driven by full pass-through of commodity price movements to direct users, with relatively little pass-through deeper into the network. Figure 5 plots the results of the regression specification given by Equation (9), 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

¹⁴Both 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.

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 pass-through, of which very little occurs on impact. Instead, pass-through phases in slowly over the course of eight months or so.¹⁵

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 modeling framework implies that changes in TFP are in the error term of these regressions. By using more exogenous sources of oil price variation that are unlikely to be correlated with movements in TFP, we aim to minimize concerns relating to omitted-variable-bias. Consequently, Panel 2 of Figure 5 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 – suggesting faster pass-through than we saw in the OLS variants of the regression.

In Appendix C, we show that these results are robust to a variety of modifications and alternative approaches. In C.1, we show that the same patterns are evident in a binscatter analysis. In C.2, we show that our results are essentially unchanged if we do not include controls for general equilibrium wage effects. In C.3, we control for gas and electricity exposure - two commodities likely to be close substitutes for oil in some contexts. In C.4, we use alternative cost shares that exclude payments to capital from the share denominator. In C.5, we conduct permutation tests on our main specifications as an alternative robust method of generating p-values from within-sample.

5.3 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

¹⁵Dynamic treatment effect heterogeneity predicted by the model suggests that, if anything, this is an underestimate of the speed of pass-through. We defer a more in-depth discussion to our structural empirics using the lag structure implied by the model.

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, 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 6 displays the regression coefficients originating from applying Equation (11) 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 per barrel that occurred in 2008, oil prices remained near all-time highs throughout the early 2010s. They averaged over \$90 per 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 6 displays the regression coefficients originating from applying Equation (11) 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 of 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, 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 while much of the productive capacity 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 6 displays the regression coefficients originating from applying Equation (11) 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.

5.4 Heterogeneity

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

$$\begin{aligned} \Delta P_{i,t} = & \lambda_t + \sum_{h=-\tilde{H}}^H \beta_h \left[(\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z \right]_i \Delta P_{Z,t-h} \\ & + \sum_{h=-\tilde{H}}^H \tilde{\beta}_h \left[(\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z \right]_i \Delta P_{Z,t-h} \times \text{heterogeneity}_i \\ & + \sum_{h=-\tilde{H}}^H \gamma_h \left[(\mathbf{I} - \Phi)^{-1} \text{diag}(s_i^L) \Delta \mathbf{w}_{t-h} \right]_i + \epsilon_{i,t} \end{aligned}$$

We first confirm that our measure of industries' frequencies of price adjustment, from Pasten et al. (2017), are indeed predictive of the extent of price pass-through in the short-run ($\tilde{H} = H = 0$, so that we are focusing on one-month pass-through). Following the discussion in our model

section, define industry i 's time to reach a fraction X of long-run pass-through in a model with no network as $-\ln(1 - X)/\phi_i$. This is of course declining in the rate of price adjustment ϕ_i . We prefer using a measurement of the pass-through "half-life," so that $X = .5$.

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.

Columns (4), (5), and (6) repeat this exercise, adding downstreamness defined in Definition (3) 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.¹⁶

Figure 7 plots the impulse response function corresponding to the specification in column (4), returning to our previous dynamics setting the number of leads to $\tilde{H} = 6$ and the number of lags to $H = 24$. 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. Because this IRF corresponds to the specification in column (4), the interpretation is that, if two industries have the same frequency of price adjustment, but one is more downstream, the more downstream industry has slower pass-through.

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.

¹⁶We may be over-controlling here if industries have a low frequency of price adjustment because they are downstream from volatile input price movements and there is measurement error in downstreamness.

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

6 Empirical Analysis using Model-derived Dynamics

Our pricing model provides impulse response functions (IRFs) measuring how much each industry changes prices, and when, 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 model IRFs into the component due to pass-through under the model under full myopia and the component due to forward-lookingness under rational expectations.

6.1 Specifications for Semi-Structural Tests and Structural Estimation

Denote by $passthrough_{i,h}(\alpha)$ the model-defined pass-through of a commodity shock h months ago to prices in industry i under a model calibration α . We describe the key parameters in α after discussing our empirical strategy. The impulse response function characterizing how prices in industry i evolve over H periods in response to an unexpected change in the commodity price is then

$$IRF_{i,H}(\alpha) = \sum_{h=0}^H passthrough_{i,h}(\alpha).$$

We can test whether pass-through in the data follows the model's predictions using the regression

$$\Delta P_{i,t} = \lambda_t + \sum_{h=0}^H \beta_h(\alpha) passthrough_{i,h}(\alpha) \Delta P_{Z,t-h} + \epsilon_{i,t}. \quad (12)$$

If the model is capable of fitting the patterns in the data and our OLS is unbiased, we should have $\beta(\alpha_0) = \mathbf{1}$ for some calibration α_0 . As in our reduced-form analysis, we still include a time fixed effect; the regression has a dynamic shift-share interpretation. The interpretation of $\beta_h(\alpha)$ is again a relative effect: how much more an industry with a high pass-through at horizon h as determined by the model changes prices than an industry with lower pass-through at horizon h . If we did not include time fixed effects in our regressions, we might worry that action by the monetary authority affecting the aggregate price level would lead our pass-through predictions in the model to be wrong, leading to estimates of $\beta_h(\alpha) \neq 1$ even under the correct calibration α_0 .

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. Though this is a meaningful concern in principle, we show in Appendix D.1 that including refineries in our structural estimation does not meaningfully change our results, revealing robustness of our procedure to a large outlier.

Our test of how close the β 's are to 1 can be visualized with

$$IRF_{i,H}^{Data}(\alpha) = \sum_{h=0}^H \hat{\beta}_h(\alpha) passthrough_{i,h}(\alpha),$$

where $\hat{\beta}_h(\alpha)$ are the estimates from regression equation (12). While we will refer to this object as “the empirical IRF,” it is crucial to note that this measure varies with α and so is not independent of the model. It is best to think about it as a conveniently weighted test of whether $\hat{\beta}_h = 1$ for all h – intuitively, we care less about some $\beta_h \neq 1$ if the pass-through coefficient from the model at that horizon is small (which typically occurs when h is large). If the model is capable of fitting the data for the calibration α_0 and our OLS estimates from regression equation (12) are unbiased, then we should have

$$m(\alpha_0, H) = E[IRF_{i,H}^{Data}(\alpha_0) - IRF_{i,H}(\alpha_0)] = 0.$$

When presenting the results of regression (12), we plot the impulse response in the model compared to its empirical variant. It is infeasible to show IRFs for every industry (this would be hundreds of plots), so we summarize our results in IRFs for weighted industry averages. We will use an artificial industry comprised of the top 10th percentile of upstream industries affected by the shock, an artificial industry comprised of the top 10th percentile of downstream industries affected by the shock, and the average industry. The downstreamness split is simply made on the median pass-through half-life due to downstreamness, as defined under Definition 3. 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.

In our GMM estimation, we use $m(\alpha_0, H) = 0$ as our moment condition for a vector of horizons H , typically between 12 and 24 months, depending on our power. The number of horizons $H + 1$ is the number of moment conditions. The GMM procedure is as follows: given α , we will estimate the model, run regression (12) with H lags, construct IRFs for each industry, and then choose the α that gets the estimated IRFs as close as possible to the model-generated IRFs. The sample average of our vector of moments is

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

We can think about this sample average as the difference in the estimated and model IRFs for an artificial industry whose IRF is the average of all industry IRFs in the data. Then our estimate is

$$\hat{\alpha} = \arg \min_{\alpha} \hat{m}(\alpha)' \hat{m}(\alpha).$$

The GMM-optimal model is denoted by $\hat{\alpha}$. While the standard errors for $\hat{\alpha}$ will be informative about how important each parameter is in enabling model fit to the data, we will also be interested in testing whether moments we did not specifically seek to match in the GMM estimation are also matched. These tests will also be of independent interest. For example, our plots showing the fit for more upstream and downstream industries may differ than the fit for the average industry, which the GMM was designed to match. Further, note that the degree of myopia, m_f , will be one element of α . We can decompose

$$passthrough_{i,h}(\hat{\alpha}) = \underbrace{passthrough_{i,h}(\hat{\alpha}_{-m_f}, m_f = 0)}_{\text{Myopic Pass-through}} + FLGap_{i,h}(\hat{\alpha}),$$

with

$$FLGap_{i,h}(\hat{\alpha}) = \underbrace{passthrough_{i,h}(\hat{\alpha}_{-m_f}, m_f = \hat{m}_f) - passthrough_{i,h}(\hat{\alpha}_{-m_f}, m_f = 0)}_{\text{Additional Pass-through due to GMM-optimal Forward-lookingness}}.$$

If the component of industry pricing due to forward-lookingness is correct, we can run the regression

$$\begin{aligned} \Delta P_{i,t} = & \lambda_t + \sum_{h=0}^H \beta_{h, \text{Myopic}} passthrough_{i,h}(\hat{\alpha}_{-m_f}, m_f = 0) \Delta P_{Oil,t-h} \\ & + \sum_{h=0}^H \beta_{h, \text{FLGap}} FLGap_{i,h}(\hat{\alpha}) \Delta P_{Oil,t-h} + \epsilon_{i,t}, \end{aligned} \quad (13)$$

to test whether $\beta_{h, \text{Myopic}} = \beta_{h, \text{FLGap}} = 1$ for all h . In particular, $\beta_{h, \text{FLGap}} \neq 0$ tells us that the forward-looking component of the New Keynesian model has predictive power for pass-through of commodity shocks. Formally, 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 forward-lookingness affects the timing of their pass-through.

6.2 Model Calibration

Recall from Proposition 2 that, to solve the model for a persistent, unexpected commodity price change, we require measurements of (1) the frequencies of price adjustment for each industry, θ , (2) the degree of forward-lookingness m_f , (3) the discount factor δ (or a transformation of the steady-state interest rate), (4) the input-output matrix Φ , and (5) the commodity cost shares s^Z .

Now, the model solutions depend on whether the commodity price change is persistent and unexpected. Fortunately, real oil prices are approximately a random walk over a horizon of one year, and the real oil price is relevant for us (rather than the nominal price) because of our use of time fixed effects. We have verified this persistence in 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)). For our analysis of other commodities, we will only include commodities for which we cannot reject the random walk model in a time series analysis.

While real oil prices follow a random walk on average over a one year horizon, oil futures prices provide some information about how a given oil price movement is expected to evolve in the future. We incorporate futures explicitly in the model and structural estimation procedure in Appendix E. We show that the results in this section are robust to incorporating futures.

As discussed in the data section, we have measures of the frequencies of price adjustment from Pasten et al. (2017). It is not obvious whether these are required or helpful for fitting the pass-through patterns in the data. We might also consider the average frequency of price adjustment in this data, $\bar{\theta}_i$. We will calibrate

$$\theta_i = \gamma \bar{\theta}_i + (1 - \gamma) \theta_i^{\text{Pasten et al.}},$$

where γ is our first specified parameter in α and denotes the amount that we bias each industry's measured θ_i in the Pasten et al. data toward the average. The second parameter in α is the degree of forward-lookingness, m_f . The third parameter will be related to the steady state markup, and we describe it now.

Recall that industry i 's intermediate input share in industry j is

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

where C_i were the total input costs in industry i . Now, we do not observe C_i . Instead, we observe output in each industry, $P_i Y_i$. Using constant returns to scale and the optimal industry price in steady state ($P_i = \mu_i M C_i$), it follows that

$$C_i = \frac{P_i Y_i}{\mu_i}, \quad \mu_i = \frac{\sigma_i}{\sigma_i - 1},$$

where $\sigma_i > 1$ was the elasticity of substitution between varieties (firm-level output) in each industry. Now, the industry markup μ_i is also unobserved. One approach would be to estimate a common μ for all industries in our GMM procedure. We can do better, however, because in addition to measuring intermediate input and labor input costs for each industry, we also observe payments to capital. Payments to capital are comprised of profit and rental payments. In

particular, because rental payments are weakly positive in every period, industry profits cannot exceed measured payments to capital in the industry. Therefore, we can estimate the profit share of payments to capital, α , leading to cost calibrations in each industry of

$$C_i(\alpha) = \sum_j P_j X_i^j + W_i L_i + (1 - \alpha) R_i K_i.$$

Then $\mu_i(\alpha) = P_i Y_i / C_i(\alpha)$. We will present the results of our estimation of α in terms of the markup for the average industry, i.e. $\bar{\mu}(\alpha) = \frac{1}{I} \sum_i \mu_i(\alpha)$, with standard errors computed using the delta method and the GMM standard errors for $\hat{\alpha}$.

Summarizing, then, we are estimating three parameters in the vector $\alpha = (\gamma, m_f, \bar{\mu})$. Now, there is no strong reason to worry that γ should vary with the commodity price movement being analyzed. It is more complicated with the other two parameters. In particular, we know that the ideal calibration of C_i for each industry uses a markup measured for each industry. Because different industries are changing prices depending on which commodity is shocked, the estimate of $\bar{\mu}$ is likely to vary across commodities. Now, we have up until now not provided a microfoundation for m_f , and it is not the goal of this paper to evaluate different microfoundations for this parameter. We do have some hypotheses worth testing, however.

Consider the example of a furniture manufacturer that is exposed to oil because it uses petroleum-based foam cushioning in its chairs or produces memory foam mattresses. Suppose there is an oil price increase just before the company sets its annual catalog of prices. Forward-lookingness, i.e. $m_f > 0$, implies that such a manufacturer will increase prices, even if it has not yet experienced an increase in the underlying foam costs, anticipating such a cost increase in the near future. We believe there are many reasons, both behavioral and non-behavioral, that such a firm may not be forward-looking in this way. Oil is quite far upstream from furniture manufacturers, so it may be unlikely that they pay attention, even if the oil shock is particularly salient. For a firm to build in upstream cost shocks into pricing before marginal cost has been affected, they need information on when and how much the shock should affect their marginal cost. To acquire this information, they would likely need to pay a consulting firm, and so it is not costless to be forward-looking. To even think about hiring consulting services, they would need to be attentive to supply chain risks. It is sensible to think that it would be quite rational for firms not to form costly forward-looking expectations about marginal cost emerging from upstream price increases, particularly when ultimate exposure to the upstream price increases is small or the commodity prices are not volatile. For these reasons, we will allow m_f to vary across commodity and assess its variation.

6.3 Results for Oil

First, we use our GMM framework to estimate α using oil price movements. We estimate $\hat{\gamma} = 0$ (SE = 0.210), $\hat{m}_f = 1$ (SE = 0.204), and $\hat{\bar{\mu}} = 1.230$ (SE = 0.093).¹⁷ Therefore, the data prefers the heterogeneous frequencies of price adjustment, $\theta_i^{\text{Pasten et al.}}$, rational expectations, and a steady state markup of 1.23. Recall again that this is the markup required to fit impulse responses of industries passing through oil shocks, which is not necessarily similar to the aggregate markup if oil-dependent industries (in the network sense) have systematically different markups than other industries. To get a sense for how good the model fit is, and to assess model fit under other parameter values, developing the intuition behind our standard errors, we now turn to plotting our results.

Panel 1 of Figure 8 shows the fit of our GMM-optimal model for the average industry, which was the fit we tried to optimize. We see that the fit is quite good for the duration of the IRF. Panel 2 uses the GMM-optimal model but sets $\theta_i = \bar{\theta}$, the average frequency of price adjustment. We see suggestive evidence that there is more pass-through in the data relative to the model initially, followed by strong evidence of less pass-through in the data relative to the model in later periods. This follows in part because the industries most exposed to oil have high frequencies of price adjustment. We also note that the IRF is essentially flat starting around month 10. This means that the β 's estimated in regression (12) are, on average, 0 starting in month 10; assuming a homogeneous θ therefore means that cross-sectional variation in the model's pass-through coefficients is broadly not useful for predicting which industries will change prices more in response to oil during those periods.

Panel 3 uses the GMM-optimal model but sets $m_f = 0$ rather than using the GMM-optimal rational expectations. Early on, pass-through is not substantially different, suggesting that very upstream industries are identifying these coefficients – recall from Proposition 3 that myopia does not affect pass-through speed for direct users of the commodity when the shock is unexpected and persistent. Later in the IRF, however, the myopic pass-through measures predict too much pass-through, and so the β 's estimated from regression (12) at higher lags are on average less than one. The intuition here is that, because pass-through is slower under myopia, particularly for downstream firms, the pass-through measures are larger at higher lags than they are under rational expectations.

Finally, Panel 4 uses the GMM-optimal model but sets $\bar{\mu} = 1$ rather than using the GMM-optimal markup. Setting a lower markup reduces industries' network exposures to the commodity, thereby reducing our prediction of the extent of long-run pass-through. We see that there is uniformly more pass-through in the data than predicted by the model. By revising the estimated markup up from 1, the GMM-optimal model is able to improve fit.

¹⁷As in our empirics in the previous section, standard errors are clustered by industry.

Our estimates are robust to instead using a 2-step GMM procedure, which uses our previous estimates to form an optimal weight matrix for the moments. In the 2-step procedure, we find $\hat{\gamma} = 0$ (SE = 0.149), $\hat{m}_f = 1$ (SE = 0.130), and $\hat{\bar{\mu}} = 1.06$ (SE = 0.041). The only difference here is a lower markup and smaller standard errors.

Now, we address the question of why pass-through seems so much slower in our structural analysis than it did in our analysis from the previous section. 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 C.6, finding pass-through for sufficiently downstream industries that is slower. This pass-through is also larger than 1, which reflects our use of $\bar{\mu} = 1$ for the reduced-form analysis and replicates the finding of Panel 4 in Figure 8 but in reduced form.

Finally, rather than looking at the IRF for the average industry, we can visualize the solution for the most upstream and downstream industries, as described after our regression specification. These IRFs let us visualize how close the β 's from regression (12) were to 1 for different lags, since upstream industries pass through the shock faster than the average industry, while downstream industries pass through the shock more gradually than the average industry. In Figure 9, we see that fit remains good for these visualizations. Moreover, we see how different the speed of pass-through can be as a result of downstreamness. Upstream industries reach 75% of long-run pass-through in 6 months, while the average industry takes around 20 months and downstream industries have not yet achieved 75% of long-run pass-through after two years.

Now, recall that regression (13) split the GMM-optimal model's solution into the myopic solution and the component due to forward-lookingness. If forward-lookingness is truly operative, it should have predictive power in the cross-section: industries that should increase prices faster due to forward-lookingness, holding myopic exposure to the shock constant, should increase prices faster in the data. 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. To

prepare for the shape of the IRF for the forward-looking gap, we note that

$$IRF_{i,H}^{FLGap} = \sum_{h=0}^H FLGap_{i,h}(\hat{\alpha}) \rightarrow 0 \quad \text{as } H \rightarrow \infty,$$

which follows from long-run pass-through of a persistent shock being the same under rational expectations and myopia (Proposition 1).

We show our results in Figure 10. 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. For the industries comprising the IRF – those most affected by rational expectations – we see that the boost to pass-through from forward-lookingness is nearly 20% of long-run pass-through 4-5 months after the shock.

The finding of rational expectations suggests that firms should be forward-looking not only about gradual pass-through of oil price movements to their marginal costs; they should also be forward-looking about future movements in oil prices, to the extent that the market possesses information about such movements. We verify that this forward-lookingness is present in Appendix E, where we use oil futures data for all monthly maturities extending out five years from Bloomberg. Indeed, firms respond to “forward-guidance” from oil futures prices, in accord with our finding that firms are forward-looking about changes in marginal cost resulting ultimately from oil price movements. We demonstrate these findings using regressions incorporating the model-implied lag structure for how industries should respond to oil futures price movements in Appendix E.3. In Appendix E.4, we confirm these findings in a reduced-form approach consistent with our empirical analysis from Section 5. In Appendix E.5, we re-run our GMM procedure and estimate an α that is not statistically different from the estimate found in this section.

6.4 Results for Non-oil Commodities

Both for power reasons, and because we do not think, *a priori*, that the parameter should vary with the commodity that is shocked, we now set $\gamma = 0$, so that $\theta = \theta^{\text{Pasten et al.}}$. We estimate $\alpha = (m_f, \bar{\mu})$ for each of fourteen commodities – all of the commodities with random walk price movements and for which we can estimate GMM for at least a horizon of 12 months. We do not have substantial power to estimate the markup for individual commodities, but we do have power to estimate the degree of myopia for most commodities. We show our results in Figure 11. For all precisely estimated cases, we find $m_f = 0$, or complete myopia. The only two outliers are (1) beef cattle ranching and farming and (2) animal production, except cattle/poultry/eggs, and in these cases the standard errors reveal we have no power to estimate myopia in the data – these two cases are excluded from the figure for readability. The relatively tight standard errors on the precisely estimated myopia coefficients imply that the model fit is meaningfully worse under

different values of m_f .¹⁸

In this sense, oil is very special among the commodities. Downstream industries, even those that do not use oil directly in production, pay attention to oil prices, passing through changes before their marginal costs fully reflect the changes in the oil price. For commodity price movements beyond oil, firms are less attentive to the cost increases of their upstream suppliers.

We now ask if it is possible to recover a role for forward-lookingness when non-oil commodity prices change. To do this, we consider price changes for composite commodities, allowing all commodity prices to move jointly and assessing how industry prices change in response to all non-oil commodity prices moving at the same time. Under this experiment, we will capture if firms are forward looking about upstream cost increases when many commodity prices move jointly, as they did during the commodity price boom of the early-mid 2000s, prior to the Great Recession. In such cases, industry prices are typically predicted to move much more on average, and across many more industries, than when a commodity price moves in isolation.

Our GMM-optimal estimates are $\hat{m}_f = 0.691$ and $\hat{\mu} = 1$. We show the optimal model fit in Figure 12. We see that the overall fit is good, regardless of whether we focus on the average industry or the most upstream and downstream industries. What is also interesting is that the standard errors span the entire possible space for the markup and myopia. Put differently, the model fit remains relatively good if we consider different values of $m_f \in [0, 1]$ and somewhat higher values of the markup. The reason for this is instructive and seems to be twofold: (1) there are many direct users with high direct exposures in the composite commodity, and coefficients are disproportionately estimated from direct users, for which Proposition 3 says there is no role for myopia to influence pass-through; and (2) our approach for estimating markups means the maximal markup is bounded above by a transformation involving all of industries' payments to capital being attributed to profits rather than rental payments, and direct users have a relatively small capital share of cost on average due to extensive use of intermediate inputs.

In contrast, all coefficients are well identified for our oil results because there is limited direct use of oil, particularly when excluding petroleum refineries, as we have done in all of our structural analysis for oil. Concerning forward-lookingness, again by Proposition 3, we only have power to identify myopia empirically when the analysis includes indirect users of the commodity. This is empirically confirmed in our robustness check that includes petroleum refineries in our oil results (Appendix D.1), where we see somewhat larger standard errors on the myopia parameter in both 1-step and 2-step GMM estimates.

Further, there is also a larger role for the markup when indirect users of a commodity are primarily used to identify our parameters. To see this, define the input shares in revenues (rather

¹⁸Standard errors are adjusted for the estimates being on the boundary of the parameter space, in which case they (asymptotically) have a halfnormal distribution (Andrews 1999).

than costs) as

$$\tilde{\Phi}_{i,j} = \frac{\Phi_{i,j}}{\bar{\mu}}, \quad \tilde{s}_i^Z = \frac{s_i^Z}{\bar{\mu}}.$$

Then

$$\begin{aligned} (\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z &= \mathbf{s}^Z + \Phi \mathbf{s}^Z + \Phi^2 \mathbf{s}^Z + \dots \\ &= \bar{\mu} \tilde{\mathbf{s}}^Z + \bar{\mu}^2 \tilde{\Phi} \tilde{\mathbf{s}}^Z + \bar{\mu}^3 \tilde{\Phi}^2 \tilde{\mathbf{s}}^Z + \dots \\ &= \bar{\mu} (\mathbf{I} - \bar{\mu} \tilde{\Phi})^{-1} \tilde{\mathbf{s}}^Z. \end{aligned}$$

Therefore, similarly to how the effect of nominal rigidities is amplified with downstreamness, as shown in Proposition 3, the effect of the markup on the extent of pass-through is amplified with downstreamness. Intuitively, if an industry's suppliers are predicted to have lower pass-through due to a lower markup, the ultimate input price increase the industry sees in response to an upstream commodity price increase is smaller, and this is an additional effect leading to smaller downstream price increases.

Therefore, we think our methodology is best suited to assessing whether and how quickly individual sectoral price changes pass through the production network. Doing otherwise means our GMM procedure disproportionately isolates pass-through of direct commodity users, minimizing the role of the network. In cases where we isolate variation in individual commodity prices, we find that rational expectations is required to explain the pace of oil price pass-through, while myopia is required to explain the pace of non-oil commodity price pass-through. A crucial caveat is that, during periods when many commodity prices are moving together, we cannot reject that downstream firms pay just as much attention as they do to oil price changes, and the role for myopia may be attenuated.

7 Application: Network Oil Inflation

As we have shown in our empirical work, oil and other commodity shocks generate inflation beyond changing prices for products in which they are directly used. An inflation measure that subtracts only the direct component of oil inflation – derived from changes in consumer gasoline prices – 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, inflation should be predictable using the network component of oil inflation.

We begin in subsection 7.1 with a practitioner's guide for how to apply our model to purge the network effects of a commodity price movement from inflation. In subsection 7.2, we then illustrate this procedure for the case of oil price changes, providing estimates of how oil prices affect aggregate inflation both directly and indirectly. We produce historical series of inflation

resulting from the direct and indirect effects of historical oil price changes. In subsection 7.3, we explore whether Core Personal Consumption Expenditures (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. In subsection 7.4, we produce a new inflation series that fully purges the effects of oil price movements from inflation. Finally, in subsection 7.5, we describe how one would fail to match the data using a representative industry model without a production network.

7.1 A User’s Guide

While we focus on the inflation arising from oil in this section, our model can be applied to purge the direct and indirect effects of any commodity price movement from aggregate inflation. This may be a desirable exercise for practitioners at private and central banks and policy institutes, among others. In particular, if one thinks a specific supply chain disruption or commodity price increase is responsible for observed inflation, our framework can be used to quantify how much inflation should result and over what time horizon. In this subsection, we provide a framework outlining how to apply our model to remove the direct and indirect effects of any sectoral price movement from inflation. The rest of our application section illustrates this methodology for oil.

