Delayed Inflation in Supply Chains: Theory and Evidence*

Robert Minton† Brian Wheaton‡

December 6, 2023

Abstract

How fully and quickly do supply chains transmit commodity price movements into inflation? In a production network with sticky prices, we show that the network delays full propagation of commodity price shocks to downstream firms. This delay from downstreamness occurs even when firms are forward-looking, and myopia amplifies it. We confirm the theory using shift-share designs exploiting firms' differential exposures to commodities through their networks. We find forward-looking responses to oil price movements but myopic responses for other commodities. Applying our model, we show that delayed network propagation of oil price movements forecasts the future path of core inflation.

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. Oil prices, for example, are thought to have contributed to many historical inflationary episodes. Beyond oil's central role in energy, everything from fabric in clothing to foam cushions in furniture to the plastic in every consumer electronic can be made indirectly from oil. If oil price movements only gradually filter through

*This paper was previously circulated as Robbie's job market paper and was titled "Hidden Inflation in Supply Chains: Theory and Evidence." We are especially grateful for invaluable feedback and advice from Robbie's advisors: Xavier Gabaix, Ed Glaeser, Andrei Shleifer, Ludwig Straub, and Adi Sunderam. We are also grateful to Maxim Alekseev, Isaiah Andrews, Pol Antras, Robert Barro, Adrien Bilal, Raj Chetty, Gabe Chodorow-Reich, Andrew Figura, Jason Furman, Jordi Galí, Paola Giuliano, Yuriy Gorodnichenko, Elhanan Helpman, Larry Katz, Greg Mankiw, Marc Melitz, Jeff Miron, Emi Nakamura, Ariel Pakes, Giacomo Ponzetto, Ken Rogoff, Emanuel Schertz, Roman Sigalov, Jón Steinsson, Jim Stock, Elie Tamer, John Taylor, Nico Voigtlander, and Romain Wacziarg for excellent comments and advice. We also thank participants of the Macroeconomics, International Economics, and Labor/Public lunch seminars at Harvard, the Macroeconomics faculty lunches at Berkeley and Stanford, the Hoover Economic Policy seminar at Stanford, the Institute for Humane Studies SuRF seminar, and the Conference in Honor of Robert Barro. Finally, we thank the participants at Robbie's job market talks. The views in this paper are those of the authors and do not represent those of the Board of Governors of the Federal Reserve System or the Federal Reserve System.

†Federal Reserve Board. Contact: robert.j.minton@frb.gov

‡UCLA, Anderson School of Management. Contact: brian.wheaton@anderson.ucla.edu

supply chains to final products, they could generate delayed inflation that – without a dynamic model of supply chain propagation – might be misattributed to other causes. We build a theoretical and empirical framework to assess how quickly and fully commodity price movements propagate through supply chains, demonstrating their delayed and predictable effects on aggregate inflation.

In our model and empirical analysis, we obtain five key results. First, we find that network pass-through is full in the long run but occurs much more gradually than one might expect. In the case of oil, the average industry has only passed through 60% of a persistent oil price increase after one year, rising to 80% after two years. Second, we show that this delay is a consequence of price rigidity interacting with the network. Formally, in a network where all industries have the same frequency of price adjustment, industries more downstream from the shock pass it through more gradually, even if they are forward-looking. Third, deviations from rational expectations can amplify the delay due to downstreamness substantially. While our empirical results find network propagation of oil price movements in line with rational expectations, propagation of other commodities' price movements is even more gradual, suggesting substantial inattention. Fourth, relative price movements throughout the network generated from oil largely pass through to aggregate inflation. Fifth, this pass-through explains the predictability of core inflation from oil price movements. We now summarize our results in greater detail.

We begin by studying supply chain propagation theoretically, using a dynamic supply-side pricing model with Calvo pricing and a production network. Our setup is quite standard, following the supply side of Rubbo (2023) and augmented to allow for myopic expectations, following Gabaix (2020). It is now well-understood theoretically that how much a price movement in sector j affects prices in sector i depends on the overall share of cost in sector i that lies in sector j, inclusive of supply chain connections. Our theoretical contribution is to show that 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 separate sector i from sector j (on average, as many supply chain links between sectors may exist). Notably, downstreamness in our setting can be due to both loops in the network, as in Basu (1995), and in-line production, as in Blanchard (1983). Finally, we show that forward-lookingness interacts strongly with downstreamness. 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, 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 empirically, using shift-share designs analyzing benchmark input-output data from the BEA and producer price index (PPI) data from the BLS. We find evidence that pass-through of commodity price movements to industry prices throughout the production network is limited in the short run but full in the long run. In the case of oil, we replicate

these results using the Kanzig (2021) series of exogenous oil price shocks obtained through high-frequency identification of the surprise effects of OPEC announcements. 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 – finding analogous results.

Performing heterogeneity analysis, we confirm our central proposition that industries further downstream from the commodity price movement experience significantly less rapid pass-through. This result holds while controlling for industries' heterogeneous frequencies of price adjustment, another key source of variation that affects speed of pass-through. Controlling for other plausible mediators – such as market concentration, firm size, inventories, and capital share – does not diminish the magnitude or significance of the downstreamness result.

We then turn to a structural estimation of our dynamic pricing model, primarily to determine the degree of firms' forward-lookingness about gradual network propagation of commodity price movements. Specifically, we compare the pace of network pass-through in the data to models calibrated with different degrees of forward-lookingness. We implement a generalized method of moments (GMM) procedure identified using a dynamic shift-share design that follows directly from the model. Because the approach exploits the cross section, it removes the need to estimate pass-through industry-by-industry.

We find that context-dependent forward-lookingness is 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 connect this result to shock salience by showing that there are more news articles and Google searches about oil than other commodities.

We also augment the model to accommodate information about shock persistence contained in commodity futures markets. We verify our result of forward-lookingness for oil, finding that firms' pricing behaviors respond to "forward guidance" about future oil prices as provided by oil futures data. We also verify our result of myopia outside of oil by performing a second study using corn futures, another fairly liquid futures market. In this context, we find no response of industry prices to forward guidance. Put differently, we cannot reject that firms pass through corn spot price changes the same way, regardless of the contemporaneous shift in the futures curve.

We argue that our empirically-validated model 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 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. To arrive at this conclusion, we first present evidence on the effect of oil price movements on industries with zero network exposure to oil. We find no evidence of price responses in such industries, suggesting an absence of price declines associated with any contractionary responses by the monetary authority or other general equilibrium effects. Put differently, we find that all of the relative price movements predicted by the model pass through to aggregate inflation.

Purging all of oil's effects from aggregate PCE inflation modifies the time pattern of the current inflationary episode. Underlying inflation is lower than the headline figure for all of 2021 and much of 2022; in the latter half of 2022, underlying inflation continues to increase despite headline inflation declining. More generally, inclusive of supply chain effects, oil price movements explain 33% 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 network 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 reduced-form empirical analysis demonstrating significant but delayed commodity price pass-through, including a variety of robustness checks and heterogeneity analysis. In Section 6, we empirically characterize pass-through speed and structurally estimate the degree of forward-lookingness, further validating our results by performing heterogeneity analyses exploiting commodity futures data. In Section 7, we focus on our application: stripping the full network effects of oil price movements from core inflation. 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 response to exogenous shocks. 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 (2019a,b,c) 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.

Three papers with findings related to some of our empirical results are Auer, Levchenko, and Saure (2019), Smets, Tielens, and Van Hove (2018), and Luo and Villar (2023). Auer et al. study

international input-output linkages, presenting evidence that global input-output linkages contribute to the synchronization of PPI inflation across countries. Smets et al. empirically 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. Akin to Acemoglu, Akcigit, and Kerr (2016) but assessing price outcomes, Luo and Villar (2023) study whether shocks propagate upstream and downstream. Most of our arguments are distinct from these papers: while we do also verify that network propagation through prices occurs – in our setting, in response to commodity price movements – we are primarily focused on whether the extent and duration of this pass-through is consistent with our model, especially the channel by which a sector's downstreamness from the shock and lack of forward-lookingness delay pass-through.

Additionally, we draw on a large literature pertaining to models of price-setting. Broadly speaking, these models fall into two categories. The first category is that of time-dependent models. In these models, in each period, firms have the opportunity to adjust their prices with some exogenous probability. Seminal examples of time-dependent models are those of Taylor (1979), Calvo (1983), and Yun (1996), amongst which Calvo pricing has been the most commonly-used. As detailed micro data on prices became more available in succeeding decades, some deficiencies of time-dependent models became apparent (documented in Bils and Klenow 2004 and Nakamura and Steinsson 2008), and so a second category of price-setting models emerged to address these issues. In these state-dependent models, price changes are endogenous, depending on the state of the firm and the broader economy. For example, firms may have menu costs, whereby changing prices in response to changed costs is only advantageous if the expected gains from shifting to the new optimal price exceed the fixed cost of changing prices. Examples of state-dependent models can be found in Golosov and Lucas (2007), Klenow and Kryvtsov (2008), Nakamura and Steinsson (2010), and Midrigan (2011). More recently, Auclert et al. (2022) has proven an equivalence result between state-dependent and time-dependent models: the Phillips curve of the canonical menu cost model is the same as a mixture of two time-dependent models.

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, and Huang and Liu (2001) extend his findings in this setting. Building on this work, Basu (1995) develops a roundabout production model arguing that this compounding of nominal rigidities magnifies productivity fluctuations and thereby contributes to the intensity of business cycles. Theoretical results in parallel work from Afrouzi and Bhattarai (2023) provide sufficient statistics for shock pass-through

in production networks, with a broader focus on GDP in addition to inflation. Their Proposition 2 is closely related to the component of our Proposition 2 associated with rational expectations. Our work is distinct in characterizing how pass-through speed in networks depends on forward-lookingness (also in Proposition 2 and further in Proposition 3), linking pass-through speed of a shock explicitly to an industry's downstreamness from the shock (in Proposition 3), and empirically confirming that downstreamness amplifies inflation predictability from past commodity price movements.

It is important to verify whether nominal rigidities compound in supply chains because this conjecture increasingly underlies policy advice. Rubbo (2023) 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 Carlstom, Fuerst, and Paustian (2012) and discussed in detail by Del Negro, Giannoni, and Patterson (2012), in a workhorse New Keynesian model (such as Smets and Wouters 2007), central bank promises with regard to policies that will be undertaken in the (sometimes-distant) future can have unreasonably large effects on present-day inflation and output. This is referred to as the "forward guidance puzzle." Gabaix (2020) discusses myopia, the notion that agents are not perfectly forward-looking in all contexts, as a solution to the forward guidance puzzle. We show a way of estimating myopia from observational data on the macroeconomy (as opposed to microeconomic experiments) and, in so doing, provide empirical evidence of 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 (2023). 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 $i \in \{1,...,I\}$. Each firm produces output denoted $Y_{i,j,t}$, and these varieties are transformed[1] 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 $\boldsymbol{X}_{i,j,t} = (X_{i,j,t}^1, ..., X_{i,j,t}^I)$ and labor $L_{i,j,t}$.[2] 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, $Y_{i,j,t} = A_{i,t} F_i(L_{i,j,t}, \boldsymbol{X}_{i,j,t}, Z_{i,j,t})$. 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.

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

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

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

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

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 $\boldsymbol{P}_t = (P_{1,t}, ..., P_{I,t})$, 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 (2015) 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 $SDF_{t,t+k}$.[3] The optimal reset price will not vary across firms within an industry, and it solves

$$\max_{P^*_{i,j,t}} \sum_{k=0}^{\infty} \theta_i^k \tilde{\mathrm{E}}_t \left[SDF_{t,t+k} Y_{i,j,t+k} \left(P^*_{i,j,t} - MC_{i,t+k} \right) \right] \tag{2}$$

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

Now, $\tilde{\mathrm{E}}_t$ 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

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

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).*[4]

[3] One 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.

[4] The 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.

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. We note that, in Appendix A.4, we provide details for a fully specified model.

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

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

with $m_f \in [0,1]$ ($m_f = 1$ denoting rational expectations) and E_t being the rational expectations operator under the period t information set. When there is myopia, firms discount expected future 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. In our empirics, we will allow m_f to vary by the excluded commodity Z to establish whether firms act as through they are more forward-looking about some commodity shocks than others. We suppress the dependence on Z for notational simplicity.

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 $\mathbf{\Phi}$ 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

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

The log-linearizations of the law of motion for industry prices and the first-order condition deter-

mining the optimal reset price are well-known in the New Keynesian theory. With the addition of myopia, these equations log-linearize to

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

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

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*

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

If wages adjust in general equilibrium,

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

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. Proposition 1 is related to many existing results, particularly the forward network propagation result in Proposition 1 of Baqaee and Rubbo (2022).

We can test Proposition 1 directly in the data. In particular, we can compute $(\boldsymbol{I} - \boldsymbol{\Phi})^{-1}\boldsymbol{s}^Z$ 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:

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

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

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

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 composition of network exposures at various orders is irrelevant for long-run pass-through, conditional on the full network exposure. As we show now, however, the composition of network exposure is important for the speed of adjustment.

3.4 Transitory Dynamics and Downstreamness Measures

Next, we seek to develop empirical measures of the speed of pass-through in an industry. We derive closed-form solutions for two special cases, allowing us to 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 $\psi \to \infty$, 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*

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

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

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

and we take the roots with all positive eigenvalues.[5] *In the myopic case,* $\boldsymbol{B} = \bar{\phi}(\boldsymbol{I} - \boldsymbol{\Phi})$. *Under rational expectations, when the interest rate* r *is* 0, *we have* $\boldsymbol{B} = \left(\bar{\phi}^2(\boldsymbol{I} - \boldsymbol{\Phi})\right)^{1/2}$.

Proposition 2 yields a general method for determining the time path of pass-through of price movements to all industry prices, given $\boldsymbol{\phi}$, $\boldsymbol{\psi}$, $\boldsymbol{\Phi}$, and $\boldsymbol{s}^Z$. In two special cases – complete myopia and rational expectations with a steady state interest rate of 0 – we actually know the functional form of the solution.

Now, we turn to further characterizations of the speed of pass-through in a sector. Denote by $\boldsymbol{e}^i$ 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 $\hat{p}_{i,\infty} = \left[(\boldsymbol{I} - \boldsymbol{\Phi})^{-1}\boldsymbol{s}^Z\right]_i \hat{p}_{Z,0}$. It follows from Proposition 2 that

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

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. In the network setting, the rate of price adjustment is insufficient to characterize pass-through duration. We now 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*

[5] If $\phi_i = \phi$ for all i, then $\psi_i = \psi$ for all i, $\boldsymbol{B}$ satisfies the equation

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

and we take the roots with all positive eigenvalues:

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

industry i is

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

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

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

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

$$D_i \approx \frac{1}{\phi}\frac{(\boldsymbol{e}^i)'(\boldsymbol{I} - \boldsymbol{\Phi})^{-3/2}\boldsymbol{s}^Z}{(\boldsymbol{e}^i)'(\boldsymbol{I} - \boldsymbol{\Phi})^{-1}\boldsymbol{s}^Z} = \frac{1}{\phi}\frac{(\boldsymbol{e}^i)'\sum_{n=1}^{\infty}\left|\binom{-3/2}{n-1}\right|\boldsymbol{\Phi}^{n-1}\boldsymbol{s}^Z}{(\boldsymbol{e}^i)'\sum_{n=1}^{\infty}\boldsymbol{\Phi}^{n-1}\boldsymbol{s}^Z}.$$[6]

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 (2022). Our result directly links this measure to pass-through duration in a network when firms are completely myopic.

Under rational expectations, the delay due to downstreamness still exists but is shortened. 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 $\left|\binom{-3/2}{n-1}\right|$ (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 $\boldsymbol{\phi}$), duration becomes

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

with equality in continuous time. In general, $D_i \neq \frac{1}{\phi_i}\frac{(\boldsymbol{e}^i)'(\boldsymbol{I}-\boldsymbol{\Phi})^{-2}\boldsymbol{s}^Z}{(\boldsymbol{e}^i)'(\boldsymbol{I}-\boldsymbol{\Phi})^{-1}\boldsymbol{s}^Z}$, though this was the case under homogeneous θ. Intuitively, two sectors with the same frequency of price adjustment may

[6] Unfortunately, the coefficient $\left|\binom{-3/2}{n-1}\right|$ does not admit a clean functional form. For values $n \in \{1, 2, ..., 5\}$, it is 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.

vary still in pass-through speed when one has more price-flexible suppliers than the other. As a result, we view Proposition 3 as a ceteris paribus result characterizing how downstreamness delays pass-through, holding frequency of price adjustment constant.

Duration is defined conveniently to connect transitory dynamics to downstreamness in closed-form. Numerically, however, we can compute the pass-through delay due to downstreamness at any given moment in time. To define this, 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.[7]

Connecting to our previous result, this no-network model would have $\phi_i D_i = 1$, so that pass-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 no-network 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

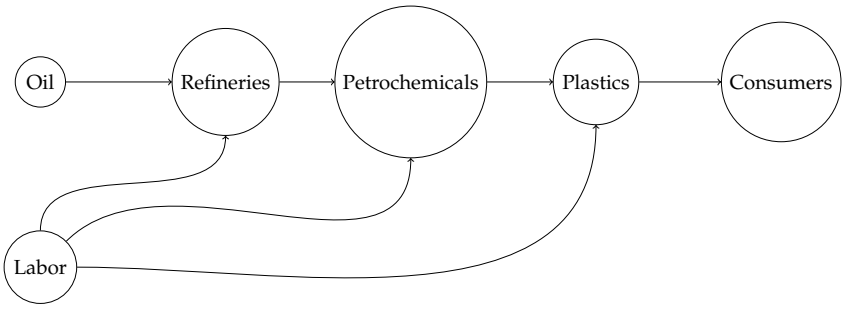
In this section, we develop intuition on how expectations influence the dynamics of commodity shock pass-through in a linear network. A linear network of length N is comprised of industries

[7] This 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 $m_f \to 1/(\theta_i \delta)$ (or as $\psi \to 0$ in continuous time). The rate of pass-through in the hyper forward-looking limit is

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

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.

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, petroleum refineries pay for labor and oil. The plastics sector pays for labor and inputs from the petrochemical sector but not from petroleum refineries directly – they only use petroleum refinery output indirectly through their use of petrochemical inputs. Ultimately, the supply chain provides plastic products to downstream consumers.

$n \in \{1, ..., N\}$, each populated by a continuum of firms that only use labor and an input from industry $n-1$. Industry 1 uses labor and an exogenous supplied commodity with price denoted $P_{0,t}$. Industry n's price is denoted $P_{n,t}$. This setup is nested in our general setup from the previous subsection. 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).

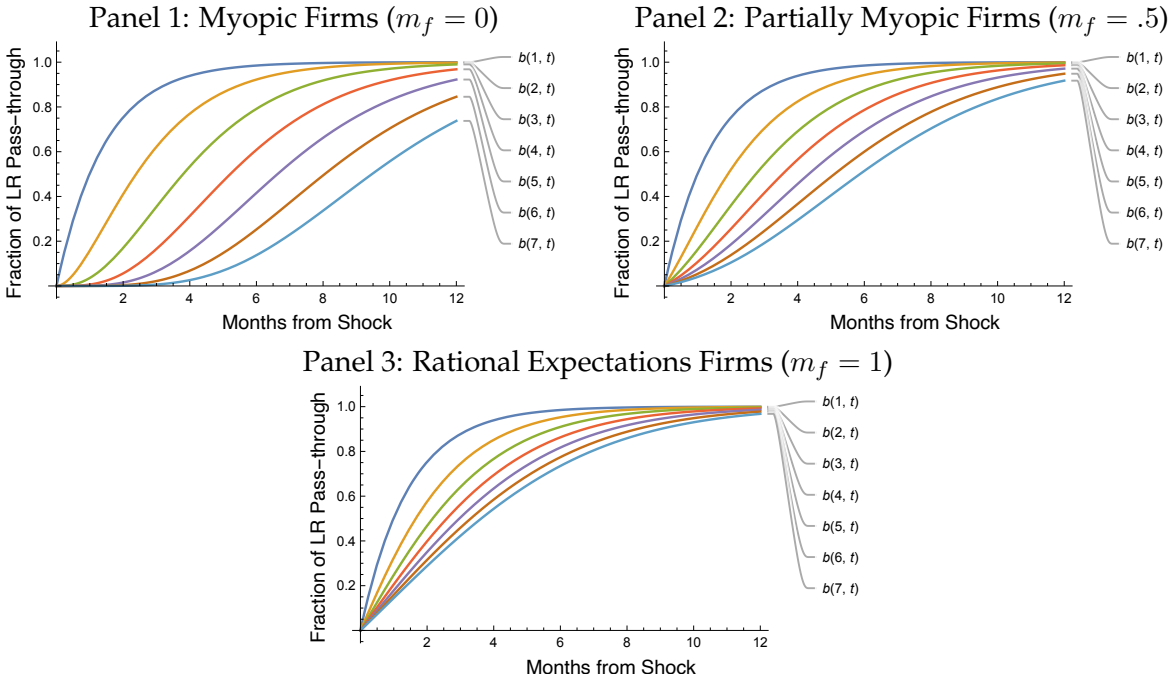
Now, we visualize how much longer pass-through takes to occur in each sector in the supply chain. We set $\delta = .96$ (a standard annual discount factor), and we set $\theta = .5$ (half of the firms in each industry getting the chance to change prices each period). Finally, we will consider rational 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

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



Note: We display the response of each of seven sectors' prices to an unexpected and persistent shock to the commodity price, or the price in sector 0. Sector n uses labor and inputs from sector $n-1$, and the fraction of long-run pass-through achieved in sector n by time t is denoted $b(n,t)$. The panels show how pass-through rates vary with the degree of firm forward-lookingness. We see that 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.

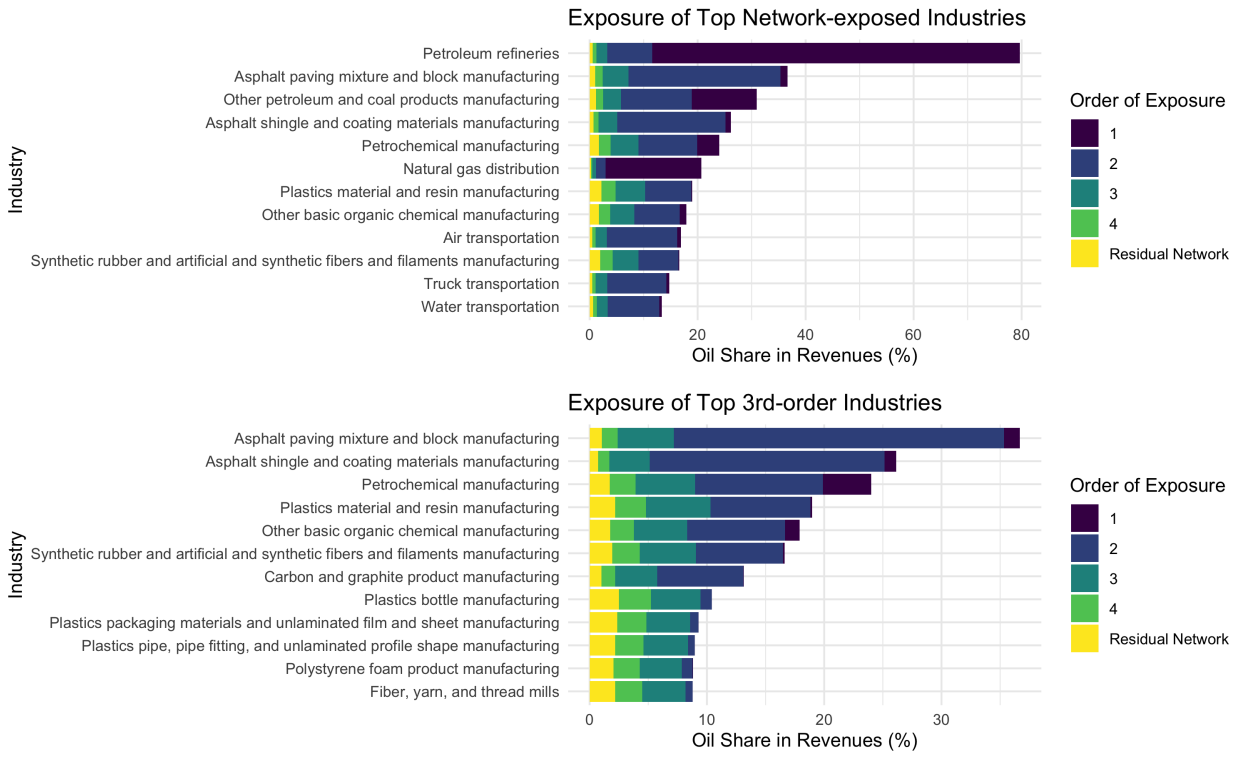
our propositions can be applied exactly. We see that $b_{1,t}$ does not vary with the degree of forward-lookingness. The rate of pass-through for every other sector, however, does vary, and the degree to which these rates vary is increasing in n, just as predicted under Proposition 3. This variation 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.

Every five years, the Bureau of Economic Analysis (BEA) publishes a detailed input-output table – approximately 400 sectors in size – which represents interdependencies between industries in the U.S. economy. Specifically, the tables display the extent to which the output of a given sector

Figure 3 – Network Oil Shares for a Selection of Industries



Note: This figure displays the total network share of oil - inclusive of both direct, first-order exposure and indirect, higher-order exposure - in each industry's revenues for a selection of industries. Panel 1 displays the twelve industries with the highest total 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%.

is used as an input by each other sector in the economy (or consumed by final demand). We outline our processing of the input-output tables in Appendix B.

To provide an example of the kind of information contained in these input-output tables, 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 procedure by which the data is produced is described in detail in BLS (2015). In short, the BLS collects price micro data on an extensive variety of individual goods. The BLS then computes industry prices in a given period by taking the average price of all transactions recorded 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. The BEA released correspondences between SIC codes and BEA codes with each input-output table through 1992, and it released correspondences between NAICS codes and BEA codes for the 1997 table onward. By utilizing these various correspondences, it is possible to merge the BLS industry price data with the BEA input-output tables.

Additionally, we obtain data 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.[8] 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 initial empirical analysis tests our theoretical results in a reduced-form way, allowing the data to speak with minimal added structure. Our reduced-form empirics (1) provide evidence that pass-through is gradual, particularly so for downstream firms, and (2) motivate the structural analysis of Section 6, which assesses how well our model can match pass-through dynamics in the data.

We first utilize a shift-share design taken directly from 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 then assess whether pass-through is gradual. Finally, we test whether industries more downstream from the commodity price increase experience less rapid pass-through.

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

5.1 Regression Specifications

We begin by testing the implication of Proposition 1 that pass-through of a commodity price movement to industry prices is governed by industries' network cost shares in the commodity. We regress the price change in an industry on the network-implied cost change due to movements in the price of a commodity or commodities of interest, the network-implied cost change due to movements in wages, and a time fixed effect. Formally, we take first-differences of Proposition 1 and add a time fixed effect to enable a shift-share interpretation:[9]

$$\Delta P_{i,t} = \lambda_t + \beta[(\boldsymbol{I} - \boldsymbol{\Phi})^{-1}\boldsymbol{s}^Z]_i \Delta P_{Z,t} + \gamma[(\boldsymbol{I} - \boldsymbol{\Phi})^{-1}\text{diag}(s_i^L)\Delta \boldsymbol{w}_t]_i + \epsilon_{i,t}, \tag{8}$$

where $P_{i,t}$ denotes the log price of industry i at time t, λ_t is a time fixed-effect, $[(\boldsymbol{I} - \boldsymbol{\Phi})^{-1}\boldsymbol{s}^Z]_i$ 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, $[(\boldsymbol{I} - \boldsymbol{\Phi})^{-1}\text{diag}(s_i^L)\Delta \boldsymbol{w}_t]_i$ 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. It is possible to vary the time horizon over which the differences Δ are taken, so we examine the extent of pass-through over both one-month and one-year horizons. If pricing were fully flexible, the model implies $\beta = 1$ no matter the first-differencing horizon.[10]

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

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.