Recall that the model can be solved to determine the dynamic pass-through of any commodity price change to all industry prices, using the formula provided by Proposition 2. As in our structural estimation section, denote by $passthrough_{i,h}(\alpha)$ the proposition’s prediction for how a unit log point increase in the chosen commodity price h periods ago affects prices in sector i under model calibration α . The optimal calibration $\hat{\alpha}$ is chosen to match the cross-sectional variation in industry price changes resulting from variation in the commodity price, as we did in our structural estimation. We recommend using rational expectations for oil price changes and complete myopia for non-oil commodity price changes. We also recommend using frequencies of price adjustment measured for each industry, as we have through our Pasten et al. (2017) data. Finally, we recommend, for oil price changes, using our GMM-estimated aggregate markup. For other commodities, we were underpowered to estimate the markup; the competitive case would provide a conservative benchmark, but we would also recommend using a calibration with measured markups in each industry.

The response of aggregate PCE inflation to a sequence of changes in the commodity price is then

$$\hat{\Pi}_t = Intercept + \sum_i PCEShare_i \sum_{h=0}^H passthrough_{i,h}(\hat{\alpha}) \Delta P_{Z,t-h}, \quad (14)$$

an intercept plus the PCE-weighted average of industry pass-through coefficients (under the optimal calibration) multiplied by the changes in the commodity price. The intercept is required

because our model is a supply-side model (not a closed model with a demand side and monetary rule) and α is estimated using cross-sectional variation (exploiting use of a time fixed effect).

The literature on aggregating cross-sectional estimates provides many methods for determining an intercept. We now discuss one such method – our favored method for oil – which requires the user to possess a commodity-price instrument valid for time-series identification.¹⁹ With such an instrument, one can estimate the intercept as the effect of the commodity price movement on a hypothetical sector with no network exposure to the commodity. Intuitively, in our supply-side model, such a sector is not predicted to experience a price change in response to the commodity price increase. Any price change experienced by the sector, therefore, must come from systematic comovement in general equilibrium effects, such as a response of the monetary authority to the commodity price change or a response of inflation expectations. This exercise can be conducted using a 2SLS variant of regression (8) without a time fixed effect or a 2SLS variant of regression (12) without a time fixed effect. We provide the full details for these procedures applied to oil in Appendix F.2 and find that they do not meaningfully differ.

Given an intercept, the PCE share of each industry, model-derived pass-through coefficients with the optimal calibration $\hat{\alpha}$, and the sequence of commodity price changes, the user can apply formula (14) to compute the effects of the price changes on aggregate inflation. Because propagation takes time to occur, these effects should predict aggregate inflation – a testable hypothesis. Moreover, subtracting these effects provides a measure of inflation purged of the direct and indirect effects of the sequence of commodity price changes; formally, the new measure considers the counterfactual scenario in which the commodity price was constant.

Many of the steps listed above can optionally be subjected to additional robustness checks. We perform some of these checks in our illustration. Furthermore, the direct and indirect effects of the commodity on inflation can be decomposed. This is particularly useful in the case of oil because its direct effects on gas prices are large and already purged from PCE inflation in the Core PCE inflation series. The Core PCE inflation series, on the other hand, does not purge the indirect effects of oil price movements.

Finally, the user may wish to incorporate expectations information from commodity futures data when computing the commodity inflation series. We outline this procedure for oil in Appendix F.1.

¹⁹One alternative, model-driven approach to the time-series procedure described here is to close the macroeconomic model with a demand side and a monetary rule. In this case, we caution that the choice of monetary rule is critically important in driving the magnitude of the intercept, and we note that the cross-sectional estimation of α should be conducted with pass-through coefficients from the closed model, not from the supply-side component of the model.

7.2 Illustration: Oil

Earlier in the paper, we derived, estimated, and measured all of the elements needed to compute network oil inflation using formula (14), except the intercept. We can estimate the intercept with the aforementioned time-series method. We use our IV specification with the Kanzig (2021) shocks, given that the time-series variation in these shocks is valid for identification. This procedure is discussed in detail in Appendix F.2. We cannot reject the hypothesis that, in response to oil shocks, a hypothetical industry with a zero network cost share in oil does not change prices. 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. Therefore, our inflation prediction from formula (14) collapses to

$$\hat{\Pi}_t = \sum_i PCEShare_i \sum_{h=0}^H passthrough_{i,h}(\hat{\alpha}) \Delta P_{Z,t-h}.$$

Now, we apply formula (14) to compute network oil inflation.²⁰ We decompose this inflation into direct and indirect sources. Consumer use of gasoline appears in the data as personal consumption expenditures from the petroleum refining sector. The effect of oil prices on gas prices paid by consumers 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 13 reports our results. In Panel 1, we see that, of the about 6 percentage points of aggregate inflation resulting from a unit shock to log oil prices, just over 2 percentage points comes from consumer gas purchases, while about 4 percentage points comes from consumer purchases from all other industries. The indirect inflationary effect is realized slowly, dominating the direct effect only after 5 months, and only 75% of the total indirect effect is realized over the first year after the oil price change. In Panel 2, we compute the predicted year-over-year inflation resulting from historical oil price movements, applying equation (14) but decomposing direct and indirect effects. The combined direct and indirect effects are not uncommonly above 2%, and they exceeded 3% in the early-mid 1970s. We note again how gradual the indirect effects can be, a result that is clearly visible in the indirect effects of the 1970s and early 1980s oil price changes.

²⁰While our GMM-optimal procedure optimized fit for the equally weighted average industry, we confirm in Appendix F.3 that the fit remains good for the PCE-weighted average industry, both including and excluding petroleum refineries.

7.3 Explaining the Predictability of Official Inflation Measures

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. This distinction is important because wholesalers and retailers must pass through producer price changes before consumers see a price change. If pass-through is disrupted at this step, or is additionally slowed through price rigidity among wholesalers and retailers, our model's predictions for aggregate producer price inflation could differ from their effects on consumer price inflation.

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. Our results are shown in Table 2. We see in Panel 1 that indirect oil inflation explains a substantial fraction of official monthly PCE inflation: the R^2 is 33%. 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 explain 34% 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 2 reveals that it is only partially successful in doing this. While direct oil inflation indeed has little predictive power for core PCE inflation, indirect oil inflation is still strongly predictive. We retain an R^2 of 16% when explaining official core PCE inflation with network oil inflation.

We can further assess whether the dynamics of inflation predictability using oil price movements are in line with those in the time series. First, we assess how predictable PCE and Core PCE inflation are from oil price movements. To do this, we first run the 2SLS time-series regression

$$\Delta P_t = \alpha + \sum_{h=0}^{24} \beta_h \Delta P_{Oil,t-h} + \epsilon_t$$

where $\Delta P_{Oil,t-h}$ is instrumented with the OPEC shock series identified in Kanzig (2021). Accumulating the β 's estimated from this time series regression allows us to determine the overall effect of oil price increases on inflation. We can use P_t equal to PCE or Core PCE inflation. We then overlay our model's predicted effects on these accumulated estimates. If our model's predictions for aggregate inflation, constructed using producer price indices, pass through immediately to official measures, constructed primarily using consumer price indices, the dynamics from our model should fit the inflation predictability using oil that we see in the time series

We show our results in Figure 14. First, from the time series impulse responses, we see that aggregate inflation is predictable using oil price movements. Even Core PCE inflation remains

predictable using oil price movements. We see that our model can explain this predictability: oil price movements take time to pass-through to industry prices through supply chains, leading to lagged effects on aggregate inflation. These effects are large in magnitude: from a 1 log point oil price increase, about 3.5 percentage points of Core inflation are realized over the 2 years following the price increase. About 5.5 percentage points of PCE inflation are realized over the 2 years following the price increase. Pass-through in the time series is somewhat slower than what the model predicts, suggesting some additional delay resulting from gradual pass-through of wholesalers and retailers.

7.4 Removing All Oil Inflation from Aggregate 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 maintains the interpretation that 1970s/80s 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 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 and whether they played a role in causing inflation expectations to become unanchored.

7.5 Alternative “Strawman” Model: A Representative Industry

Rather than using our network model, we consider whether it is possible to match the time-series predictability of aggregate inflation using a model with just one rational expectations industry that uses labor and oil. We continue using the same model setup, so the firms in the industry can adjust prices only with some probability. Because there is no network, long-run pass-through from Proposition 1 is just s_i^Z . The rate of pass-through is approximately $(1 - e^{-\phi t})$, with equality in continuous time. We use an average price duration of 5 months, the average duration of industry prices in our Pasten et al. (2017) data.

The calibration of s_i^Z is not straightforward in the context of a representative industry. In the 2012 input-output table, the most recent table available, s_i^Z cannot be the share of oil sector GDP in aggregate GDP, since oil sector GDP is negative due to extensive oil imports. A different approach is to use the total intermediate use of oil (including use of domestic and imported oil) across all sectors in the economy as a fraction of aggregate output. Using aggregate output as the denominator instead of GDP solves the usual double-counting of GDP concerns in national accounting – total intermediate use of oil should be scaled by aggregate output, which is about 1.8 times larger than GDP. This calculation of long-run pass-through is also inconsistent with our results, even when using very high assumed values of the aggregate markup, as the production network amplifies oil shocks to have larger effects. Instead, we will plot results where we scale pass-through in each period by long-run pass-through and determine whether we can at least fit the duration of oil pass-through using either of aforementioned price adjustment frequencies.

We show our results in Figure 16. Pass-through under an average price duration of five months is faster than we found from our time series analysis of inflation predictability, depicted in Figure 14. In both the model and the time series, less than 50% of long-run pass-through has occurred after 5 months in Figure 14, while about 65% has already occurred under the representative industry model.

Our model, which used heterogeneous frequencies of price adjustment with an average price duration of five months, was able to fit the data because the network amplifies nominal rigidities. It is also able to fit the extent of pass-through, unlike the representative industry model. We are also matching the pass-through patterns in the time series with a markup and frequencies of price adjustment that are consistent with the cross-sectional pass-through of oil shocks to industry prices.

8 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 occurs primarily in the month of the cost shock's impact, whereas pass-through to indirectly-exposed industries is substantially more gradual. 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

our pricing model, which features Calvo pricing and a production network. 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 price adjustment. Empirically, we show that price pass-through responses to oil shocks are consistent with fully rational expectations, whereas responses to non-oil commodity price movements 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 ($R^2 = 15\%$) predictive power for official Core PCE inflation, as well as even greater predictive power ($R^2 = 33\%$) for total PCE inflation. As such, we create a revised measure of PCE inflation which removes all inflation resulting from oil. This results in a noticeably lower level inflation in 2021 and 2022 (a peak of 5%) but does not overturn the result that inflation in those years was well above the normal range of recent decades.

We think there are many interesting avenues for future work. While we have focused on commodity price movements, our model can be applied to price shocks in any sector. It could also be applied to assess the network pass-through of exchange rate movements. We have also neglected the distinction between labor and capital, but there may be interesting heterogeneity to explore. Further, our analysis was first-order, and higher-order effects may be important for certain shocks and in certain sectors. We also think the network setting provides a powerful empirical lab for testing different ways of endogenizing industries’ frequencies of price adjustment and the degree of forward-lookingness in expectations.

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.
- Afrouzi, Hassan and Saroj Bhattarai**, “Inflation and GDP Dynamics in Production Networks: A Sufficient Statistics Approach,” Working Paper March 2022.
- Alquist, R., L. Kilian, and R.J. Vigfusson**, “Forecasting the Price of Oil,” *Handbook of Economic Forecasting*, 2013, 2, 427–507.
- Andrews, Donald W. K.**, “Estimation When a Parameter is on a Boundary,” *Econometrica*, 1999, 67 (6), 1341–1383.
- Angeletos, George-Marios and Zhen Huo**, “Myopia and Anchoring,” *American Economic Review*, April 2021, 111 (4), 1166–1200.
- Antràs, Pol and Davin Chor**, “Global Value Chains,” Working Paper 28549, National Bureau of Economic Research March 2021.
- 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.
- Baqae, David R.**, “Cascading Failures in Production Networks,” *Econometrica*, 2018, 86 (5), 1819–1838.
- **and Emmanuel Farhi**, “The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten’s Theorem,” *Econometrica*, 2019, 87 (4), 1155–1203.
- **and –**, “Networks, Barriers, and Trade,” Working Paper 26108, National Bureau of Economic Research July 2019.
- **and –**, “Productivity and Misallocation in General Equilibrium,” *The Quarterly Journal of Economics*, 09 2019, 135 (1), 105–163.
- 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.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer**, “Diagnostic Expectations and Credit Cycles,” *The Journal of Finance*, 2018, 73 (1), 199–227.
- , – , and – , “Overreaction and Diagnostic Expectations in Macroeconomics,” *Journal of Economic Perspectives*, August 2022, 36 (3), 223–44.
- Borusyak, Kirill, Peter Hull, and Xavier Jaravel**, “Quasi-Experimental Shift-Share Research Designs,” *The Review of Economic Studies*, 2022, 89 (1), 181–213.
- 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**, “The Granular Origins of Aggregate Fluctuations,” *Econometrica*, 2011, 79 (3), 733–772.
- , “Chapter 4 - Behavioral inattention,” in B. Douglas Bernheim, Stefano Della Vigna, and David Laibson, eds., *Handbook of Behavioral Economics - Foundations and Applications 2*, Vol. 2 of *Handbook of Behavioral Economics: Applications and Foundations 1*, North-Holland, 2019, pp. 261–343.
- , “A Behavioral New Keynesian Model,” *American Economic Review*, 2020, 110 (8), 2271–2327.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift**, “Bartik Instruments: What, When, Why, and How,” *American Economic Review*, 2020, 110 (8), 2586–2624.
- Gordon, Robert J.**, “Output Fluctuations and Gradual Price Adjustment,” *Journal of Economic Literature*, 1981, 19 (2), 493–530.
- Horvath, M.**, “Cyclicalities and Sectoral Linkages: Aggregate Fluctuations 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.
- La’O, J. and A. Tahbaz-Salehi**, “Optimal Monetary Policy in Production Networks,” *Econometrica*, 2022, 90 (3), 1295–1336.
- Liu, E.**, “Industrial Policies in Production Networks,” *Quarterly Journal of Economics*, 2021, 134 (4), 1883–1948.
- Mankiw, N. Gregory and Ricardo Reis**, “Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve*,” *The Quarterly Journal of Economics*, 11 2002, 117 (4), 1295–1328.
- Nakamura, Emi and Jón Steinsson**, “Five Facts about Prices: A Reevaluation of Menu Cost Models*,” *The Quarterly Journal of Economics*, 11 2008, 123 (4), 1415–1464.
- and – , “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.
- Rubbo, E.**, “Networks, Phillips Curves and Monetary Policy,” Working Paper 2020.

Sims, Christopher A., “Implications of rational inattention,” *Journal of Monetary Economics*, 2003, 50 (3), 665–690. Swiss National Bank/Study Center Gerzensee Conference on Monetary Policy under Incomplete Information.

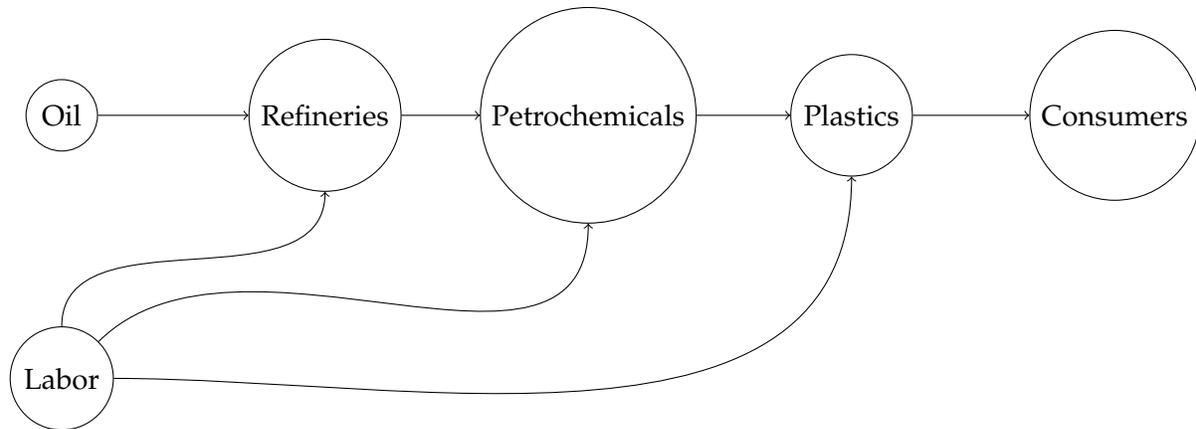
Smets, F. and R. Wouters, “Shocks and Frictions in US Business Cycles: A Bayesian DSGE approach,” *American Economic Review*, 2007, 97 (3), 586–606.

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

Woodford, Michael, “Imperfect Common Knowledge and the Effects of Monetary Policy,” Working Paper 8673, National Bureau of Economic Research December 2001.

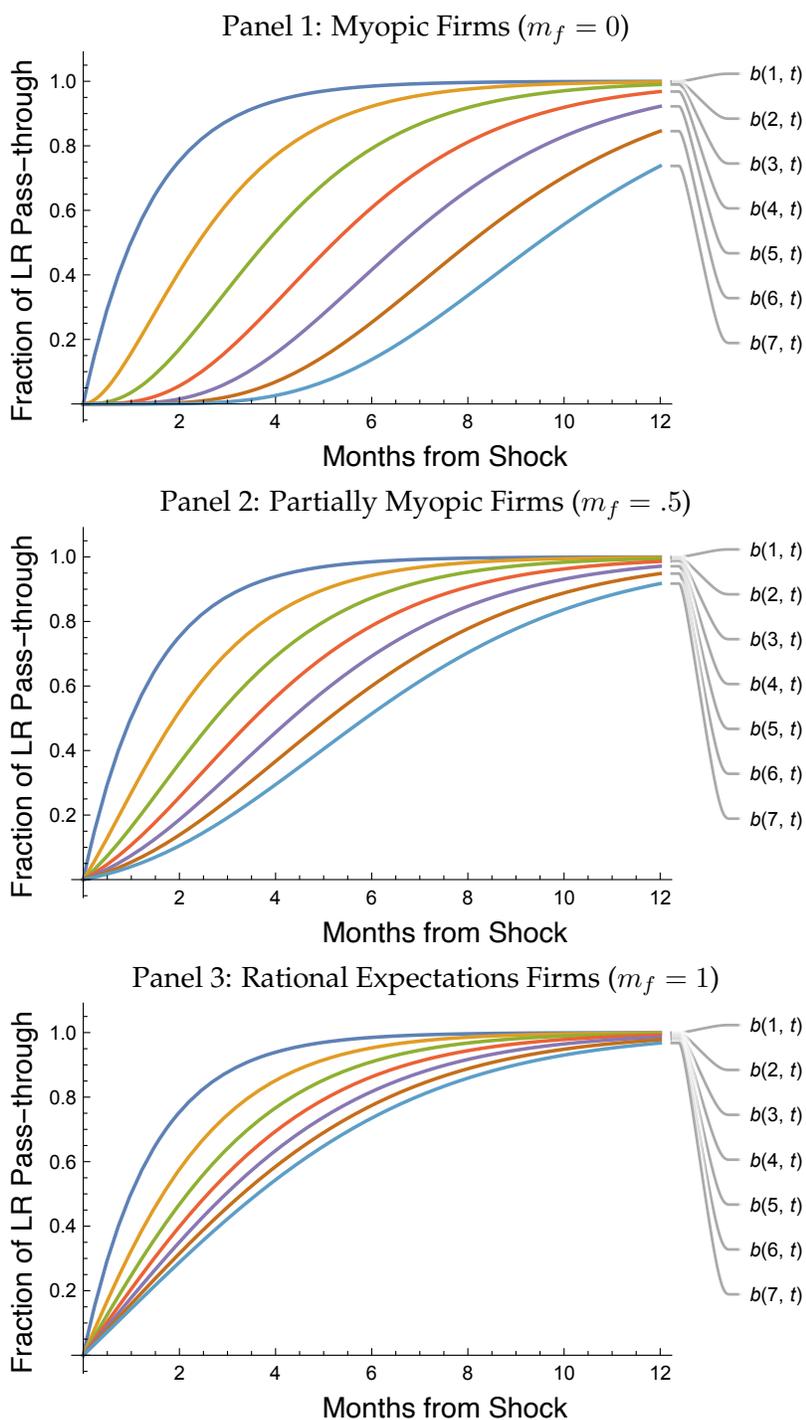
9 Figures and Tables

Figure 1 – A Linear Network



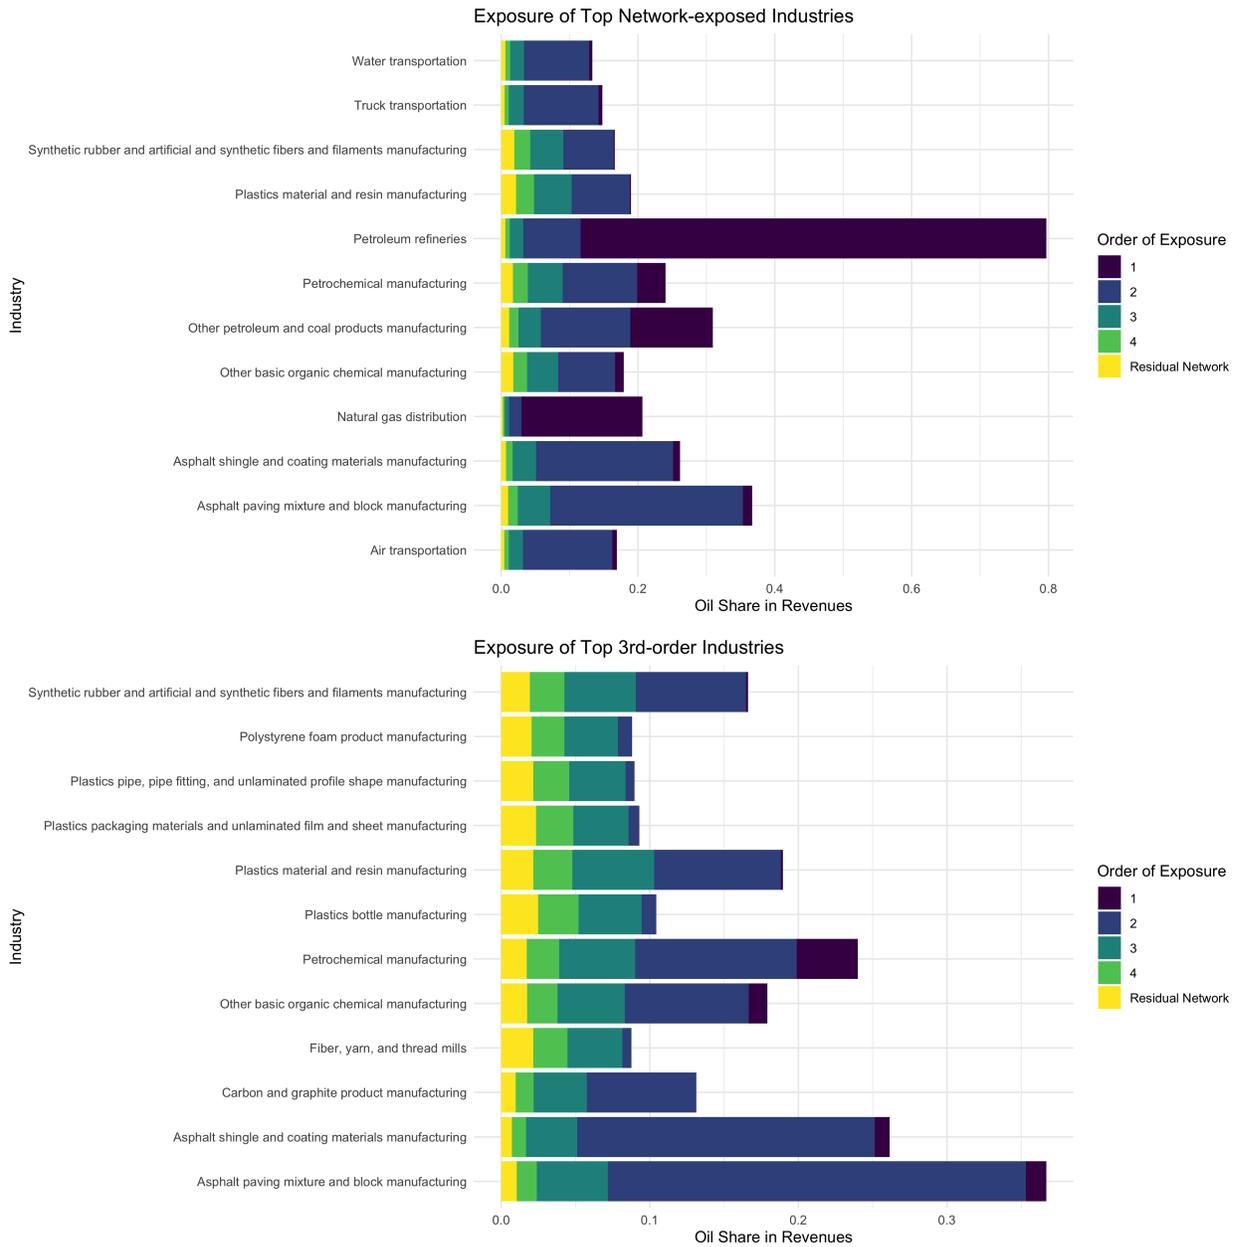
Note: In our linear network illustration of the general model, each industry uses labor and intermediate inputs only from the industry immediately preceding it in the supply chain. In the illustration, the plastics sector purchases inputs from the petrochemical sector but not from petroleum refineries directly – they only use petroleum refinery output indirectly through their use of petrochemical inputs.

Figure 2 – The Influence of Expectations on Pass-through Dynamics



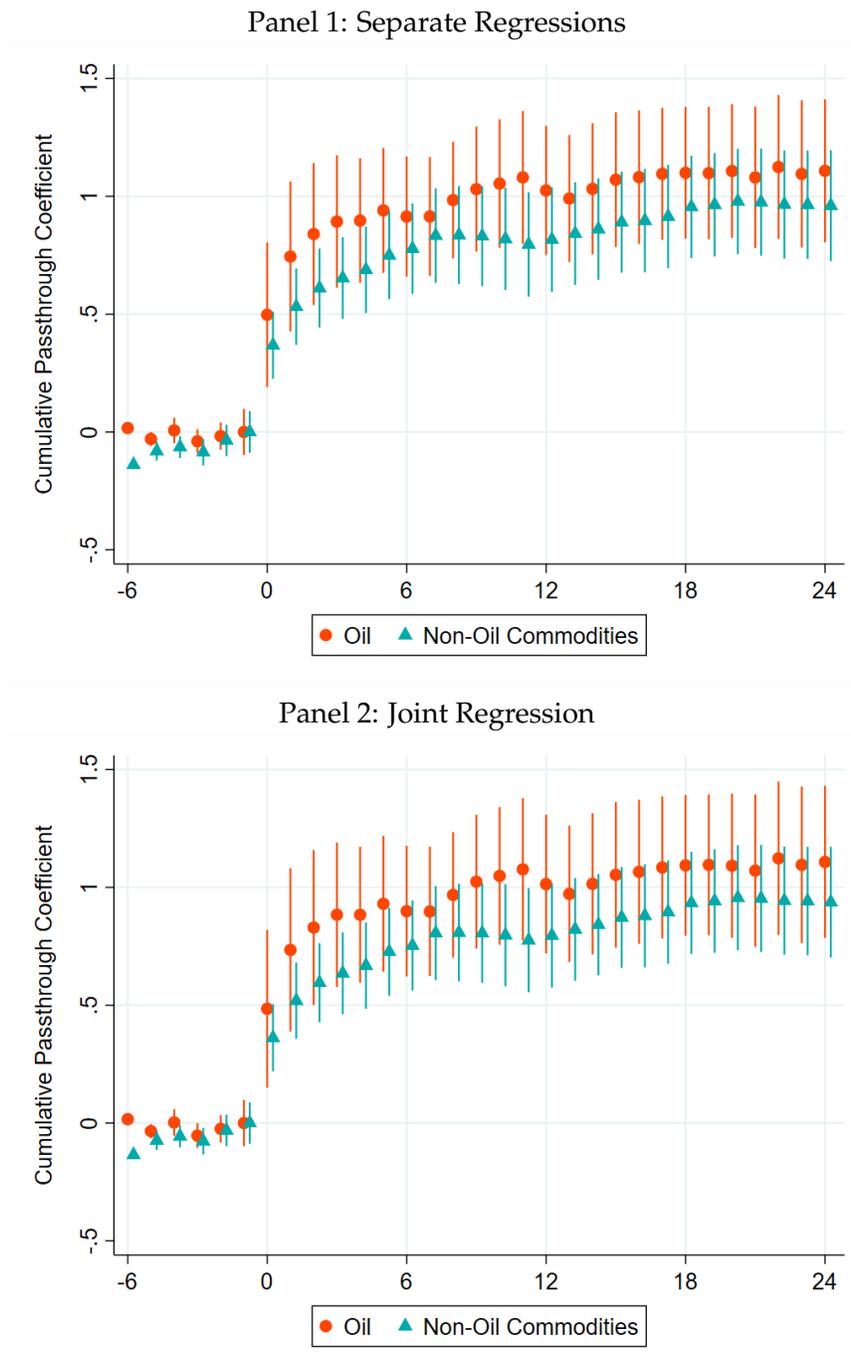
Note: We display the time it takes each of seven sectors in the linear network to pass through a persistent and unexpected shock to the upstream price. More downstream sectors take more time to pass through the shock, and the amount of additional time is increasing in the extent of myopia. The rate of price adjustment is calibrated so that half of the firms in each sector have the opportunity to change prices in each month; the results are plotted for the continuous-time variant of the model, so there is no jump in prices on impact. The y-axis abbreviation LR stands for long-run, so we are plotting the fraction of long-run pass-through achieved at each month under the myopia calibration in the panel description.

Figure 3 – Network Oil Shares for a Selection of Industries



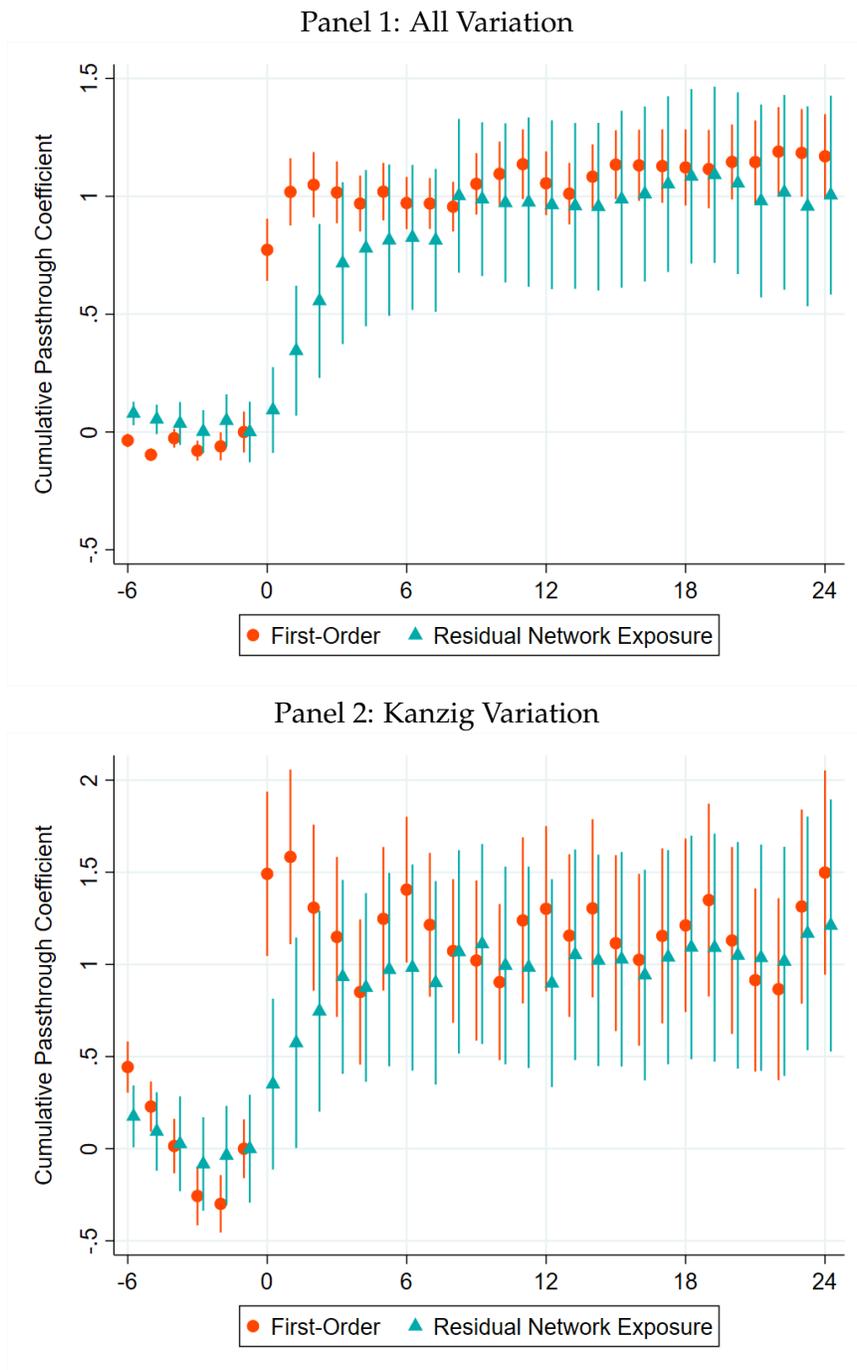
Note: This figure displays the total network share of oil in each industry’s revenues for a selection of industries. Panel 1 displays the twelve industries with the highest network oil shares. Panel 2 displays the twelve industries with the highest third-order oil cost shares (i.e., indirect exposure 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. We see that many sectors primarily exposed to oil only through suppliers have exposures as high as 10%.

Figure 4 – Month-by-Month pass-through of Oil & Non-Oil Commodity Price Changes



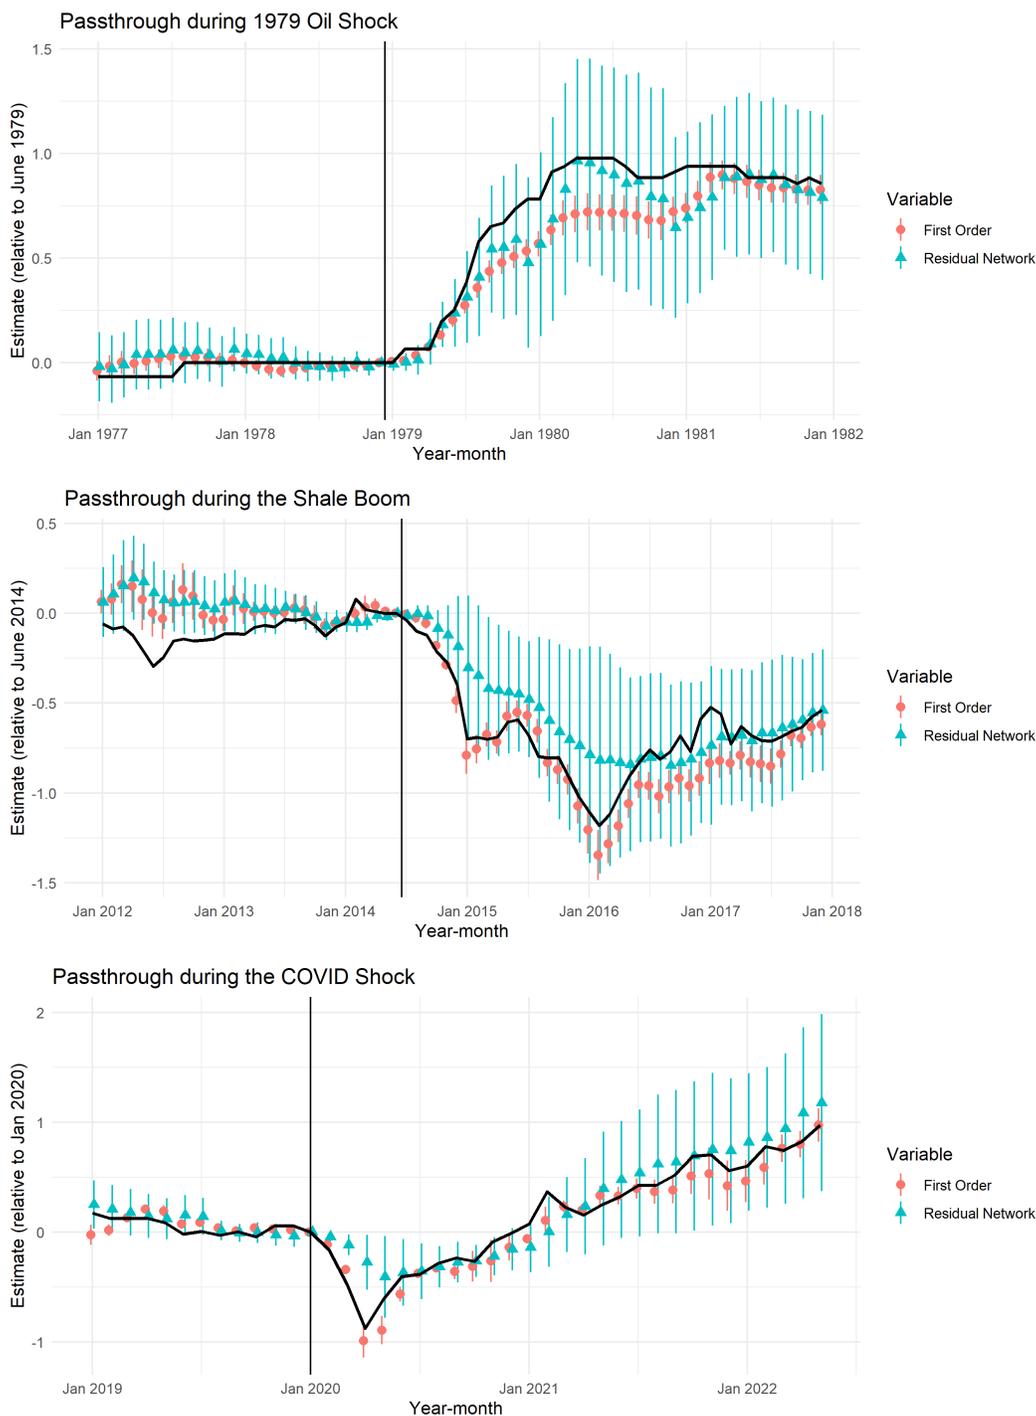
Note: The top panel plots two separate regression specifications, both corresponding to Equation (8). Red coefficients plot monthly price pass-through of network exposure to oil price changes; blue coefficients plot monthly price pass-through of network exposure to non-oil commodity price changes. The 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. Both plots are consistent with full but gradual pass-through of commodity price movements to industry prices.

Figure 5 – Month-by-Month pass-through of Oil Price Changes: Direct vs. Indirect



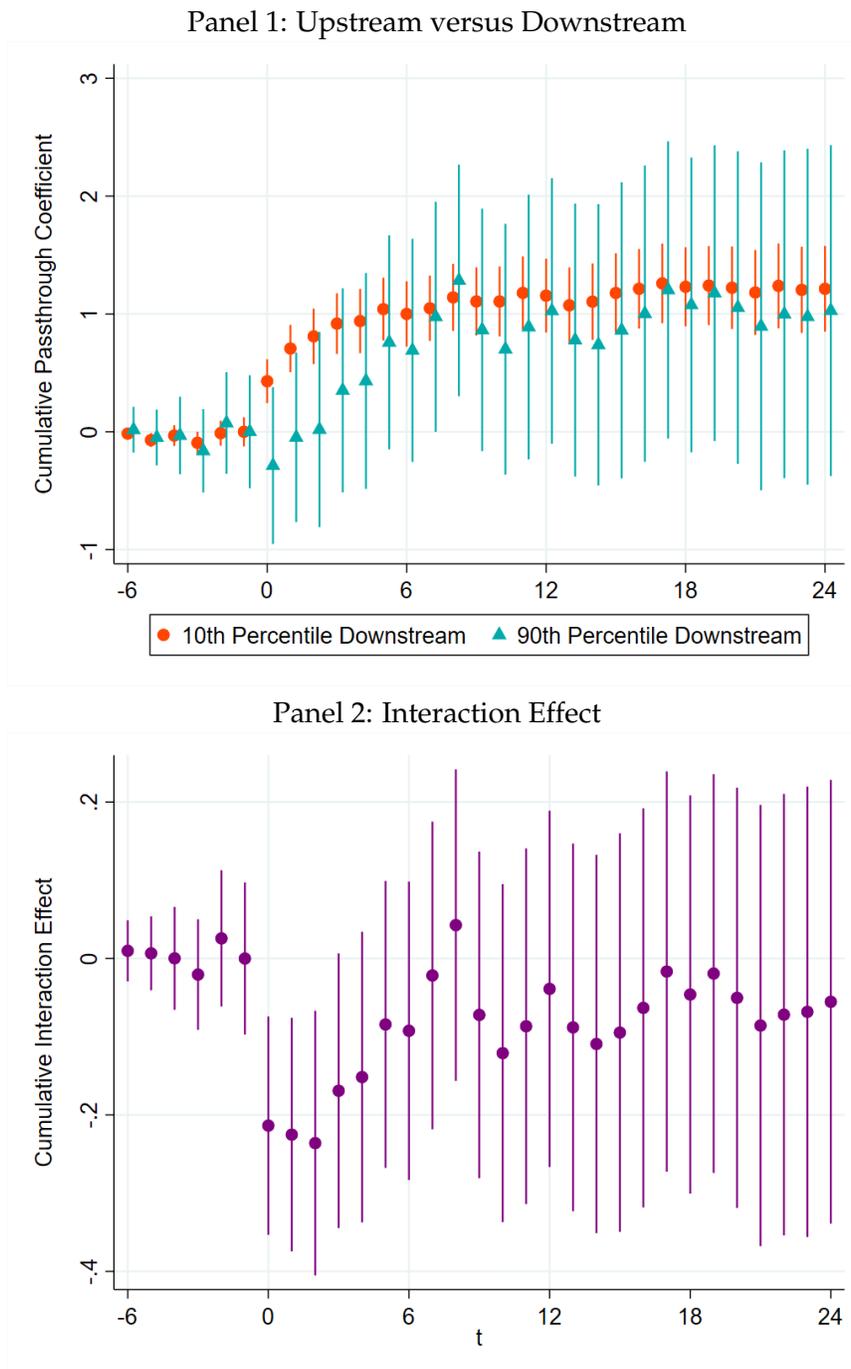
Note: Regression specifications correspond to Equation (9). Red coefficients plot monthly price pass-through of first-order exposure to crude oil price shocks; blue coefficients plot monthly price pass-through of second- and higher-order network exposure to crude oil price shocks. Both plots are consistent with full pass-through of oil price movements to industry prices, with slower pass-through when exposure is indirect.

Figure 6 – Case Studies



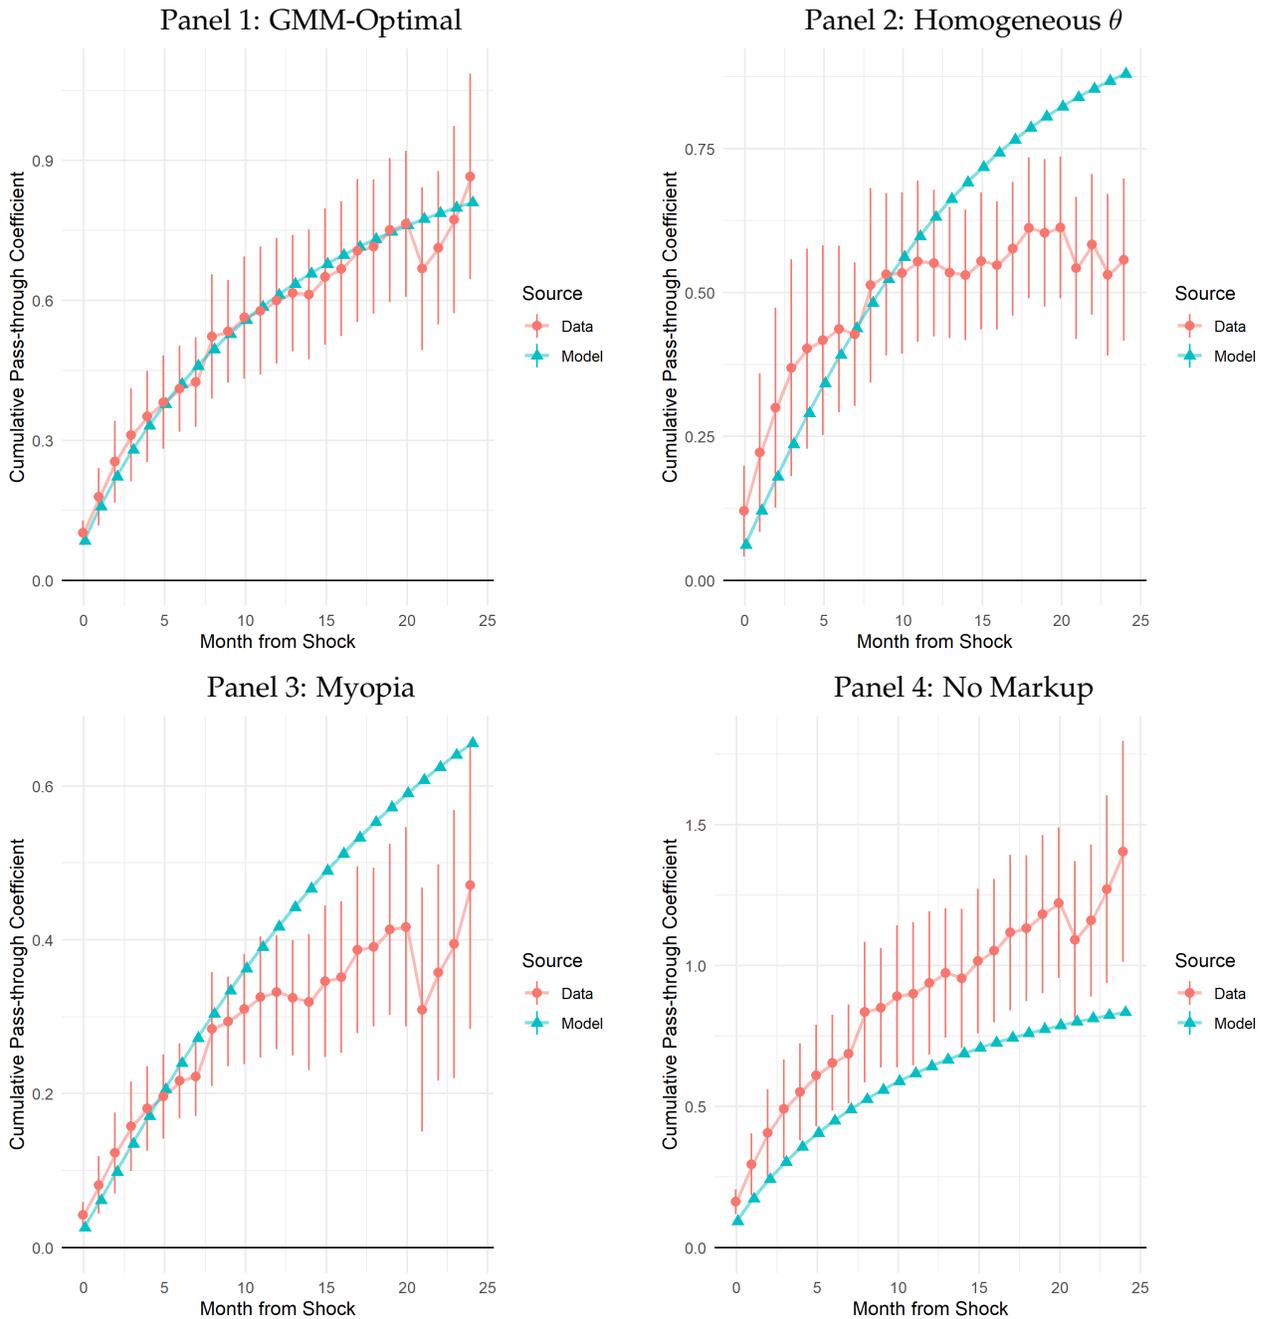
Note: Regression specification corresponds to Equation (11). Black line plots West Texas Intermediate (WTI) crude oil price, red coefficients plot cumulative industry price pass-through of direct/first-order exposure to crude oil price changes, and blue coefficients plot cumulative industry price pass-through of indirect/residual network exposure to crude oil price changes. We confirm our findings of oil price pass-through to industry prices as predicted by the network model. Therefore, our pass-through findings do not vary substantially across large oil shock episodes.

Figure 7 – Month-by-Month pass-through of Oil Price Changes: Upstream vs. Downstream Industries



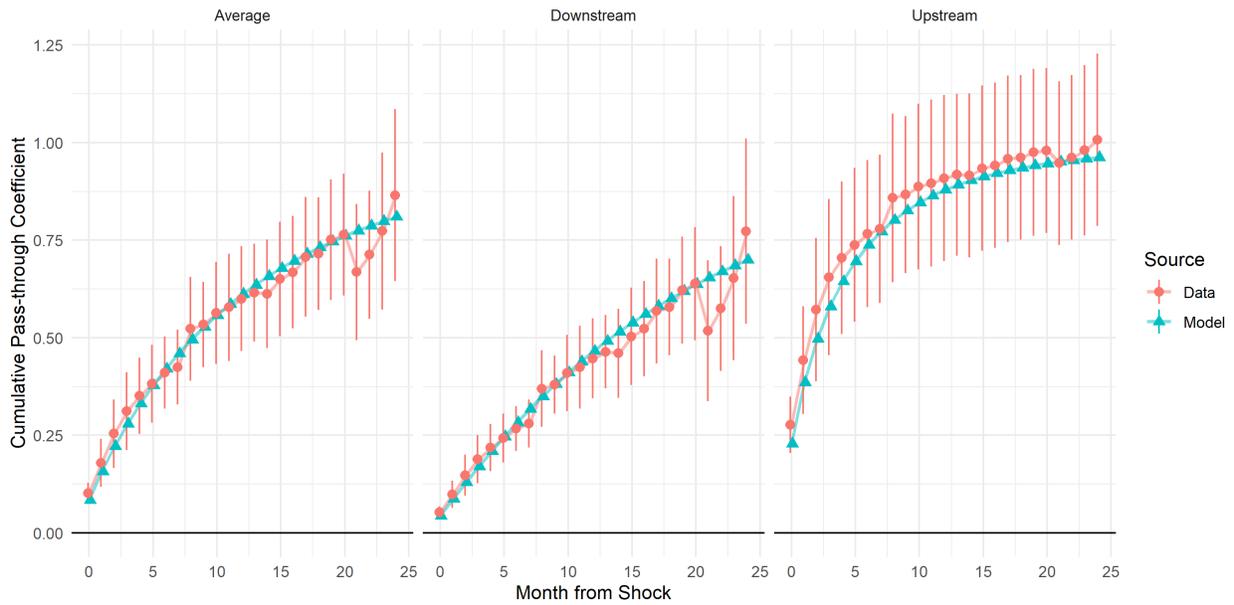
Note: Regression specifications correspond to column (4) of Table 1, with shock terms interacted with our measure of downstreamness. Consequently, the top panel plots monthly cumulative price pass-through of crude oil price shocks for both upstream (10th percentile downstreamness) and downstream (90th percentile downstreamness) firms. The bottom panel plots the cumulated interaction coefficient itself. Downstream firms have delayed pass-through by several months relative to upstream firms, even conditional on frequency of price adjustment – which is typically lower for firms more downstream from oil.