[9] Recall that Proposition 1 assumed no shocks to TFP. The time fixed effect will not completely address comovements in aggregate TFP with oil price changes because industries have somewhat heterogeneous loadings on aggregate TFP in a production network model. In appendix C.1, we show that interacting a time fixed effect with industries' model-implied loadings on aggregate TFP does not meaningfully change our results.

[10] The model also implies that $\gamma = 1$. Pass-through of wages to prices is not the focus of this paper, however, and so we will interpret the term due to general equilibrium wage effects as a control. In Appendix Tables G.2 and G.3, we show that omitting this control does not meaningfully affect our results and that the general equilibrium wage effects are uncorrelated with the treatment of interest over both one-month and one-year horizons. We therefore argue that there is no loss in ignoring this term for the purposes of our analysis.

Consequently, the identification assumption for plim $\hat{\beta} = \beta$ is given by

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

where $\mu_t \equiv E[[(\boldsymbol{I} - \boldsymbol{\Phi})^{-1}\boldsymbol{s}^Z]_i\epsilon_{i,t}]$ 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 differentially for omitted reasons ($\Rightarrow \mu_t \neq 0$) in periods when shocks to the price of commodity Z also tend to occur ($\Delta P_{Z,t} \neq 0$) 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.

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.[11] We note that in all of our regressions, the shocked industries will be excluded from the regression analysis.[12]

Table 1 shows the results of the preceding regression specification. Panel 1 shows the re-

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

[12] For example, when we consider shocks to the oil price, we will only be interested in how that shock propagates to non-oil industry prices. Including oil in the regression introduces a source of mechanical dependence in the analysis; most obviously, 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.

Table 1 – Pass-through Regressions: The Effects of Cost Changes on Industry Price Changes

Panel 1: One Month Horizon

Dependent Variable: ΔPPI (Month-over-Month)	(1) Oil	(2) Oil (No Refineries)	(3) Kanzig	(4) Kanzig (No Refineries)	(5) Non-Oil
Total Network Cost Change	0.503***	0.252***	0.774***	0.266**	0.340***
	(0.145)	(0.093)	(0.295)	(0.125)	(0.071)
R-Squared	0.1188	0.1121	0.1211	0.1197	0.1561
Observations	103,480	103,186	97,324	97,048	97,974

Panel 2: One Year Horizon

Dependent Variable: ΔPPI (Year-over-Year)	(1) Oil	(2) Oil (No Refineries)	(3) Kanzig	(4) Kanzig (No Refineries)	(5) Non-Oil
Total Network Cost Change	0.853***	0.639***	0.974***	0.761***	0.737***
	(0.153)	(0.185)	(0.150)	(0.178)	(0.089)
R-Squared	0.1626	0.1378	0.1572	0.1361	0.1764
Observations	96,479	96,197	90,323	90,059	91,643

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. This table shows the results of regressions corresponding to Equation (8). That is, we regress industry price changes on cost changes resulting from the network effects of commodity price movements. Columns (1) and (2) focus on all oil price variation from 1997-2022. Columns (3) and (4) focus on oil price variation induced by Kanzig's (2021) OPEC shock series. Column (5) focuses on non-oil commodity price variation from 1997-2022. Standard errors are clustered by industry. We find that pass-through is far below the full pass-through benchmark of 1 over a one-month horizon. More pass-through accumulates over a one-year horizon.

sults for month-over-month changes, whereas Panel 2 shows the results for year-over-year long differences. Standard errors are clustered by industry in all cases.

Beginning with column (1) of Panel 1, we find much less than full pass-through of an oil price shock over a one-month horizon. The extent of pass-through into industry prices is 0.503 - just over half of full pass-through. Column (1) of Panel 2 reveals that, after a year, the extent of pass-through increases to 0.853 - much closer to full pass-through. Recalling the finding from 3 that the petroleum refinery industry is an extreme outlier in terms of both its overall network oil share and the proportion of this share that is due to direct, first-order exposure to oil, in column (2) we make one simple change: dropping petroleum refineries from our sample. This allows us to focus more on industries that are indirectly exposed to oil. Indeed, here we see that, over a one-month horizon, the extent of pass-through of an oil price shock into industry prices is 0.252 - about a quarter of full pass-through. Over the course of a year, this number increases greatly to 0.639 but still falls short of full pass-through. Differences in coefficients between the one-month and one-year horizon regressions are statistically significant in all cases.

In columns (3) and (4), we repeat the preceding exercise using a narrower but more plausibly exogenous source of variation. Specifically, we use 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. This leads to a qualitatively similar conclusion: full pass-through of an oil price shock into industry prices does not occur on impact. Some pass-through occurs on impact and more pass-through occurs over subsequent months.

In column (5), we repeat the exercise using variation in non-oil commodity prices. That is, we pool all commodities apart from oil and compute the network share in all these commodities by industry. The result is again qualitatively similar: some pass-through – but less than full pass-through – of the shocks in the month of impact, and substantially more pass-through after a year.

In Appendix C, we show that these results are robust to a variety of modifications and alternative approaches. In C.1, we repeat the analysis in Table 1 (i) without the time fixed effects, (ii) without the wage control variable, (iii) with an added TFP control variable, (iv) with an added control for network gas/electricity cost changes (two commodities likely to be close substitutes for oil), and (v) with cost shares that exclude payments to capital from the denominator. All of these exercises yield very similar results. In C.3, we modify our preceding regression specification by adding leads and lags to study the month-by-month dynamics of pass-through in a reduced-form manner. In C.4, we repeat this analysis using a local projections approach, finding similar results. In C.5, we show that the same patterns are again evident in a binscatter analysis. Finally, in C.6, we take a different approach to isolating exogenous variation. We focus on a few case studies likely to be highly exogenous - the 1979 oil shock, the 2014 oil shale boom, and the 2020 COVID shock - and show that these settings yield the same results as our pooled analysis.

5.3 Heterogeneity

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

$$\begin{aligned}\Delta P_{i,t} =& \lambda_t + \beta \left[(\boldsymbol{I} - \boldsymbol{\Phi})^{-1} \boldsymbol{s}^Z \right]_i \Delta P_{Z,t} \\ &+ \tilde{\beta} \left[(\boldsymbol{I} - \boldsymbol{\Phi})^{-1} \boldsymbol{s}^Z \right]_i \Delta P_{Z,t} \times heterogeneity_i \\ &+ \gamma [(\boldsymbol{I} - \boldsymbol{\Phi})^{-1} \mathrm{diag}(s_i^L) \Delta \boldsymbol{w}_t]_i + \epsilon_{i,t}\end{aligned}$$

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. 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

Table 2 – Heterogeneity Analysis

Dependent Variable: ΔPPI	(1) Oil	(2) Kanzig	(3) Non-Oil	(4) Oil	(5) Kanzig	(6) Non-Oil
Total Network Cost Change	0.673***	0.956***	0.397***	0.689***	0.981***	0.425***
	(0.061)	(0.238)	(0.099)	(0.048)	(0.223)	(0.105)
Total Network Cost Change	-0.141***	-0.156***	-0.067***	-0.120***	-0.130***	-0.055***
*Price Adjust. Half-Life	(0.018)	(0.039)	(0.018)	(0.017)	(0.034)	(0.016)
Total Network Cost Change				-0.212***	-0.274**	-0.126*
*Downstreamness				(0.054)	(0.113)	(0.065)
Total Network Cost Change						
*SD[Marginal Cost]						
R-Squared	0.1100	0.1103	0.0705	0.1111	0.1108	0.0728
Observations	85,667	79,511	1,507,249	85,667	79,511	1,507,249

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 over a one-month horizon. 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.

$X = .5$.

In Table 2, 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 excess pass-through time due to downstreamness (defined in Definition 3) as another interaction term. In the case of all oil price variation, the Kanzig IV variation, and the non-commodity price variation, the downstreamness interaction term is negative and statistically significant. Industries further downstream have lower levels of price pass-through over a one-month horizon. Because we have controlled for potential differences in rates 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.[13]

In Appendix C.2, we investigate heterogeneity on some other factors that could plausibly affect pass-through: concentration, firm size, inventories, and the capital share (all measured at

[13]We 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.

the sector level). Adding these additional heterogeneity terms to the regressions in the preceding table does not change the result that pass-through is slower for downstream firms.

In Appendix C.5, we plot binscatters of industry price change on network cost change over varying time horizons. This allows us to investigate some additional dimensions of potential heterogeneity: size and sign of shock, both of which would exist if second-order effects were operative. What we find is little to no evidence of heterogeneity in either of these dimensions. Specifically, after a year, we do not find evidence that the extent of pass-through is lesser for small cost changes than for large ones; this would entail a flat slope around the origin, which is not the case. Nor do we find an asymmetry on the sign of the shock; both negative and positive cost shocks are passed through into industry prices to the same extent. The slope of the line of best fit in these binscatter plots is also increasing in the time horizon over which the industry price change is examined. In other words, the binscatters once again reveal that pass-through tends to accumulate over time rather than be realized immediately upon shock impact.

Taken together, these initial empirical results suggest that long-run pass-through implied in Proposition 1 is achieved gradually, rather than on shock impact. They also suggest that a careful examination of the dynamics of pass-through - and how they are mediated by the production network structure - is in order.

6 Empirical Analysis using Model-derived Dynamics

Our model does not just determine the long-run pass-through of commodity price movements into industry prices. It also characterizes the transition dynamics to that long run, as seen in Proposition 2. Compared to the previous section, which focused largely on whether long-run pass-through occurred and whether it was gradual, this section treats the pass-through dynamics of the model seriously and tests whether they are empirically accurate.

We begin by describing a minimum distance procedure that would require estimating industry price responses to commodity price movements for each industry individually using time series variation. Motivated by some challenges inherent in this methodology – among them the noise in estimates, particularly for downstream firms – we develop a broadly useful generalized method of moments (GMM) approach for testing dynamic macroeconomic models using a shift-share technique in which the shares are time- and horizon-specific and the shifts are dynamic. This approach, we think, jointly removes the need to estimate pass-through industry-by-industry and enables convincing cross-sectional identification using commodity prices movements directly.

This section also contains our structural estimation of the degree of forward-lookingness, a parameter that can substantially affect the dynamics of price adjustment, especially for downstream industries, as shown in Proposition 3. We further validate these estimates by augmenting the model and structural analysis to handle expected deviations of shocks from complete persis-

tence as measured in futures markets—deviations that would be neglected by myopic agents.

6.1 Minimum Distance

To motivate our dynamic shift-share identification strategy, we first fix ideas through a straightforward minimum distance exercise. Simply put, we estimate the pass-through of oil price movements to industries' prices and compare those empirical impulse response functions to those we find when solving the model.

Formally, we estimate empirical pass-through of a persistent commodity shock by using the regression

$$\Delta P_{i,t} = \alpha_i + \sum_{h=0}^{24} \beta_{ih} \Delta P_{Z,t-h} + \epsilon_{i,t}, \tag{9}$$

where $\Delta P_{i,t}$ is the month-over-month change in industry i's log price, α_i is an industry-specific constant, $\Delta P_{Z,t-h}$ is the month-over-month change in the commodity price h periods ago, and $\epsilon_{i,t}$ is an error term. Aggregating the estimated coefficients, we form impulse response functions for each industry and horizon:

$$IRF_{i,H}^{Data} = \sum_{h=0}^{H} \hat{\beta}_{ih}.$$

Stacking across industries and horizons, we form the vector of impulse responses $\boldsymbol{IRF}^{Data}$. Note that these impulse responses are specific to the commodity chosen in regression equation (9), but we suppress this dependence to minimize additional notation.

After solving the model, it is straightforward to compute impulse response functions for each industry and horizon. Solving the model numerically requires some calibration, which we denote by $\boldsymbol{\alpha}$, the vector of all parameters we would like to estimate. We specify these parameters in more detail later, but they include, for example, the degree of forward-lookingness m_f. Stacking impulse responses from the model's solution across industries and horizons in the same order we did for the empirical IRFs, we form the empirical moment

$$\hat{\boldsymbol{m}}(\boldsymbol{\alpha}) \equiv \boldsymbol{IRF}^{Data} - \boldsymbol{IRF}^{Model}(\boldsymbol{\alpha}).$$

The parameter vector $\boldsymbol{\alpha}$ is then selected to minimize squared error. The moments can be weighted so that the procedure is efficient by using an estimate of inverse variance of the moments. This estimate can be constructed using the variance matrix estimated in the panel regression (9). The optimal parameters solve

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

While we set up this problem in the space of impulse response functions, it would have been equivalent to match the β's estimated in regression (9) with their model variants, which we term

the pass-through coefficient

$$passthrough_{i,h}(\boldsymbol{\alpha}) = IRF_{i,h}^{Model}(\boldsymbol{\alpha}) - IRF_{i,h-1}^{Model}(\boldsymbol{\alpha}),$$

where $IRF_{i,-1}^{Model}(\boldsymbol{\alpha})$ is set to 0. Formally, it is also true that

$$\hat{\boldsymbol{\alpha}} = \arg\min_{\boldsymbol{\alpha}} \left(\hat{\boldsymbol{\beta}} - \boldsymbol{passthrough}(\boldsymbol{\alpha})\right)' \hat{\text{Var}}(\hat{\boldsymbol{\beta}})^{-1} \left(\hat{\boldsymbol{\beta}} - \boldsymbol{passthrough}(\boldsymbol{\alpha})\right).$$

6.2 Shift-share GMM Motivation

There are several challenges introduced by the minimum distance design. Foremost is whether the estimates of $\boldsymbol{\beta}$ in regression (9) are unbiased. Even oil price changes are often correlated with, for example, movements in aggregate TFP or business cycle developments that also influence prices. Movements in other commodity prices are typically seen as even less exogenous than oil price movements, making this problem worse when Z is not set to oil. Beyond bias, the industry-specific estimates from regression 9 reveal substantial noise outside of industries such as petroleum refineries. We specifically want to focus on the timing and extent of pass-through to downstream firms, where it is difficult to estimate pass-through of commodity price movements using a time series approach, and where the degree of forward-lookingness matters substantially.

Second, the coefficients in regression (9) also capture general equilibrium effects of oil price movements, such as reactions in inflation expectations and responses by the monetary authority. Such effects are thought to generate movements that affect all industries similarly much more than than magnifying or attenuating relative price movements. For this reason, results we present using the minimum distance approach will use a fully specified macroeconomic model that includes sticky wages and reactions by the monetary authority,[14] while results we present using the approach below can use the supply-side results presented in Proposition 2.

To address these issues jointly, we propose using a shift-share design with the model's pass-through coefficients $passthrough_{i,h}(\boldsymbol{\alpha})$ as industry- and horizon-specific shares and the commodity price movement as the shift. We can test whether pass-through in the data follows the model's predictions using the regression

$$\Delta P_{i,t} = \lambda_t + \sum_{h=0}^{H} \beta_h(\boldsymbol{\alpha}) passthrough_{i,h}(\boldsymbol{\alpha}) \Delta P_{Z,t-h} + \epsilon_{i,t}, \tag{10}$$

where $\Delta P_{i,t}$ is the month-over-month change in industry i's log price, λ_t is a time fixed effect, $\Delta P_{Z,t-h}$ is the month-over-month change in the commodity's log price h periods ago, and $\epsilon_{i,t}$ is an error term. If the model is capable of fitting the patterns in the data and our OLS is unbiased, we should have the vector $\boldsymbol{\beta}(\boldsymbol{\alpha}_0) = \mathbf{1}$ for some calibration $\boldsymbol{\alpha}_0$.

[14] The details for the fully specified model can be found in Appendix A.4.

Compared to regression equation (8), our shift-share design from the reduced-form section, we now have a serious treatment of the horizon-composition of the industry's network cost share.[15] Formally, since $\sum_{h=0}^{\infty} passthrough_{i,h}(\boldsymbol{\alpha}) = [(\boldsymbol{I} - \boldsymbol{\Phi})^{-1}\boldsymbol{s}^{Z}]_i$, a unit shock to the commodity price has the same long-run effect as in the prior regression model. Because shocks take time to propagate throughout the network, there is now a role for lagged commodity price changes to predict current industry price changes through the component of network propagation that has yet to occur.

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

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

where $\hat{\beta}_h(\boldsymbol{\alpha})$ are the estimates from regression equation (10), to the model-derived IRFs, $IRF_{i,H}^{Model}(\boldsymbol{\alpha})$. While we will still refer to $IRF_{i,H}^{Data}(\boldsymbol{\alpha})$ as "the empirical IRF," it is crucial to note that this measure varies with $\boldsymbol{\alpha}$ and so is not independent of the model. We find this empirical IRF to be useful as it does not require estimating empirical IRFs for each industry using time-series variation; instead, it is identified using a dynamic shift-share design. Importantly, this means our GMM procedure is no longer minimum distance but rather a "true" GMM in which the empirical moments are model-dependent.

Why do we not consider a minimum distance procedure instead using our shift-share estimates from the previous section? This approach is flawed because the model-consistent calculation of the shift variable used to estimate the empirical moments changes with the model parameters being estimated in our setting. Our approach correctly adjusts the empirical moments as the underlying model changes.

If the model is capable of fitting the data for the calibration $\boldsymbol{\alpha}_0$ and our OLS estimates from regression equation (10) are unbiased, then our empirical and model-derived IRFs should be identical, and we have the moment equation

$$m(\boldsymbol{\alpha}_0, H) = \mathrm{E}[IRF_{i,H}^{Data}(\boldsymbol{\alpha}_0) - IRF_{i,H}(\boldsymbol{\alpha}_0)] = 0.$$

In our GMM estimation, we use $m(\boldsymbol{\alpha}_0, H) = 0$ 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 $\boldsymbol{\alpha}$, we will estimate the model, run regression (10) with H lags, construct IRFs for each industry, and generate the quadratic error measuring how far the moment conditions are from 0. Then we choose the $\boldsymbol{\alpha}$

[15] We have also now completely dropped the general equilibrium wage control. As discussed previously, in Appendix Tables G.2 and G.3, we showed that omitting this control does not meaningfully affect our results and that the general equilibrium wage effects are uncorrelated with the treatment of interest over both short and long horizons. Any worries that this result changes meaningfully if a dynamic path of sticky wages is considered are handled by our robustness check with the minimum distance procedure, which uses a fully closed model incorporating gradual adjustment of sticky wages. The results from this fully specified model are not meaningfully different.

that minimizes the error. Standard errors are clustered at the industry level.

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.

6.2.1 Specifications Testing Forward-lookingness about Network Dynamics

Now, we preview that one element of $\boldsymbol{\alpha}$ is the myopia parameter, m_f. For any estimate $\hat{\boldsymbol{\alpha}} = (\hat{\boldsymbol{\alpha}}_{-m_f}, \hat{m}_f)$ such that $\hat{m}_f \neq 0$, we can measure how much industries are increasing prices due to forward-lookingness using

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

with

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

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

$$\begin{aligned} \Delta P_{i,t} = \lambda_t &+ \sum_{h=0}^{H} \beta_{h,\text{Myopic}} passthrough_{i,h}(\hat{\boldsymbol{\alpha}}_{-m_f}, m_f = 0) \Delta P_{Z,t-h} \\ &+ \sum_{h=0}^{H} \beta_{h,\text{FLGap}} FLGap_{i,h}(\hat{\boldsymbol{\alpha}}) \Delta P_{Z,t-h} + \epsilon_{i,t}, \end{aligned} \tag{11}$$

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.

6.2.2 Specifications Testing Forward-lookingness about Shock Persistence

We have so far assumed that all commodity price changes are fully persistent. In appendix E, we show that futures data can be used to measure expected movements in future commodity prices. The expanded model advances upon our pass-through object $passthrough_{i,h}(\boldsymbol{\alpha})$, delivering predictions $passthrough_{i,h,m}(\boldsymbol{\alpha})$: how much industry i should change prices in response to a unit shock to commodity prices h periods ago at maturity m. This is a rich object: fix $h = 0$ and con-

sider $m = 0$ and $m = 1$. For $m = 0$, the object is the predicted effect of increasing commodity 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 commodity 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 $passthrough_{i,h}(\boldsymbol{\alpha}) = \sum_{m \geq 0} passthrough_{i,h,m}(\boldsymbol{\alpha})$, i.e. pass-through of a level shift up in the commodity futures curve.

To test whether agents pass through changes in the futures curve, holding changes in the spot price constant, we consider the regression

$$\begin{aligned} \Delta P_{i,t} =& \lambda_t + \sum_{h=0}^{H} \beta_{h,Spot}(\boldsymbol{\alpha}) passthrough_{i,h,m=0}(\boldsymbol{\alpha}) \Delta P_{Oil,t-h} \\ &+ \sum_{h=0}^{H} \beta_{h,Futures}(\boldsymbol{\alpha}) \sum_{m=1}^{M} passthrough_{i,h,m}(\boldsymbol{\alpha}) \Delta f_{Z,t-h,m} + \epsilon_{i,t}. \end{aligned} \tag{12}$$

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

6.3 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, $\boldsymbol{\theta}$, (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 $\boldsymbol{\Phi}$, and (5) the commodity cost shares $\boldsymbol{s}^Z$.

We take the frequencies of price adjustment from the Pasten et al. (2017) data. We estimate the myopia parameter, m_f. We calibrate a standard annualized discount factor, $\delta = 0.96$. Finally, we use the input-output data as much as possible to calibrate the input-output matrix and commodity cost shares, but these data lack measures of markups required to convert input shares in sales to input shares in cost; the key difference between sales and cost are capital payments, which are comprised of rental payments (part of cost) and profits (not part of costs).

To overcome this issue, we estimate how much payments to capital are profits, on average. In particular, because rental payments are weakly positive in every period, industry profits cannot exceed measured payments to capital in the industry. We also impose that steady-state profits are

weakly positive. Therefore, we can estimate the profit share of payments to capital, $\gamma \in [0, 1]$, leading to cost calibrations in each industry of

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

Then the input-output matrix is calibrated using $\Phi_{i,j} = P_i X_i^j / C_i(\gamma)$, and the commodity cost shares are calibrated using $s_i^Z = P_Z Z_i / C_i(\gamma)$.

Note that, though we estimate a common γ for all industries, the industry markup is heterogeneous: $\mu_i(\gamma) = P_i Y_i / C_i(\gamma)$. We will present the results of our estimation of γ in terms of the markup for the average industry, i.e. $\overline{\mu}(\gamma) = \frac{1}{I} \sum_i \mu_i(\gamma)$, with standard errors computed using the delta method and the GMM standard errors for $\hat{\gamma}$.

The model solutions also 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.[16] 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. Finally, futures prices may provide some information about how a given commodity price movement is expected to evolve in the future. We incorporate futures explicitly in two subsections here and provide additional details in Appendix E.

Model identification requires that, for some calibration $\boldsymbol{\alpha}_0 = (m_{f,0}, \gamma_0)$, the model's impulse response functions are expected to be the same as their empirical variants ($m(\boldsymbol{\alpha}_0, H) = 0$). We want to note that this is not mechanically the case, because cumulative pass-through in an industry is, at every horizon, bounded below by that coming from complete myopia, $m_f = 0$, and no industry profits $\gamma = 0$, and bounded above by that coming from rational expectations, $m_f = 1$, and no rental payments, $\gamma = 1$. The model may not be identified, for example, if there is no pass-through to indirect users of the commodity or if pricing is completely flexible. Fortunately, we have already rejected these alternative hypotheses in Section 5.

Now, we conduct our analysis for both oil and non-oil commodities. Because the industries using these commodities in their supply chains are different, we estimate a different γ for each case to capture the potential for heterogeneity in profit rates across industries. Though we have not microfounded the degree of myopia, we will also allow it to vary with the commodity being shocked. This will allow us to establish facts about whether firms pass through some shocks as if they are forward looking and some shocks as if they are myopic.[17]

[16] 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)).

[17] 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

6.4 Results for Oil

First, we use our shift-share GMM framework to estimate $\boldsymbol{\alpha}$ using oil price movements. When equally weighting the moments (1-step GMM), we find $\hat{m}_f = 1$ (SE = 0.224) and $\hat{\overline{\mu}} = 1.231$ (SE = 0.098).[18] Therefore, the data prefers 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 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 plot our results.

Panel 1 of Figure 4 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 adequate for the duration of the IRF, and the average industry requires about 20 months to reach 75% of long-run pass-through.

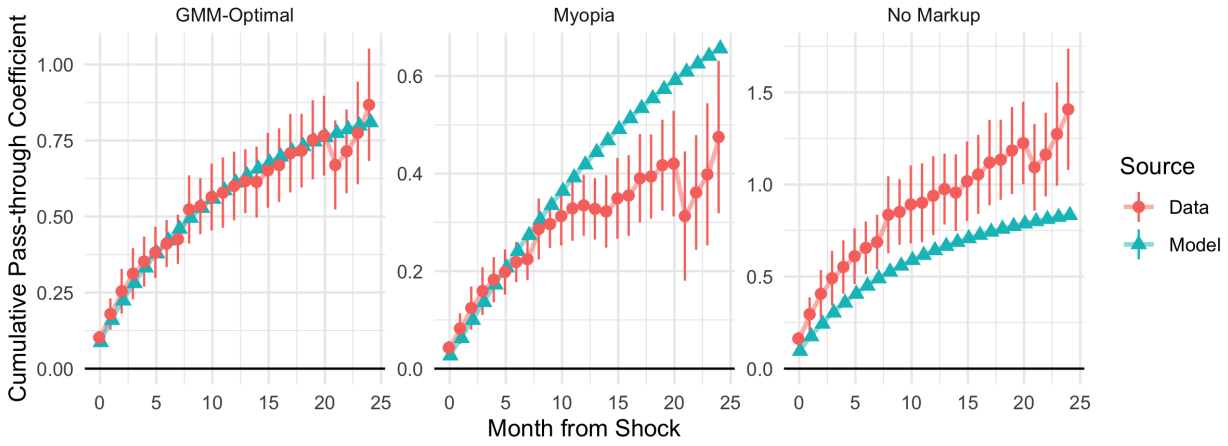
Panel 2 uses the GMM-optimal model but sets $m_f = 0$ rather than using the GMM-optimal rational expectations. Early on, pass-though is not substantially different, suggesting that very upstream industries are identifying these coefficients – recall from Proposition 3 that myopia does not affect pass-through speed for direct users of the commodity when the shock is unexpected and fully persistent. Later in the IRF, however, the myopic pass-through measures predict too much pass-through, and so the β's estimated from regression (10) at higher lags are on average less than one.

Finally, Panel 3 uses the GMM-optimal model but sets $\overline{\mu} = 1$ 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 more pass-through in the data than predicted by the model under a calibration of no markup. By revising the estimated markup up from 1, the GMM-optimal model is able to improve fit.

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.

[18] As in our empirics in the previous section, standard errors are clustered by industry.

Figure 4 – 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 and 3. Under alternative calibrations, model fit is meaningfully worse. Panel 2 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 3 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.

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 $\hat{m}_f = 1$ (SE = 0.101), and $\hat{\bar{\mu}} = 1.048$ (SE = 0.043). In other words, the rational expectations result is unchanged. The only difference here is a lower markup and smaller standard errors.