Figure 8 – GMM Results for Industry Pass-through of Oil Price Changes



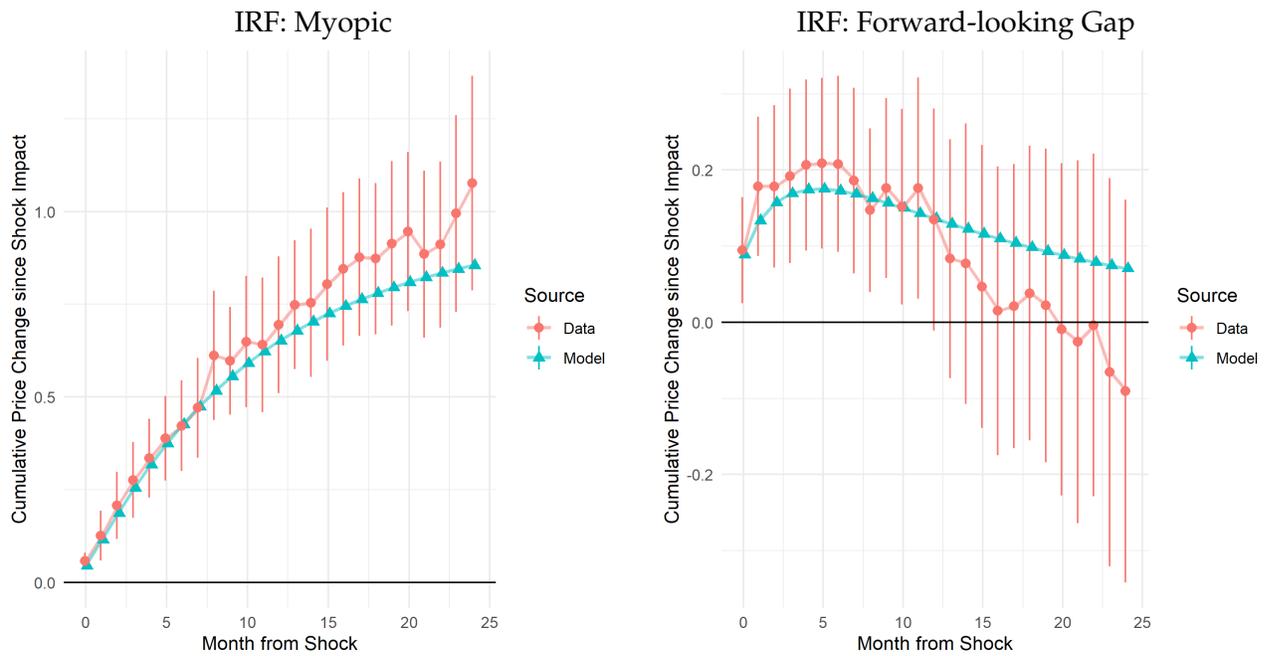
Note: Model fit is good for the GMM-optimal model, as shown in Panel 1. We provide intuition for how parameters are identified and show results for alternative calibrations in Panels 2 - 4. Under alternative calibrations, model fit is meaningfully worse. Panel 2 assesses the experiment in which all industries have the average frequency of price adjustment; the data clearly prefer the model in which more upstream industries have faster frequencies of price adjustment, as we observe in the data. Panel 3 analyzes the case where firms respond myopically to oil price increases; for higher lags, the myopic model predicts more pass-through than the data appears to warrant. Panel 4 assesses the case where firms are competitive instead of pricing with some markup over marginal cost. The competitive case yields too little pass-through at all horizons to be consistent with the data.

Figure 9 – Industry Pass-through of Oil Price Changes, Upstream and Downstream



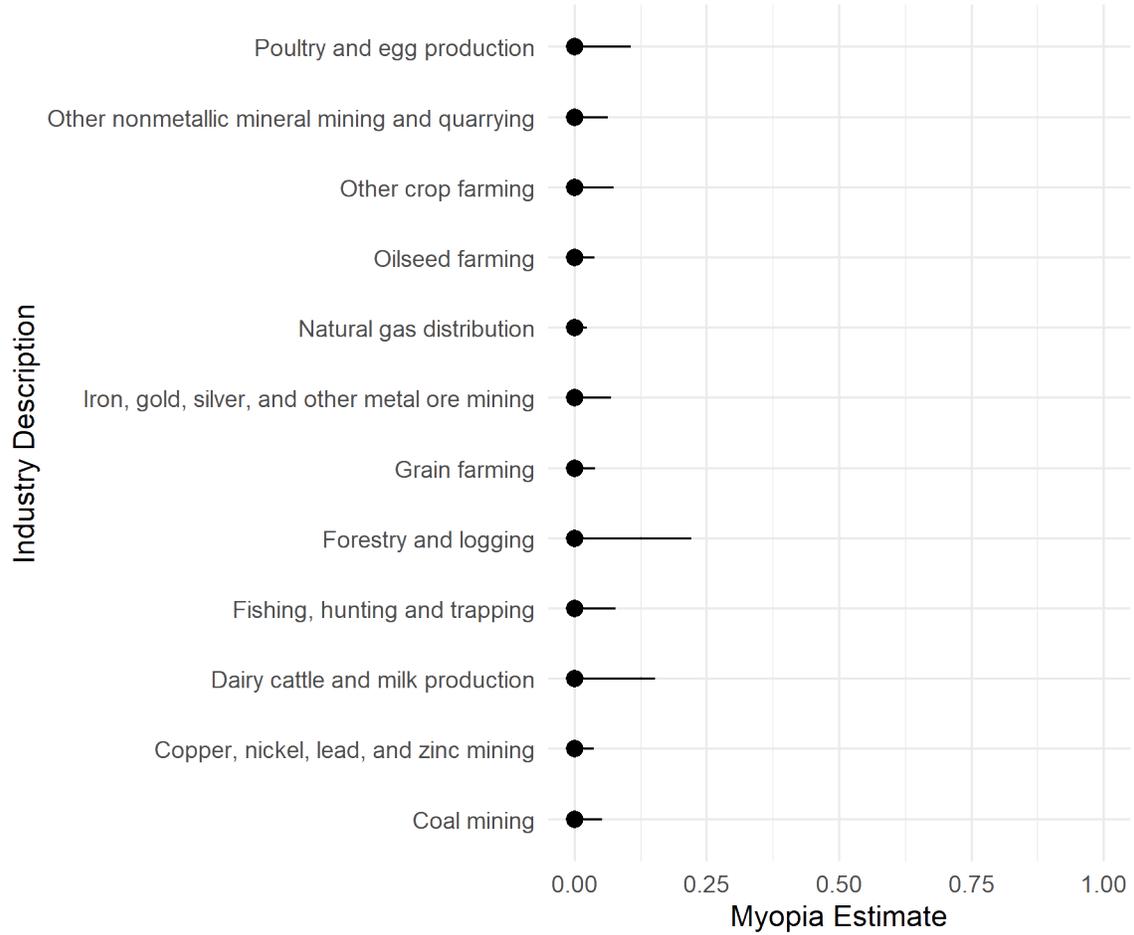
Note: The GMM procedure optimized fit for the average industry (leftmost plot). We visualize how well the fit is for the most upstream and downstream industries in the middle and rightmost plots, finding that model fit remains good. Moreover, we see how much pass-through speed varies with downstreamness. Upstream industries achieve 75% of long-run pass-through in just 5 months, while the average industry requires 20 months, and the most downstream industries have not reached 75% of long-run pass-through even in two years.

Figure 10 – Tests of Forward-lookingness



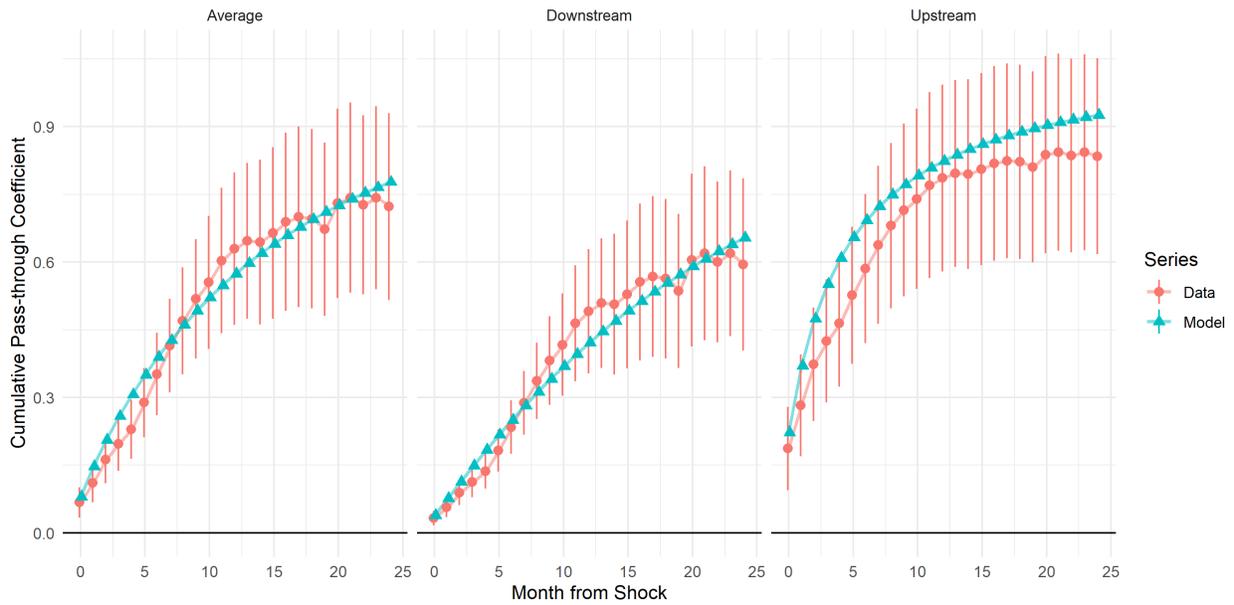
Note: We plot the results of our test that pass-through due to forward-lookingness under rational expectations is present in the data. Statistical significance in the right panel implies we cannot reject that firms are forward-looking about the gradual pass-through of upstream shocks to their marginal costs. The fact that the model lies within the standard error bars in the right figure visualizes that we cannot reject a degree of forward-lookingness consistent with rational expectations. Moreover, we see that in particularly exposed industries, rational expectations provides a pass-through boost of more than 15% of long-run pass-through five months after the oil price change.

Figure 11 – Estimates of Myopia for Non-oil Commodities



Note: We find that complete myopia is required to fit the pass-through of specific non-oil commodity price movements. The standard errors are relatively tight and certainly allow us to reject rational expectations, or a myopia estimate of 1.

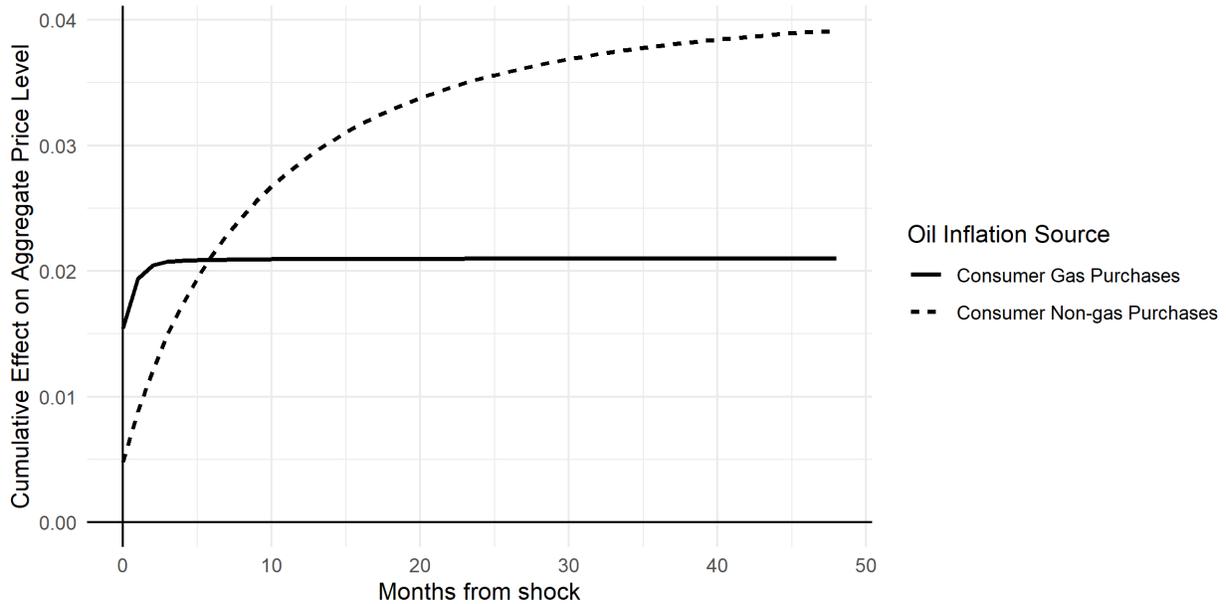
Figure 12 – GMM-Optimal Model, Pooled Non-oil Commodity Price Movements



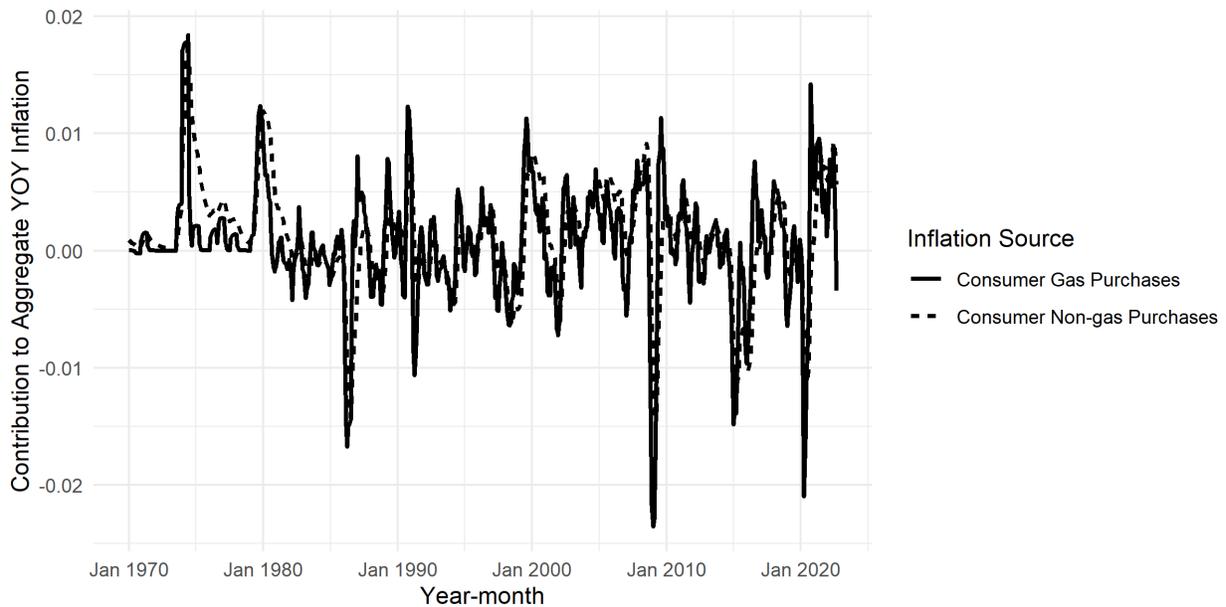
Note: Model fit for pooled commodity price increases is good for the average and most upstream and downstream industries. There is suggestive, but not statistically significant, evidence of slightly lower upstream pass-through in the data relative to the model.

Figure 13 – Inflationary Effects of Oil

Panel 1: Components of Inflation due to Oil, Predicted by the Model

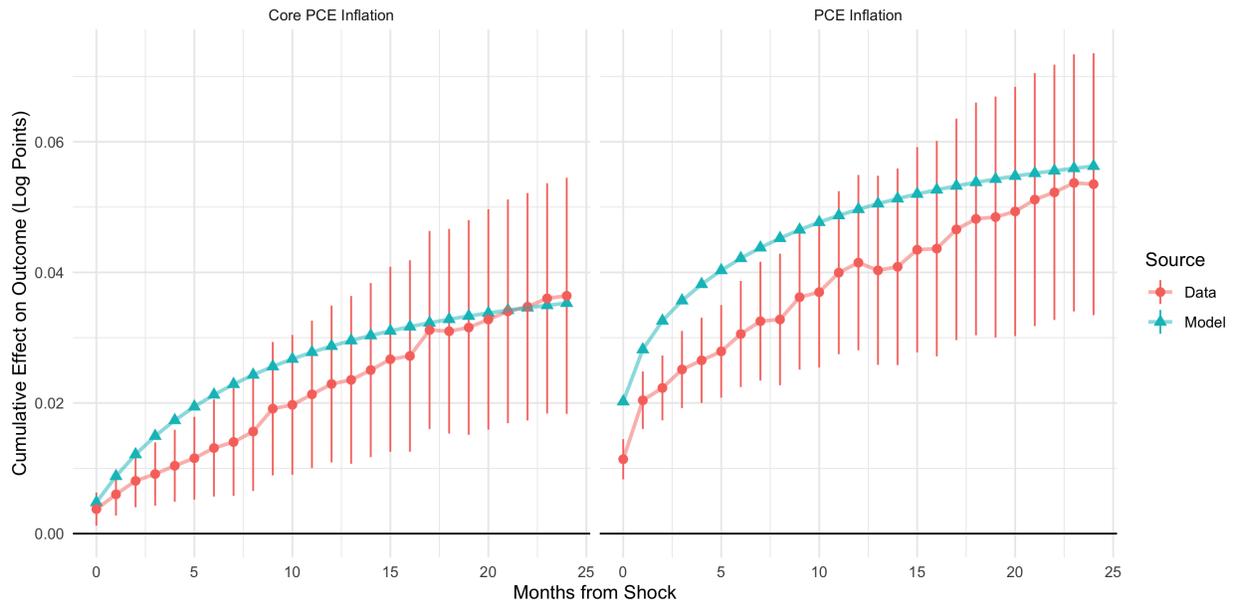


Panel 2: Inflation due to Oil since 1970



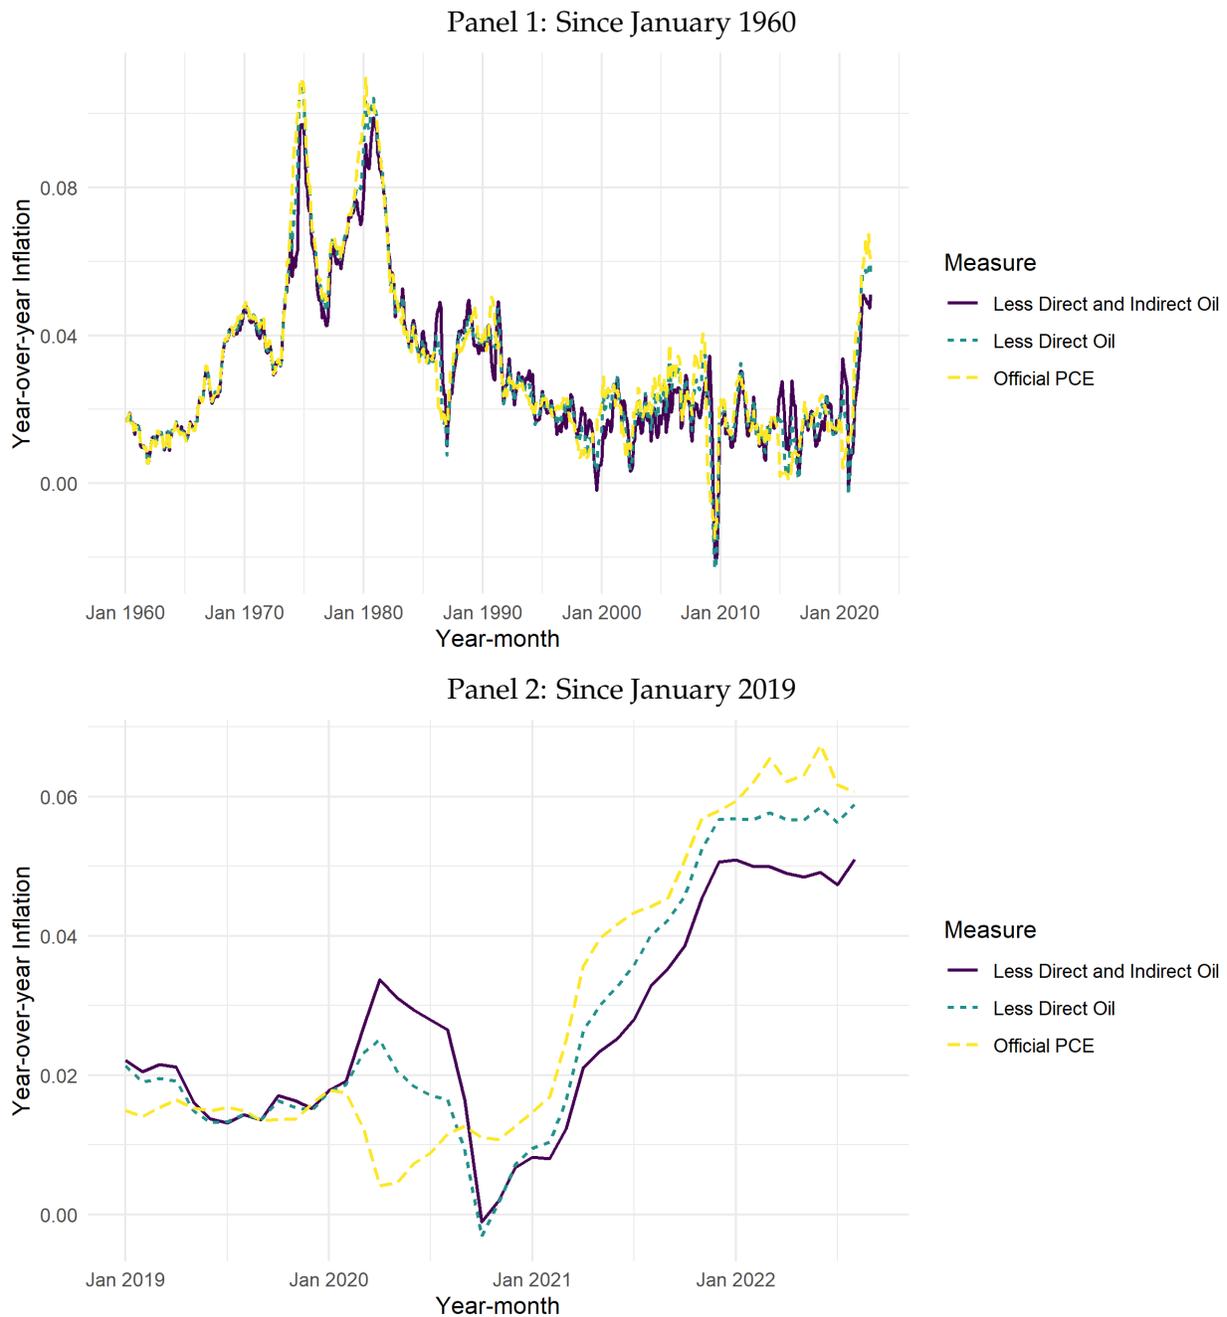
Note: Panel 1 is the model-predicted IRF of aggregate inflation to a log point shock to oil prices. The total inflation response due to a log point oil price increase is 6%, with most of the direct effect realized on impact. The indirect effect is realized slowly and only dominates the direct effect after about 6 months. Panel 2 shows how both components of oil inflation have affected aggregate inflation historically. In the 1970s and 1980s, it is particularly clear that indirect oil inflation has lasting effects on the aggregate price level. In the 1970s, direct and indirect effects of oil contributed to more than 3 percentage points of aggregate, year-over-year inflation.

Figure 14 – Explaining the Predictability of Official Inflation Series



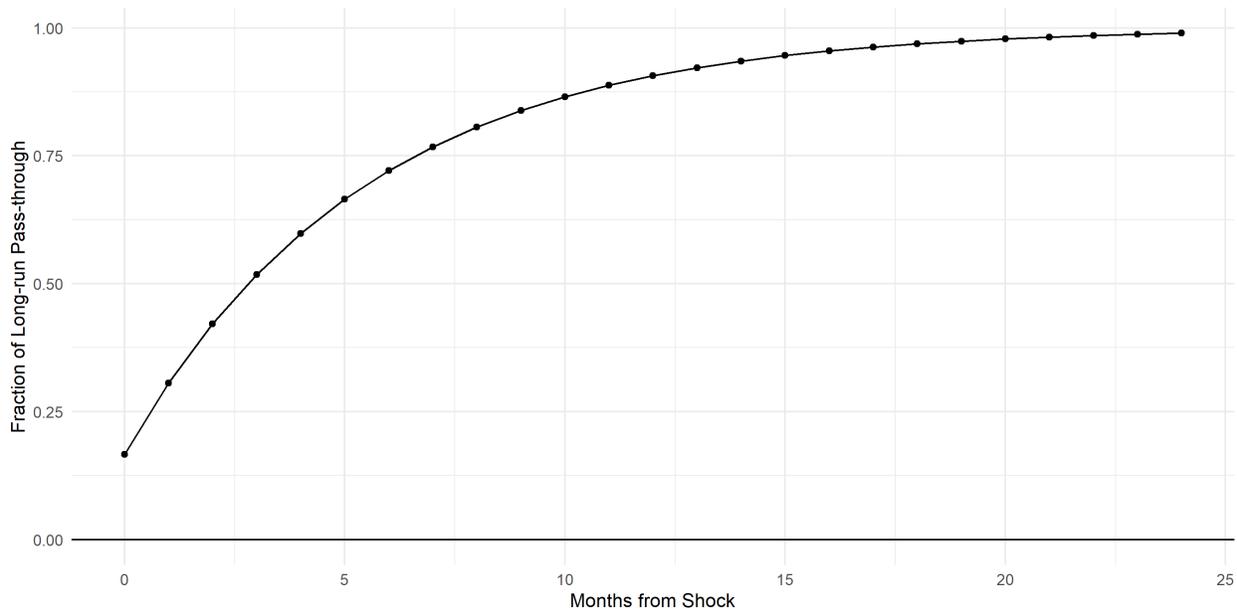
Note: Time series impulse response functions estimated in the data are plotted against the model's predictions for how aggregate inflation, measured in PPIs, should respond to oil price movements. Faster pass-through in the model, combined with the long-run effect being correct, likely implies that it takes some time for producer price (PPI) movements to pass through wholesale and retail prices to consumer prices (CPIs).