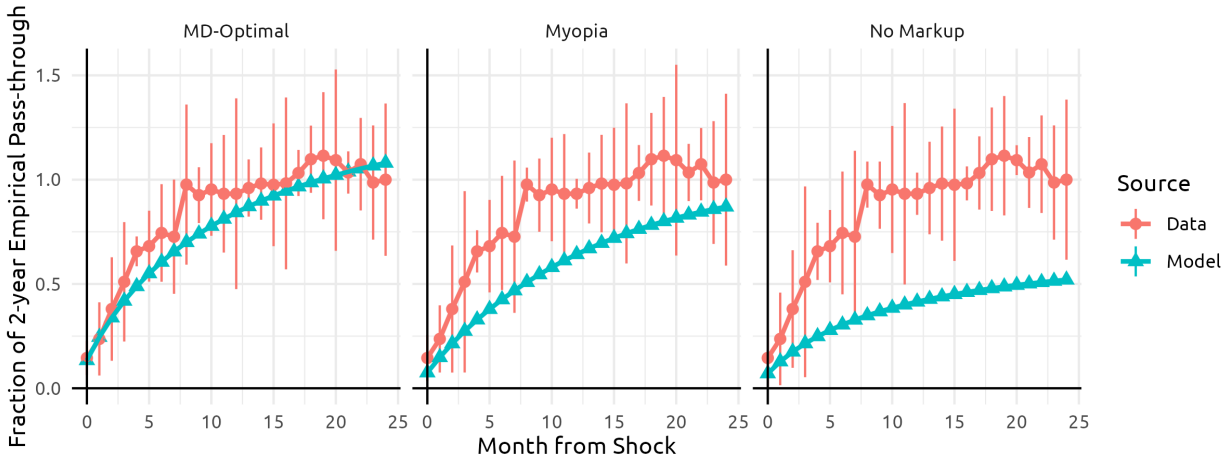
6.4.1 Minimum Distance Results

While the time series estimates of empirical impulse response functions for each industry are very noisy, even using all oil price variation, the minimum distance procedure does yield similar results. Using minimum distance with the efficient weight matrix, we find rational expectations, $m_f = 1$ (SE = 0.092), and a markup over marginal cost, $\bar{\mu} = 1.312$ (SE = 0.052).[19] The minimum distance procedure tries to fit every industry's impulse response at every horizon, and weights these responses by the inverse variance matrix. It is not quite possible to show model fit in impulse response form as a result — the set of industries identifying coefficients changes at each horizon, and such a plot, even with standard errors, does not show how covariances between these estimates at each horizon affect the result.

We can and do, however, show model fit for the equally weighted average industry, the same industry shown in in Panel 1 of Figure 4. To keep the empirical moment independent of the model, the key benefit of this approach, we scale all series by the cumulative pass-through achieved in two years according to the empirical pass-through estimates. A complex covariance

[19]This procedure was performed on the same sample as the shift-share design. For example, petroleum refineries were excluded from the analysis.

Figure 5 – Minimum Distance Results for Industry Pass-through of Oil Price Changes



Note: The figure depicts empirical and model derived pass-through of a persistent oil price increase for the average industry. As in our shift-share approach, we see that minimum distance estimation favors rational expectations and a markup over marginal cost to fit oil pass-through patterns in the data. Note that this figure scales all series by 2-year empirical pass-through to keep the empirical moment independent of long-run pass-through in the model, which depends on the estimated markup.

pattern is evident in the standard errors for the empirical moments; strong negative covariance at certain horizons means we are only relatively certain about the degree of pass-through at a few horizons (e.g., 7 months, 13 months, etc.). Despite this noise, we retain reasonable standard errors on parameter estimates.

We also point out that the relationship between the IRFs from the myopic model and the data reverses between the GMM procedure and the minimum distance procedure. In Figure 5, the data series is independent of the model series, and the model series clearly shows less pass-through than is present in the data, i.e. it is below the data series. In Figure 4, the data series depends on the model. For longer horizons, when myopia is relatively more important for pass-through, the model predicts far more pass-through than is present in the data, leading to estimated β's below 1 and a resulting "empirical IRF" that lies below the model series.

Recall that because the empirical impulse responses in the minimum distance approach include general equilibrium effects that affect all industries, such as movements in the aggregate wage or the response of the monetary authority, we use a fully closed model for the results shown in Figure 5. Summarizing, the fully closed model uses a standard household setup with sticky wages (e.g., Galí 2015), nested CES production in each industry (e.g., Rubbo 2020), and a Taylor rule. Full details are contained in appendix A.4.

6.4.2 Cross-sectional Fit Compared to Section 5

In section 5, we found that one-month and one-year network pass-through of commodity price movements was more limited than long-run pass-through, which would be reached instanta-

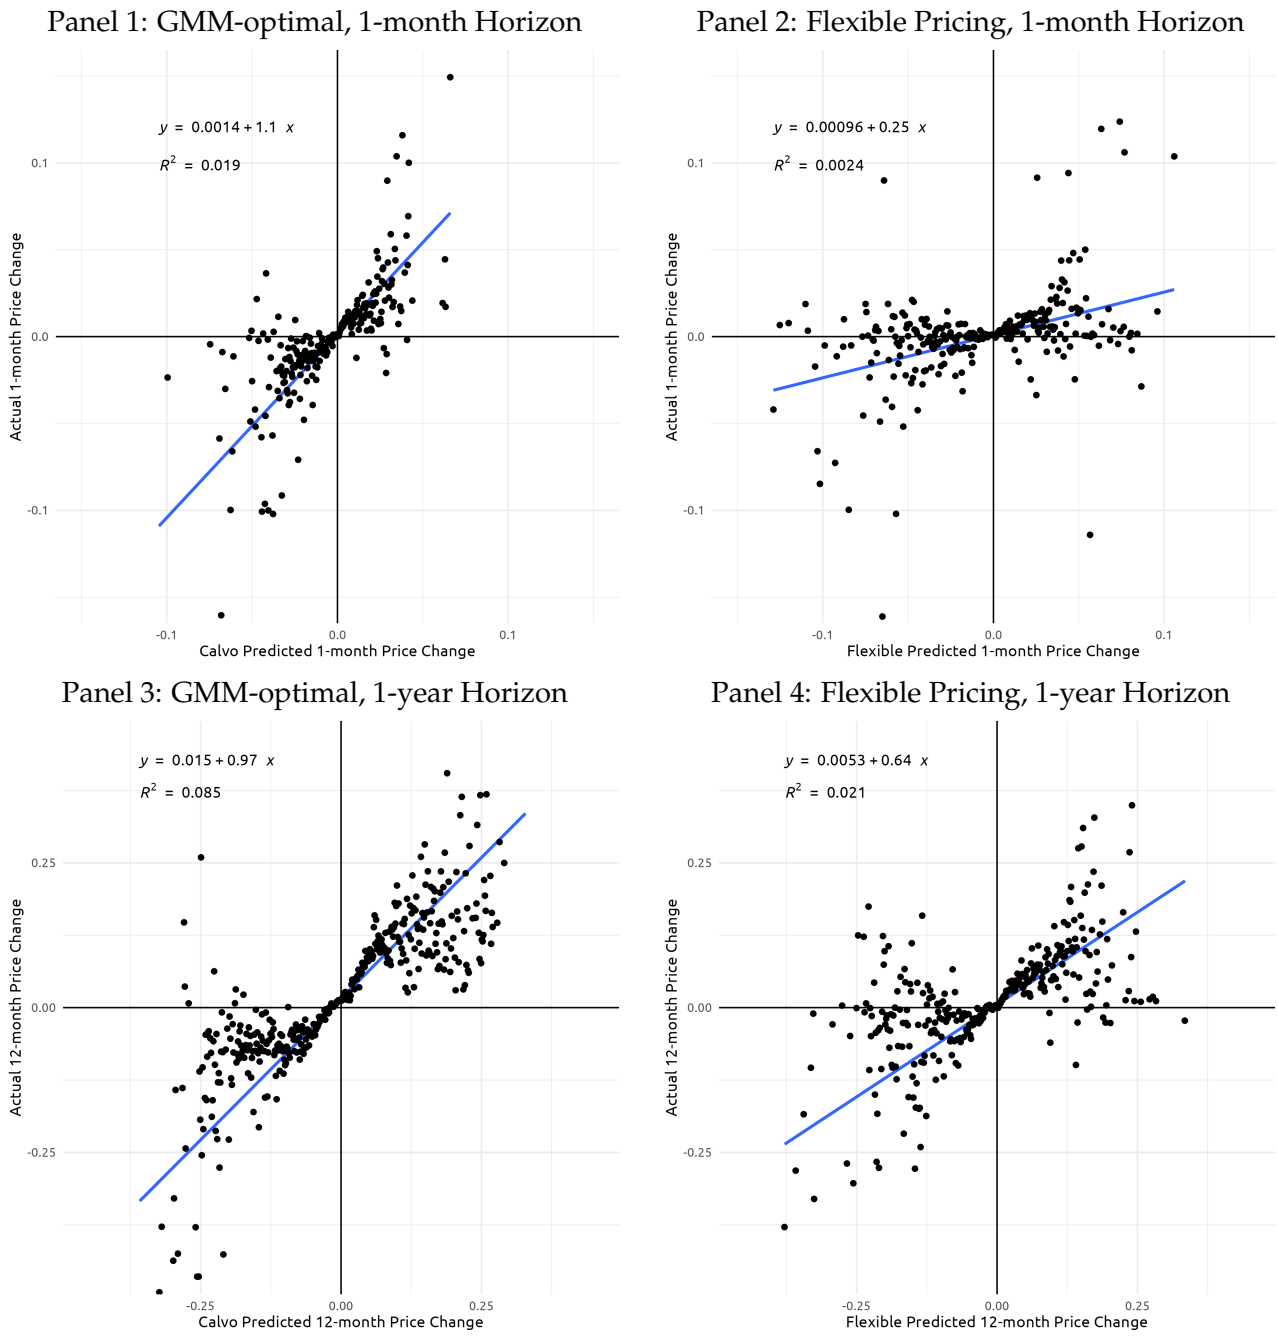
neously if pricing was fully flexible. In Figure 6, we show the fit of that model's predictions in the cross section, using binned scatter-plots, compared to the fit of the GMM-optimal dynamic model. As in section 5, we analyze model fit over both one-month and one-year horizons. To let the data speak as clearly as possible, we do not residualize first for time fixed effects. Each point in the plot can be interpreted as showing how much the x-axis delineated model predicts an industry should change prices in response to an oil price movement, compared to the industry's realized price change. These points are averages, as each binned scatter-plot contains 500 bins.

Comparing the GMM-optimal fit (Panels 1 and 3) to fit under the model that would be correct if pricing were fully flexible (Panels 2 and 4), we first notice that the GMM-optimal model has pass-through coefficients that are much closer to 1 at both horizons. This means that the GMM-optimal model is accurately predicting empirical pass-through at both one-month and one-year horizons, as opposed to the flexible pricing model that predicts too much pass-through at both one-month and one-year horizons. While the GMM coefficients are partially a result of the moment targeting, we reiterate here that the model did not mechanically have the ability to fit any pass-through data, and we used just two parameters to match average pass-through at 25 distinct horizons. We also highlight that, with just two parameters, we did not mechanically have the ability to explain the large heterogeneity that exists in the data. Despite this, the GMM-optimal model has a much higher R^2 in the panel of PPI changes. In particular, the GMM-optimal model explains 8.5% of the variation in one-year PPI changes, compared to just 2.1% in the model that would be correct under flexible pricing. The R^2 of the GMM-optimal model in one-year PPI changes is comparable to the R^2 of a time fixed effect in one-year PPI changes.

6.4.3 Upstream and Downstream Industries

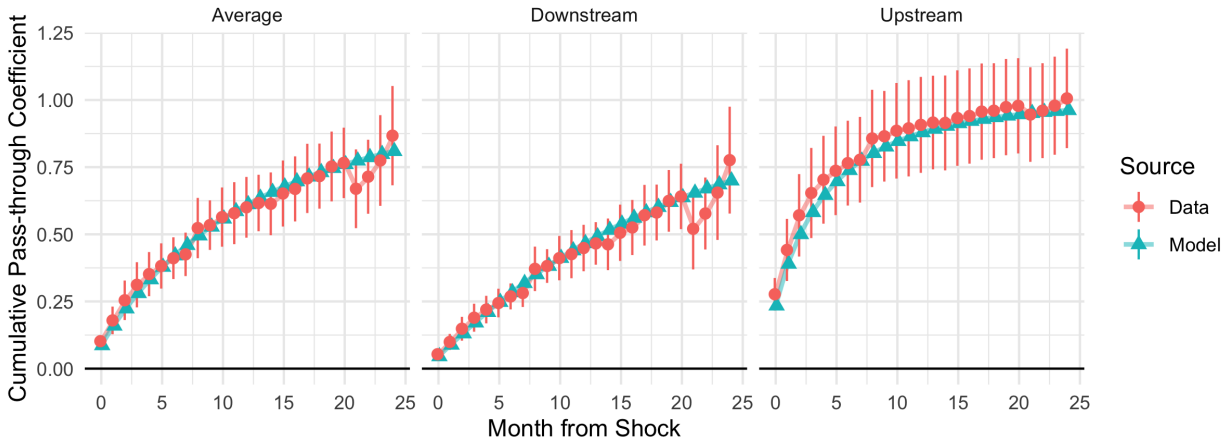
Rather than looking at the IRF for the average industry, we can visualize the solution for the 10% most upstream and downstream industries. These IRFs let us examine how close the β's from regression (10) are to 1 for different lags, since upstream industries pass through the shock faster than the average industry, while downstream industries pass through the shock more gradually than the average industry. In Figure 7, 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.

Figure 6 – Cross-sectional Fit of the GMM-optimal Model for Oil Price Changes



Note: These plots show fit of the x-axis designated model predictions on realized changes in PPIs. Each binned scatterplot contains 500 bins. Compared to the right column panels using the model from our reduced-form section 5, we see that the GMM-optimal model has a pass-through coefficient closer to one and a higher R^2 at both one-year and one-month horizons.

Figure 7 – 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 good 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 6 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.

6.4.4 The Role of Forward-lookingness about Network Dynamics

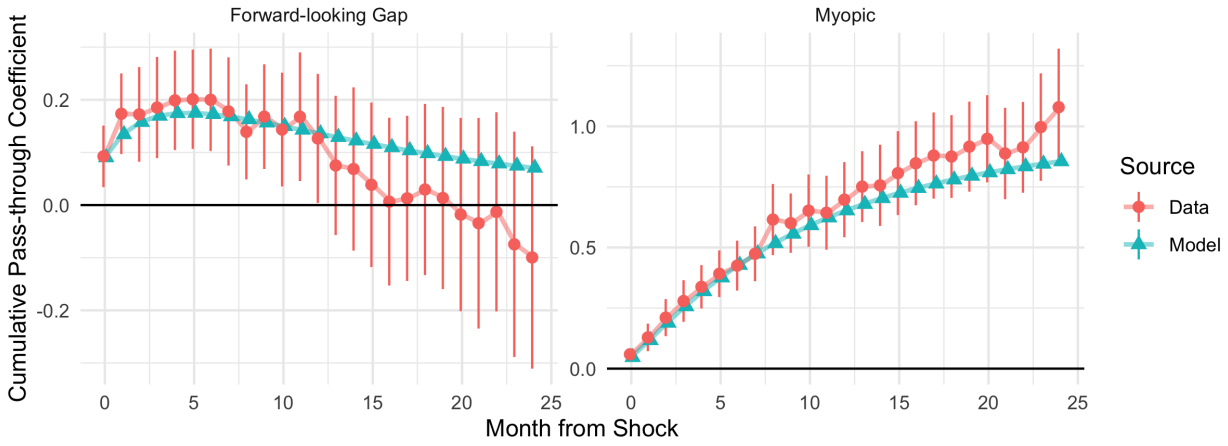
Now, recall that regression (11) split the GMM-optimal model's solution into the myopic solution and the component due to forward-lookingness about network dynamics. If forward-lookingness is truly operative, it should have predictive power in the cross-section: industries that should increase prices 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 notional 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

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

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

We show our results in Figure 8. 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.

Figure 8 – Tests of Forward-lookingness about Network Dynamics



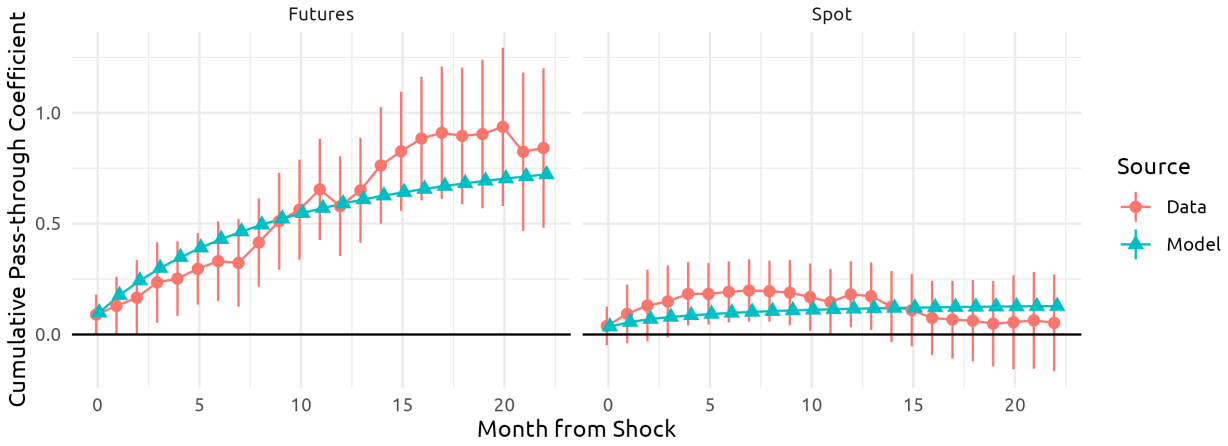
Note: We plot the results of our test that pass-through due to forward-lookingness about network dynamics under rational expectations is present in the data. Statistical significance in the left panel implies we cannot reject that firms are forward-looking about the gradual pass-through of upstream shocks to their marginal costs. The fact that the model lies within the standard error bars in the left figure, moreover, visualizes that we cannot reject a degree of forward-lookingness consistent with rational expectations. Further, we see that in the top 10% of industries affected by forward-lookingness, rational expectations provides a pass-through boost nearing 20% of long-run pass-through five months after the oil price change.

6.4.5 The Role of Forward-lookingness about Shock Persistence

The finding of rational expectations suggests that firms should be forward-looking not only about gradual pass-through of oil price movements through supply chains 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. The futures test we described represents an experiment asking whether firm pass-through differs under different expected paths for oil prices, measured with oil futures, holding the price change constant in the spot market.

We show our results in Figure 9, which decomposes the impulse response of a level shift in the entire futures curve into the pass-through component due to the change in the spot price and the component due to the change in futures prices. Our results are consistent with the model's predictions, and rational expectations in particular, that pass-through of spot price changes differs substantially with the change in the futures curve. In Appendix E.3, we confirm these findings in a reduced-form approach consistent with our empirical analysis from Section 5.

Figure 9 – Tests of Forward-lookingness about Oil Shock Persistence



Note: We plot the results of our test that pass-through due to forward-lookingness about oil shock persistence under rational expectations is present in the data. Statistical significance in the left panel implies we cannot reject that firms are forward-looking about the persistence of an oil price shock; formally, firms pass-through the same shock to the oil spot price differently under different shifts in the oil futures curve. The fact that the model lies within the standard error bars in the left panel visualizes that we cannot reject a degree of forward-lookingness consistent with rational expectations. We are only powered to perform this analysis out to 22 lags rather than 24.

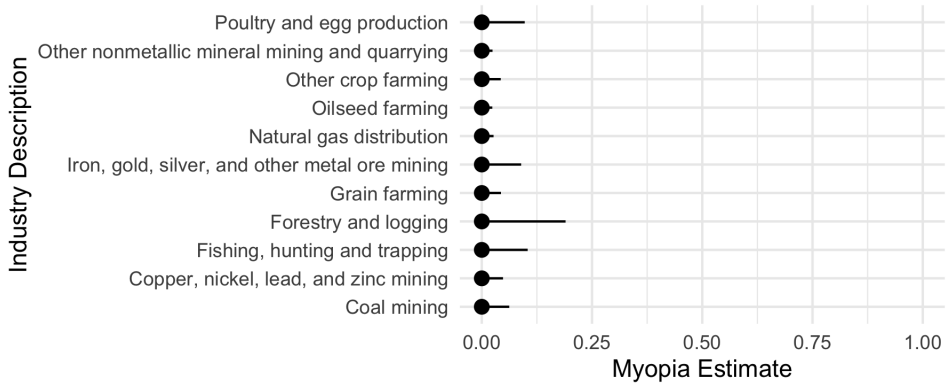
6.4.6 Homogeneous Frequency of Price Adjustment

We assess whether using the heterogeneous frequencies of price adjustment from Pasten et al. (2017) are needed to fit the data. We show in Appendix D.2 that, if we instead use the same frequency of price adjustment for all industries, set equal to the equal-weighted average frequency of price adjustment in the Pasten et al. (2017) data, model fit is substantially worse than the fit of the GMM-optimal model. The price duration associated with this average frequency is five months. Note that in this case, industries still have different rates of pass-through because they are differentially downstream from the shock (as predicted by Proposition (3)).

6.4.7 Large versus Small Shocks

While we have assumed time-dependent pricing, a model with state-dependent pricing might find that industries pass through large movements in commodity prices faster than small movements in commodity prices. We partition monthly oil price movements into two sets: those whose absolute size is larger than the median absolute size of oil price movements starting in 1997, and those whose absolute size is smaller. By a result from Auclert et al. (2022), if state-dependent pricing is operative, we should find evidence of faster pass-through for large shocks than predicted under our GMM-optimal time-dependent pricing model. In appendix D.4, we show that the GMM-optimal model retains good fit for both large and small oil price movements, and there is no statistically detectable difference in pass-through of small versus large shocks.

Figure 10 – Estimates of Myopia for Non-oil Commodities



Note: Unlike the case of oil, where the GMM-optimal model used a myopia estimate of 1, we find that complete myopia (an estimate of 0) is required to fit the pass-through of non-oil commodity price movements. The standard errors are relatively tight, suggesting meaningfully worse model fit for higher levels of forward-lookingness, and certainly allow us to reject rational expectations.

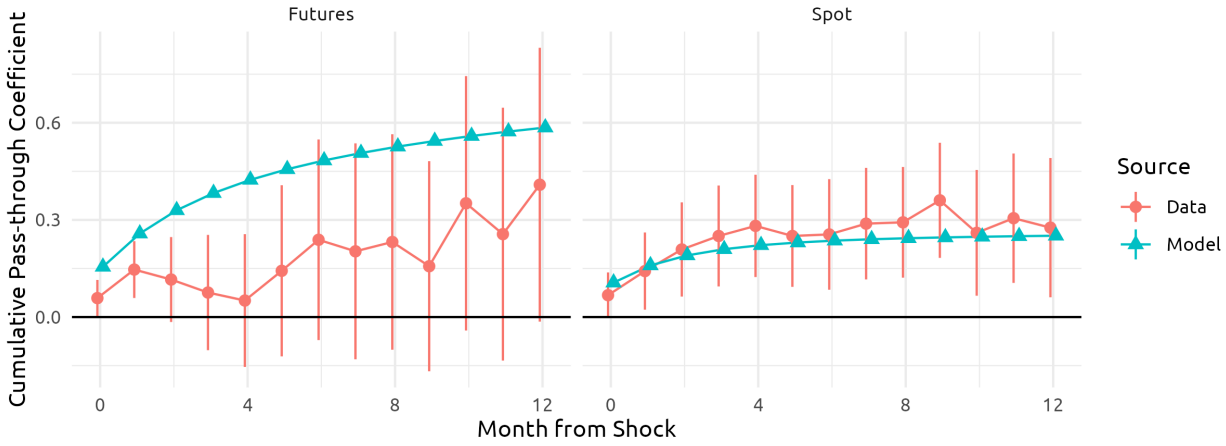
6.5 Non-oil Commodities

Using the same procedure we used for oil but for each non-oil commodity separately, we estimate $\boldsymbol{\alpha} = (m_f, \gamma)$ for each of fourteen commodities – all of the commodities with price movements for which we cannot reject the random walk model. We pool pass-through for shocks at lags 13-24 into a single treatment term as we possess substantially less power for these commodities than we do with oil. We find that we do not have substantial power to estimate the markup precisely for most non-oil commodities, but we do have power to estimate the degree of myopia for most commodities. We show our results in Figure 10. For all precisely estimated cases, we find $m_f = 0$, or complete myopia. The only outliers are (1) beef cattle ranching and farming, (2) animal production, except cattle/poultry/eggs, and (3) dairy cattle and milk production, and in these cases the standard errors reveal we have no power to estimate myopia in the data – they span the whole space for the parameter. 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.[20]

In this sense, oil is very special among the commodities. Downstream industries, even those that do not use oil directly in production, act as though they 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 act as though they are less attentive to the cost increases of their upstream suppliers. This underlines how the delay in pass-through due to downstreamness can be even more important than one would expect under rational expectations. We show in Appendix D.5 that our results are again robust to a 2-step GMM procedure, with identical point estimates and smaller standard errors.

[20]Standard 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).

Figure 11 – Tests of Forward-lookingness about Grain Shock Persistence



Note: We plot the results of our test that pass-through due to forward-lookingness about grain shock persistence under rational expectations is present in the data. Statistical insignificance in the left panel implies we can reject that firms are forward-looking about the persistence of a grain price shock; formally, firm pass-through is the same for all changes in grain spot prices, regardless of the shift in the grain futures curve. Point estimates above 0 in the left panel suggest that there may be some forward-lookingness, but this evidence is much weaker than the evidence we found for oil. We are only powered to perform this analysis out to 12 lags rather than 24.

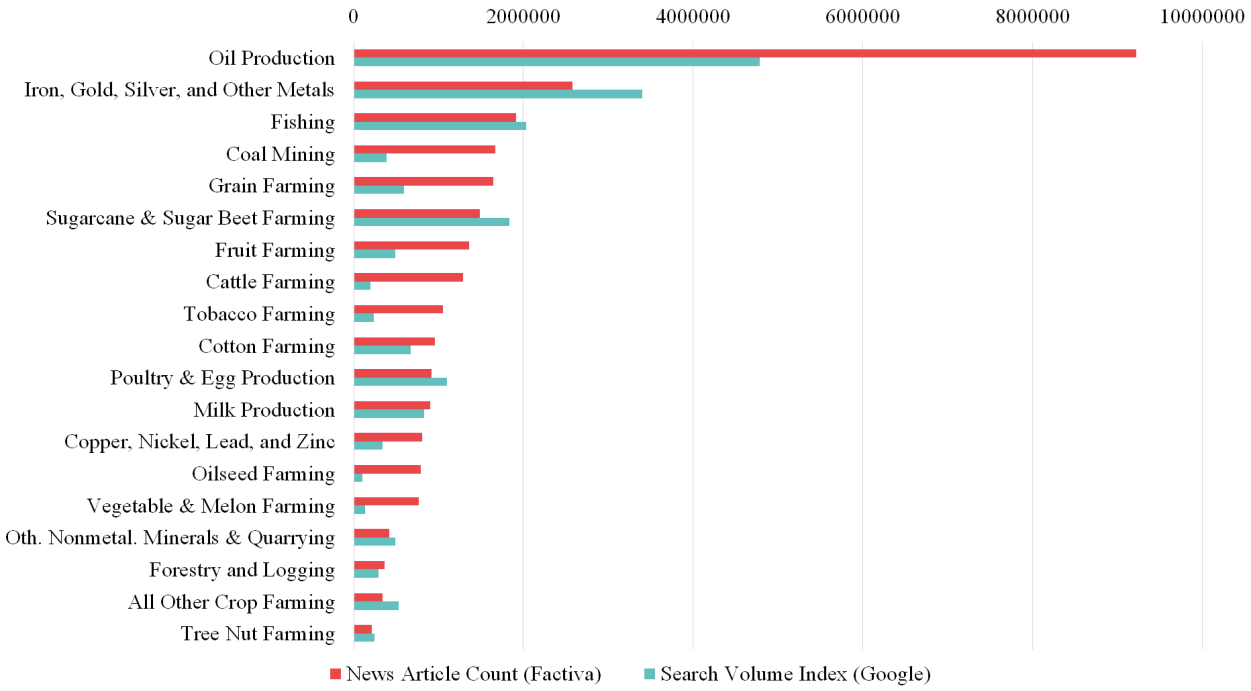
6.5.1 The Role of Forward-lookingness about Shock Persistence

If firms behave myopically concerning network dynamics when passing through shocks to non-oil commodity prices, then there may also not be detectable differences in pass-through of the same spot price change when there have been different movements in commodity futures curves. More formally, the futures terms in Figure 9 should be indistinguishable from 0 when performed for non-oil commodities if firms act as myopically about differences in shock persistence as they do about network dynamics.

Because commodity futures outside of oil and corn do not merge cleanly with the commodity categories in our input-output data, we test the myopic hypothesis for pass-through of non-oil commodity price movements using corn. The results are shown in Figure 11 using the GMM-optimal model for corn but assuming rational expectations. While there is one period of statistical significance of the futures pricing terms, most futures terms are statistically indistinguishable from zero, and evidence of forward-lookingness is substantially weaker than we saw for oil in Figure 9. For corn, movements in the spot price lead to statistically similar pass-through regardless of the movement in the futures curve.

We replicate this finding using our reduced form approach pooling all non-oil commodities (and their future prices) in Appendix E.3. In this case, evidence of myopia for pass-through of non-oil commodity price movements, compared with rational expectations for oil price movements, is even starker.

Figure 12 – News Coverage and Search Volume by Commodity



Note: Data on number of news articles is from Factiva for the period 1997-2023. Commodity categories are constructed to mirror BEA commodity categories from the input-output tables. For commodity categories made up of multiple constituent commodities, a weighted average is taken in order to reflect the average level of attention paid to that commodity category. For example, for "Poultry and Egg Production," the average number of articles making reference to each type of poultry and to eggs is computed (weighted by employment, per the QCEW, which has granular data for each constituent commodity). Data on Search Volume Index (SVI) is from Google Trends for the period 2008-2023. SVI is a relative measure of search frequency; here, "Oilseed Farming" is indexed to 1. Commodity categories with multiple constituent commodities are again averaged as above.

6.5.2 Commodity Salience

As discussed above, we find evidence of rational expectations in the context of oil and myopia in the context of non-oil commodities. But why is behavior different in the context of oil? We suggest that this relates to the uniquely pervasive salience of oil relative to other commodities. Virtually every American receives a daily signal about the changing price of oil in the form of fluctuating prices at the gas station around the corner. And swings in oil prices are invariably accompanied by a bevy of news reports discussing the swings and their causes.

We provide some suggestive evidence to this effect in Figure 12, where we show that oil is an outlier relative to the other commodities in terms of both the number of news articles mentioning oil and the volume of Google searches for oil compared to other commodities. While this is not a causal analysis, it underlines one key way - which feeds into attentiveness - in which oil stands out from all other commodities.

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 may 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.

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 $passthrough_{i,h}(\boldsymbol{\alpha})$ 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 $\boldsymbol{\alpha}$. The optimal calibration $\hat{\boldsymbol{\alpha}}$ 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

$$\hat{\Pi}_t = Intercept + \sum_i PCEShare_i \sum_{h=0}^{H} passthrough_{i,h}(\hat{\boldsymbol{\alpha}}) \Delta P_{Z,t-h}, \tag{13}$$

an intercept plus the personal consumption expenditures 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 $\boldsymbol{\alpha}$ 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. 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 (10) without a time fixed effect. We provide the full details of this procedure applied to oil in Appendix F.2, finding no statistically-significant evidence of an intercept. This is consistent with all relative price movements generated from oil passing through to aggregate inflation. We note that our fully closed model, which uses standard parameters and is described in Appendix A.4, does not lead to a meaningfully different result.[21]

Given an intercept, the PCE share of each industry, model-derived pass-through coefficients with the optimal calibration $\hat{\boldsymbol{\alpha}}$, and the sequence of commodity price changes, the user can apply formula (13) 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

[21] It is easy to write a model where relative price movements generated from oil do not pass through to aggregate inflation. Consider a monetary authority assumed to change the money supply so that the aggregate price level always remains constant. In this model, changes in the oil price would not affect inflation. This is why we prefer the empirical approach but are reassured that closing our supply-side model in a standard way leads to a similar result.

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.

7.2 Illustration: Oil

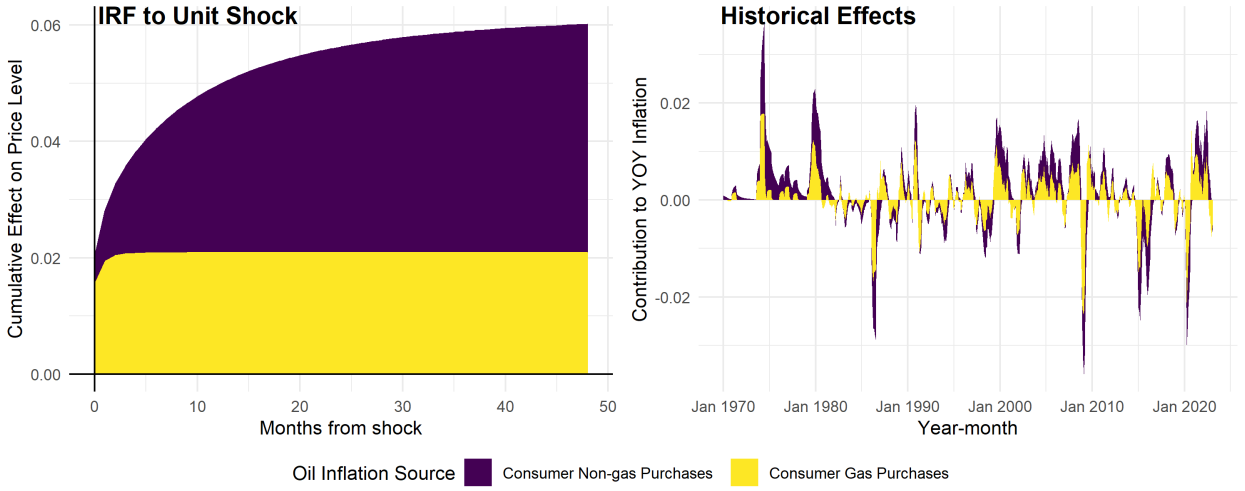
Now, we apply our guide from the preceding section to compute network oil inflation.[22] We decompose this inflation into direct and indirect sources. Consumer use of gasoline appears in the data as personal consumption expenditures from the petroleum refining sector. The effect of oil prices on gas 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 0.06 percentage points of aggregate inflation resulting from 1% shock to oil prices, just over 0.02 percentage points comes from consumer gas purchases, while about 0.04 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 (13) but decomposing direct and indirect effects. The magnitude of the combined direct and indirect effects is not uncommonly above 2%, and it 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.

In Appendix F.6, we subtract these historical effects from aggregate PCE inflation. Underlying inflation is lower than the headline figure for all of 2021 and much of 2022; in the latter half of 2022, underlying inflation continues to increase despite headline inflation declining.

[22] While our GMM-optimal procedure optimized fit for the equally weighted average industry, we confirm in Appendix F.5 that the fit remains good for the PCE-weighted average industry, both including and excluding petroleum refineries.

Figure 13 – Inflationary Effects of Oil



Note: Panel 1 is the model-predicted response of the aggregate price level (in log points) to a log point shock to oil prices. The total price level response due to a log point oil price increase is 0.06 log points, with most of the effect due to gas prices increasing being realized on impact. The indirect effect due to all other price increases is realized slowly and only dominates the direct effect after about 6 months. Eventually, the indirect effect is about twice as large. Panel 2 shows how both components of oil inflation have affected aggregate, year-over-year 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.

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 measurements from the producer price indices. This distinction is important because wholesalers and retailers must past 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 3. 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 approximately the same if we include direct oil inflation, the component of oil inflation due to consumer gas purchases; together, direct and indirect oil inflation explain 33% of the variation in official PCE inflation. We noted earlier that official core PCE inflation tries to purge the effects of oil inflation from official PCE inflation. Panel 2 of Table 3 reveals that it is only partially successful in doing this; indirect oil inflation is still rather strongly predictive. We retain an R^2 of 16% when explaining official core PCE inflation with network oil inflation.

Table 3 – Predicting Inflation with Network Oil Inflation

Panel 1: Total Inflation

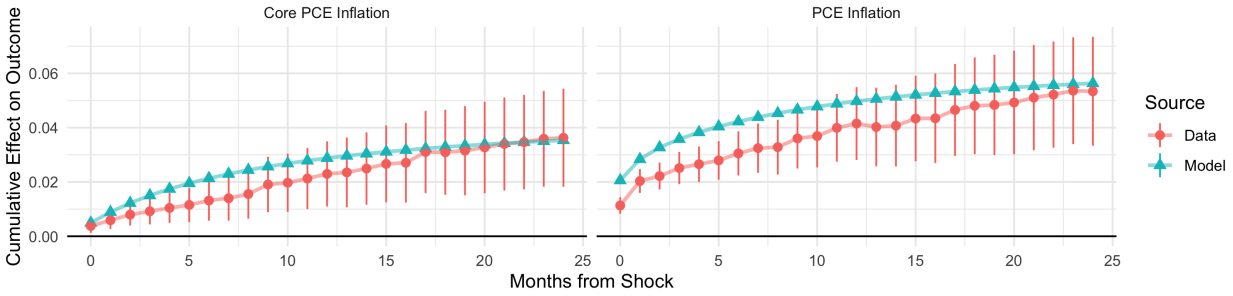
	(1)	(2)	(3)	(4)
Dependent Variable:	All	Kanzig	All	Kanzig
Total PCE Inflation	Variation	Variation	Variation	Variation
Direct Oil Inflation			0.106	0.319**
			(0.079)	(0.138)
Indirect Oil Inflation	1.775***	1.355***	1.654***	0.965***
	(0.112)	(0.200)	(0.138)	(0.284)
R-Squared	0.3326	0.3140	0.3347	0.3054
Observations	768	768	768	768

Panel 2: Core Inflation

	(1)	(2)	(3)	(4)
Dependent Variable:	All	Kanzig	All	Kanzig
Core PCE Inflation	Variation	Variation	Variation	Variation
Direct Oil Inflation			-0.219***	-0.084
			(0.620)	(0.111)
Indirect Oil Inflation	0.921***	0.424**	1.172***	0.527**
	(0.092)	(0.181)	(0.126)	(0.253)
R-Squared	0.1437	0.1019	0.1581	0.1119
Observations	768	768	768	768

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. We find that indirect oil inflation predicts both PCE inflation and Core PCE inflation, even conditional on direct oil inflation. Moreover, network oil inflation explains 16% of the variation in Core PCE inflation, even though Core PCE inflation is designed to factor out oil price movements.

Figure 14 – Explaining the Predictability of Official Inflation Series



Note: The red-circle series show the cumulative effect of oil price movements, instrumented using the Kanzig (2021) variation, on Core PCE inflation (Panel 1) and PCE inflation (Panel 2). Core PCE inflation should not be predictable by oil price movements. Our model rationalizes its predictability: inflation resulting from indirect use of oil is shown in the blue-triangle series in Panel 1. The model-based series in Panel 2 additionally includes inflation due to gasoline 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).

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 0.035 log points (3.5% of the oil shock magnitude) of Core PCE inflation are realized over the 2 years following the price increase. About 0.055 log 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.

8 Conclusion

We study how much and how quickly supply chains transmit commodity price movements throughout the production network, finding statistically-significant evidence of full but delayed pass-through. Price rigidity interacting with the network is a major source of this delay – industries more downstream from the shock pass it through more gradually, even if they are forward-looking.

Our model suggests that the delay due to downstreamness is intensified when firms are inattentive. 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 rational expectations, whereas responses to non-oil commodity price movements are consistent with more myopic behavior.

We then apply our model in the specific case of oil, showing that relative price movements throughout the network generated from oil largely pass through to aggregate inflation. We show that network oil inflation has statistically-significant and non-trivial (R^2 = 16%) predictive power for official Core PCE inflation, as well as even greater predictive power (R^2 = 33%) for total PCE inflation.

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. Finally, we 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, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, "Microeconomic Origins of Macroeconomic Tail Risks," *American Economic Review*, 2017, *107* (1), 54–108.

_ , Ufuk Akcigit, and William Kerr, "Networks and the Macroeconomy: An Empirical Exploration," *NBER Macroeconomics Annual*, 2016, *30* (1), 275–335.

_ , Vasco Carvalho, and Alireza 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 31218, National Bureau of Economic Research 2023.

Alquist, Ron, Lutz Kilian, and Robert 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," *Handbook of International Economics*, 2022, *5*, 297–376.

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

Auclert, Adrien, Rodolfo Rigato, Matthew Rognlie, and Ludwig Straub, "New Pricing Models, Same Old Phillips Curves?," Working Paper 30264, National Bureau of Economic Research 2022.

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

Baqaee, David and Elisa Rubbo, "Micro Propagation and Macro Aggregation," Working Paper 30538, National Bureau of Economic Research 2022.

Baqaee, 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, Jean-Noel and Julien Sauvagnat, "Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks," *Quarterly Journal of Economics*, 2016, *131* (3), 1543–1592.

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

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

Bils, Mark and Peter J. Klenow, "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 2004, *112* (5), 947–985.

Blanchard, Olivier, "Price Asychronization and Price-level Inertia," in Rudiger Dornbush and Maria Henrique Simonson, eds., *Inflation, Debt, and Indexation*, MIT Press, 1983.

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

Blinder, Alan, Elie R.D. Canetti, David E. Lebow, and Jeremy B. Rudd, *Asking about Prices: A New Approach to Understanding Price Stickiness*, Russell Sage Foundation, 1998.

BLS, *Handbook of Methods*, Division of BLS Publishing, Office of Publications and Special Studies, 2015.

Boehm, Christoph E., Aaron Flaaen, and Nitya Pandalai-Nayar, "Input Linkages and the Transmission of Shocks: Firm Level Evidence from the 2011 Tohoku Earthquake," *Review of Economics and Statistics*, 2019, *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, Lorenzo, Fernando Parro, and Aleh Tsyvinski, "Distortions and the Structure of the World Economy," Technical Report 4 2022.

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

Carlstrom, Charles T., Timothy S. Fuerst, and Matthias 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, Vasco M. and Xavier Gabaix, "The Great Diversification and its Undoing," *American Economic Review*, 2013, *103* (5), 1697–1727.

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

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

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

Gabaix, Xavier, "The Granular Origins of Aggregate Fluctuations," *Econometrica*, 2011, *79* (3), 733–772.

_ , "Chapter 4 - Behavioral inattention," in B. Douglas Bernheim, Stefano DellaVigna, 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.

Galí, Jordi, *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and its Applications*, Princeton University Press, 2015.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 2020, *110* (8), 2586–2624.

Golosov, Mikhail and Robert E. Lucas, "Menu Costs and Phillips Curves," *American Economic Review*, 2007, *115* (2), 171–199.

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

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

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

Huang, Kevin X.D. and Zheng Liu, "Production chains and general equilibrium aggregate dynamics," *Journal of Monetary Economics*, 2001, *48* (2), 437–462.

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

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

Klenow, Peter J. and Oleksiy Kryvtsov, "State-Dependent or Time-Dependent Pricing: Does it Matter for Recent U.S. Inflation?," *Quarterly Journal of Economics*, 2008, *123* (3), 863–904.

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

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

Luo, Shaowen and Daniel Villar, "Propagation of Shocks in an Input-output Economy: Evidence from Disaggregated Prices," *Journal of Monetary Economics*, 2023, *137*, 26–46.

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.

Midrigan, Virgiliu, "Menu Costs, Multiproduct Firms, and Aggregate Fluctuations," *Econometrica*, 2011, *79* (4), 1139–1180.

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, Ernesto, Raphael Schoenle, and Michael Weber, "Price Rigidity and the Origins of Aggregate Fluctuations," Working Paper 23750, National Bureau of Economic Research 2017.

Rubbo, Elisa, "Networks, Phillips Curves and Monetary Policy," Technical Report 4 2023.

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

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

Taylor, John B., "Staggered Wage Setting in a Macro Model," *American Economic Review*, 1979, *69* (2), 108–113.

Woodford, Michael, "Imperfect Common Knowledge and the Effects of Monetary Policy," in Philippe Aghion, Roman Frydman, Joseph Stiglitz, and Michael Woodford, eds., *Knowledge, Information, and Expectations in Modern Macroeconomics: in Honour of Edmund S. Phelps*, Princeton University Press, 2003, pp. 25–58.

Yun, Tack, "Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles," *Journal of Monetary Economics*, 1996, *37* (2), 345–370.

[For Online Publication]

A Appendix: Model

Appendix A.1 proves the propositions in the model section of the paper. A.2 describes the numerical solution of the model in discrete time, which is used for our empirical work. We use this discrete time solution at a daily frequency in our empirics for two reasons: (1) to improve comparability of the model with our continuous time results and (2) so that our code can easily allow forecasters to exploit that commodity prices are usually measured at least at a daily frequency – enabling much better eve-of-release forecasts of PPI, for example, than using the previous month's PPI data. A.3 recasts the model in continuous time and is used for our main results in the paper. A.4 fully closes the model and describes assumptions about household consumption, sticky wages, and the Taylor rule.

A.1 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 $\hat{\boldsymbol{p}}^*_\infty = \hat{\boldsymbol{p}}_\infty$. The reset price equation simplifies to $\hat{\boldsymbol{p}}_\infty = \widehat{\boldsymbol{mc}}_\infty$, and the marginal cost equation, using that there are no TFP shocks, simplifies to $\widehat{\boldsymbol{mc}}_\infty = \text{diag}(s_i^L)\hat{\boldsymbol{w}}_\infty + \boldsymbol{s}^Z\hat{p}_Z + \boldsymbol{\Phi}\hat{\boldsymbol{p}}_\infty$. Eliminating marginal cost, we have

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

which simplifies to

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

This was our desired result. □

Proof of Proposition 2:

Proof. First, see subsection A.3, 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:

$$\hat{p}^*_{i,t} = \psi_i \text{E}_t \int_{\tau=t}^{\infty} e^{-\psi_i(\tau-t)}\widehat{mc}_{i,\tau}d\tau$$
$$\dot{\hat{p}}_{i,t} = \phi_i(\hat{p}^*_{i,t} - \hat{p}_{i,t}).$$

The log-linearization of marginal cost is the same as in the discrete time model. We begin with the myopic case, $\psi_i \to \infty$ for all i, which is more straightforward than when firms are forward

looking. In this case, $\hat{p}^*_{i,t} = \widehat{mc}_{i,t}$. Holding wages and TFP constant, then $\hat{\boldsymbol{p}}^*_t = \boldsymbol{s}^Z \hat{p}_{Z,0} + \boldsymbol{\Phi}\hat{\boldsymbol{p}}_t$. So

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

This differential equation solves to

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

where $\boldsymbol{B} = \bar{\phi}(\boldsymbol{I} - \boldsymbol{\Phi})$. 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

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

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

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

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

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

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

$$\begin{aligned}\bar{\psi}\left[(\boldsymbol{I} - \boldsymbol{\Phi})\hat{\boldsymbol{p}}_t - z\right] &= \bar{\psi}\left[(\boldsymbol{I} - \boldsymbol{\Phi})(\boldsymbol{I} - \boldsymbol{\Phi})^{-1}z - (\boldsymbol{I} - \boldsymbol{\Phi})e^{-\boldsymbol{B}t}(\boldsymbol{I} - \boldsymbol{\Phi})^{-1}z - z\right] \\ &= -\bar{\psi}(\boldsymbol{I} - \boldsymbol{\Phi})e^{-\boldsymbol{B}t}(\boldsymbol{I} - \boldsymbol{\Phi})^{-1}z.\end{aligned}$$

The LHS is

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

Therefore, we have

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

We take derivatives on the LHS:

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

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

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

So the conjectured solution works if

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

Under rational expectations and a steady state interest rate of 0, $\bar{\phi} = \bar{\psi}$, so $\boldsymbol{B} = \left(\bar{\phi}^2(\boldsymbol{I} - \boldsymbol{\Phi})\right)^{1/2}$. When the frequency of price adjustment does not vary by sectors, $\bar{\phi} = \phi\boldsymbol{I}$ and $\bar{\psi} = \psi\boldsymbol{I}$ commute with all matrices, and so the matrix equation can be rewritten as

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

or

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

where we take the roots with positive eigenvalues. □

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

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

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

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

Now, just as in scalar case, we have

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

which was the desired result.

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

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

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

A.2 Eigenvalue Solution of the Model (with Decomposition)

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

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

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 $m_f > 0$,

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

Stacking across industries,

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

Recall that

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

Therefore,

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

Now, define

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

Then

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

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

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

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

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

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

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

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

which can be forward solved to get

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

The required transverality condition is

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

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

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

is finite. Then

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

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

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

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

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

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

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

Then

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

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

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

The long-run pass-through is

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

which, as we have already shown, is

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

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

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

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

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

so that

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

Putting everything together, we have

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

with

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

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

A.3 Model in Continuous Time

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

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

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

subject to the demand conditions

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

The first-order condition log-linearizes to

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

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

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

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

Now, for a standard household problem, we have $SDF^{SS}_{t,\tau} = e^{-\rho(\tau-t)}$, where ρ is the discount rate. Therefore, the above equation simplifies to

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

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

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

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.4 Fully Closed Model

A.4.1 Households and Government

The setup for the representative household is completely standard in the NK theory of sticky wages, following, e.g., Woodford (2003) or Galí (2015). Each firm $j \in [0,1]$ uses a CES aggregate of individual labor types $h \in [0,1]$ to produce:

$$L_t(j) = \left(\int_0^1 L_t(j,h)^{\frac{\sigma_L-1}{\sigma_L}}dh\right)^{\frac{\sigma_L}{\sigma_L-1}}.$$

Firm j's demand for individual h's labor is thus

$$L_t(j,h) = L_t(j)\left(\frac{W_t(h)}{W_t}\right)^{-\sigma_L},$$

and total demand for individual h's labor is (with an abuse of notation)

$$L_t(h) = \int_0^1 L_t(j,h)dj = L_t\left(\frac{W_t(h)}{W_t}\right)^{-\sigma_L},$$

with $L_t = \int_0^1 L_t(j)dj$ and the usual wage index. Denote the discount factor by β with $0 < \beta < 1$ and the probability an individual cannot adjust their wage in any period by θ_W with $0 \leq \theta_W < 1$. Consumption utility is given by

$$U(C_t) = \frac{C_t^{1-\gamma} - 1}{1-\gamma},$$

so that the marginal utility of nominal income is $C_t^{-\gamma}/P_t$, where P_t is the price of consumption. Labor disutility is given by

$$V(L_t) = \frac{L_t^{1+1/\epsilon}}{1+1/\epsilon}.$$

A wage $W_t(h)$ that is adjusted in period t solves

$$\max \tilde{\mathrm{E}}_t \sum_{k=0}^{\infty} (\beta\theta_W)^k \left[\frac{C_{t+k}^{-\gamma}}{P_{t+k}} W_t(h) L_{t+k}(h) - \frac{L_{t+k}(h)^{1+1/\epsilon}}{1+1/\epsilon}\right],$$

which maximizes the expected (potentially myopic) discounted value of nominal income, weighed by the marginal utility of nominal income, net of the disutility of labor. The first-order condition is

$$\tilde{\mathrm{E}}_t \sum_{k=0}^{\infty} (\beta\theta_W)^k L_{t+k}(h) \left[W_t(h) - \frac{\sigma_L}{\sigma_L - 1} \frac{L_{t+k}(h)^{1/\epsilon}}{C_{t+k}^{-\gamma}/P_{t+k}}\right] = 0,$$

which sets the expected discounted value of the gap between the wage and a markup over the marginal rate of substitution (marginal labor disutility over the marginal utility of nominal income) equal to 0.

This equation loglinearizes to

$$w_t(h) = \frac{1-\beta m\theta_W}{1+\sigma_L/\epsilon} \mathrm{E}_t \sum_{k=0}^{\infty} (\beta m\theta_W)^k \left(\gamma c_{t+k} + p_{t+k} + \frac{1}{\epsilon} l_{t+k} + \frac{\sigma_L}{\epsilon} w_{t+k}\right), \tag{14}$$

where m is the degree of consumer myopia. Clearly, anyone who resets their wage does so in the same way, i.e. $w_t(h) = w_t^*$ for all h. The loglinearized wage index is

$$w_t = \theta_W w_{t-1} + (1-\theta_W) w_t^*.$$

These equations simplify to

$$\mathrm{E}_t w_{t+1} = \left(\frac{1}{\beta m \theta_W} + \theta_W - \frac{K}{\beta m \theta_W} \frac{\sigma_L}{\epsilon} \right) w_t - \frac{1}{\beta m} w_{t-1} - \frac{K}{\beta m \theta_W} (\gamma c_t + p_t + \frac{1}{\epsilon} l_t),$$

with

$$K = \frac{(1-\theta_W)(1-\beta m \theta_W)}{1+\sigma_L/\epsilon}.$$

The Euler equation is as usual, augmented for myopia:

$$m\mathrm{E}_t[c_{t+1}] = c_t + \frac{1}{\gamma}(i_t - m\mathrm{E}_t[p_{t+1}] + p_t).$$

Consumption is a CES bundle of industry-specific consumption bundles (with elasticity σ_C), so that

$$c_{it} = \sigma_C(\boldsymbol{s}'\boldsymbol{p}_t - p_{it}) + c_t,$$

or, stacking across industries,

$$\boldsymbol{c}_t = \sigma_C(\mathbf{1}\boldsymbol{s}' - \boldsymbol{I})\boldsymbol{p}_t + \mathbf{1}c_t. \tag{15}$$

Finally, there is a standard Taylor rule with price and wage inflation coefficients ψ_p and ψ_w.