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



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. During the current inflationary episode, year-over-year inflation falls to a peak of 5%, down from a peak above 6.5%, when removing all the elements of oil inflation from official PCE inflation. The results also suggest there was inflation early during COVID that was hidden by the large decline in oil prices in the early part of the pandemic; this finding is particularly driven by removing oil's indirect contributions to aggregate inflation.

Figure 16 – The Effect of an Oil Price Increase on the Price Level: “Strawman” Model



Note: We see that the representative industry model calibrated with an average price duration of 5 months generates faster aggregate inflation resulting from oil than is present in our time series analysis depicted in Figure 14. In both the model and the time series, less than 50% of long-run pass-through has occurred after 5 months in Figure 14, while about 65% has already occurred under the representative industry model.

Table 1 – Heterogeneity Analysis

Dependent Variable: Δ PPI	(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*** (0.063)	0.912*** (0.241)	0.349*** (0.097)	0.654*** (0.051)	0.937*** (0.227)	0.365*** (0.097)	0.380*** (0.116)	0.071 (0.224)	0.452** (0.190)
Cost Change*Price Adjust. Half-Life	-0.128*** (0.017)	-0.147*** (0.038)	-0.051*** (0.015)	-0.110*** (0.016)	-0.122*** (0.033)	-0.044*** (0.013)	-0.108*** (0.019)	-0.084** (0.034)	-0.059*** (0.017)
Cost Change*Downstreamness				-0.188*** (0.050)	-0.273** (0.109)	-0.070 (0.055)			
Cost Change*SD[Marginal Cost]							5.152*** (1.702)	17.468*** (3.174)	-3.566 (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.

Table 2 – Predicting Inflation with Network Oil Inflation

Panel 1: Total Inflation				
	(1)	(2)	(3)	(4)
Dependent Variable:	All	Kanzig	All	Kanzig
Total PCE Inflation	Variation	Variation	Variation	Variation
Direct Oil Inflation			0.139*	0.341***
			(0.081)	(0.133)
Indirect Oil Inflation	1.780***	1.362***	1.623***	0.951***
	(0.112)	(0.201)	(0.140)	(0.284)
R-Squared	0.3324	0.3140	0.3359	0.3078
Observations	762	762	762	762

Panel 2: Core Inflation				
	(1)	(2)	(3)	(4)
Dependent Variable:	All	Kanzig	All	Kanzig
Core PCE Inflation	Variation	Variation	Variation	Variation
Direct Oil Inflation			-0.185***	-0.083
			(0.066)	(0.111)
Indirect Oil Inflation	0.929***	0.431**	1.139***	0.530**
	(0.093)	(0.182)	(0.130)	(0.254)
R-Squared	0.1455	0.1036	0.1556	0.1117
Observations	762	762	762	762

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.

A Appendix: Proofs

Proof of Proposition 1:

Proof. Start with the 3 equation model given by equations (5 – log-linear marginal cost), (7 – log-linear reset price), and (6 – log-linear law of motion for prices). In the new steady state, the law of motion implies $\hat{p}_\infty^* = \hat{p}_\infty$. The reset price equation simplifies to $\hat{p}_\infty = \widehat{m}c_\infty$, and the marginal cost equation, using that there are no TFP shocks, simplifies to $\widehat{m}c_\infty = \text{diag}(s_i^L)\hat{w}_\infty + s^Z\hat{p}_Z + \Phi\hat{p}_\infty$. Eliminating marginal cost, we have

$$\hat{p}_\infty = \text{diag}(s_i^L)\hat{w}_\infty + s^Z\hat{p}_Z + \Phi\hat{p}_\infty,$$

which simplifies to

$$\hat{p}_\infty = (\mathbf{I} - \Phi)^{-1}\text{diag}(s_i^L)\hat{w}_\infty + (\mathbf{I} - \Phi)^{-1}s^Z\hat{p}_Z.$$

This was our desired result. □

Proof of Proposition 2:

Proof. First, see subsection A.2, which recasts the model in continuous time. We start with the log-linearized equations for the optimal reset price and the law of motion for industry prices:

$$\begin{aligned}\hat{p}_{i,t}^* &= \psi_i \mathbb{E}_t \int_{\tau=t}^{\infty} e^{-\psi_i(\tau-t)} \widehat{m}c_{i,\tau} d\tau \\ \dot{\hat{p}}_{i,t} &= \phi_i(\hat{p}_{i,t}^* - \hat{p}_{i,t}).\end{aligned}$$

The log-linearization of marginal cost is the same as in the discrete time model. We begin with the myopic case, $\psi_i \rightarrow \infty$ for all i , which is more straightforward than when firms are forward looking. In this case, $\hat{p}_{i,t}^* = \widehat{m}c_{i,t}$. Holding wages and TFP constant, then $\hat{p}_t^* = s^Z\hat{p}_{Z,0} + \Phi\hat{p}_t$. So

$$\dot{\hat{p}}_t = \bar{\phi}(s^Z\hat{p}_{Z,0} - (\mathbf{I} - \Phi)\hat{p}_t).$$

This differential equation solves to

$$\hat{p}_t = (\mathbf{I} - e^{-\mathbf{B}t})(\mathbf{I} - \Phi)^{-1}s^Z\hat{p}_{Z,0},$$

where $\mathbf{B} = \bar{\phi}(\mathbf{I} - \Phi)$. This gives the desired results for myopia under heterogeneous and homogeneous frequencies of price adjustment.

Now we turn to the case of (partially) forward-looking firms. Denote by D the time derivative operator. Rewrite the log-linearized model equations in partial equilibrium (holding wages and

TFP constant, which simplifies the marginal cost equation) as

$$D\hat{\mathbf{p}}_t = \bar{\phi}(\hat{\mathbf{p}}_t^* - \hat{\mathbf{p}}_t), \quad D\hat{\mathbf{p}}_t^* = \bar{\psi}(-\Phi\hat{\mathbf{p}}_t - \mathbf{s}^Z\hat{p}_{Z,0} + \hat{\mathbf{p}}_t^*).$$

The first equation can be rewritten as $\hat{\mathbf{p}}_t^* = \bar{\phi}^{-1}(D\mathbf{I} + \bar{\phi})\hat{\mathbf{p}}_t = (\bar{\phi}^{-1}D + \mathbf{I})\hat{\mathbf{p}}_t$. The second equation can be rewritten as $(D\mathbf{I} - \bar{\psi})\hat{\mathbf{p}}_t^* = -\bar{\psi}(\Phi\hat{\mathbf{p}}_t + z)$, where $z \equiv \mathbf{s}^Z\hat{p}_{Z,0}$. Plugging in the solution for $\hat{\mathbf{p}}_t^*$, we have

$$(D\mathbf{I} - \bar{\psi})(\bar{\phi}^{-1}D + \mathbf{I})\hat{\mathbf{p}}_t = -\bar{\psi}(\Phi\hat{\mathbf{p}}_t + z).$$

Rewrite this by grouping the non time derivative terms on the RHS:

$$(\bar{\phi}^{-1}D^2 + D\mathbf{I} - \bar{\psi}\bar{\phi}^{-1}D)\hat{\mathbf{p}}_t = \bar{\psi}\hat{\mathbf{p}}_t - \bar{\psi}(\Phi\hat{\mathbf{p}}_t + z) = \bar{\psi}[(\mathbf{I} - \Phi)\hat{\mathbf{p}}_t - z].$$

Now, we plug in our conjecture that $\hat{\mathbf{p}}_t = (\mathbf{I} - e^{-Bt})(\mathbf{I} - \Phi)^{-1}z$. First, note that the RHS simplifies nicely:

$$\begin{aligned} \bar{\psi}[(\mathbf{I} - \Phi)\hat{\mathbf{p}}_t - z] &= \bar{\psi}[(\mathbf{I} - \Phi)(\mathbf{I} - \Phi)^{-1}z - (\mathbf{I} - \Phi)e^{-Bt}(\mathbf{I} - \Phi)^{-1}z - z] \\ &= -\bar{\psi}(\mathbf{I} - \Phi)e^{-Bt}(\mathbf{I} - \Phi)^{-1}z. \end{aligned}$$

The LHS is

$$(\bar{\phi}^{-1}D^2 + D\mathbf{I} - \bar{\psi}\bar{\phi}^{-1}D)(\mathbf{I} - e^{-Bt})(\mathbf{I} - \Phi)^{-1}z = -(\bar{\phi}^{-1}D^2 + D\mathbf{I} - \bar{\psi}\bar{\phi}^{-1}D)e^{-Bt}(\mathbf{I} - \Phi)^{-1}z.$$

Therefore, we have

$$(\bar{\phi}^{-1}D^2 + D\mathbf{I} - \bar{\psi}\bar{\phi}^{-1}D)e^{-Bt}(\mathbf{I} - \Phi)^{-1}z = \bar{\psi}(\mathbf{I} - \Phi)e^{-Bt}(\mathbf{I} - \Phi)^{-1}z.$$

We take derivatives on the LHS:

$$(\bar{\phi}^{-1}D^2 + D\mathbf{I} - \bar{\psi}\bar{\phi}^{-1}D)e^{-Bt} = \bar{\phi}^{-1}e^{-Bt}\mathbf{B}^2 - e^{-Bt}\mathbf{B} + \bar{\psi}\bar{\phi}^{-1}e^{-Bt}\mathbf{B}.$$

Now, $e^{-Bt} = \sum_{k=0}^{\infty} \frac{1}{k!}\mathbf{B}^k$ by definition, and so it is clear that \mathbf{B} and \mathbf{B}^2 commute with e^{-Bt} . Therefore, we have

$$(\bar{\phi}^{-1}\mathbf{B} - (\mathbf{I} - \bar{\psi}\bar{\phi}^{-1}))\mathbf{B}e^{-Bt}(\mathbf{I} - \Phi)^{-1}z = \bar{\psi}(\mathbf{I} - \Phi)e^{-Bt}(\mathbf{I} - \Phi)^{-1}z.$$

So the conjectured solution works if

$$(\bar{\phi}^{-1}\mathbf{B} - (\mathbf{I} - \bar{\psi}\bar{\phi}^{-1}))\mathbf{B} = \bar{\psi}(\mathbf{I} - \Phi).$$

Under rational expectations and a steady state interest rate of 0, $\bar{\phi} = \bar{\psi}$, so $\mathbf{B} = (\bar{\phi}^2(\mathbf{I} - \Phi))^{1/2}$. When the frequency of price adjustment does not vary by sectors, $\bar{\phi} = \phi\mathbf{I}$ and $\bar{\psi} = \psi\mathbf{I}$ commute

with all matrices, and so the matrix equation can be rewritten as

$$\left(\mathbf{B} - \frac{\psi - \phi}{2} \right)^2 = \left(\frac{\psi + \phi}{2} \right)^2 - \phi\psi\Phi,$$

or

$$\mathbf{B} = \frac{\psi - \phi}{2} \mathbf{I} + \left(\left(\frac{\psi + \phi}{2} \right)^2 \mathbf{I} - \phi\psi\Phi \right)^{1/2},$$

where we take the roots with positive eigenvalues. \square

Proof of Proposition 3:

Proof. Start with the duration definition, $D_i = \int_0^\infty (1 - b_{i,s}) ds$, and the result from Proposition 2 that $b_{i,t} \approx 1 - \frac{(e^i)' e^{-\mathbf{B}t} \hat{\mathbf{p}}_\infty}{(e^i)' \hat{\mathbf{p}}_\infty}$. Then

$$D_i \approx \int_0^\infty \frac{(e^i)' e^{-\mathbf{B}s} \hat{\mathbf{p}}_\infty}{(e^i)' \hat{\mathbf{p}}_\infty} ds = \frac{(e^i)' \left(\int_0^\infty e^{-\mathbf{B}s} ds \right) \hat{\mathbf{p}}_\infty}{(e^i)' \hat{\mathbf{p}}_\infty} = \frac{(e^i)' \mathbf{B}^{-1} \hat{\mathbf{p}}_\infty}{(e^i)' \hat{\mathbf{p}}_\infty}.$$

Plug in the result for myopia under homogeneous θ , $\mathbf{B} = \phi(\mathbf{I} - \Phi)$, and use that $\hat{\mathbf{p}}_\infty = (\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z$. Then, under myopia,

$$D_i \approx \frac{1}{\phi} \frac{(e^i)' (\mathbf{I} - \Phi)^{-2} \mathbf{s}^Z}{(e^i)' (\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z}.$$

Now, just as in scalar case, we have

$$\begin{aligned} (\mathbf{I} - \Phi)^{-2} &= \mathbf{I} + 2\Phi + 3\Phi^2 + \dots \\ &= \sum_{n=1}^{\infty} n\Phi^{n-1}, \end{aligned}$$

which was the desired result.

Next, plug in the result for rational expectations and a 0 interest rate under homogeneous θ , $\mathbf{B} = \phi(\mathbf{I} - \Phi)^{1/2}$. Then

$$D_i \approx \frac{1}{\phi} \frac{(e^i)' (\mathbf{I} - \Phi)^{-3/2} \mathbf{s}^Z}{(e^i)' (\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z}.$$

In this case, term n of $(\mathbf{I} - \Phi)^{-3/2}$ is $(-1)^{n-1} \binom{-3/2}{n-1} \Phi^{n-1}$, which is less than n for $n > 1$. \square

A.1 Eigenvalue Solution of the Model (with Decomposition)

The compact version of the log-linearized optimal reset price equation (7) is

$$\hat{p}_{i,t}^* = \theta_i \beta m_f E_t[\hat{p}_{i,t+1}^*] + (1 - \theta_i \beta m_f) \widehat{m}c_{i,t}.$$

The log-linearized law of motion for prices, equation (6), was

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

Combining and restricting to $m_f > 0$,

$$E_t[\hat{p}_{i,t+1}] = \frac{1 + \theta_i^2 \beta m_f}{\theta_i \beta m_f} \hat{p}_{i,t} - \frac{1}{\beta m_f} \hat{p}_{i,t-1} - \frac{(1 - \theta_i)(1 - \theta_i \beta m_f)}{\theta_i \beta m_f} \widehat{m}c_{i,t}.$$

Stacking across industries,

$$E_t[\hat{\mathbf{p}}_{t+1}] = \text{diag} \left(\frac{1 + \theta_i^2 \beta}{\theta_i \beta m_f} \right) \hat{\mathbf{p}}_t - \frac{1}{\beta m_f} \hat{\mathbf{p}}_{t-1} - \text{diag} \left(\frac{(1 - \theta_i)(1 - \theta_i \beta m_f)}{\theta_i \beta m_f} \right) \widehat{\mathbf{m}}c_t.$$

Recall that

$$\widehat{\mathbf{m}}c_t = \Phi \hat{\mathbf{p}}_t + \mathbf{s}^Z \hat{p}_{Z,t} + \text{diag}(s_i^L) \hat{\mathbf{w}}_t - \hat{\mathbf{a}}_t.$$

Therefore,

$$E_t[\hat{\mathbf{p}}_{t+1}] = \left(\text{diag} \left(\frac{1 + \theta_i^2 \beta m_f}{\theta_i \beta m_f} \right) - \text{diag} \left(\frac{(1 - \theta_i)(1 - \theta_i \beta m_f)}{\theta_i \beta m_f} \right) \Phi \right) \hat{\mathbf{p}}_t - \frac{1}{\beta m_f} \hat{\mathbf{p}}_{t-1} \\ - \text{diag} \left(\frac{(1 - \theta_i)(1 - \theta_i \beta m_f)}{\theta_i \beta m_f} \right) (\mathbf{s}^Z \hat{p}_{Z,t} + \text{diag}(s_i^L) \hat{\mathbf{w}}_t - \hat{\mathbf{a}}_t).$$

Now, define

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

Then

$$E_t \hat{\mathbf{x}}_{t+1} = \mathbf{B}_x \hat{\mathbf{x}}_t + \mathbf{B}_e \hat{\mathbf{e}}_t,$$

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

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

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

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

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{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

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

which can be forward solved to get

$$\tilde{\mathbf{x}}_{1,t} = - \sum_{j=0}^{\infty} (\Lambda_1^{-1})^{j+1} \tilde{\mathbf{B}}_1 E_t[\hat{e}_{t+j}] + \lim_{j \rightarrow \infty} (\Lambda_1^{-1})^j E_t[\tilde{\mathbf{x}}_{1,t+j}].$$

The required transversality condition is

$$\lim_{j \rightarrow \infty} (\Lambda_1^{-1})^j E_t[\tilde{\mathbf{x}}_{1,t+j}] = 0.$$

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

$$\sum_{j=0}^{\infty} (\Lambda_1^{-1})^{j+1} \tilde{\mathbf{B}}_1 E_t[\hat{e}_{t+j}]$$

is finite. Then

$$\tilde{\mathbf{x}}_{1,t} = - \sum_{j=0}^{\infty} (\Lambda_1^{-1})^{j+1} \tilde{\mathbf{B}}_1 E_t[\hat{e}_{t+j}].$$

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

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

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

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

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, $\hat{e}_t = \hat{e}_0$ 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

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

when the commodity price is shocked by 1 log point. This is long-run pass-through from Proposition 1. We will focus on the case of shocking the commodity price while leaving TFP and wages constant. In this case,

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

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

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

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

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

with

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

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.2 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

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

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

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

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

defined in our setting as

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

where \mathbb{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, we have $SDF_{t,\tau}^{SS} = e^{-\rho(\tau-t)}$, where ρ is the discount rate. Therefore, the above equation simplifies to

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

or (setting $\psi_i = \phi_i + \rho + \tilde{m}_f$)

$$\hat{p}_{i,t}^* = \psi_i \mathbb{E}_t \int_{\tau=t}^{\infty} e^{-\psi_i(\tau-t)} \widehat{m}c_{i,\tau} d\tau.$$

The industry price index satisfies

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

where the dot notation denotes a time derivative.

A.3 Additional Characterizations in the 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

$$p_{n,t}^* = \psi \mathbb{E}_t \int_{\tau=t}^{\infty} e^{-\psi(\tau-t)} [sp_{n-1,t}] d\tau$$

$$\dot{p}_{n,t} = \phi (p_{n,t}^* - p_{n,t}),$$

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. In the case of $\psi \rightarrow \infty$, we get

$$p_{n,t}^* = sp_{n-1,t}$$

$$\dot{p}_{n,t} = \phi (p_{n,t}^* - p_{n,t}),$$

or just

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

We will be able to find a closed-form solution. 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 $X \rightarrow \infty$,

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

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 so $t_n(X)$ is characterized by

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

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

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

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

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

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,

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

etc. Again, we define $t_n^F(X)$ implicitly as satisfying

$$b_{n,t_n^F(X)} = 1 - e^{-X}.$$

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

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

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 com-

modity 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, $(I - \Phi)^{-1}$. 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 to maximize consistency with our regressions, which include a time fixed effect. We then split the data into 100 quantiles of industry oil cost changes and show the results for a variety of time horizons. Figure G.1 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 G.2, 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.

C.2 Excluding Wage Controls

It is possible to replicate our main results after estimating regression specifications that exclude our controls for general equilibrium wage changes. 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 include the wage changes. However, these wage changes are largely uncorrelated with price changes, and results are very similar if they are excluded. 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 G.3 shows that this scarcely changes the findings.

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 shock exposure to show that the main findings on oil are not driven by these related sectors. These findings are displayed in Figure G.4; again, the main results are scarcely changed.

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 G.5. Once again, the pace of pass-through is little changed from baseline. The extent of indirect passthrough when all oil price variation is used is slightly depressed relative to baseline. When using the Kanzig variation, this is no longer the case, and both the extent and pace of passthrough are highly consistent with the baseline regressions including capital in the denominator.

C.5 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 permute 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 G.6 correspond to the regression specification given by Equation (8) 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 (9) 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.

C.6 Higher-order pass-through

In our structural work, we noted that different assumptions about the aggregate markup meaningfully affected overall pass-through of oil price movements as predicted by our model. We also noted that our reduced-form results were much less influenced 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 G.7, we confirm this hypothesis. When we include payments to capital in industries' costs, third- and higher-order pass-through of oil price movements maxes out around 2, and fourth- and higher-order pass-through of oil price movements maxes out around 2.5. This means that industries 3 and higher or 4 and higher links away from oil in the supply chain increase prices 2 and 2.5 times more, respectively, than predicted under Proposition 1. These predictions are brought back into a reasonable range in the structural model,

where a markup higher than 1 is estimated.

D Appendix: Robustness of Structural Estimation

D.1 Including Petroleum Refineries

Our step 1 estimates including petroleum refineries are $\gamma = 0$ (SE = 0.134), $m_f = 1$ (SE = 0.248), and $\hat{\mu} = 1.158$ (SE = 0.104). These are very similar to our step 1 estimates without petroleum refineries. The results from 2-step GMM, where we use the first step estimates to form an optimal weight matrix, are also quite similar regardless whether we include or exclude petroleum refineries: $\gamma = 0$ (SE = 0.062), $m_f = 1$ (SE = 0.213), and $\hat{\mu} = 1.019$ (SE = 0.042). Just as before, the only difference in estimates is a smaller value for the markup and smaller standard errors.

Now, the visualization of our results changes with petroleum refineries because they are such an outlier. So we reproduce our main figures displaying our oil results in Figures [G.8](#), [G.9](#), and [G.10](#).

E Appendix: Incorporating Oil Futures

In this appendix, we first extend our model solution to incorporate use of commodity futures to measure expected commodity spot prices in the future. The analysis will be robust to time-invariant but maturity-varying deviations from the expectations hypothesis for futures, which says in its simplest form that a commodity futures price at time t and maturity m equals the date t expectation of the commodity spot price at period $t + m$.

We then empirically test whether our GMM-optimal model in the main body of the paper, which uses rational expectations (i.e., no myopia at all in the degree of forward-lookingness) and assumes oil prices follow a random walk, maintains an adequate fit of the data when expectations are instead measured directly using oil futures. We also test additional moments of interest, such as whether firms pass-through expected changes in futures prices, holding the spot price constant. We view this as a method of asking whether “forward-guidance” has predictive power for firm pricing.

Finally, we return to our analysis permitting an arbitrary lag structure, using a regression equation like those in section [5](#). Again, we confirm that firms pass through changes in expected future oil prices, holding the spot price constant.

E.1 Model Solution with Futures

As we showed in detail in Appendix A.1, the model can be written as

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

where

$$\hat{\mathbf{x}}_t = \begin{bmatrix} \hat{\mathbf{p}}_{t-1} \\ \hat{\mathbf{p}}_t \end{bmatrix}, \quad \hat{\mathbf{e}}_t = \begin{bmatrix} \hat{\mathbf{a}}_t \\ \hat{\mathbf{w}}_t \\ \hat{p}_{Z,t} \end{bmatrix},$$

and the matrices \mathbf{B} come from the equation²¹

$$m_f \mathbb{E}_t \hat{\mathbf{p}}_{t+1} = \text{diag} \left(\frac{1 + \theta_i^2 \delta}{\theta_i \delta} \right) \hat{\mathbf{p}}_t - \frac{1}{\delta} \hat{\mathbf{p}}_{i,t} - \text{diag} \left(\frac{(1 - \theta_i \delta m_f)(1 - \theta_i)}{\theta_i \delta} \right) (\Phi \hat{\mathbf{p}}_t + \mathbf{s}^Z \hat{p}_{Z,t} + \text{diag}(s_i^L) \hat{\mathbf{w}}_t - \hat{\mathbf{a}}_t)$$

and using the typical method to create a first-order difference equation from the above second-order equation. If the eigendecomposition of \mathbf{B}_x is $\mathbf{B}_x = \mathbf{V} \Lambda \mathbf{V}^{-1}$, and we define $\tilde{\mathbf{x}}_t = \mathbf{V}^{-1} \hat{\mathbf{x}}_t$ and $\tilde{\mathbf{B}}_e = \mathbf{V}^{-1} \mathbf{B}_e$, then

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

where the elements of the diagonal matrix Λ_1 are all greater than 1. Denote

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

Proposition 4. *The solution for prices as a function of the state variable (lagged prices) and the shock vector is*

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

where

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

Now, when we shock oil (and oil futures) we will be setting all wage and TFP dimensions of the shock $\hat{\mathbf{e}}_{t+j}$ equal to 0. In this case,

$$\tilde{\mathbf{x}}_{1,t} = \sum_{j=0}^{\infty} (\Lambda_1^{-1})^{j+1} \tilde{\mathbf{B}}_{1,Z} \mathbb{E}_t[\hat{p}_{Z,t+j}],$$

where the column vector $\tilde{\mathbf{B}}_{1,Z}$ is the appropriate subset of the matrix $\tilde{\mathbf{B}}_1$, namely all rows and the last column, given the placement of $\hat{p}_{Z,t}$ in $\hat{\mathbf{e}}_t$.

²¹We restrict here to $m_f \in (0, 1]$ so that we can divide by it. This isn't necessary if we instead wrote the model without expectations of $\hat{\mathbf{x}}_{t+1}$ on the left-hand side, but I did it this way first, and it doesn't affect our ability to solve or estimate the model, since the solution is still well-defined for very small values of m_f .

Now, it is instructive to write the solution in first-differences form. The price component is straightforward, but the first-difference of $\tilde{x}_{1,t}$ is more interesting. We have

$$\tilde{x}_{1,t} - \tilde{x}_{1,t-1} = \sum_{j=0}^{\infty} (\Lambda_1^{-1})^{j+1} \tilde{B}_{1,Z} (E_t[p_{Z,t+j}] - E_{t-1}[p_{Z,t-1+j}]),$$

where we can write p_Z without a hat because the steady-state component is not time-varying. The difference in expectations can be decomposed as follows:

$$E_t[\hat{p}_{Z,t+j}] - E_{t-1}[\hat{p}_{Z,t-1+j}] = E_t[\Delta P_{Z,t+j}] + (\Delta E_t)[\hat{p}_{Z,t-1+j}].$$

In this formulation, we see that the relevant oil shocks for pricing behavior include the current and expected future changes in oil prices, $E_t[\Delta P_{Z,t+j}]$, as well as the news received about oil prices, measured by the changes in expectations $(\Delta E_t)[\hat{p}_{Z,t-1+j}]$.