A.4.2 Firms

Firms in each industry produce using a bundle of intermediate goods and labor (which could be construed as a value added composite). The intermediate bundle is a CES aggregator of intermediate inputs (with elasticity σ_X):

$$X_{it}(j) = \tilde{F}_i(X_{it}(j)^1, ..., X_{it}(j)^I).$$

To a first order, firm-level demands aggregate, and we can consider the loglinearized demand of industry i for intermediate input j:

$$x_{it}^j = \sum_{i'=1}^{I} \epsilon_{i'}^{X_i^j} p_{i't} + x_{it},$$

where x_i is industry production of the intermediate bundle and the ϵ's are log partial derivatives of the superscript with respect to the subscript argument. Using properties of the factor demand elasticities and defining s_i^X to be the share of intermediate inputs in cost, we have

$$x_{it}^j = \sigma_X \left(\frac{[\Phi]_{i,\cdot}}{s_i^X} \boldsymbol{p}_t - p_{jt} \right) + x_{it}. \tag{16}$$

Firms then combine the intermediate bundle and labor according to

$$Y_{it}(j) = A_{it}F_i(, L_{it}(j)).$$

By a similar argument, again assuming a CES aggregator with elasticity σ_V, we have

$$x_{it} = \sigma_V s_i^L \left(w_t - \frac{[\Phi]_{i,\cdot}}{s_i^X} \boldsymbol{p}_t \right) + y_{it} - a_{it}. \tag{17}$$

Combining equations 16 and 17, we have

$$x_{it}^j = \sigma_X \left(\frac{[\Phi]_{i,\cdot}}{s_i^X} \boldsymbol{p}_t - p_{jt} \right) + \sigma_V s_i^L \left(w_t - \frac{[\Phi]_{i,\cdot}}{s_i^X} \boldsymbol{p}_t \right) + y_{it} - a_{it}. \tag{18}$$

Now, market clearing is

$$Y_{it} = \sum_j X_{jt}^i + C_{it},$$

which log-linearizes to

$$y_{it} = \frac{\left(\boldsymbol{X}^i\right)'}{Y_i} \boldsymbol{x}_t^i + \frac{C_i}{Y_i} c_{it}.$$

Now, we can stack equation 18 across using industries to obtain

$$\boldsymbol{x}_t^i = (\sigma_X \boldsymbol{I} - \sigma_V \text{diag}(s_i^L)) diag(s_i^X)^{-1} \boldsymbol{\Phi} \boldsymbol{p}_t - \sigma_X \mathbf{1} p_{it} + \sigma_V \boldsymbol{s}^L w_t + \boldsymbol{y}_t - \boldsymbol{a}_t.$$

Combining this with the market clearing equation, we have

$$y_{it} = \frac{\left(\boldsymbol{X}^i\right)'}{Y_i} (\sigma_X \boldsymbol{I} - \sigma_V \text{diag}(s_i^L)) diag(s_i^X)^{-1} \boldsymbol{\Phi} \boldsymbol{p}_t - \sigma_X \left(1 - \frac{C_i}{Y_i}\right) p_{it} + \sigma_V \frac{\left(\boldsymbol{X}^i\right)'}{Y_i} \boldsymbol{s}^L w_t + \frac{\left(\boldsymbol{X}^i\right)'}{Y_i} (\boldsymbol{y}_t - \boldsymbol{a}_t) + \frac{C_i}{Y_i} c_{it}.$$

This form is convenient because it can also be readily stacked across industries. Denote the demand-side IO matrix by

$$\boldsymbol{\Omega} = \begin{pmatrix} \frac{X_1^1}{Y_1} & \frac{X_2^1}{Y_1} & \cdots & \frac{X_I^1}{Y_1} \\ \frac{X_1^2}{Y_2} & \frac{X_2^2}{Y_2} & \ddots & \frac{X_I^2}{Y_2} \\ \vdots & \ddots & \ddots & \vdots \\ \frac{X_1^I}{Y_I} & \frac{X_2^I}{Y_I} & \cdots & \frac{X_I^I}{Y_I} \end{pmatrix}.$$

Then

$$\begin{aligned} \boldsymbol{y}_t = {} & (\boldsymbol{I} - \boldsymbol{\Omega})^{-1} \boldsymbol{\Omega} (\sigma_X \boldsymbol{I} - \sigma_V \text{diag}(s_i^L)) diag(s_i^X)^{-1} \boldsymbol{\Phi} \boldsymbol{p}_t - \sigma_X (\boldsymbol{I} - \boldsymbol{\Omega})^{-1} \text{diag} \left(1 - \frac{C_i}{Y_i}\right) \boldsymbol{p}_t \\ & + \sigma_V (\boldsymbol{I} - \boldsymbol{\Omega})^{-1} \boldsymbol{\Omega} \boldsymbol{s}^L w_t - (\boldsymbol{I} - \boldsymbol{\Omega})^{-1} \boldsymbol{\Omega} \boldsymbol{a}_t + (\boldsymbol{I} - \boldsymbol{\Omega})^{-1} \text{diag} \left(\frac{C_i}{Y_i}\right) \boldsymbol{c}_t. \end{aligned}$$

Now recall that the vector of consumption was already pinned down in terms of prices and aggregate consumption in equation 15.

Armed with this solution for output, we can finally eliminate labor supply from the NK wage equation 14. We start with the log-linearized factor demand:

$$\boldsymbol{l}_t = \sigma_V(\boldsymbol{\Phi}\boldsymbol{p}_t - \boldsymbol{s}^X w_t) + \boldsymbol{y}_t - \boldsymbol{a}_t.$$

Then $l_t = \frac{\boldsymbol{L}'}{L}\boldsymbol{l}_t$.

The pricing equation, which we have derived elsewhere, is

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

A.4.3 Calibration

We use a standard calibration for all parameters. We set $\sigma_X = 0.1$, $\sigma_V = 1$, $\sigma_C = 1$, $\sigma_L = 3$, $\epsilon = 1/4$, $\gamma = 2$, $\theta_W = 0.42$ (annualized rate), $\psi_p = 1.2$, and $\psi_w = 0$. Additional parameters are calibrated as described in the main body of the paper. In our structural work, we always constrain consumer and firm myopia to have the same value, i.e. $m = m_f$.

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

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

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

both the Make and Use tables.

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

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

C Appendix: Additional Robustness of Reduced-form Empirics

In this appendix, we show that these results are robust to a variety of modifications and alternative approaches. In C.1, we repeat the analysis in Table 1 (i) without the time fixed effects, (ii) without the wage control variable, (iii) with an added TFP control variable, (iv) with an added control for network gas/electricity cost changes (two commodities likely to be close substiutes for oil), and (v) with cost shares that exclude payments to capital from the denominator. All of these exercises yield very similar results. In Appendix C.2, we investigate heterogeneity on some

other factors that could plausibly affect pass-through: concentration, firm size, inventories, and the capital share (all measured at the sector level). Adding these additional heterogeneity terms to the regressions in the preceding table does not ameliorate the statistical significance of the results. In C.3, we modify our preceding regression specification by adding leads and lags to study the month-by-month dynamics of pass-through in a reduced-form manner. In C.4, we repeat this analysis using a local projections approach, finding similar results. In C.5, we show that the same patterns are again evident in a binscatter analysis. The binscatters also allow us to investigate some additional dimensions of potential heterogeneity: size and sign of shock. We find no evidence of such heterogeneity. Finally, in C.6, we take a different approach to isolating exogenous variation. We focus on a few case studies likely to be highly exogenous - the 1979 oil shock, the 2014 oil shale boom, and the 2020 COVID shock - and show that these settings yield the same results as our pooled analysis.

C.1 Main Table Robustness

In this section, we modify the regression specification used in our main analysis slightly. This regression specification is given by Equation (8).

First, we drop the time fixed effects. The time fixed effect is not present in Proposition 1, so this specification is arguably more directly linked to Proposition 1. However, as Table G.1 reveals, this modification leads to virtually no change relative to Table 1.

Second, we drop the wage control variable. 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, as shown in G.2.

Third, in a related exercise, instead of the change in industry prices, we place the change in industry wages on the right-hand-side of our regression. This allows for directly assessing the effects of network oil cost changes on industry wages. As shown in G.3, we find no evidence of a relation between these variables. This is consistent with the fact that including or excluding the wage control variable was found in the previous paragraph to have no meaningful effect on the results.

Fourth, we add a TFP control variable that accounts for changing aggregate TFP. Such changes are another factor that can feed into price changes, and if one is worried that they are correlated with oil cost shocks, this could potentially create bias in our estimates. In a network model, industries have heterogeneous loadings on aggregate TFP, and we can control for these effects by interacting industries' heterogeneous loadings with a time fixed effect. Nevertheless, we find that adding the TFP control variable scarcely changes the results, as seen in G.4.

Fifth, we add a control variable for gas and electricity. Industries with high exposure to oil may have the ability to substitute away to gas and/or electricity when the cost of oil rises

relative to the cost 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. The control variable for gas/electric network cost changes is constructed analogously to our main oil network cost change variable. The findings are displayed in G.5, where the results are little changed relative to the baseline table.

Sixth and finally, we note that 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 Table G.6. The total amount of pass-through over both the one-month horizon and the one-year horizon decreases, and the qualitative take-away - that pass-through is incomplete on impact and more thorough but still incomplete after a year - is strongly sustained.

C.2 Heterogeneity Table Robustness

In this section, we add a variety of other interaction terms to our heterogeneity table. Essentially, in order to ensure downstreamness is not simply correlated with some other variable that could plausibly affect the extent or pace of pass-through, we add such variables to the regression. We obtain information on inventories by sector from the Census Bureau's Manufacturing and Trade Inventories and Sales data. We obtain information on market concentration by sector (marketshare of the top 8 firms in the sector) from the Economic Census. We obtain information on average firm size (measured by sales) by sector from Dun and Bradstreet. And we obtain information on the capital share by sector from the BEA. Table G.7 reveals that the addition of these interaction terms does not ameliorate the significance of the downstreamness interaction term. While some of these other terms are statistically significant as well, downstreamness remains a crucial determinant of pass-through. In fact, the magnitude of the downstreamness coefficient actually increases in case of the Kanzig variation and the case of all non-oil commodity price variation. In the case of all oil variation, it is scarcely changed.

C.3 Reduced-Form Empirics with Dynamics

C.3.1 Dynamic Reduced-Form Regression Specifications

We slightly modify our main reduced-form regression specification by adding leads and lags to flexibly study the month-by-month dynamics of pass-through. The specification is again inspired by Proposition 1, albeit with an arbitrary lag structure. 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,

$$\Delta P_{i,t} = \lambda_t + \sum_{h=-6}^{24} \beta_h [(\boldsymbol{I} - \boldsymbol{\Phi})^{-1} \boldsymbol{s}^Z]_i \Delta P_{Z,t-h} + \sum_{h=-6}^{24} \gamma_h [(\boldsymbol{I} - \boldsymbol{\Phi})^{-1} \text{diag}(s_i^L) \Delta \boldsymbol{w}_{t-h}]_i + \epsilon_{i,t}, \quad (19)$$

where, as before, $P_{i,t}$ denotes the log price of industry i at time t, λ_t is a time fixed-effect, $[(\boldsymbol{I} - \boldsymbol{\Phi})^{-1} \boldsymbol{s}^Z]_i$ 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, $[(\boldsymbol{I} - \boldsymbol{\Phi})^{-1} \text{diag}(s_i^L) \Delta \boldsymbol{w}_t]_i$ represents input cost changes due to wage movements in various sectors whose output industry i utilizes, and $\epsilon_{i,t}$ is an error term. In this context, $\sum_h \beta_h = 1$ 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.

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:

$$\begin{aligned} \Delta P_{i,t} = & \lambda_t + \sum_{h=-6}^{24} \eta_h [\boldsymbol{s}^Z]_i \Delta P_{Z,t-h} + \sum_{h=-6}^{24} \rho_h [(\boldsymbol{I} - \boldsymbol{\Phi})^{-1} \boldsymbol{s}^Z - \boldsymbol{s}^Z]_i \Delta P_{Z,t-h} \\ & + \sum_{h=-6}^{24} \gamma_h [(\boldsymbol{I} - \boldsymbol{\Phi})^{-1} \text{diag}(s_i^L) \Delta \boldsymbol{w}_{t-h}]_i + \epsilon_{i,t}. \end{aligned} \quad (20)$$

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.

C.3.2 Results

We begin by using all variation in oil prices from 1997 onward, running the regression specification given by Equation (19). 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.[23]

The red coefficients in Panel 1 of Figure G.1 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%.

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

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 G.1 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 G.1 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 G.2 plots the results of the regression specification given by Equation (20), decomposing total network exposure to oil price variation into direct and indirect exposure. Panel 1 again focuses on all oil price variation. Here, the red coefficients correspond to pass-through of direct exposure to oil shocks, whereas the blue coefficients correspond to pass-through of indirect exposure through the network to oil shocks (i.e., total network exposure minus direct exposure). The results reveal no evidence of substantial pre-trends. That is, in the months prior to an oil price movement, coefficients are not substantially positive or negative. In month 0, on impact of the shock, a high degree of direct pass-through occurs (approximately 75%). One month after impact of the shock, additional direct pass-through occurs (approximately 25%). At this point, after just a couple months, full pass-through has already occurred. This contrasts with the pattern of indirect pass-through, of which very little occurs on impact. Instead, pass-through phases in slowly over the course of eight months or so.[24]

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 G.2 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 – instru-

[24]Dynamic 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.

menting 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.

C.3.3 Heterogeneity

We similarly modify our heterogeneity specification by adding leads and lags to study the time path of downstreamness' effects on pass-through. Specifically,

$$\begin{aligned}\Delta P_{i,t} =& \lambda_t + \sum_{h=-6}^{24} \beta_h \left[(\boldsymbol{I} - \boldsymbol{\Phi})^{-1} \boldsymbol{s}^Z \right]_i \Delta P_{Z,t-h} \\ &+ \sum_{h=-6}^{24} \tilde{\beta}_h \left[(\boldsymbol{I} - \boldsymbol{\Phi})^{-1} \boldsymbol{s}^Z \right]_i \Delta P_{Z,t-h} \times heterogeneity_i \\ &+ \sum_{h=-6}^{24} \gamma_h [(\boldsymbol{I} - \boldsymbol{\Phi})^{-1} \text{diag}(s_i^L) \Delta \boldsymbol{w}_{t-h}]_i + \epsilon_{i,t}\end{aligned}$$

Figure G.3 plots the impulse response function corresponding to the specification in column (4) of table 2. 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.

C.3.4 Robustness

As we did for the main regression table, it is possible to replicate these dynamic results after excluding the wage control, controlling for gas/electricity exposure, and using cost shares that exclude capital from the denominator. This is done in Figures G.4, G.5, and G.6, respectively.

C.3.5 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 permutate treatment across industries 1000 times. We run our main specifications on these placebo variations and compare the magnitude of the coefficients resulting from these regressions with the magnitude of the actual coefficients, yielding information on the likelihood with which the actual coefficients resulted from pure chance. In particular, the permutation tests are performed on the cumulative one-year pass-through coefficient. Panels 1a and 1b of Figure G.7 correspond to the re-

gression specification given by Equation (19) 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 (20) 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.4 Reduced-Form Empirics using Local Projections

We slightly modify our dynamic specifications from the preceding section to run them as local projections instead and confirm that the results are similar. In particular, we run

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

repeatedly for each k from -6 through 24, plotting the coefficient β_0 in each case.

Figures G.8 and G.9 exhibit qualitatively similar behavior to the parallel Figures G.1 and G.2 from the preceding section. The only difference is that they are somewhat more ill-behaved after month 12, exhibiting slight decline in the cumulative pass-through coefficient.

C.5 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.10 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.11, 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.6 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. To do so, we slightly modify our main regression specifications to make them 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,

$$\sum_{j=0}^{t} \Delta P_{i,j} = \lambda_t + \beta_t [(\boldsymbol{I} - \boldsymbol{\Phi})^{-\mathbf{1}} \boldsymbol{s}^Z]_i + \epsilon_{i,t} \tag{21}$$

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, $[(\boldsymbol{I} - \boldsymbol{\Phi})^{-1} \boldsymbol{s}^Z]_i$ represents the network cost share of commodity or commodities $\mathbf{Z}$ in industry i, and $\epsilon_{i,t}$ is an error term.[25] 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:

$$\sum_{j=0}^{t} \Delta p_{i,j} = \lambda_t + \eta_t [\boldsymbol{s}^Z]_i + \rho_t [(\boldsymbol{I} - \boldsymbol{\Phi})^{-1} \boldsymbol{s}^Z - \boldsymbol{s}^Z]_i + \epsilon_{i,t} \tag{22}$$

We begin our case studies with the 1979 oil crisis. BLS PPI data for the Petroleum Refineries industry – one of the most crucial links in the oil supply chain – does not exist for the earlier 1970s, nor does PPI data for many other oil-consuming sectors; consequently, the 1979 shock is the earli-

[25] To ease the additional notational burden, we exclude general equilibrium wage controls in the case study specifications. As we discuss in Appendices C.1 and C.3, there are no meaningful differences in our results if we exclude the wage controls from the regression.

est 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 G.12 displays the regression coefficients originating from applying Equation (22) 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 G.12 displays the regression coefficients originating from applying Equation (22) 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 G.12 displays the regression coefficients originating from applying Equation (22) 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.

D Appendix: Robustness of Structural Estimation

In this appendix, we complete several robustness checks for our structural estimation. In D.1, we include petroleum refineries in the analysis. In D.2, we use the average frequency of price adjustment. In D.3, we compare pass-through speed to that found in our reduced-form dynamic analysis. In D.4, we assess whether pass-through of large and small shocks is different. In D.5, we report the results of our 2-step GMM procedure for non-oil commodities.

D.1 Including Petroleum Refineries

Our step 1 estimates including petroleum refineries are $m_f = 1$ (SE = 0.248) and $\hat{\bar{\mu}} = 1.162$ (SE = 0.095). 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: $m_f = 1$ (SE = 0.095), and $\hat{\bar{\mu}} = 1.019$ (SE = 0.040). 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 may change with petroleum refineries because they are such an outlier. So we reproduce our main figures displaying our oil results in Figures G.13, G.14, G.15, and G.16. In all four cases, we see similar findings to those discussed in the paper, where

refineries were excluded.

D.2 Using the Average Frequency of Price Adjustment

This section uses the GMM-optimal model but sets $\theta_i = \bar{\theta}$, the average frequency of price adjustment – associated with an average duration of industry prices of 5 months. We plot our results in Figure G.17.

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 (10) 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.

D.3 Pass-through Speed Compared to Reduced Form Dynamics

Now, we address the question of why pass-through seems so much slower in our structural analysis than it did in our analysis from our dynamic reduced-form section C.3, where our shift-share did not use horizon-specific shares. 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 (10), the main shift-share regression equation underlying our GMM procedure. 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 downstream industries are not as responsible for identifying any reduced-form coefficients in section 5, 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 demonstrate this formally by simulating data from the GMM-optimal model. Formally, we feed observed oil price movements starting in January 1997 into regression equation (10), setting $\epsilon_{i,t} = 0$ for all i, t and $\lambda_t = 0$ for all t, generating simulated price changes for every industry, $\Delta P_{i,t}^{Sim}$. We then re-run our reduced form regression (20) using these simulated outcomes. The results are shown in Panel 2 of Figure G.18. We see that we nearly replicate the result found earlier in the paper, shown again for convenience in Panel 1. This formalizes the intuition that the dynamics in reduced-form are heavily biased towards showing pass-through that is too fast.

D.4 Pass-through of Large versus Small Shocks

While we have assumed time-dependent pricing, a model with state-dependent pricing might find that industries pass through large movements in commodity prices faster than small movements in commodity prices. We partition monthly oil price movements into two sets: those whose absolute size is larger than the median absolute size of oil price movements starting in 1997, and those whose absolute size is smaller. Define the corresponding dummy variable $Large_t = \mathbf{1}\{|\Delta P_{Oil,t}| > Med(|\Delta P_{Oil,t}|)\}$. Our regression of interest is a simple heterogeneity analysis on our main structural regression equation (10):

$$\begin{aligned}\Delta P_{i,t} = \lambda_t &+ \sum_{h=0}^{H} \beta_{large,h}(\boldsymbol{\alpha}) Large_t passthrough_{i,h}(\boldsymbol{\alpha}) \Delta P_{Z,t-h} \\ &+ \sum_{h=0}^{H} \beta_{small,h}(\boldsymbol{\alpha})(1 - Large_t) passthrough_{i,h}(\boldsymbol{\alpha}) \Delta P_{Z,t-h} + \epsilon_{i,t}.\end{aligned} \tag{23}$$

By a result from Auclert et al. (2022), if state-dependent pricing is operative, we should find evidence of faster pass-through for large shocks than predicted under our GMM-optimal time-dependent pricing model. In Figure G.19, we show that the GMM-optimal model retains good fit for both large and small oil price movements, and there is no statistically detectable difference in pass-through of small versus large shocks.

D.5 2-Step GMM for Non-oil Commodities

We show that our results for non-oil commodities are robust to a 2-step GMM procedure, where we use the results from our 1-step GMM procedure to form an optimal weight matrix and re-run GMM re-weighting the moments using this matrix. The results are shown in Figure G.20.

E Appendix: Incorporating 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. We then 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 Appendix C.3. Again, we confirm that firms pass through changes in expected future oil prices, holding changes in the spot price constant.

E.1 Model Solution with Futures

As we showed in detail in Appendix A.2, the model can be written as

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

where

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

and the matrices $\boldsymbol{B}$ come from the equation[26]

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

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

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

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

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

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

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

where

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

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

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

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

[26] We 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 $\hat{x}_{t+1}$ 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 $\tilde{\boldsymbol{x}}_{1,t}$ is more interesting. We have

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

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:

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

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

Now we will assume log commodity futures prices satisfy $f_{Z,t,m} = -c_m + \mathrm{E}_t[p_{Z,t+m}]$, 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.[27] 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

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

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

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

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

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

where the time difference operator holds the maturity fixed.

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

[27] This assumption neglects time-varying risk free rates, which are not a major driver of the variation over this period, and embeds some limits to arbitrage.

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

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

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

$$\begin{aligned}
\Delta \boldsymbol{p}_t &= A_1\Delta \boldsymbol{p}_{t-1} + A_2 \sum_{m=0}^{\infty} \left(\Lambda_1^{-1}\right)^{m+1} \tilde{\boldsymbol{B}}_{1,Z}\Delta f_{Z,t,m} \\
&= A_1^2\Delta \boldsymbol{p}_{t-2} + A_1A_2 \sum_{m=0}^{\infty} \left(\Lambda_1^{-1}\right)^{m+1} \tilde{\boldsymbol{B}}_{1,Z}\Delta f_{Z,t-1,m} \\
&+ A_2 \sum_{m=0}^{\infty} \left(\Lambda_1^{-1}\right)^{m+1} \tilde{\boldsymbol{B}}_{1,Z}\Delta f_{Z,t,m},
\end{aligned}$$

and, more generally,

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

We have found that, for H large enough, the state variable $\Delta \boldsymbol{p}_{t-(H+1)}$ should not matter much for 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 $\boldsymbol{\alpha}$, prices in industry i react to a change in the commodity future h periods ago at maturity m according to*

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

The following corollary results:

Corollary 2. *For large H,*

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

In pratice, we will let the highest maturity be 60 months. It is also the case that $\left(\Lambda_1(\boldsymbol{\alpha})^{-1}\right)^{61} \approx 0$, 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 $passthrough_{i,h}(\boldsymbol{\alpha})$, where $\boldsymbol{\alpha}$ was the vector of calibration parameters used in the model. We tested these model predictions using the regression

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

If the model under calibration $\boldsymbol{\alpha}$ 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 $\hat{\beta}_h(\boldsymbol{\alpha}) = 1$ for all h. Our GMM estimated $\boldsymbol{\alpha}$ to get these β's as close as possible to 1.

After incorporating futures, the model instead delivers predictions $passthrough_{i,h,m}(\boldsymbol{\alpha})$: 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 $passthrough_{i,h}(\boldsymbol{\alpha}) = \sum_{m \geq 0} passthrough_{i,h,m}(\boldsymbol{\alpha})$.

A full test of the model incorporating futures would be

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

where $\Delta f_{Z,t-h,m}$ is the change in the commodity 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 a commodity 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.

This regression is infeasible for a typical number of lags and maturities because it yields $(H + 1) \times (M + 1)$ β's to estimate. We now refer back to the main body of the paper for how we collapse the large number of β's above into a smaller set more feasible to estimate.

E.3 Reduced-form Test of Forward-lookingness about Shock Persistence

We now incorporate futures exposure into our regression setup from Section C.3 of the paper. This section used an arbitrary lag structure with industries' long-run exposures to commodity shocks, rather than using the dynamic model's implied lag structure for industry pass-through. We alter regression equation (19) to include an additional set of terms that should be predictive of industry price changes if industries incorporate commodity futures data when forming expectations about future commodity prices. This regression is most closely related to the fully structural regression test captured in regression 12.

$$\Delta P_{i,t} = \lambda_t + \sum_{h=-6}^{24} \beta_{h,Impact}[(\boldsymbol{I} - \boldsymbol{\Phi})^{-1}\boldsymbol{s}^Z]_i \Delta P_{Z,t-h} + \sum_{h=-6}^{24} \beta_{h,Forward}[(\boldsymbol{I} - \boldsymbol{\Phi})^{-1}\boldsymbol{s}^Z]_i \Delta f_{Z,t-h,m} + \epsilon_{i,t}.$$

We use the change in the one-year maturity future, so $m = 12$. When pooling non-oil commodities, we weight changes in commodity futures by the commodity weights underlying the S&P commodity index.

We plot our results for oil in the left panel of Figure G.21. We see very similar results to those found in Figure 9, with oil price movements only translating into (lasting) price increases further down the supply chain when the oil price movements are expected to persist. We also plot our results for non-oil commodities in the right panel of Figure G.21. In stark contrast, we find very little evidence of differential pass-through for different movements in the non-oil commodity future curve, holding movements in spot prices constant.

F Appendix: Additional Empirics for Application

In this section, we discuss further details for the application of our model. In F.1, we outline how to incorporate commodity futures into measured network oil inflation. In F.2, we estimate the aggregation intercept. In F.3, we show that oil price movements have no effect on the Federal Funds Rate. In F.4, we show how oil price movements pass through to aggregate inflation in a fully specified model. In F.5, we show the fit of our GMM-optimal model for the PCE-weighted average industry. In F.6, we subtract network oil inflation from aggregate PCE inflation.

F.1 Incorporating Oil Futures Data

The formula for predicting inflation changes from formula 13 to

$$\hat{\Pi}_t = Intercept + \sum_i PCEShare_i \sum_{h=0}^{H} \sum_{m=0}^{M} passthrough_{i,h,m}(\hat{\boldsymbol{\alpha}}) \Delta f_{Z,t-h,m}, \tag{24}$$

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

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

with

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

for H and M suitably large, say $H = 24$ and $M = 60$, which is feasible using futures data.

F.2 Appendix: Regressions for Computing Aggregation Intercept

Recall that regression 10 tested our GMM-optimal model and was given by

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

where $\hat{\boldsymbol{\alpha}}$ 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 merely removes the time fixed-effect and adds interaction terms as required:[28]

$$\Delta P_{i,t} = \alpha + \sum_{h=0}^{24} \delta_{1,h} passthrough_{i,h}(\hat{\boldsymbol{\alpha}}) + \sum_{h=0}^{24} \delta_{2,h} passthrough_{i,h}(\hat{\boldsymbol{\alpha}})\Delta P_{Z,t-h} + \sum_{h=0}^{24} \tilde{\beta}_h \Delta P_{Z,t-h} + \epsilon_{i,t}. \tag{25}$$

Then $\sum_h \tilde{\beta}_h$ is the effect of a unit oil price increase on prices in a sector with no network exposure to the commodity, i.e. the aggregation intercept.

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 exposures and the oil price changes with the interaction between the network exposures and the Kanzig instruments at the appropriate horizons. Our results are depicted in Panel 1 of Figure G.22. We see that, as before, with the exception of two values for h, we cannot reject that $\sum_h \tilde{\beta}_h = 0$.

We can alternatively use the dynamic forms of our reduced-form regressions to compute the intercept, though we prefer the approach using the GMM-optimal dynamics from above. Recall

[28] The δ_1 terms are optional.

that regression 19 was

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

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

$$\Delta P_{i,t} = \alpha + \delta_1 [(\boldsymbol{I} - \boldsymbol{\Phi})^{-1} \boldsymbol{s}^Z]_i + \sum_{h=-6}^{24} \delta_{2,h} [(\boldsymbol{I} - \boldsymbol{\Phi})^{-1} \boldsymbol{s}^Z]_i \Delta P_{Z,t-h} + \sum_{h=-6}^{24} \tilde{\beta}_h \Delta P_{Z,t-h} + \epsilon_{i,t}. \quad (26)$$

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 2 of Figure G.22. We see that, with the exception of a couple pre-periods, we cannot reject that $\sum_h \tilde{\beta}_h = 0$.