Now we will assume log commodity futures prices satisfy $f_{Z,t,m} = -c_m + E_t[p_{Z,t+m}]$, so that the log price of an oil future at period t with maturity m is the period t expectation of log oil prices m months in the future, minus a composite risk-premium/opportunity cost, cost of carry, and convenience yield component c_m which may vary with the maturity m but not the time period t . We say $c_0 = 0$ because the futures price at a maturity of 0 is known and equal to the spot price. It follows from these assumptions that

$$f_{Z,t,m-1} = f_{Z,t-1,m} + c_m - c_{m-1} + (\Delta E_t)[p_{Z,t+m-1}],$$

so that the futures price at maturity $m - 1$ today is what the market thought the price would be last period, plus a change due to the maturity evolving (which alters the total cost of carry, convenience yield, etc.), plus any news received about the oil price. We can apply our assumption to derive

$$\begin{aligned} E_t[\hat{p}_{Z,t+m}] - E_{t-1}[\hat{p}_{Z,t-1+m}] &= f_{Z,t,m} + c_m - (f_{Z,t-1,m} + c_m) \\ &= f_{Z,t,m} - f_{Z,t-1,m}. \end{aligned}$$

Therefore, under our assumption that c_m is not time-varying, the change in futures prices at fixed maturities can be used to measure our difference of interest. So

$$\Delta \tilde{x}_{1,t} = \sum_{m=0}^{\infty} (\Lambda_1^{-1})^{m+1} \tilde{B}_{1,Z} \Delta f_{Z,t,m},$$

where the time difference operator holds the maturity fixed.

Proposition 5. *Under our assumption about how commodity futures prices relate to expectations of future*

commodity prices, the solution in Proposition 4 can be written in first-differences as

$$\Delta \mathbf{p}_t = V_{22} V_{12}^{-1} \Delta \mathbf{p}_{t-1} - (V_{21} - V_{22} V_{12}^{-1} V_{11}) \sum_{m=0}^{\infty} (\Lambda_1^{-1})^{m+1} \tilde{\mathbf{B}}_{1,Z} \Delta f_{Z,t,m}.$$

Now, we can add a lag structure when we perform our empirical tests. Define $A_1 = V_{22} V_{12}^{-1}$ and $A_2 = -(V_{21} - V_{22} V_{12}^{-1} V_{11})$. Then

$$\begin{aligned} \Delta \mathbf{p}_t &= A_1 \Delta \mathbf{p}_{t-1} + A_2 \sum_{m=0}^{\infty} (\Lambda_1^{-1})^{m+1} \tilde{\mathbf{B}}_{1,Z} \Delta f_{Z,t,m} \\ &= A_1^2 \Delta \mathbf{p}_{t-2} + A_1 A_2 \sum_{m=0}^{\infty} (\Lambda_1^{-1})^{m+1} \tilde{\mathbf{B}}_{1,Z} \Delta f_{Z,t-1,m} \\ &\quad + A_2 \sum_{m=0}^{\infty} (\Lambda_1^{-1})^{m+1} \tilde{\mathbf{B}}_{1,Z} \Delta f_{Z,t,m}, \end{aligned}$$

and, more generally,

$$\Delta \mathbf{p}_t = A_1^{H+1} \Delta \mathbf{p}_{t-(H+1)} + \sum_{h=0}^H A_1^h A_2 \sum_{m=0}^{\infty} (\Lambda_1^{-1})^{m+1} \tilde{\mathbf{B}}_{1,Z} \Delta f_{Z,t-h,m}.$$

We have found that, for H large enough, the state variable $\Delta \mathbf{p}_{t-(H+1)}$ should not matter much for pricing in period t (formally, $A_1^{H+1} \approx 0$), unless it is explosively large (which is not the case in the data). Therefore our main pass-through object for empirical work will be

Definition 4. Under model calibration α , prices in industry i react to a change in the commodity future h periods ago at maturity m according to

$$passthrough_{i,h,m}(\alpha) = \left[A_1(\alpha)^h A_2(\alpha) (\Lambda_1(\alpha)^{-1})^{m+1} \tilde{\mathbf{B}}_{1,Z}(\alpha) \right]_i.$$

The following corollary results:

Corollary 2. For large H ,

$$\Delta p_{i,t} \approx \sum_{h=0}^H \sum_{m=0}^{\infty} passthrough_{i,h,m}(\alpha) \Delta f_{Z,t-h,m}.$$

In practice, we will let the highest maturity be 60 months. It is also the case that $(\Lambda_1(\alpha)^{-1})^{61} \approx 0$, so there is likely no loss from this cut-off in practice, unless the changes in oil futures prices were exploding with the maturity – this is not the case in the data.

E.2 Empirical Setup

Before incorporating futures, we assumed that all oil price changes were fully persistent. The model gave us predictions, for each industry i , about how much the price should change in response to a unit oil price change h periods ago, which we denoted $passthrough_{i,h}(\alpha)$, where α was the vector of calibration parameters used in the model. We tested these model predictions using the regression

$$\Delta P_{i,t} = \lambda_t + \sum_{h=0}^H \beta_h(\alpha) passthrough_{i,h}(\alpha) \Delta P_{Oil,t-h} + \epsilon_{i,t}.$$

If the model under calibration α was correctly specified, and the OLS was unbiased (the argument for which relies on shift-share identification assumptions for an exogenous shifter), then we should have $\hat{\beta}_h(\alpha) = 1$ for all h . Our GMM estimated α to get these β 's as close as possible to 1.

After incorporating futures, the model instead delivers predictions $passthrough_{i,h,m}(\alpha)$: how much industry i should change prices in response to a unit shock to oil prices h periods ago at maturity m . This is a rich object: fix $h = 0$ and consider $m = 0$ and $m = 1$. For $m = 0$, the object is the predicted effect of increasing oil prices by one unit today, holding future prices constant (i.e., an immediately and fully mean-reverting shock). For $m = 1$, the object is the predicted effect of increasing oil futures at a 1 month horizon by one unit, holding current and other future prices constant (i.e., an immediately and fully mean-reverting shock expected to occur next month). So the old pass-through prediction is $passthrough_{i,h}(\alpha) = \sum_{m \geq 0} passthrough_{i,h,m}(\alpha)$.

A full test of the model incorporating futures would be

$$\Delta P_{i,t} = \lambda_t + \sum_{h=0}^H \sum_{m=0}^M \beta_{h,m}(\alpha) passthrough_{i,h,m}(\alpha) \Delta f_{Z,t-h,m} + \epsilon_{i,t},$$

where $\Delta f_{Z,t-h,m}$ is the change in the oil futures price in period $t - h$ holding the maturity m fixed ($m = 0$ denotes the change in the spot price, whereas $m > 0$ is the change in an oil future). If we include 24 lags and 60 months of futures data, this is 25×61 treatments, yielding far too many β 's to estimate. But thinking about why this is so complicated is instructive: the model predicts that, in period t , prices are still responding to, e.g., movements in the 40 months oil future 10 months ago.

We proceed with three different summations to simplify the regression analysis: (1) re-running the baseline regression but incorporating futures data, (2) jointly testing whether pass-through of actual movements in the oil price differs from pass-through of expected future movements in the oil price, and (3) testing whether pass-through differs from what should occur under a random walk assumption for oil prices if oil futures price changes suggest the random walk assumption is

wrong.

As in our structural empirics in the main body, we exclude petroleum refineries from this analysis. It is very likely that refineries make use of oil futures prices when setting prices, and we are more interested if more downstream industries make use of oil futures in forming expectations of future oil prices.

E.3 Regression Tests Incorporating Futures

A alternative baseline regression to our baseline analysis in the paper is just to run

$$\Delta P_{i,t} = \lambda_t + \sum_{h=0}^H \beta_h(\alpha) \sum_{m=0}^M \text{passthrough}_{i,h,m}(\alpha) \Delta f_{Z,t-h,m} + \epsilon_{i,t}. \quad (17)$$

Just as in our baseline analysis, we estimate one β at each horizon, and the treatment is the current and expected future changes in the oil price, scaled by how the industry should respond to those changes as predicted by the model. We find that the GMM-optimal model from the paper maintains an adequate fit, though there is evidence of slightly more/faster pass-through in the data than what the model predicts. We show the result in figure G.11 by taking the rolling average of the estimated β s at each horizon to increase power. For example, for $h = 2$, we plot $(1/3)(\hat{\beta}_0 + \hat{\beta}_1 + \hat{\beta}_2)$. If the GMM-optimal model is correct, we should find that these rolling averages are always equal to 1. We find that while point estimates are often above one, the distinction is statistically significant only at relatively few horizons. The model from the paper therefore maintains adequate fit, which is unsurprising given the result mentioned in the body of the paper that real oil prices follow essentially a random walk (on average) over a one-year horizon.

Advancing on this approach, we can double the number of estimated parameters, instead considering the regression

$$\begin{aligned} \Delta P_{i,t} = & \lambda_t + \sum_{h=0}^H \beta_{h,Current}(\alpha) \text{passthrough}_{i,h,m=0}(\alpha) \Delta P_{Oil,t-h} \\ & + \sum_{h=0}^H \beta_{h,Forward}(\alpha) \sum_{m=1}^M \text{passthrough}_{i,h,m}(\alpha) \Delta f_{Z,t-h,m} + \epsilon_{i,t}. \end{aligned}$$

Again, all of these β 's should be 1 if the model is correct and the OLS is unbiased. In particular, $\beta_{h,Forward}(\alpha) > 0$ represents that industries are passing through changes in expected future oil price changes, holding changes in the current oil price constant. It is vital that this regression is performed jointly to estimate $\beta_{h,Current}$ and $\beta_{h,Forward}$ because changes in the current oil price are highly correlated with changes in the expected oil prices in the future, so a regression separately estimating these parameters would suffer from severe omitted variables bias.

The GMM-optimal model retains adequate fit, which we visualize in Figure G.12. We are underpowered to perform this analysis out to 24 lags.

We can instead nest our empirical tests in the main body of the paper, where we assumed firms viewed all oil price changes as fully persistent. We can then ask whether using the futures data provides any additional power for predicting industry pass-through of oil price changes.

$$\begin{aligned} \Delta P_{i,t} = & \lambda_t + \sum_{h=0}^H \beta_{h,Persistent}(\boldsymbol{\alpha}) passthrough_{i,h}(\boldsymbol{\alpha}) \Delta P_{Oil,t-h} \\ & + \sum_{h=0}^H \beta_{h,FuturesEvolution}(\boldsymbol{\alpha}) \sum_{m=0}^M passthrough_{i,h,m}(\boldsymbol{\alpha}) (\Delta f_{Z,t-h,m} - \Delta P_{Oil,t-h}) \\ & + \epsilon_{i,t}. \end{aligned} \quad (18)$$

The result is shown in Figure G.13. We see again that the GMM-optimal model retains adequate fit, and there is evidence that deviations from the random walk model are useful for forecasting industry price changes in response to oil price movements. There is evidence of more pass-through of deviations from the random walk assumption than is warranted by the model, which is consistent with the evidence from Figure G.12 that futures price movements do pass through but with point estimates somewhat below 1.

E.4 Reduced-form Test of Forward-lookingness

We now incorporate futures exposure into our regression setup from Section 5 of the paper. This section used an arbitrary lag structure with industries' long-run exposures to oil shocks, rather than using the dynamic model's implied lag structure for industry pass-through. We alter regression 8 to include an additional set of terms that should be predictive of industry price changes if industries incorporate oil futures data when forming expectations about future oil prices. This regression is most closely related to the fully structural regression test captured in regression 18.

$$\begin{aligned} \Delta P_{i,t} = & \lambda_t + \sum_{h=-6}^{24} \beta_{h,Persistent} [(\mathbf{I} - \boldsymbol{\Phi})^{-1} \mathbf{s}^Z]_i \Delta P_{Z,t-h} \\ & + \sum_{h=-6}^{24} \beta_{h,FuturesEvolution} [(\mathbf{I} - \boldsymbol{\Phi})^{-1} \mathbf{s}^Z]_i (\Delta f_{Z,t-h,m} - \Delta P_{Oil,t-h}) + \epsilon_{i,t}. \end{aligned}$$

We use the change in the one-year maturity future, so $m = 12$.

We plot our results in Figure G.14. We cannot reject full pass-through in the long-run of both changes in the spot price and changes in futures prices that deviate from the random walk model for oil prices. We see evidence of more delayed pass-through of expected deviations in

spot price changes from the random walk model than pass-through of changes in the spot price. These results provide more evidence that firms are forward-looking about changes in future oil prices.

E.5 GMM Estimation incorporating Futures

We can re-estimate the model parameters in α incorporating futures data. As in the paper, the GMM procedure is performing a weighted match of the β 's estimated from regression 17 to a vector of 1's. In this case, we get $\hat{m}_f = 0.54$ (SE = .42), $\hat{\mu} = 1.24$ (SE = 0.17), and $\hat{\gamma} = 0$ (SE = 0.39). The model again prefers the heterogeneous frequencies of price adjustment ($\gamma = 0$) and a similar value of the aggregate markup. But the model prefers an estimate of myopia equal to 0.54 instead of complete rational expectations; though this estimate is less than 1, it still suggests forward-lookingness. We note that the standard errors are wider, which could reflect heterogeneity in industries' responses to movements in oil futures prices or heterogeneity in forward-lookingness about the network component of delayed pass-through relative to changes in expectations about future oil prices. The standard errors here include the estimates we found in the main body of the paper.

F Appendix: Additional Empirics for Application

F.1 Incorporating Oil Futures Data

The formula for predicting inflation changes from formula 14 to

$$\hat{\Pi}_t = Intercept + \sum_i PCEShare_i \sum_{h=0}^H \sum_{m=0}^M passthrough_{i,h,m}(\hat{\alpha}) \Delta f_{Z,t-h,m}, \quad (19)$$

which is an intercept plus the prediction arising from our analysis in Appendix E. Recall that the notation $\Delta f_{Z,t-h,m}$ is the change in the futures price at a fixed maturity. The procedure determining the intercept, outlined below in Appendix F.2, can analogously be updated for use of futures. Formally, just replace the industry shock

$$\sum_{h=0}^H passthrough_{i,h}(\hat{\alpha}) \Delta P_{Z,t-h}$$

with

$$\sum_{h=0}^H \sum_{m=0}^M passthrough_{i,h,m}(\hat{\alpha}) \Delta f_{Z,t-h,m}.$$

F.2 Appendix: Regressions for Computing Aggregation Intercept

Recall that regression 8 was

$$\Delta P_{i,t} = \lambda_t + \sum_{h=-6}^{24} \beta_h [(\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z]_i \Delta P_{Z,t-h} + \sum_{h=-6}^{24} \gamma_h [(\mathbf{I} - \Phi)^{-1} \text{diag}(\mathbf{s}_i^L) \Delta \mathbf{w}_{t-h}]_i + \epsilon_{i,t}.$$

The variant of this regression that allows us to assess the effects of oil price increases on a hypothetical industry with no network exposure to oil is therefore

$$\Delta P_{i,t} = \delta_1 [(\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z]_i + \sum_{h=-6}^{24} \delta_{2,h} [(\mathbf{I} - \Phi)^{-1} \mathbf{s}^Z]_i \Delta P_{Z,t-h} + \sum_{h=-6}^{24} \tilde{\beta}_h \Delta P_{Z,t-h} + \epsilon_{i,t}.$$

Then $\sum_h \tilde{\beta}_h$ is the effect of a sequence of oil price increases on prices in a sector with no network exposure to the commodity. Now, because this regression uses time-series variation, an instrument with valid identification for the time series is required. We therefore instrument oil price changes with the Kanzig instrument and the interaction between the network exposure and the oil price changes with the interaction between the network exposure and the Kanzig instrument. Our results are depicted in Panel 1 of Figure G.15. We see that, with the exception of a couple pre-periods, we cannot reject that $\sum_h \tilde{\beta}_h = 0$.

An alternative approach is to use the regression with model-derived dynamics for pass-through. Recall that the variant of regression 12 testing the GMM-optimal model was

$$\Delta P_{i,t} = \lambda_t + \sum_{h=0}^H \beta_h \text{passthrough}_{i,h}(\hat{\alpha}) \Delta P_{Z,t-h} + \epsilon_{i,t},$$

where $\hat{\alpha}$ was the GMM-optimal calibration. The variant of this regression that allows us to assess the effects of oil price increases on a hypothetical industry with no network exposure to oil is

$$\Delta P_{i,t} = \sum_{h=0}^{24} \delta_{1,h} \text{passthrough}_{i,h}(\hat{\alpha}) + \sum_{h=-6}^{24} \delta_{2,h} \text{passthrough}_{i,h}(\hat{\alpha}) \Delta P_{Z,t-h} + \sum_{h=-6}^{24} \tilde{\beta}_h \Delta P_{Z,t-h} + \epsilon_{i,t}.$$

We apply our instrumental variables procedure in the same way as above, except that network exposure at a given lag is now $\text{passthrough}_{i,h}(\hat{\alpha})$. Our results are depicted in Panel 2 of Figure G.15. We see that, as before, with the exception of two values for h , we cannot reject that $\sum_h \tilde{\beta}_h = 0$.

F.3 Robustness: Model Fit for the PCE-weighted Average Industry

Before computing network oil inflation, we apply an important robustness check: is the model fit for our GMM-optimal calibration $\hat{\alpha}$ still good for the PCE-weighted average industry? Our GMM

aimed at optimizing fit for the unweighted average industry. To assess this, we can compare the IRF of PCE inflation to an oil shock in the model,

$$IRF_H = \sum_i PCEShare_i \sum_{h=0}^H passthrough_{i,h}(\hat{\alpha}).$$

with the IRF in the data,

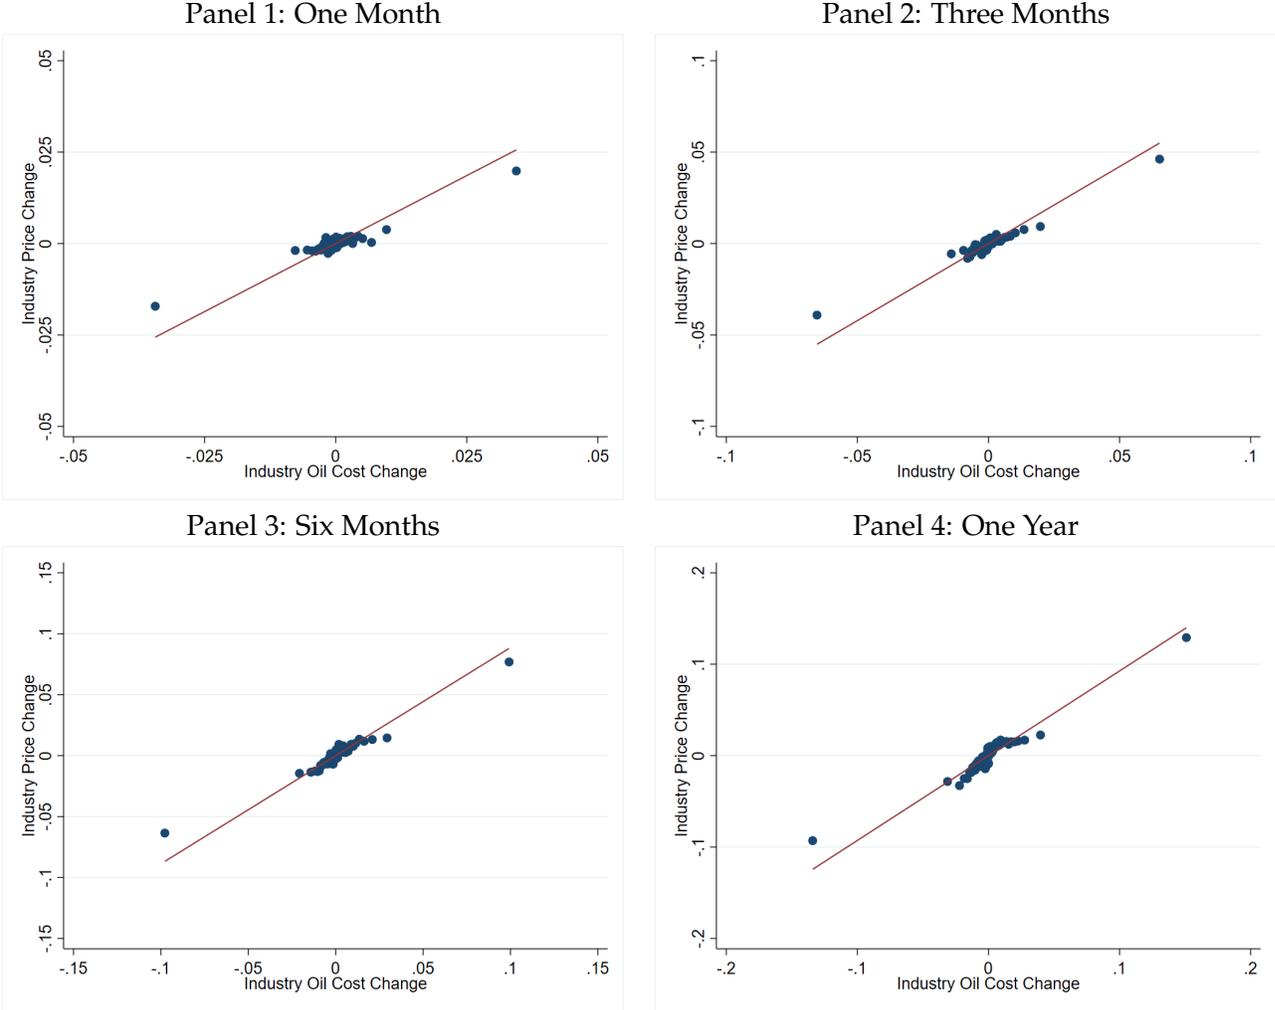
$$IRF_H^{Data} = \sum_i PCEShare_i \sum_{h=0}^H \hat{\beta}_h(\hat{\alpha}) passthrough_{i,h}(\hat{\alpha}),$$

where $\hat{\beta}_h(\hat{\alpha})$ was computed as in our structural estimation, using regression equation 12.

The results are shown in Figure G.16. We see that model fit remains good for the PCE-weighted average industry, even though fit was optimized for the equally weighted average industry.