F.3 No Effect of Oil Price Movements on the Federal Funds Rate

We regress changes in the federal funds rate on distributed lags of the Kanzig shocks to establish whether exogenous oil price increases (decreases) lead to increases (decreases) in the federal funds rate. We find no statistically significant evidence of this effect and negative point estimates. The results are shown in Figure G.23. We also show results for a regression of changes in the FFR on distributed lags of oil price changes, instrumented with the Kanzig shocks. This approach leads to the same result.

F.4 Fully Closed Model's Implications for the Intercept

In appendix section A.4, we outlined the assumptions for our fully closed macroeconomic model that includes a household and a Taylor rule. Rather than computing the intercept empirically, we can simply plot the effect of an oil shock on aggregate inflation in that model. It includes many of the potentially correlated general equilibrium effects we discussed as potentially problematic for our supply-side model: movements in the aggregate wage or in monetary policy that would primarily appear in the time fixed effect our of reduced-form and structural analyses.

We plot the results in Figure G.24. The model leads to results very similar to our findings from Figure 13. If anything, they are slightly higher effects than we found in Figure 13. The reason is that, while the Taylor rule implies contractionary monetary policy in response to the inflation resulting from the oil shock (we have no output gap in our Taylor rule), enough substitution to

labor from the intermediate input bundle is possible throughout the network that aggregate wages rise by more than enough to offset the contractionary monetary policy.

F.5 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 $\hat{\boldsymbol{\alpha}}$ 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,

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

with the IRF in the data,

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

where $\hat{\beta}_h(\hat{\boldsymbol{\alpha}})$ was computed as in our structural estimation, using regression equation 10.

The results are shown in Figure G.25. We see that model fit remains good for the PCE-weighted average industry, even though fit was optimized for the equally weighted average industry.

F.6 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 gasoline prices and non-gas prices from official PCE inflation, showing our results in Figure G.26.

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 summer 2022 to 5 percentage

points from above 6 percentage points. Thereafter, official PCE inflation and oil prices began to decline, but our measure of underlying inflation continues to increase through the latter half of 2022, suggesting that the apparent decrease in inflation in that period was entirely driven by the network effects of the oil price decreases. 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.

G Appendix Tables and Figures

Table G.1 – Pass-through Regressions: Dropping Time Fixed Effects

Panel 1: One Month Horizon

	(1)	(2)	(3)	(4)	(5)
Dependent Variable: ΔPPI (Month-over-Month)	Oil	Oil (No Refineries)	Kanzig	Kanzig (No Refineries)	Non-Oil
Total Network Cost Change	0.439***	0.247***	0.665***	0.302***	0.366***
	(0.130)	(0.070)	(0.249)	(0.095)	(0.072)
R-Squared	0.0118	0.0025	0.0057	0.0016	0.0094
Observations	103,480	103,186	97,324	97,048	97,974

Panel 2: One Year Horizon

	(1)	(2)	(3)	(4)	(5)
Dependent Variable: ΔPPI (Year-over-Year)	Oil	Oil (No Refineries)	Kanzig	Kanzig (No Refineries)	Non-Oil
Total Network Cost Change	0.809***	0.651***	0.905***	0.739***	0.793***
	(0.134)	(0.137)	(0.137)	(0.138)	(0.087)
R-Squared	0.0532	0.0244	0.0431	0.0186	0.0423
Observations	96,479	96,197	90,323	90,059	91,643

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. This table shows the results of regressions corresponding to Equation (8), except with the time fixed effects dropped. Columns (1) and (2) focus on all oil price variation from 1997-2022. Columns (3) and (4) focus on oil price variation induced by Kanzig's (2021) OPEC shock series. Column (5) focuses on non-oil commodity price variation from 1997-2022. Standard errors are clustered by industry. We find that pass-through is far below the full pass-through benchmark of 1 over a one-month horizon. More pass-through accumulates over a one-year horizon.

Table G.2 – Pass-through Regressions: Dropping Wage Control Variable

Panel 1: One Month Horizon

Dependent Variable: ΔPPI (Month-over-Month)	(1) Oil	(2) Oil (No Refineries)	(3) Kanzig	(4) Kanzig (No Refineries)	(5) Non-Oil
Total Network Cost Change	0.506***	0.263***	0.774***	0.266**	0.337***
	(0.141)	(0.093)	(0.295)	(0.125)	(0.071)
R-Squared	0.1190	0.1126	0.1211	0.1197	0.1557
Observations	107,242	106,937	97,324	97,048	101,516

Panel 2: One Year Horizon

Dependent Variable: ΔPPI (Year-over-Year)	(1) Oil	(2) Oil (No Refineries)	(3) Kanzig	(4) Kanzig (No Refineries)	(5) Non-Oil
Total Network Cost Change	0.880***	0.704***	0.974***	0.761***	0.718***
	(0.141)	(0.191)	(0.150)	(0.178)	(0.086)
R-Squared	0.1949	0.1700	0.1572	0.1361	0.2048
Observations	100,241	99,948	90,323	90,059	95,185

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. This table shows the results of regressions corresponding to Equation (8), except with the wage control variable dropped. Columns (1) and (2) focus on all oil price variation from 1997-2022. Columns (3) and (4) focus on oil price variation induced by Kanzig's (2021) OPEC shock series. Column (5) focuses on non-oil commodity price variation from 1997-2022. Standard errors are clustered by industry. We find that pass-through is far below the full pass-through benchmark of 1 over a one-month horizon. More pass-through accumulates over a one-year horizon.

Table G.3 – Pass-through Regressions: Pass-through into Wages

Panel 1: One Month Horizon

Dependent Variable: ΔWage (Month-over-Month)	(1) Oil	(2) Oil (No Refineries)	(3) Kanzig	(4) Kanzig (No Refineries)	(5) Non-Oil
Total Network Cost Change	0.006	0.007	0.013	0.005	-0.009
	(0.005)	(0.010)	(0.012)	(0.020)	(0.007)
R-Squared	0.0534	0.0532	0.0503	0.0501	0.0535
Observations	98,335	98,041	92,611	92,335	94,990

Panel 2: One Year Horizon

Dependent Variable: ΔWage (Year-over-Year)	(1) Oil	(2) Oil (No Refineries)	(3) Kanzig	(4) Kanzig (No Refineries)	(5) Non-Oil
Total Network Cost Change	0.012	0.018	0.016	0.022	-0.045**
	(0.020)	(0.038)	(0.019)	(0.035)	(0.022)
R-Squared	0.1048	0.1046	0.1011	0.1008	0.1046
Observations	91,532	91,250	85,808	85,544	88,886

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. This table shows the results of regressions corresponding to Equation (8), except with the change in industry wages rather than the change in industry prices on the left-hand-side and the wage control variable dropped. Columns (1) and (2) focus on all oil price variation from 1997-2022. Columns (3) and (4) focus on oil price variation induced by Kanzig's (2021) OPEC shock series. Column (5) focuses on non-oil commodity price variation from 1997-2022. Standard errors are clustered by industry. We find no consistent evidence of an effect of network oil cost changes on industry wages.

Table G.4 – Pass-through Regressions: Adding TFP Control Variable

Panel 1: One Month Horizon

Dependent Variable: ΔPPI (Month-over-Month)	(1) Oil	(2) Oil (No Refineries)	(3) Kanzig	(4) Kanzig (No Refineries)	(5) Non-Oil
Total Network Cost Change	0.528***	0.262***	0.793***	0.236*	0.328***
	(0.147)	(0.101)	(0.308)	(0.133)	(0.072)
R-Squared	0.1333	0.1269	0.1359	0.1346	0.1733
Observations	103,480	103,186	97,324	97,048	97,974

Panel 2: One Year Horizon

Dependent Variable: ΔPPI (Year-over-Year)	(1) Oil	(2) Oil (No Refineries)	(3) Kanzig	(4) Kanzig (No Refineries)	(5) Non-Oil
Total Network Cost Change	0.847***	0.595***	0.970***	0.725***	0.700***
	(0.166)	(0.197)	(0.159)	(0.186)	(0.093)
R-Squared	0.1741	0.1504	0.1692	0.1490	0.1894
Observations	96,479	96,197	90,323	90,059	91,643

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. This table shows the results of regressions corresponding to Equation (8), except with a TFP control variable added. Columns (1) and (2) focus on all oil price variation from 1997-2022. Columns (3) and (4) focus on oil price variation induced by Kanzig's (2021) OPEC shock series. Column (5) focuses on non-oil commodity price variation from 1997-2022. Standard errors are clustered by industry. We find that pass-through is far below the full pass-through benchmark of 1 over a one-month horizon. More pass-through accumulates over a one-year horizon.

Table G.5 – Pass-through Regressions: Adding Network Gas & Electricity Cost Change Control

Panel 1: One Month Horizon

Dependent Variable: ΔPPI (Month-over-Month)	(1) Oil	(2) Oil (No Refineries)	(3) Kanzig	(4) Kanzig (No Refineries)
Total Network Cost Change	0.571***	0.317***	0.759***	0.311**
	(0.156)	(0.121)	(0.275)	(0.153)
R-Squared	0.1360	0.1295	0.1433	0.1404
Observations	72,598	72,387	66,442	66,249

Panel 2: One Year Horizon

Dependent Variable: ΔPPI (Year-over-Year)	(1) Oil	(2) Oil (No Refineries)	(3) Kanzig	(4) Kanzig (No Refineries)
Total Network Cost Change	0.848***	0.604***	0.943***	0.796***
	(0.183)	(0.225)	(0.135)	(0.202)
R-Squared	0.1488	0.1147	0.1405	0.1110
Observations	67,621	67,421	61,465	61,283

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. This table shows the results of regressions corresponding to Equation (8), except with a control for network gas/electricity cost change added. Columns (1) and (2) focus on all oil price variation from 1997-2022. Columns (3) and (4) focus on oil price variation induced by Kanzig's (2021) OPEC shock series. Column (5) focuses on non-oil commodity price variation from 1997-2022. Standard errors are clustered by industry. We find that pass-through is far below the full pass-through benchmark of 1 over a one-month horizon. More pass-through accumulates over a one-year horizon.

Table G.6 – Pass-through Regressions: Cost Shares without Capital in Denominator

Panel 1: One Month Horizon

Dependent Variable: ΔPPI (Month-over-Month)	(1) Oil	(2) Oil (No Refineries)	(3) Kanzig	(4) Kanzig (No Refineries)	(5) Non-Oil
Total Network Cost Change	0.358***	0.198***	0.502**	0.220**	0.340***
	(0.119)	(0.061)	(0.214)	(0.086)	(0.071)
R-Squared	0.1183	0.1129	0.1226	0.1203	0.1561
Observations	103,480	103,186	97,324	97,048	97,974

Panel 2: One Year Horizon

Dependent Variable: ΔPPI (Year-over-Year)	(1) Oil	(2) Oil (No Refineries)	(3) Kanzig	(4) Kanzig (No Refineries)	(5) Non-Oil
Total Network Cost Change	0.606***	0.449***	0.750***	0.596***	0.737***
	(0.139)	(0.120)	(0.130)	(0.109)	(0.089)
R-Squared	0.1628	0.1414	0.1557	0.1377	0.1763
Observations	96,479	96,197	90,323	90,059	91,643

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. This table shows the results of regressions corresponding to Equation (8), except with cost shares computed without capital in the denominator. Columns (1) and (2) focus on all oil price variation from 1997-2022. Columns (3) and (4) focus on oil price variation induced by Kanzig's (2021) OPEC shock series. Column (5) focuses on non-oil commodity price variation from 1997-2022. Standard errors are clustered by industry. We find that pass-through is far below the full pass-through benchmark of 1 over a one-month horizon. More pass-through accumulates over a one-year horizon.

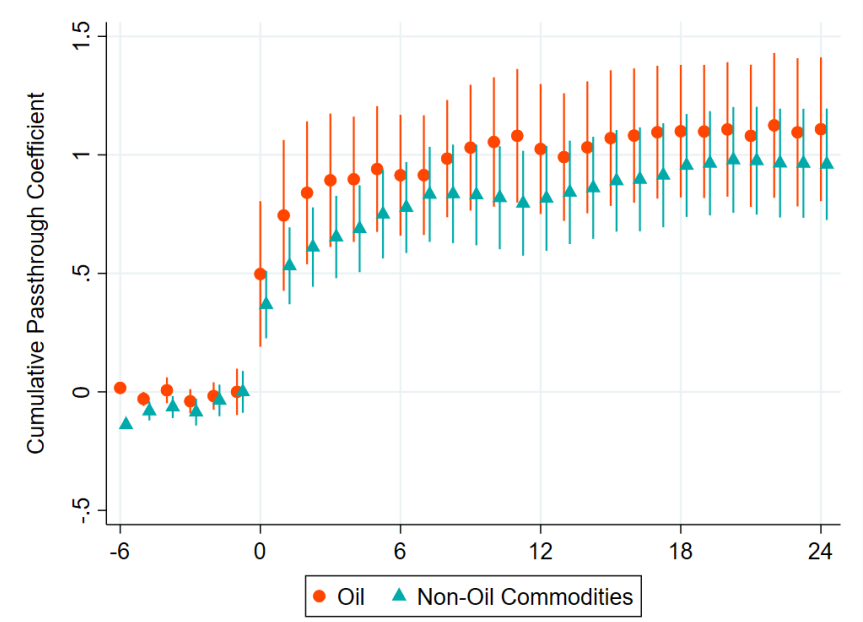
Table G.7 – Heterogeneity on Other Plausible Correlates of Price Pass-through

Dependent Variable: ΔPPI	(1) Oil	(2) Kanzig	(3) Non-Oil
Total Network Cost Change	0.535***	0.809***	1.010***
	(0.117)	(0.263)	(0.225)
Total Network Cost Change	-0.198***	-0.364***	-0.236*
*Downstreamness	(0.068)	(0.124)	(0.127)
Total Network Cost Change	-0.108***	-0.053*	-0.042
*Price Adjust. Half-Life	(0.027)	(0.031)	(0.031)
Total Network Cost Change	0.210**	0.269**	0.013
*Concentration	(0.092)	(0.127)	(0.139)
Total Network Cost Change	0.038	0.135***	-0.060*
*Sales	(0.024)	(0.031)	(0.033)
Total Network Cost Change	0.308***	0.371***	0.442***
*Inventory	(0.094)	(0.157)	(0.160)
Total Network Cost Change	0.413	-1.740	-2.913
*Capital Share	(0.794)	(1.353)	(1.878)
R-Squared	0.1941	0.1809	0.1286
Observations	27,110	27,110	492,601

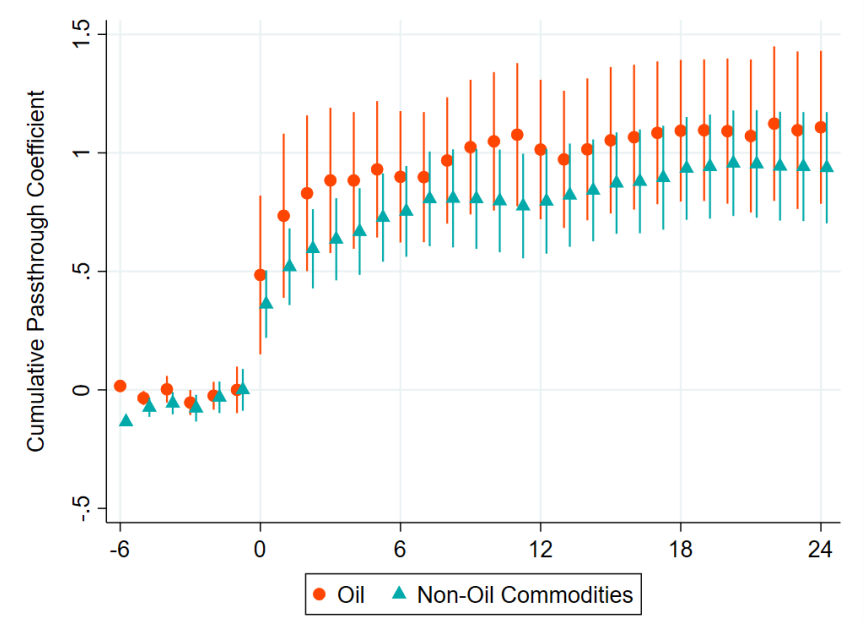
Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The outcome variable of this table is the month-over-month change in industry prices. Column (1) uses all oil price variation. Column (2) is an IV specification using the Kanzig oil shock series. Column (3) uses all non-oil commodity price variation. Standard errors are clustered by industry. We find that the interaction with downstreamness remains significant - and of similar or greater magnitude - after the addition of other factors that could plausibly affect pass-through.

Figure G.1 – Month-by-Month pass-through of Oil & Non-Oil Commodity Price Changes

Panel 1: Separate Regressions



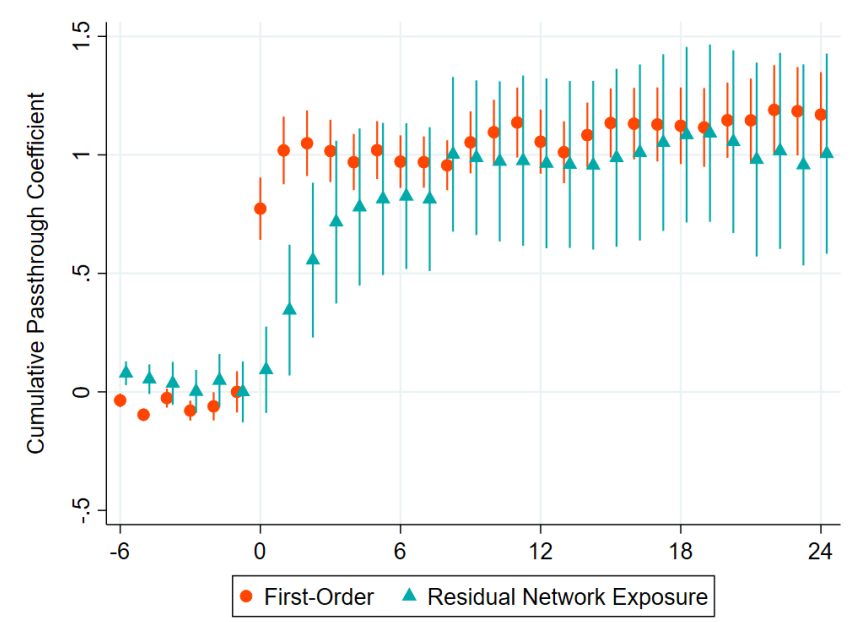
Panel 2: Joint Regression



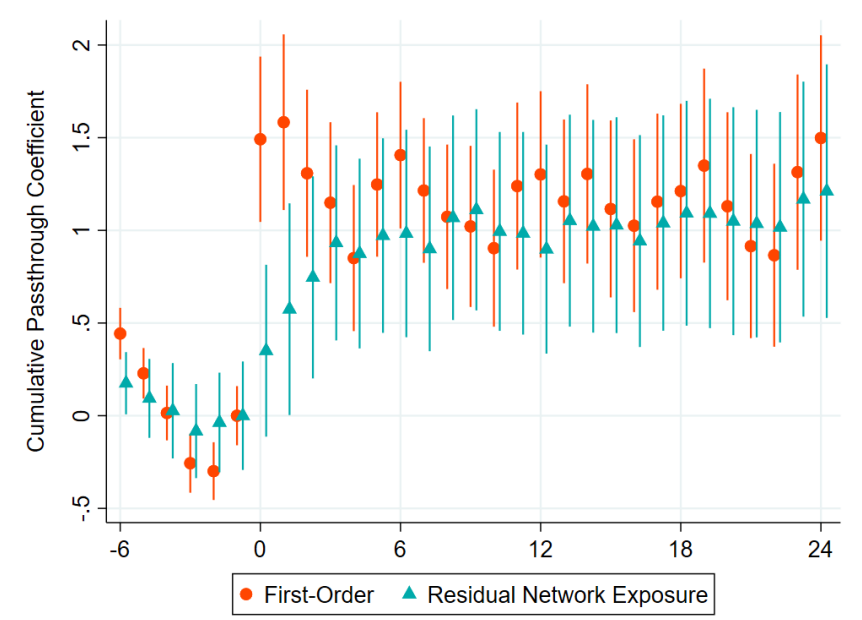
Note: The left panel plots two separate regression specifications, both corresponding to Equation (19), which entails regressing an industry's price changes on that industry's cost changes. Red (circle) coefficients plot monthly price pass-through of network exposure to oil price changes; blue (triangle) coefficients plot monthly price pass-through of network exposure to non-oil commodity price changes. The right panel includes these two separate sets of regressors in the same regression. Both plots are consistent with full but gradual pass-through of commodity price movements to industry prices.

Figure G.2 – Month-by-Month pass-through of Oil Price Changes: Direct vs. Indirect

Panel 1: All Variation



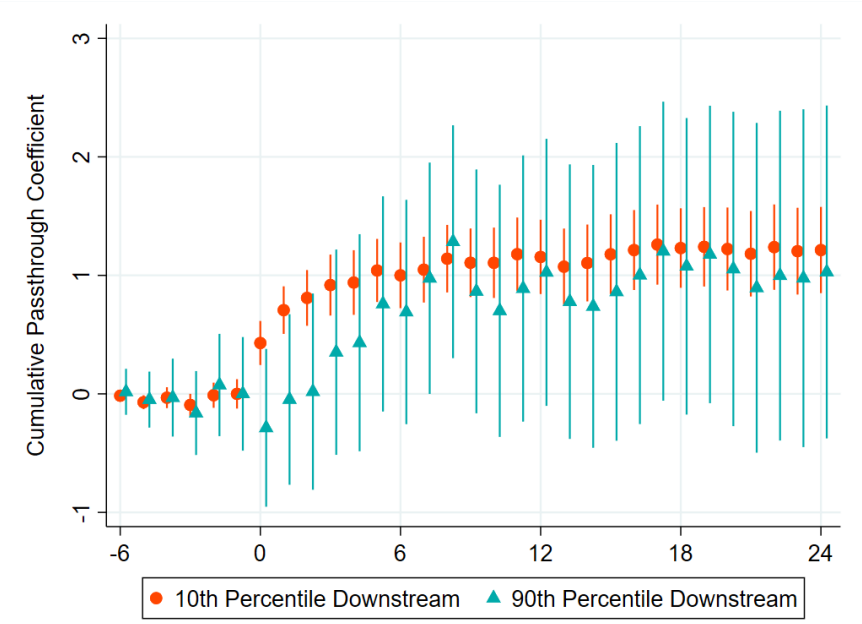
Panel 2: Kanzig Variation



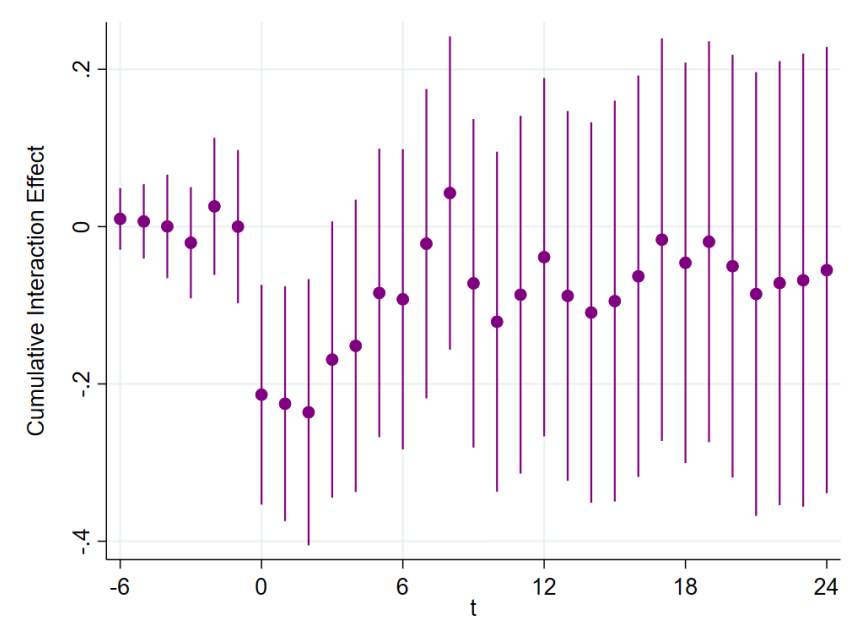
Note: Regression specifications correspond to Equation (20), which entails regressing an industry's price changes on that industry's oil cost changes - separated into direct, first-order oil cost changes and indirect, higher-order oil cost changes that filter through the production network. Red (circle) coefficients plot monthly price pass-through of first-order exposure to crude oil price shocks; blue (triangle) coefficients plot monthly price pass-through of second- and higher-order network exposure to crude oil price shocks. The regression in the left panel uses all variation in oil prices; the regression in the right panel uses the Kanzig (2021) variation stemming from high-frequency identification of the effects of OPEC announcements on oil prices. Both plots are consistent with full pass-through of oil price movements to industry prices, with slower pass-through when exposure is indirect.

Figure G.3 – Month-by-Month pass-through of Oil Price Changes: Upstream vs. Downstream Industries

Panel 1: Upstream versus Downstream



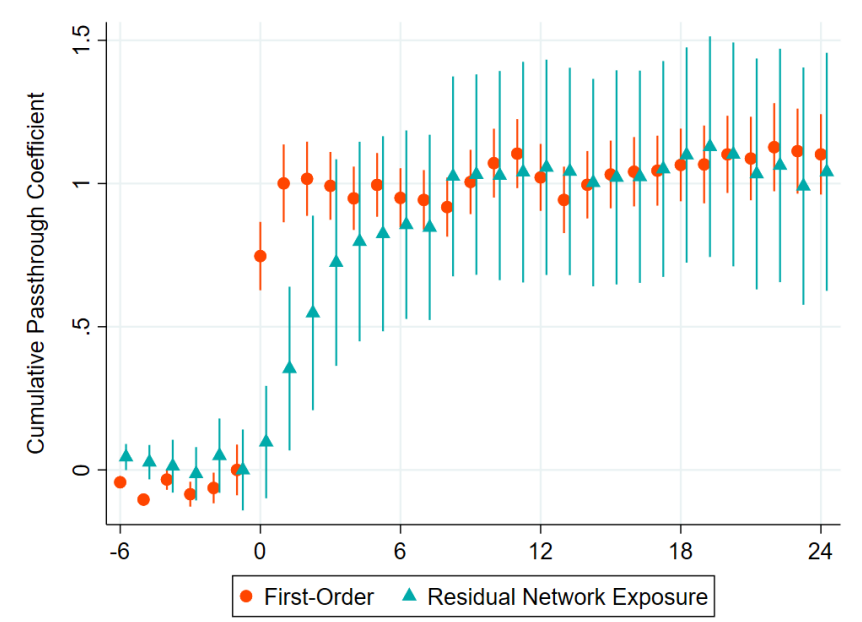
Panel 2: Interaction Effect



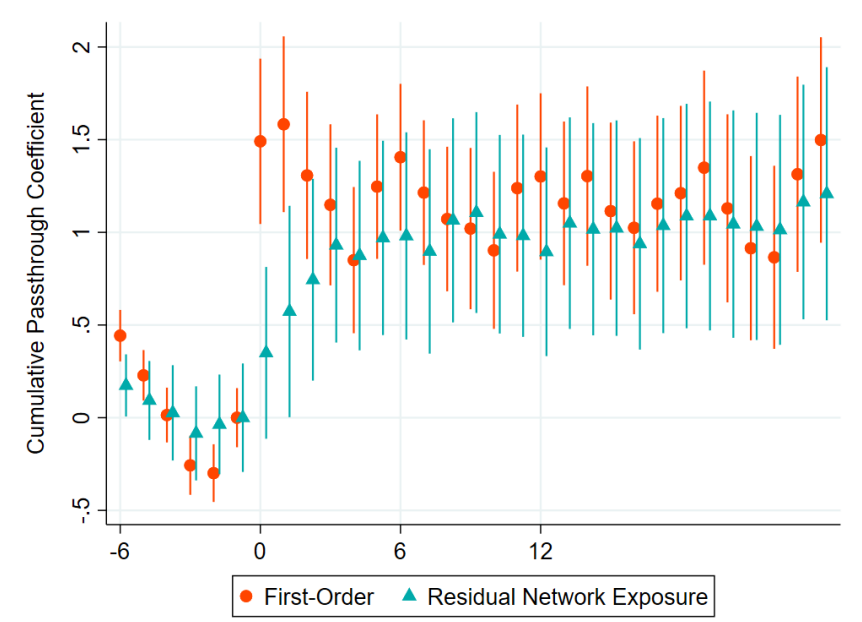
Note: Regression specifications correspond to column (4) of Table 2, with shock terms interacted with our measure of downstreamness. Consequently, the left panel plots monthly cumulative price pass-through of crude oil price shocks for both upstream firms (10th percentile downstreamness) relatively close to commodity production/extraction and downstream firms (90th percentile downstreamness) at or near the consumer side of the economy. The right panel plots the cumulated interaction coefficient to establish that the difference is statistically-significant. 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 G.4 – Month-by-Month pass-through of Oil Price Changes

Panel 1: All Variation



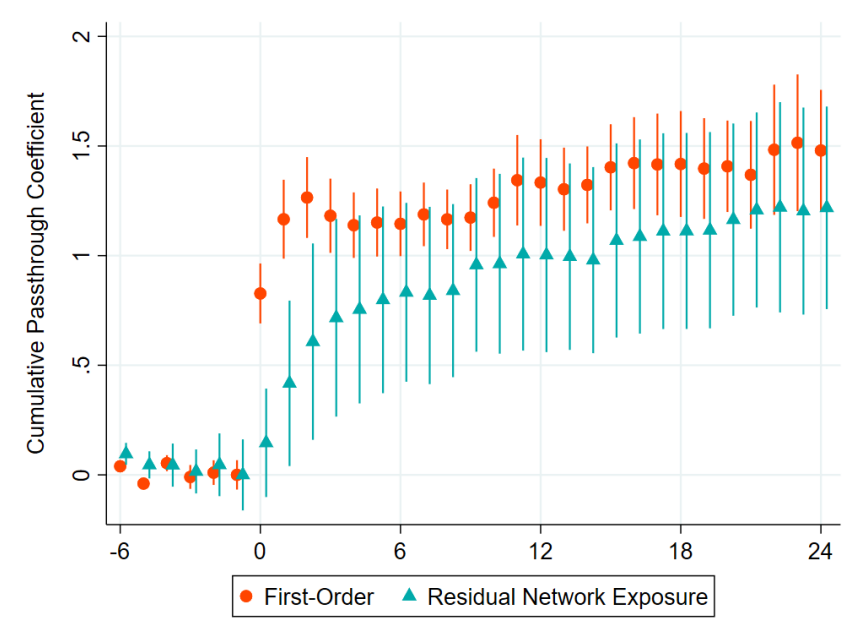
Panel 2: Kanzig IV Variation



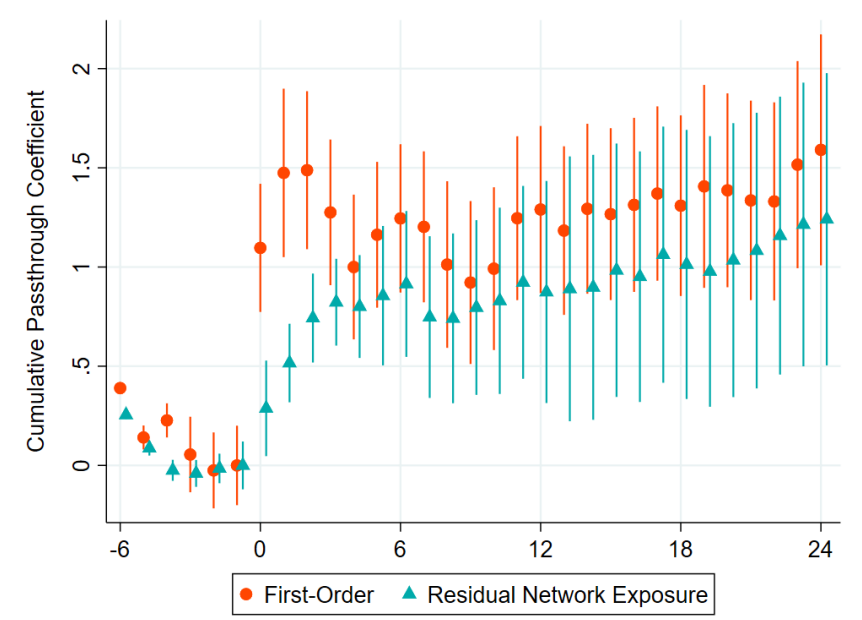
Note: Regression specifications correspond to Equation (20), which entails regressing an industry's price changes on that industry's oil cost changes - separated into direct, first-order oil cost changes and indirect, higher-order oil cost changes that filter through the production network. Red coefficients 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.5 – Month-by-Month pass-through of Oil Price Changes

Panel 1: All Variation



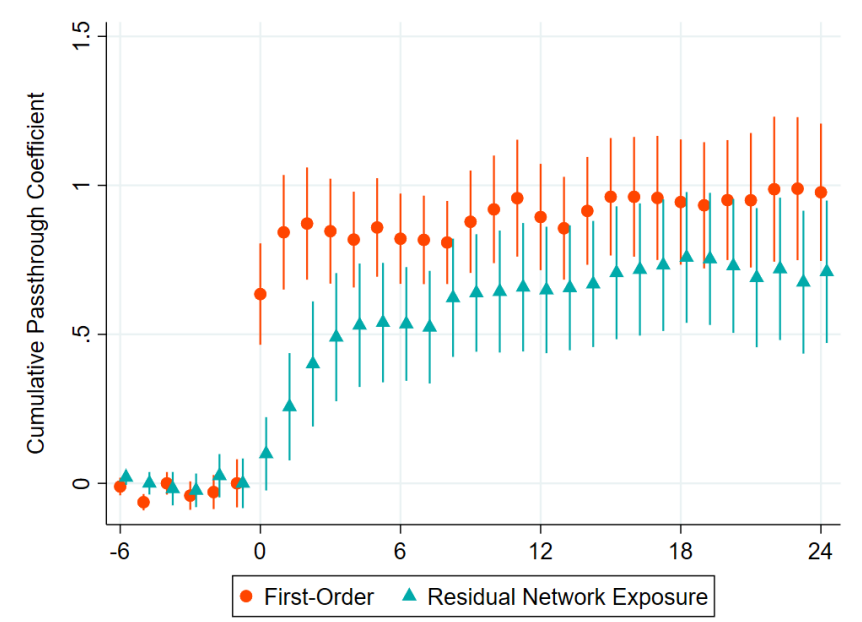
Panel 2: Kanzig IV Variation



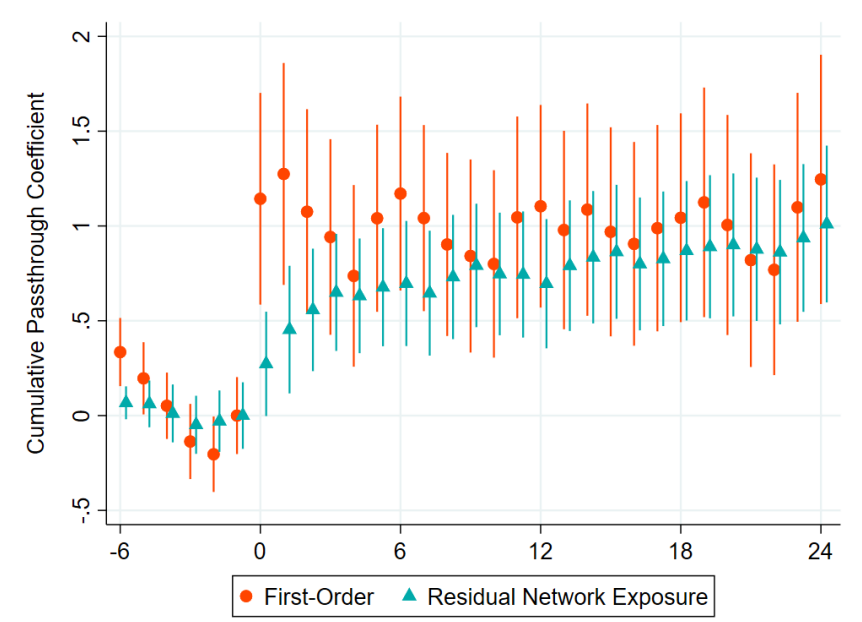
Note: Regression specifications correspond to Equation (20) with controls for gas and electricity shock exposure. Red coefficients plot monthly price pass-through of first-order exposure to crude oil price shocks; blue coefficients plot monthly price pass-through of second- and higher-order network exposure to crude oil price shocks.

Figure G.6 – Month-by-Month pass-through of Oil Price Changes

Panel 1: All Variation



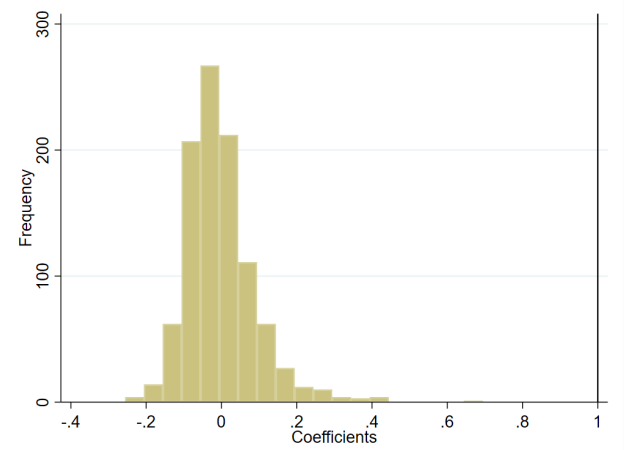
Panel 2: Kanzig IV Variation



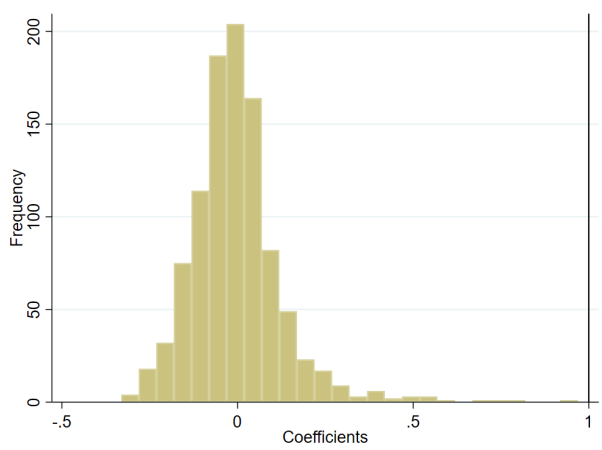
Note: Regression specifications correspond to Equation (20), albeit with alternative oil cost shares that do not include capital in the denominator. Red coefficients plot monthly price pass-through of first-order exposure to crude oil price shocks; blue coefficients plot monthly price pass-through of second- and higher-order network exposure to crude oil price shocks.

Figure G.7 – Results of Permutation Tests

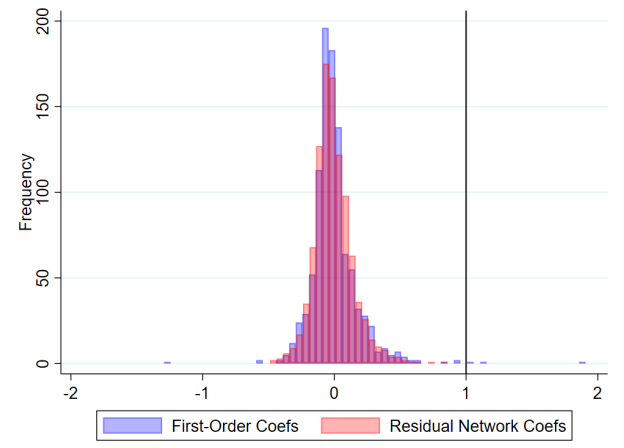
Panel 1a: All Oil Variation, Total Network



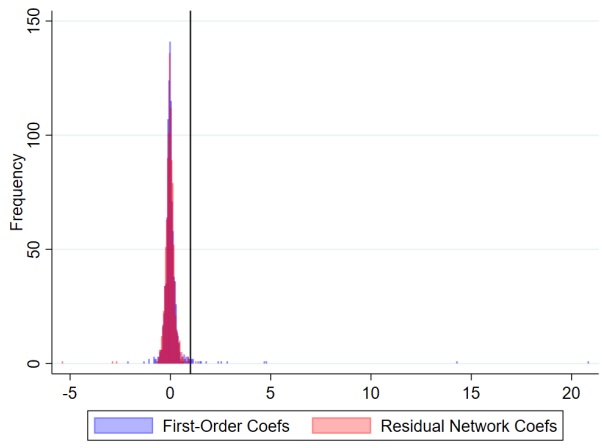
Panel 1b: Kanzig IV Variation, Total Network



Panel 2a: All Oil Variation, Direct/Indirect



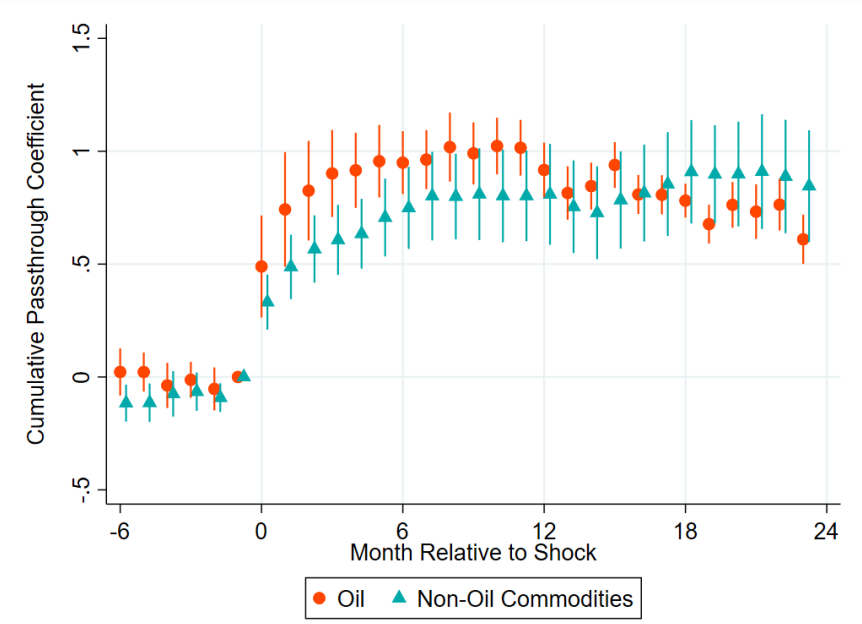
Panel 2b: Kanzig IV Variation, Direct/Indirect



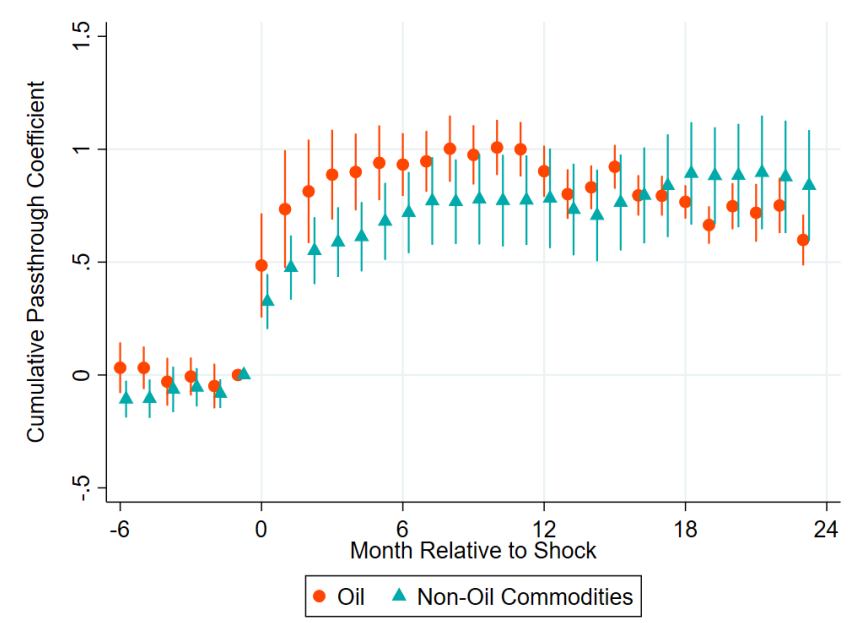
Note: These plots display the results of 1000-repetition permutation tests on the 12-month cumulative pass-through coefficient from the main specifications, as given by Equation (19) in the top panel and (20) in the bottom panel.

Figure G.8 – Local Projections: Month-by-Month pass-through of Oil & Non-Oil Commodity Price Changes

Panel 1: Separate Regressions



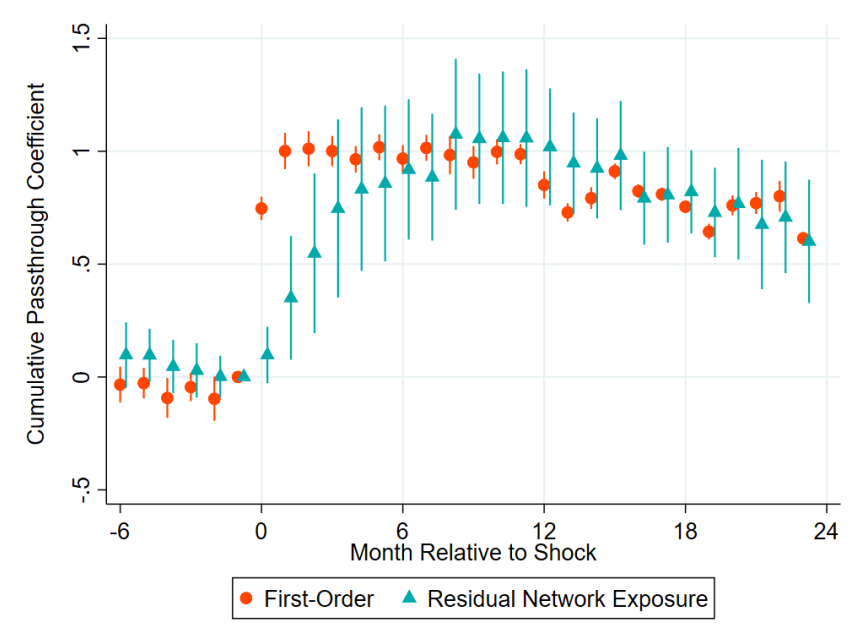
Panel 2: Joint Regression



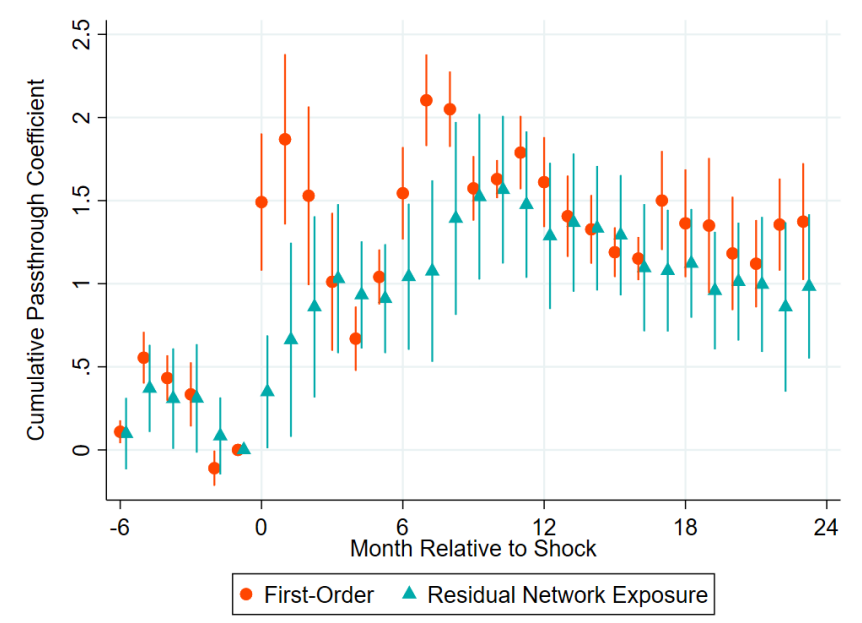
Note: The left panel plots two separate regression specifications, both corresponding to local projections regressions of an industry's cumulative price changes on that industry's network cost changes. Red (circle) coefficients plot monthly price pass-through of network exposure to oil price changes; blue (triangle) coefficients plot monthly price pass-through of network exposure to non-oil commodity price changes. The right panel includes these two separate sets of regressors in the same regression. Both plots are consistent with full but gradual pass-through of commodity price movements to industry prices.

Figure G.9 – Local Projections: Month-by-Month pass-through of Oil Price Changes: Direct vs. Indirect

Panel 1: All Variation

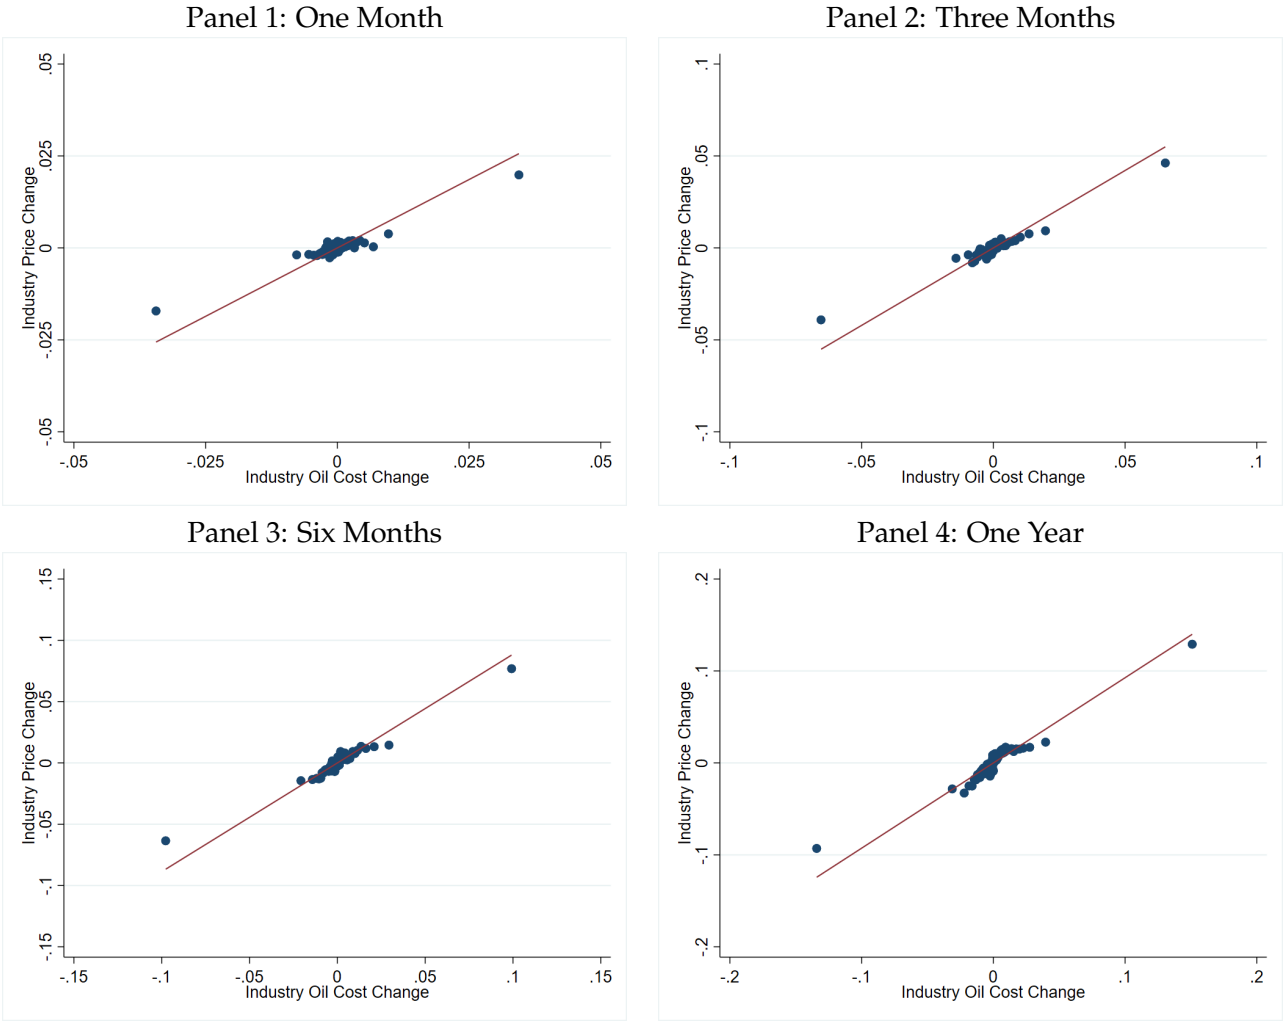


Panel 2: Kanzig Variation



Note: The left panel plots two separate regression specifications, both corresponding to local projections regressions of an industry's cumulative price changes on that industry's oil cost changes - separated into direct, first-order oil cost changes and indirect, higher-order oil cost changes that filter through the production network. Red (circle) coefficients plot monthly price pass-through of first-order exposure to crude oil price shocks; blue (triangle) coefficients plot monthly price pass-through of second- and higher-order network exposure to crude oil price shocks. The regression in the left panel uses all variation in oil prices; the regression in the right panel uses the Kanzig (2021) variation stemming from high-frequency identification of the effects of OPEC announcements on oil prices. Both plots are consistent with full pass-through of oil price movements to industry prices, with slower pass-through when exposure is indirect.