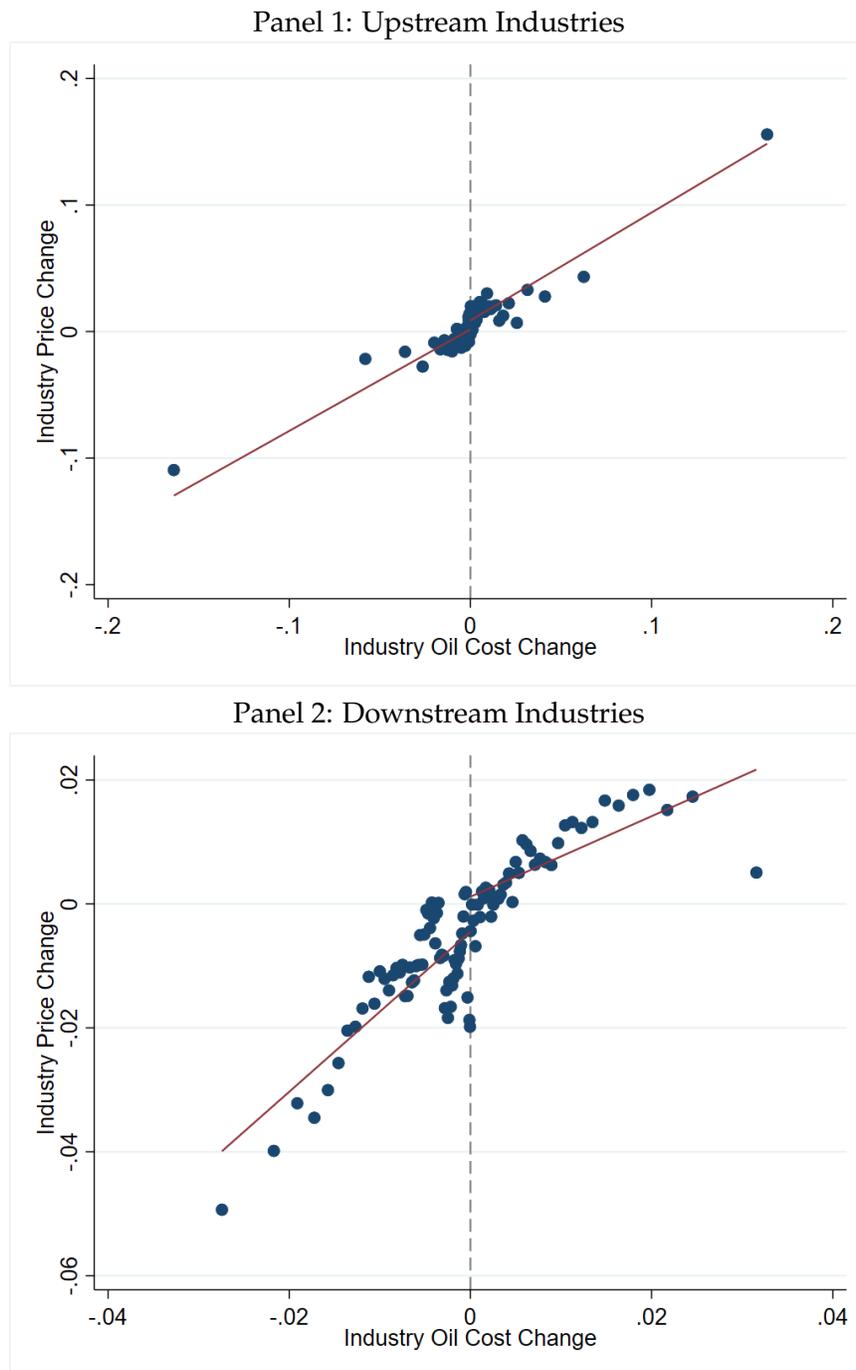
G Appendix Figures

Figure G.1 – Binscatters at Various Time Horizons



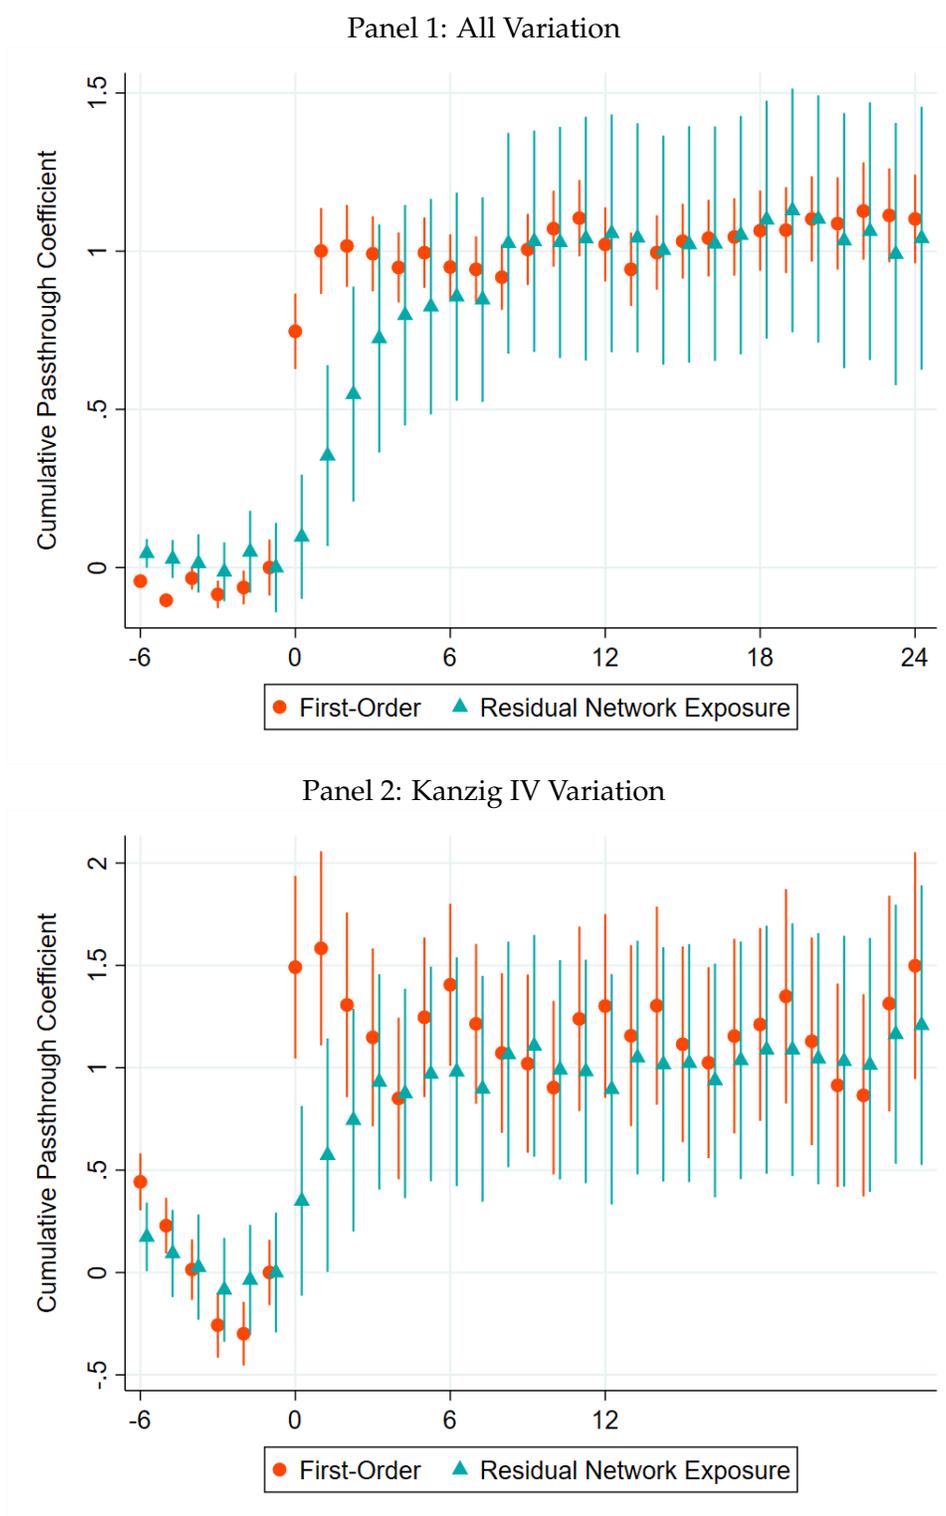
Note: These plots are 100-quantile binscatters displaying how (de-meanded) industry price changes vary with (de-meanded) industry oil cost changes. Lines of best fit are included. The slope of these lines can be interpreted as the fraction of cost increases that are passed through into prices over the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

Figure G.2 – Binscatters: Upstream vs. Downstream Industries



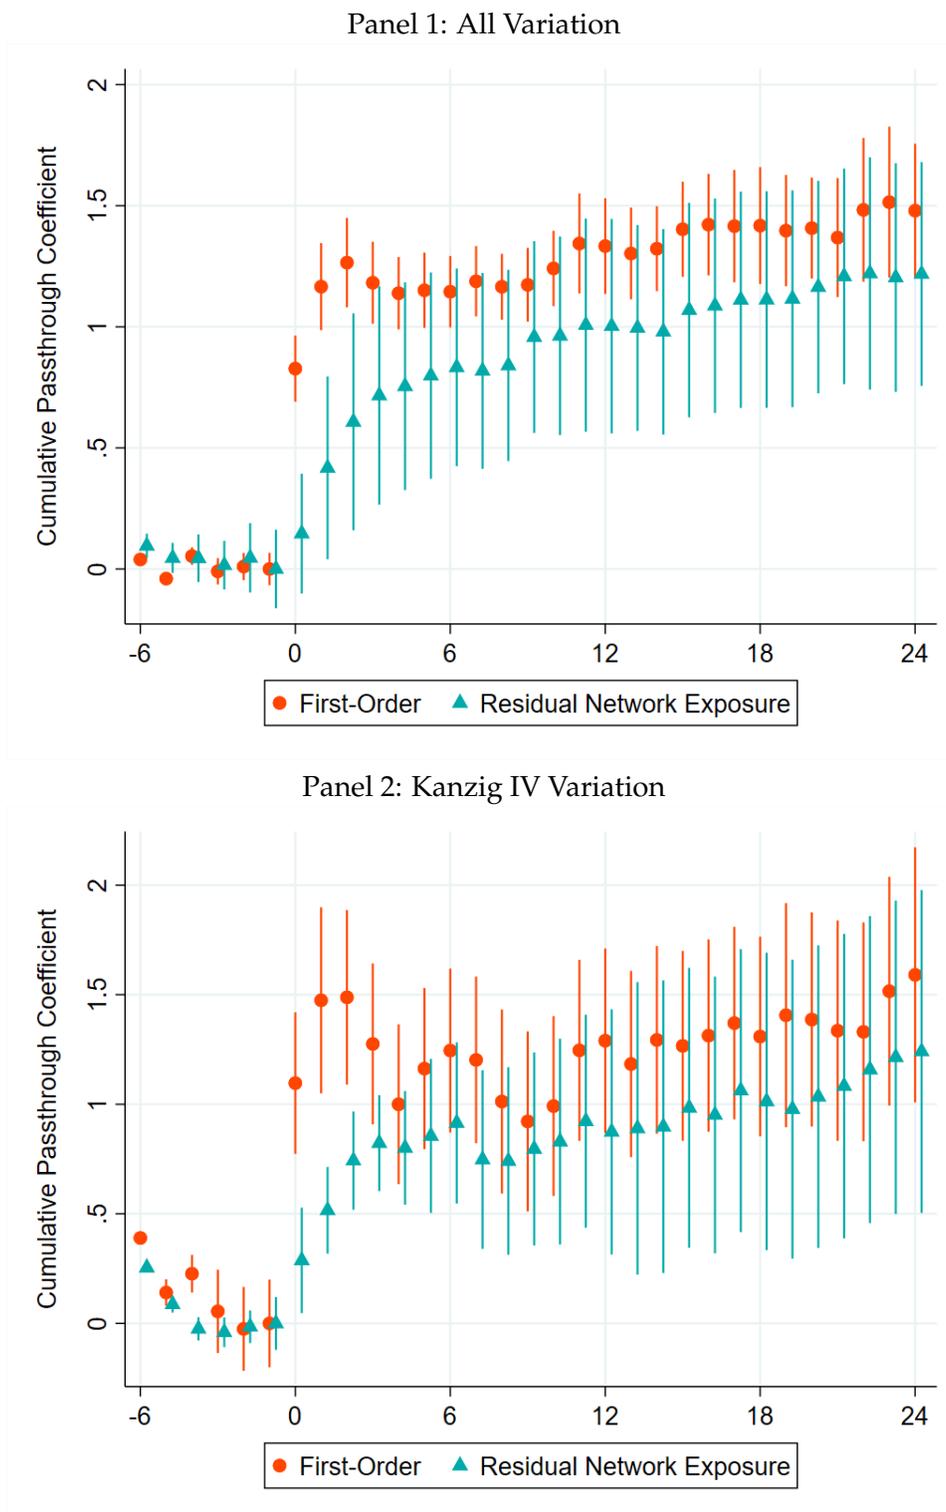
Note: These plots are 100-quantile binscatters displaying how (de-meanned) industry price changes vary with (de-meanned) industry oil cost changes. Lines of best fit with a regression discontinuity at zero are included. For upstream industries (industries with below-median downstreamness), the slope 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.

Figure G.3 – Month-by-Month pass-through of Oil Price Changes



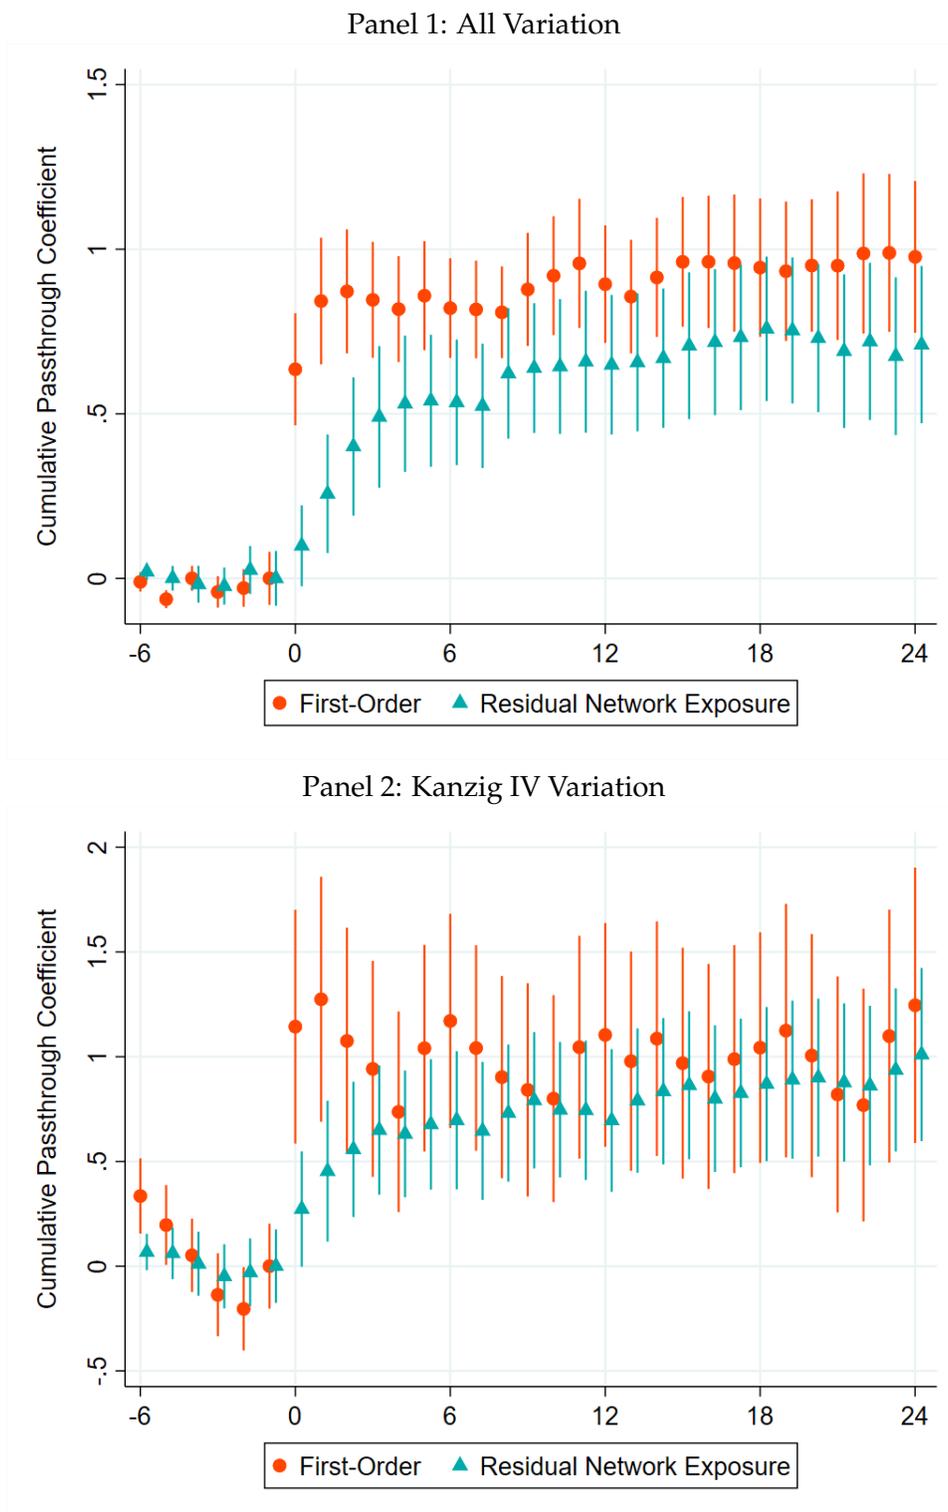
Note: Regression specifications correspond to Equation (9). Red coefficients plot monthly price pass-through of first-order exposure to crude oil shocks; blue coefficients plot monthly price pass-through of second- and higher-order network exposure to crude oil shocks. The specifications exclude the control for general equilibrium wage changes.

Figure G.4 – Month-by-Month pass-through of Oil Price Changes



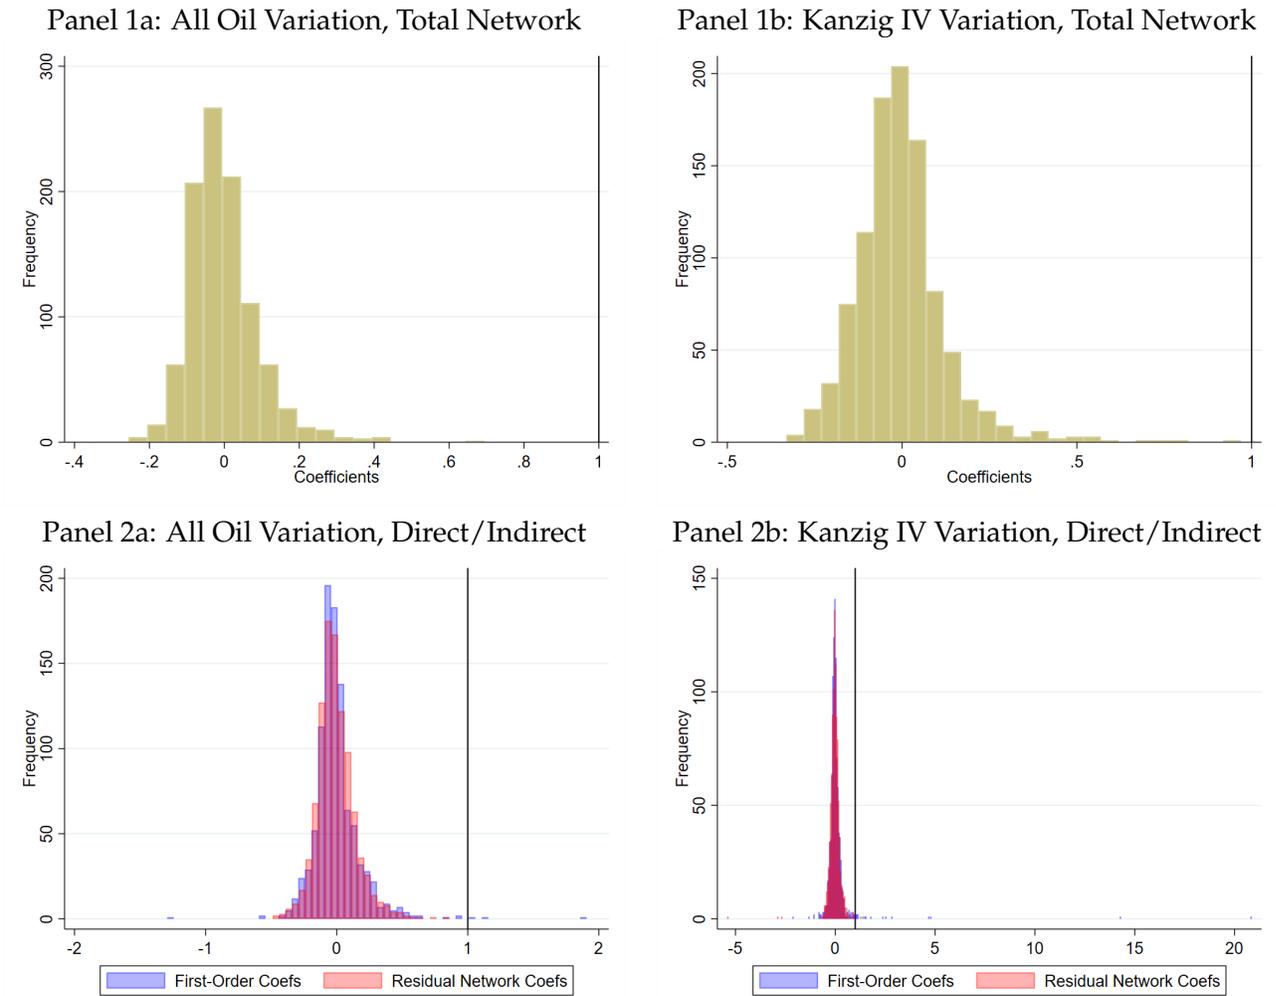
Note: Regression specifications correspond to Equation (9) with controls for gas and electricity shock exposure. Red coefficients plot monthly price pass-through of first-order exposure to crude oil price shocks; blue coefficients plot monthly price pass-through of second- and higher-order network exposure to crude oil price shocks.

Figure G.5 – Month-by-Month pass-through of Oil Price Changes



Note: Regression specifications correspond to Equation (9), 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.

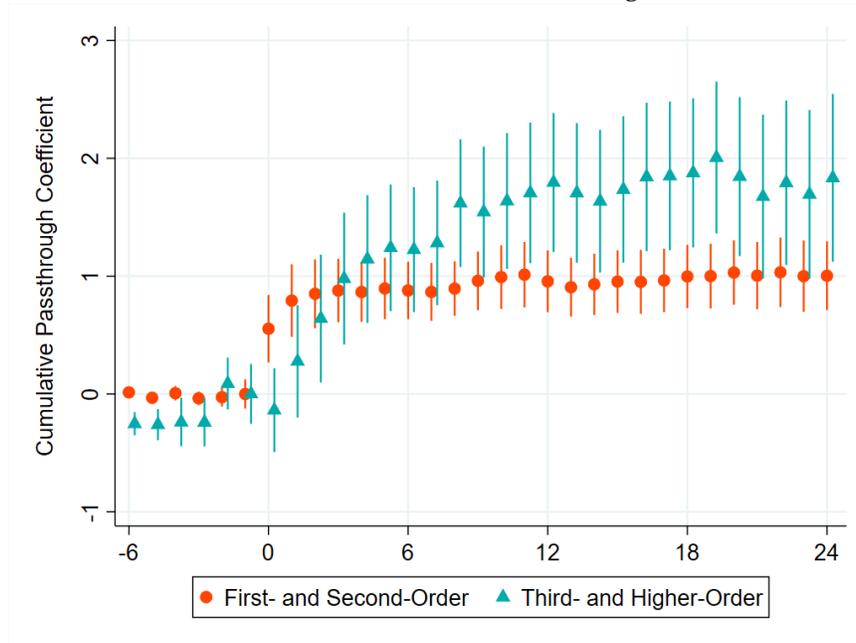
Figure G.6 – Results of Permutation Tests



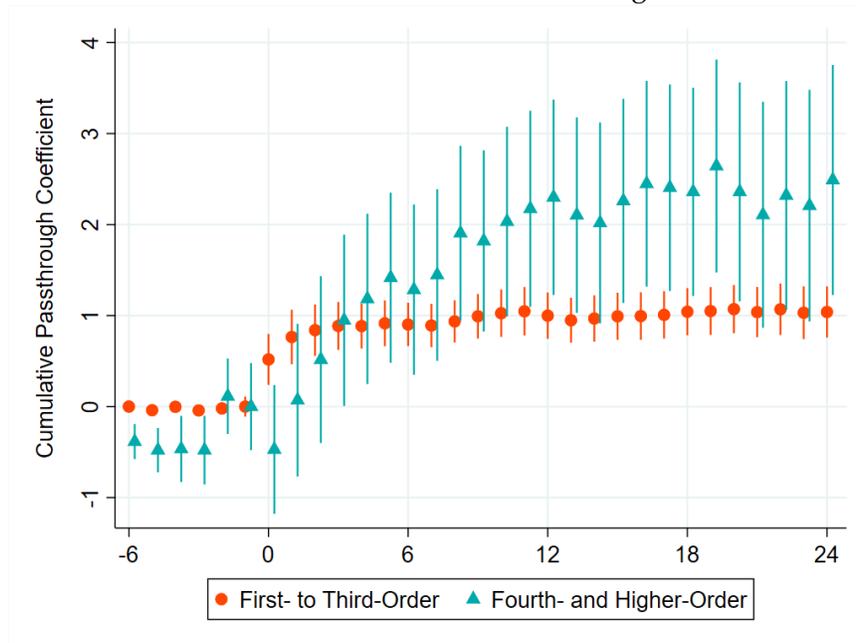
Note: These plots display the results of 1000-repetition permutation tests on the 12-month cumulative pass-through coefficient from the main specifications, as given by Equation (8) in the top panel and (9) in the bottom panel.

Figure G.7 – Higher-order Pass-through in Reduced-form

Panel 1: Third-order Pass-through

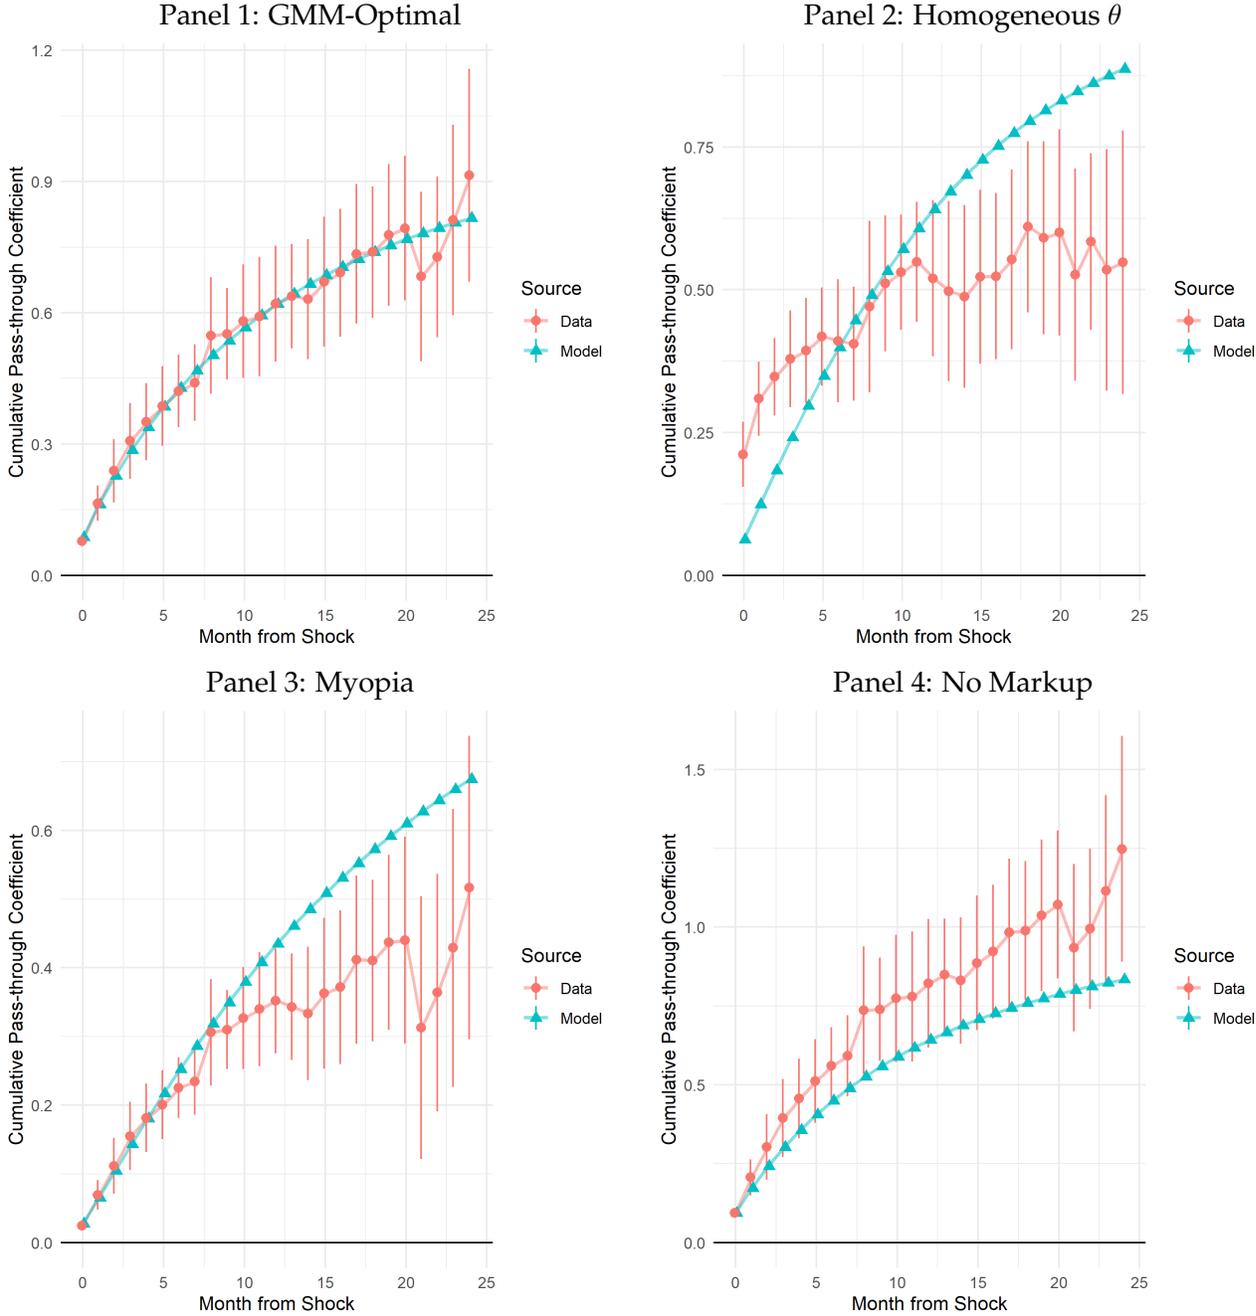


Panel 2: Fourth-order Pass-through



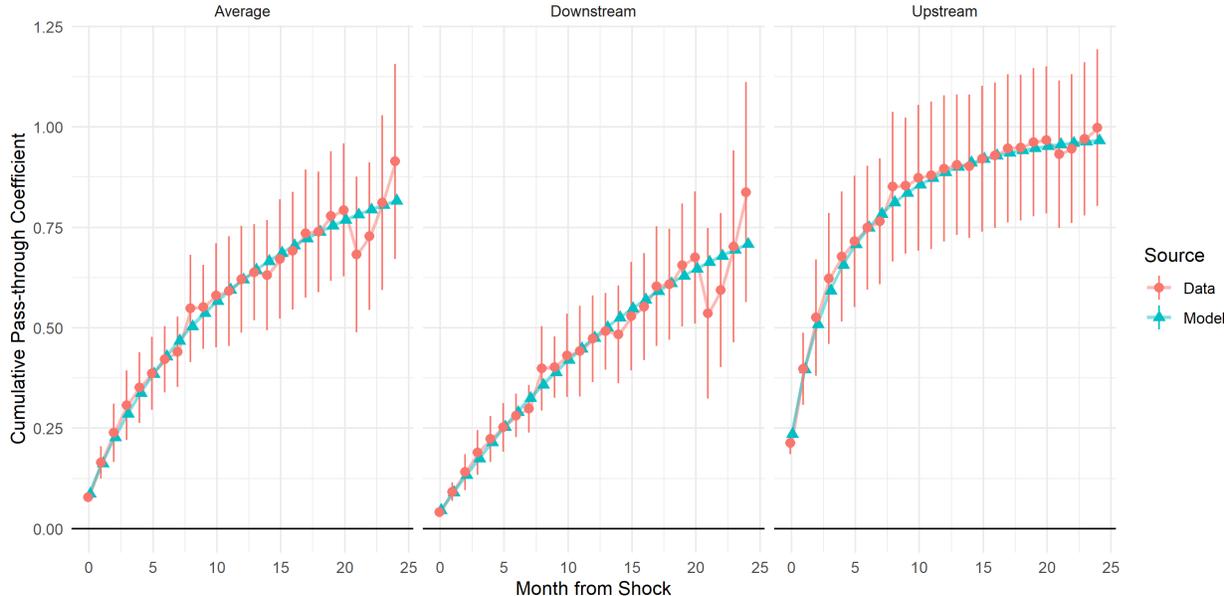
Note: These figures plot the cumulative pass-through for third-order and higher (in Panel 1) and fourth-order and higher (in Panel 2) exposure to oil price changes. The results are consistent with needing a markup larger than 1 for downstream long-run pass-through to be consistent with the predictions of Proposition 1.

Figure G.8 – GMM Results for Industry Pass-through of Oil Price Changes (including Refineries)



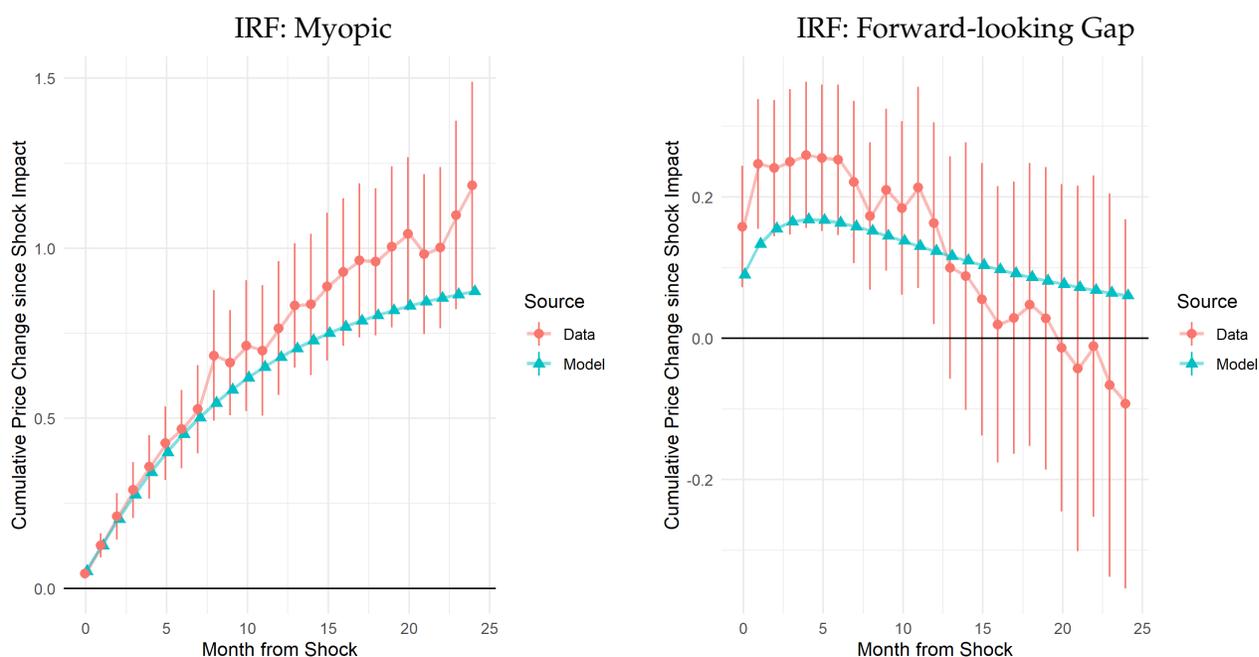
Note: Model fit is good for the GMM-optimal model when estimation includes petroleum refineries, as shown in Panel 1. Our ability to identify parameters, as visualized in panels 2-4 and captured in the tightness of our standard errors, also appears similar to the case where refineries are excluded.

Figure G.9 – Industry Pass-through of Oil Price Changes, Upstream and Downstream (including Refineries)



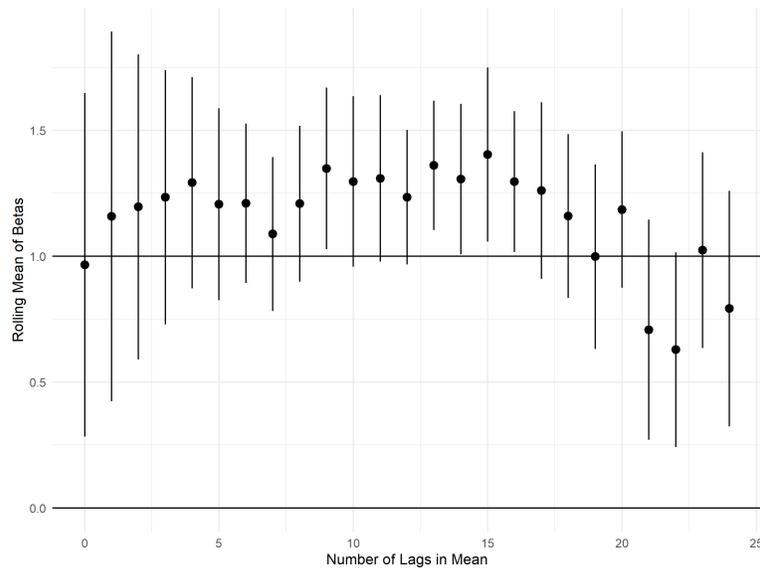
Note: The GMM procedure optimized fit for the average industry (leftmost plot). We visualize how well the fit is for the most upstream and downstream industries in the middle and rightmost plots, finding that model fit remains good. These results are similar to those we found excluding refineries.

Figure G.10 – Tests of Forward-lookingness (including Refineries)



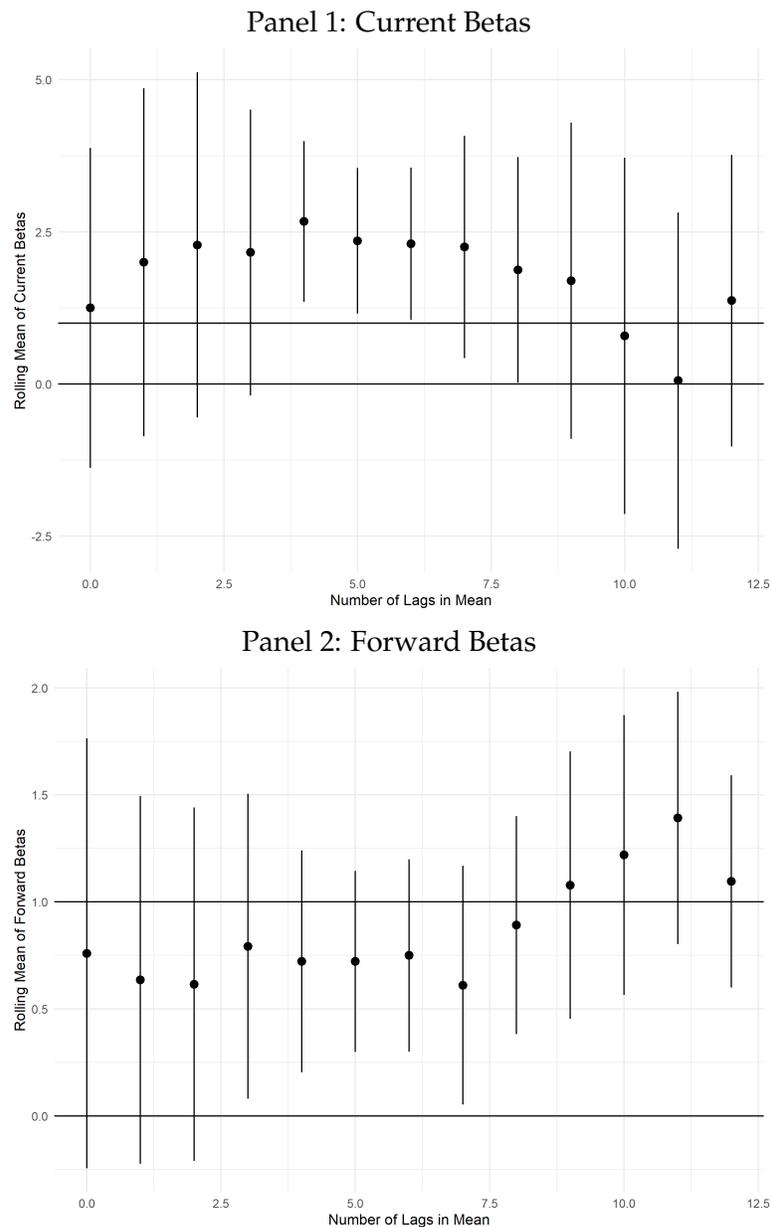
Note: We plot the results of our test that pass-through due to forward-lookingness under rational expectations is present in the data. Statistical significance in the right panel implies we cannot reject that firms are forward-looking about the gradual pass-through of upstream shocks to their marginal costs. The fact that the model lies within the standard error bars in the right figure visualizes that we cannot reject a degree of forward-lookingness consistent with rational expectations. These results are similar to those we found when excluding refineries.

Figure G.11 – Baseline Result Incorporating Futures



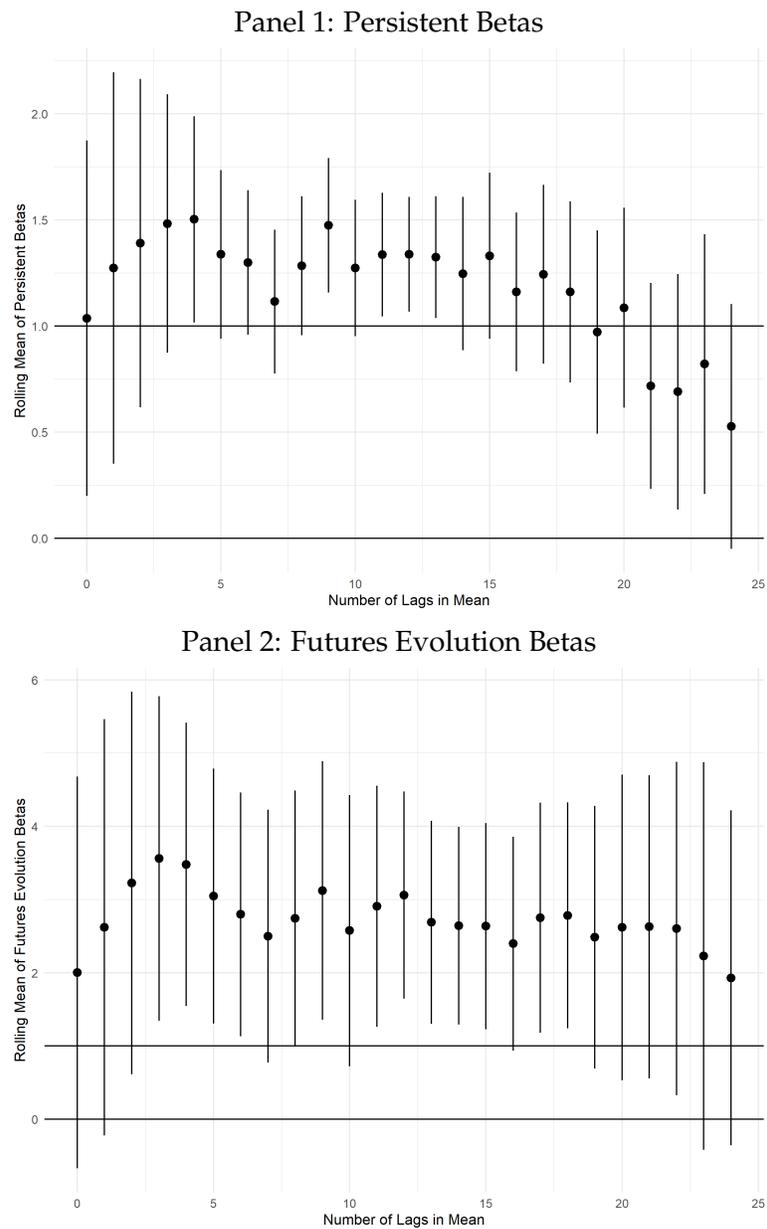
Note: Results testing whether the GMM-optimal model retains good fit of pass-through in the data when assuming that firms form expectations about future oil prices using oil futures. The rolling mean of betas should be 1 if the model is correct and our regression estimates are unbiased. We find that that overall fit is good, with some evidence of more pass-through at short horizons and less pass-through at longer horizons.

Figure G.12 – Results Testing Incorporation of Futures Price Movements



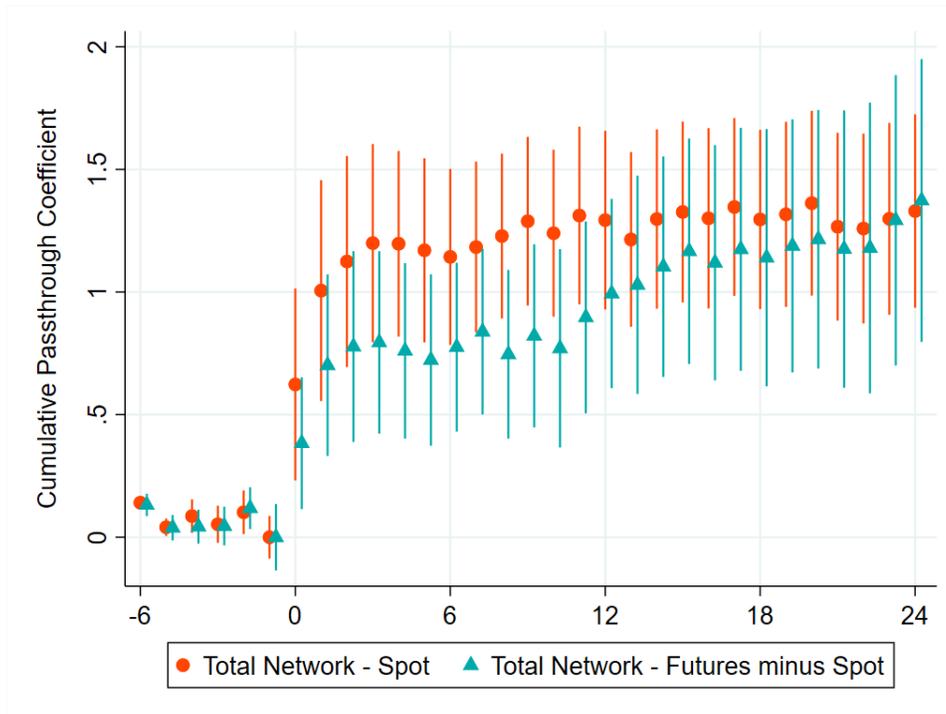
Note: Results testing whether firms respond to movements in oil futures in addition to movements in the spot price. The rolling average of betas should be 1 in both panels if the GMM-optimal model is correct and our estimates are unbiased. We find evidence of pass-through of both current and expected future movements in oil prices. Standard errors are quite large for pass-through of current oil price movements, likely because futures price movements are highly correlated. This result is therefore best paired with our result in Figure G.13. We are underpowered to run this test out to a horizon of 24 months and therefore present results for a horizon of 12 months.

Figure G.13 – Results for Deviations from the Random Walk Model



Note: Results testing whether firms respond to (1) movements in oil prices under the assumption of a random walk for future oil price movements and (2) deviations in the random walk assumption informed by oil futures price changes. We find evidence that firms respond to both (1) and (2). Results are in line with the model's predictions about pass-through of (1) and suggest somewhat more pass-through of (2), consistent with the evidence from Figure G.12 that there is less pass-through of movements in oil futures prices than pass-through of current movements in oil prices.

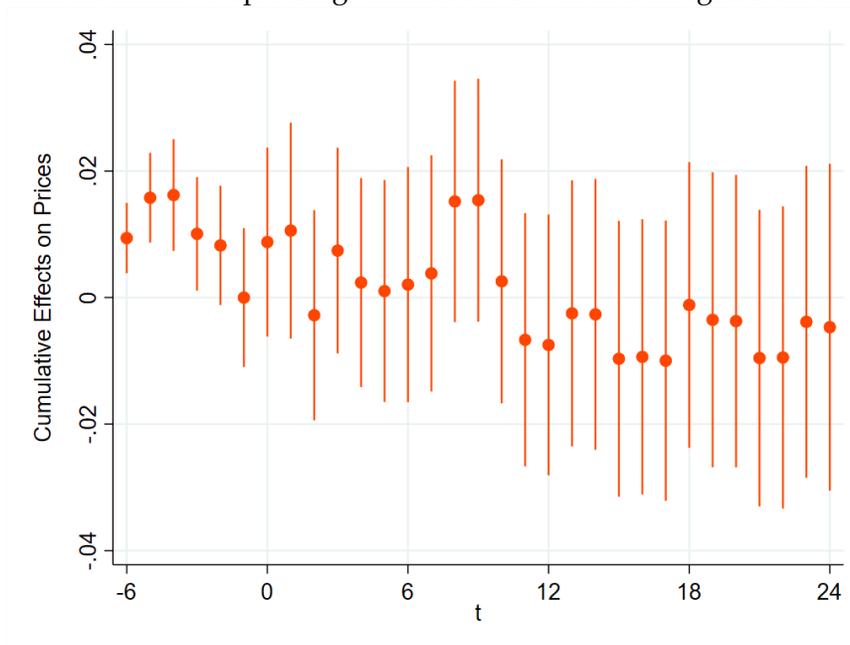
Figure G.14 – Baseline Result Incorporating Futures



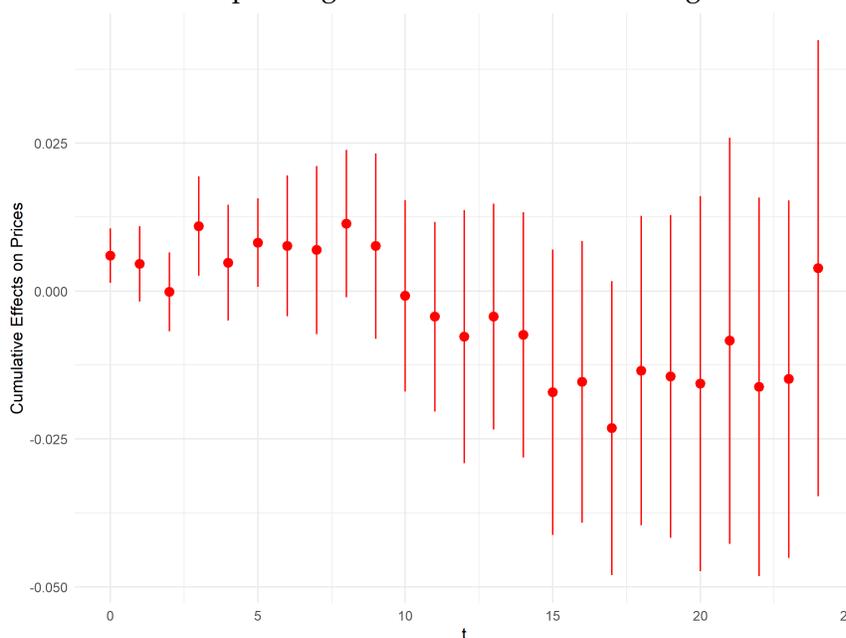
Note: Reduced-form test of whether firms respond to movements in oil futures prices in addition to movements in the oil spot price. We confirm that firms respond to both movements, with full pass-through in the long-run (a long-run point estimate of 1). There is evidence of faster pass-through of current price movements relative to futures price movements, in accord with the model's predictions.

Figure G.15 – Estimation of the Intercept Required for Aggregating Industry Price Responses

Panel 1: Intercept using a Time-series Variant of Regression 8

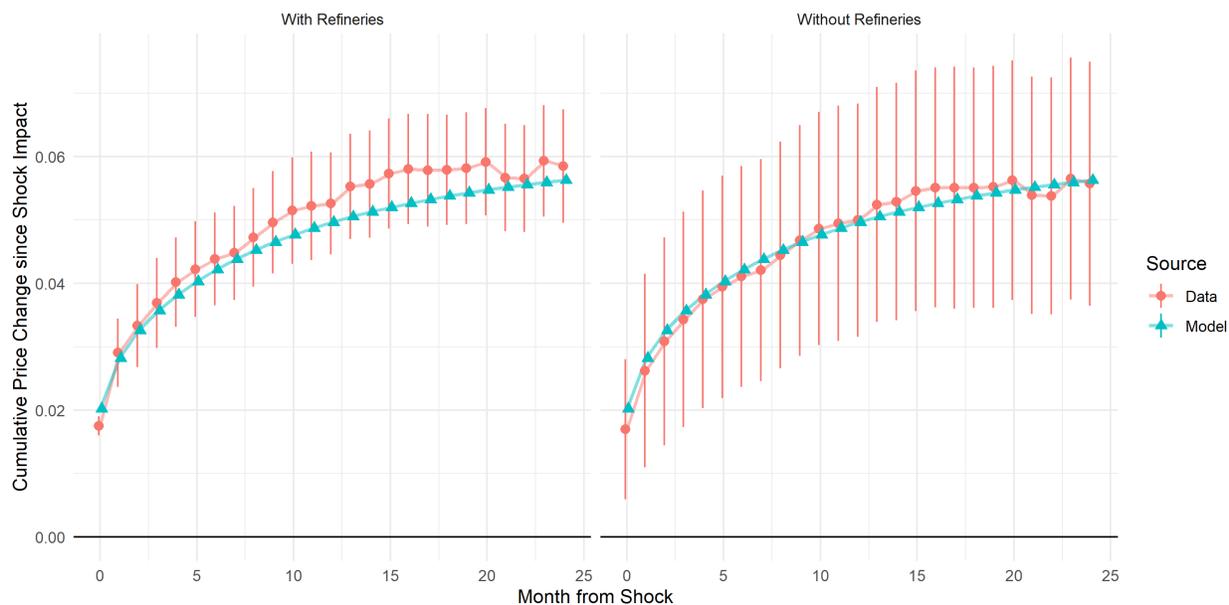


Panel 2: Intercept using a Time-series Variant of Regression 12



Note: We plot the intercept required for aggregation in our Application section using the 2SLS procedure described in appendix F.2. We see that, with the exception of a few periods, we cannot reject that the required intercept is 0. This suggests that sectors with no network exposure to oil largely do not change prices in response to oil price changes.

Figure G.16 – Model Tests for Aggregate PCE Inflation



Note: Our GMM-optimal model fit remains good when we plot fit for the PCE-weighted average industry, which up-weights industries from which consumers directly make purchases. Standard errors are much tighter when estimating model fit using petroleum refineries, highlighting the increased precision that petroleum refineries provide in assessing on-impact pass-through in the same month as the shock.