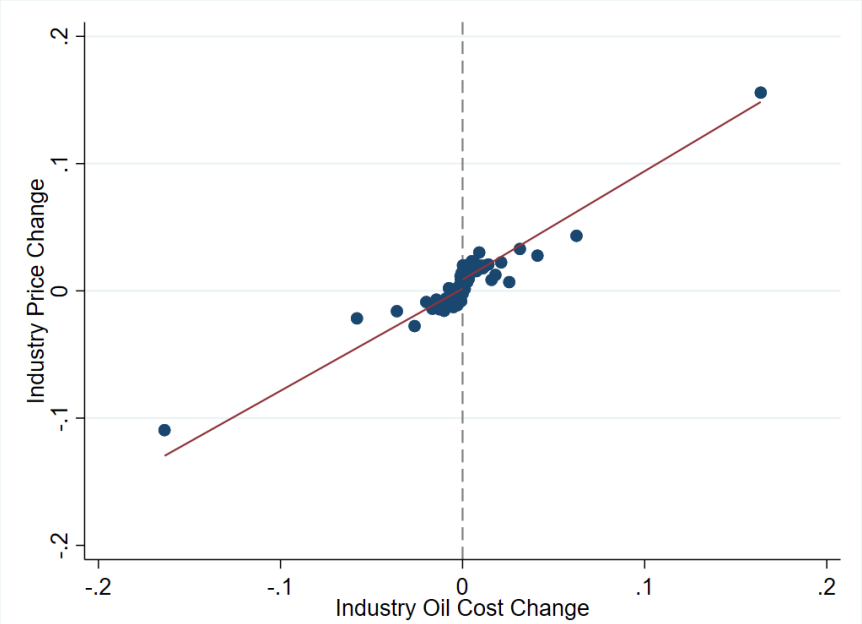
Figure G.10 – Binscatters at Various Time Horizons



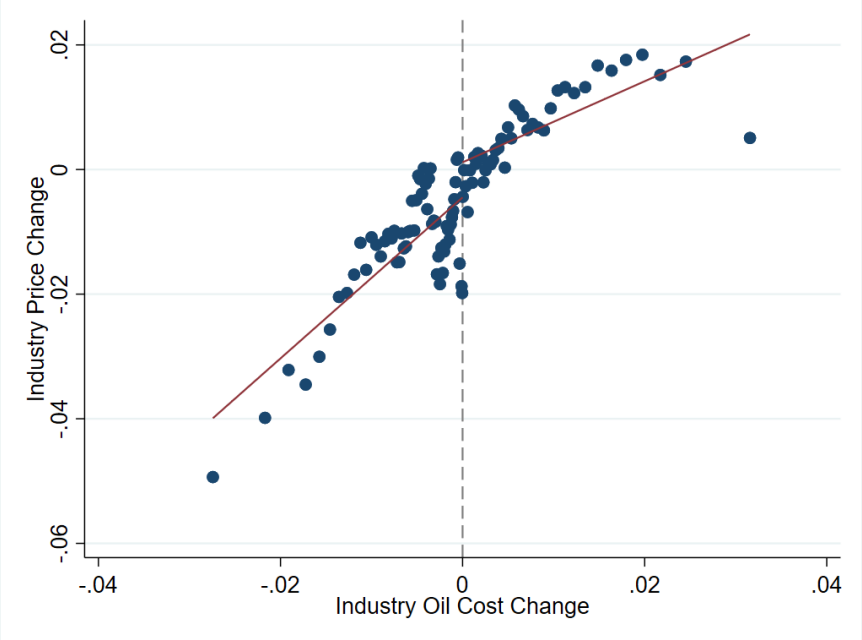
Note: These plots are 100-quantile binscatters displaying how (de-meaned) industry price changes vary with (de-meaned) industry oil cost changes over different time horizons. Lines of best fit are included. The slope of these lines can be interpreted as the fraction of cost increases that are passed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

Figure G.11 – Binscatters: Upstream vs. Downstream Industries

Panel 1: Upstream Industries

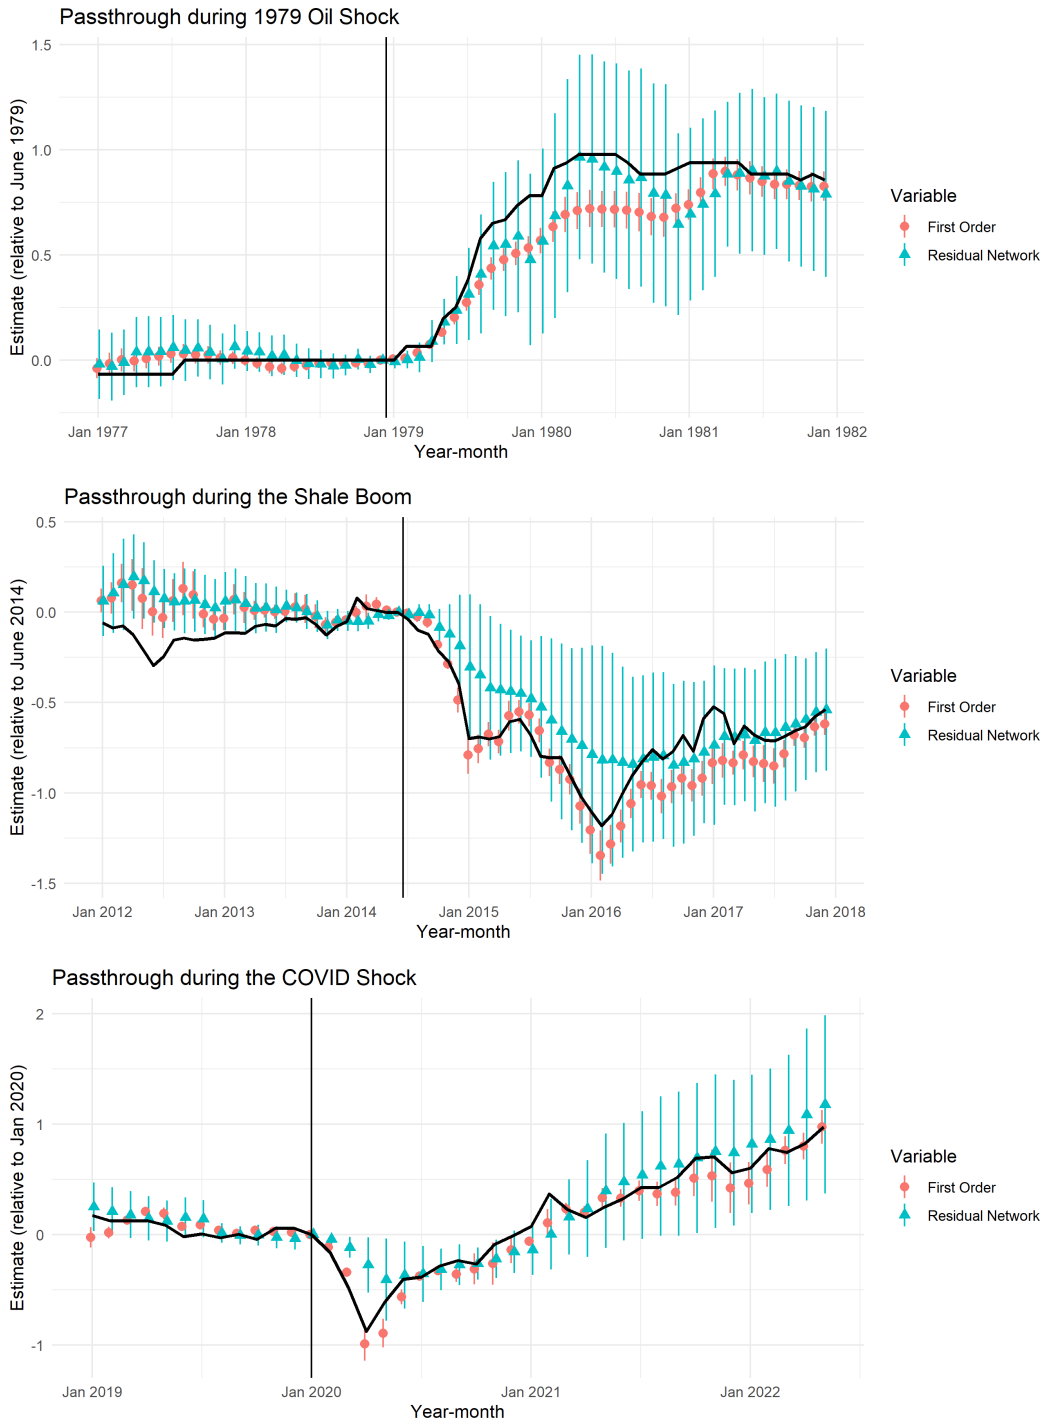


Panel 2: Downstream Industries



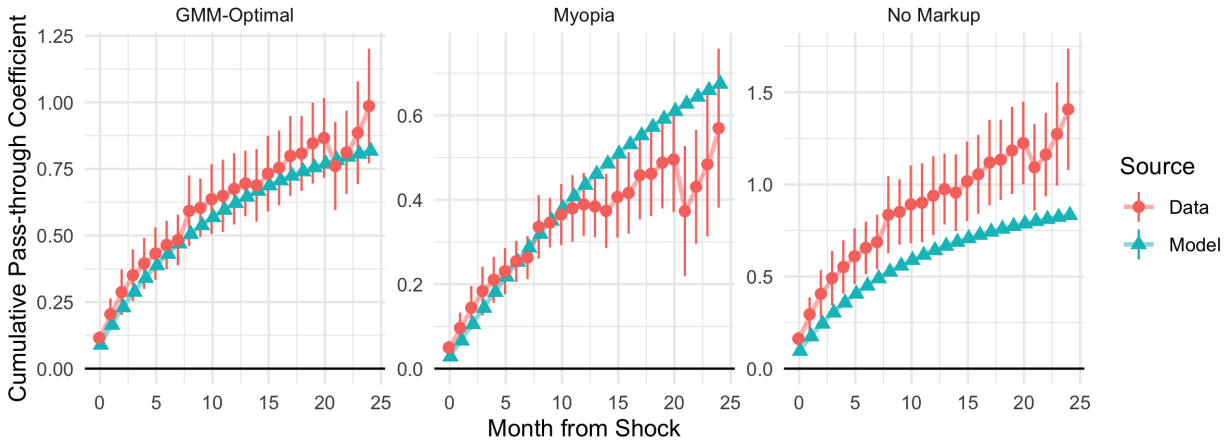
Note: These plots are 100-quantile binscatters displaying how (de-meaned) industry price changes vary with (de-meaned) industry oil cost changes for upstream versus downstream industries. Lines of best fit with a regression discontinuity at zero are included. For upstream industries (industries with 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.12 – Case Studies



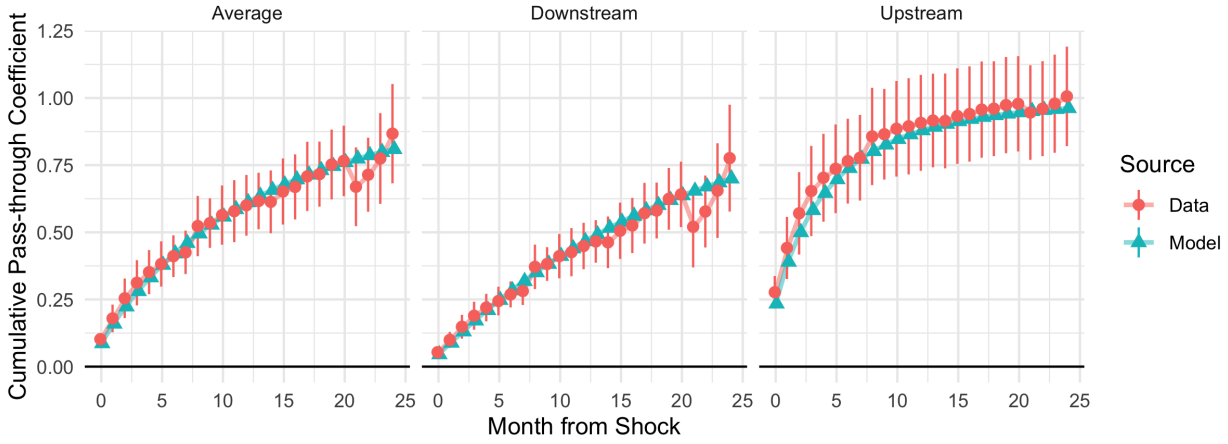
Note: Regression specification corresponds to Equation (22), which entails regressing the level of an industry's prices relative to some base period on the level of that industry's oil costs - separated into direct, first-order oil cost changes and indirect, higher-order oil cost changes that filter through the production network. The 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 G.13 – GMM Results for Industry Pass-through of Oil Price Changes (Including Refineries)



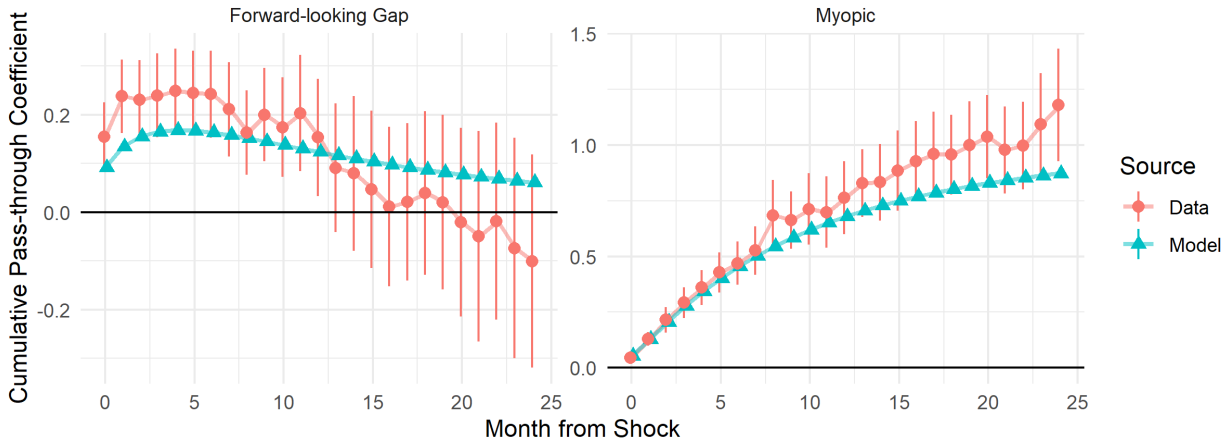
Note: Model fit is good for the GMM-optimal model, as shown in Panel 1. We provide intuition for how parameters are identified and show results for alternative calibrations in Panels 2 and 3. Under alternative calibrations, model fit is meaningfully worse. Panel 2 analyzes the case where firms respond myopically to oil price increases; for higher lags, the myopic model predicts more pass-through than the data appears to warrant. Panel 3 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 G.14 – Industry Pass-through of Oil Price Changes, Upstream and Downstream (Including Refineries)



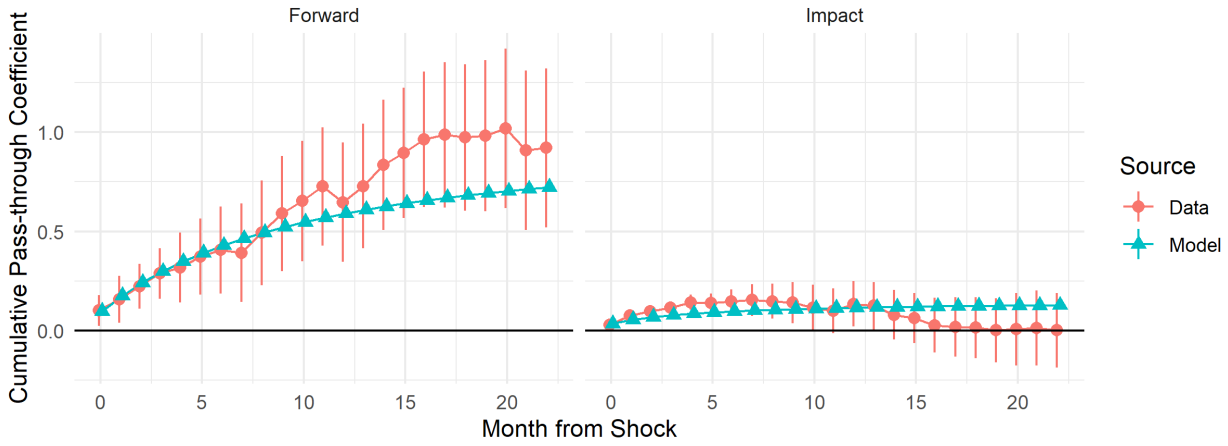
Note: The GMM procedure optimized fit for the average industry (leftmost plot). We visualize how well the fit is for the most upstream and downstream industries in the middle and rightmost plots, finding that model fit remains good. Moreover, we see how much pass-through speed varies with downstreamness. Upstream industries achieve 75% of long-run pass-through in just 6 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 G.15 – Tests of Forward-lookingness about Network Dynamics (Including Refineries)



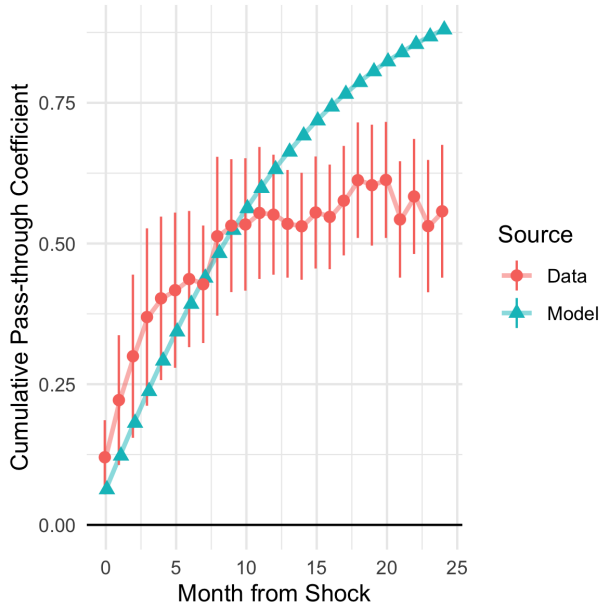
Note: We plot the results of our test that pass-through due to forward-lookingness about network dynamics under rational expectations is present in the data. Statistical significance in the right panel implies we cannot reject that firms are forward-looking about the gradual pass-through of upstream 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 G.16 – Tests of Forward-lookingness about Shock Persistence (Including Refineries)



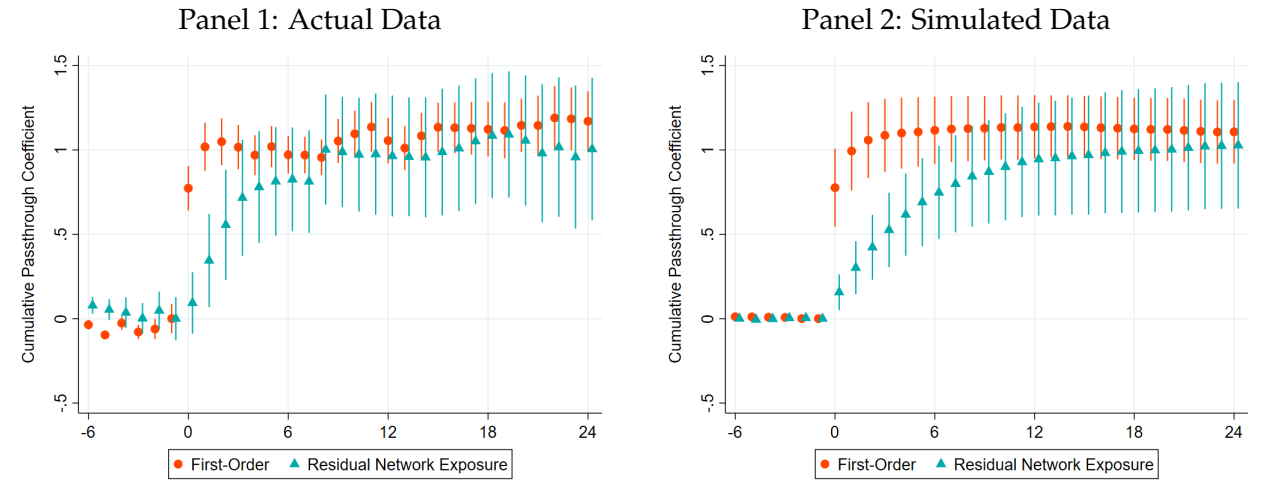
Note: We plot the results of our test that pass-through due to forward-lookingness about shock persistence under rational expectations is present in the data. Statistical significance in the left panel implies we cannot reject that firms are forward-looking about the persistence of an oil price shock. 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. We are only powered to perform this analysis out to 22 lags rather than 24.

Figure G.17 – GMM Results for Industry Pass-through of Oil Price Changes (Homogeneous θ)



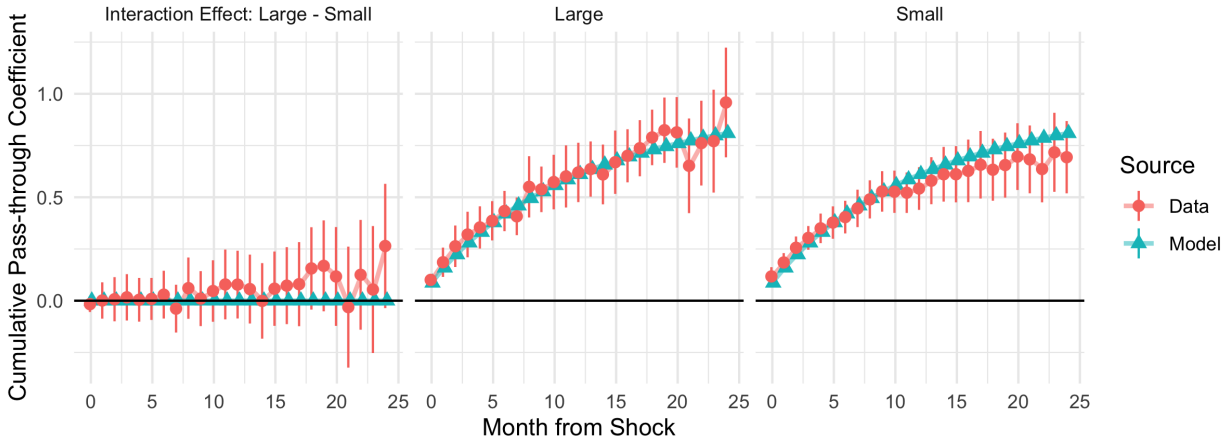
Note: The figure assesses the experiment in which all industries have the average frequency of price adjustment, associated with a price duration of 5 months; the data clearly prefer the model in which more upstream industries have higher frequencies of price adjustment and more downstream industries have lower frequencies of price adjustment, as we observe in the data.

Figure G.18 – Month-by-Month pass-through of Oil Price Changes: Actual vs. Simulated Data



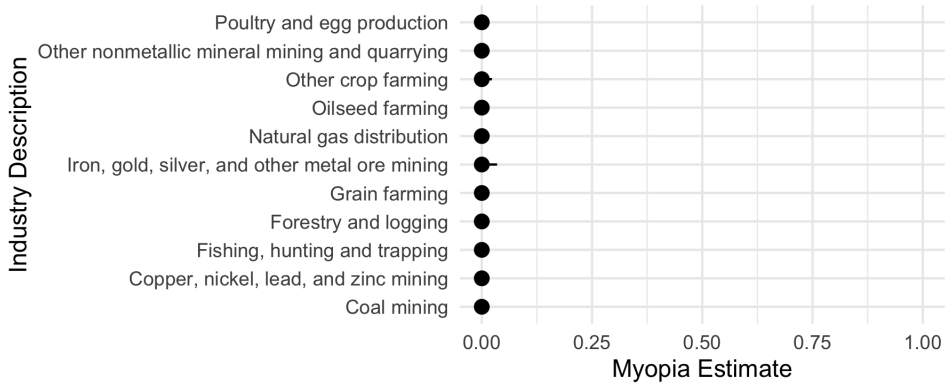
Note: In Figure 4, we showed that the GMM-optimal model yields that the average industry reaches 75% of long-run pass-through of an oil price movement only after 20 months have passed. Our reduced form findings (replicated in Panel 1), however, show 75% of long-run pass-through due to indirect use of oil occurring in as few as 4 months, and 75% of long-run pass-through due to direct use occurring immediately on shock impact. In Panel 2, we show the results of running our reduced form regression on data simulated from the GMM-optimal model. The results nearly replicate our findings in actual data, showing that the reduced form methodology is severely biased towards pass through that is too fast. The reason is that sectors making greater use of oil, both directly and indirectly, have much higher frequencies of price adjustment on average.

Figure G.19 – GMM Results for Small versus Large Oil Price Changes



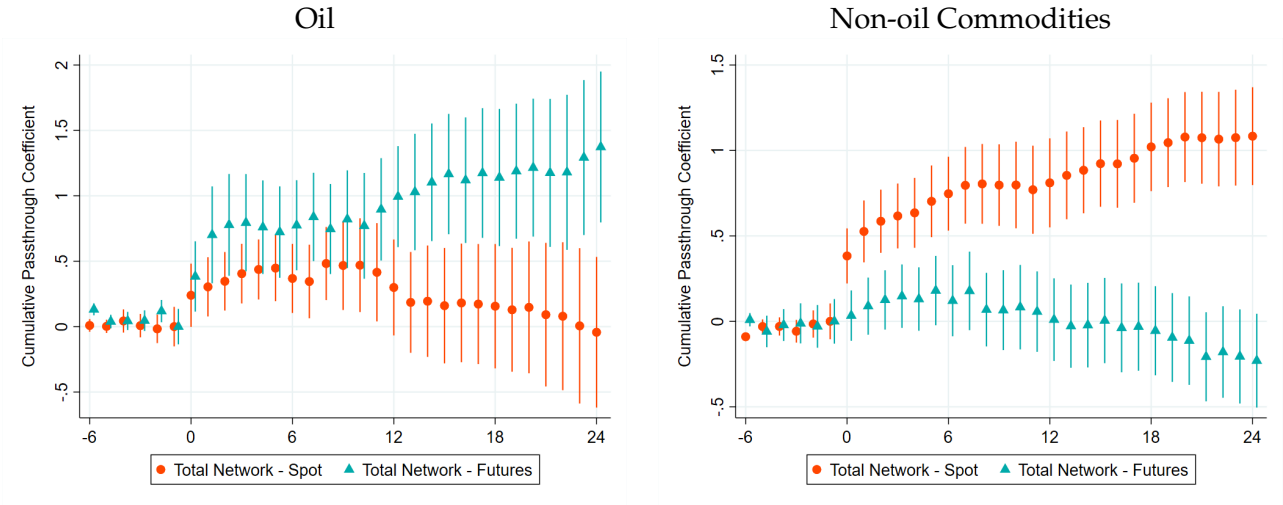
Note: The figure plots cumulative pass-through in the data compared to the GMM optimal model for small and large oil shocks separately. The standard errors for pass-through differences are quantified under the interaction effect panel. A price change is defined as large if the absolute size of the oil price change in a given month is larger than the median absolute size of oil price changes starting in 1997.

Figure G.20 – Estimates of Myopia for Non-oil Commodities (2-step GMM)



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.

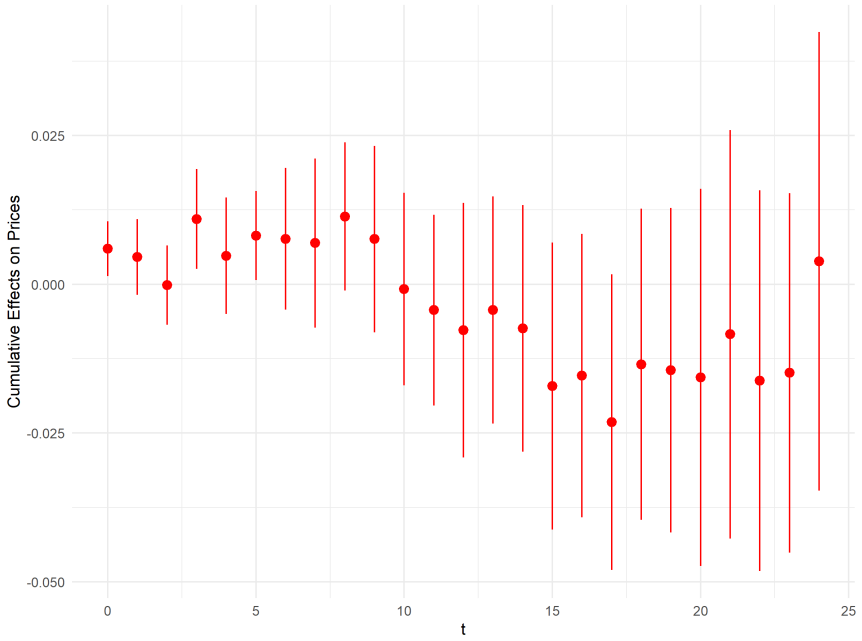
Figure G.21 – Baseline Result Incorporating Futures in Reduced Form



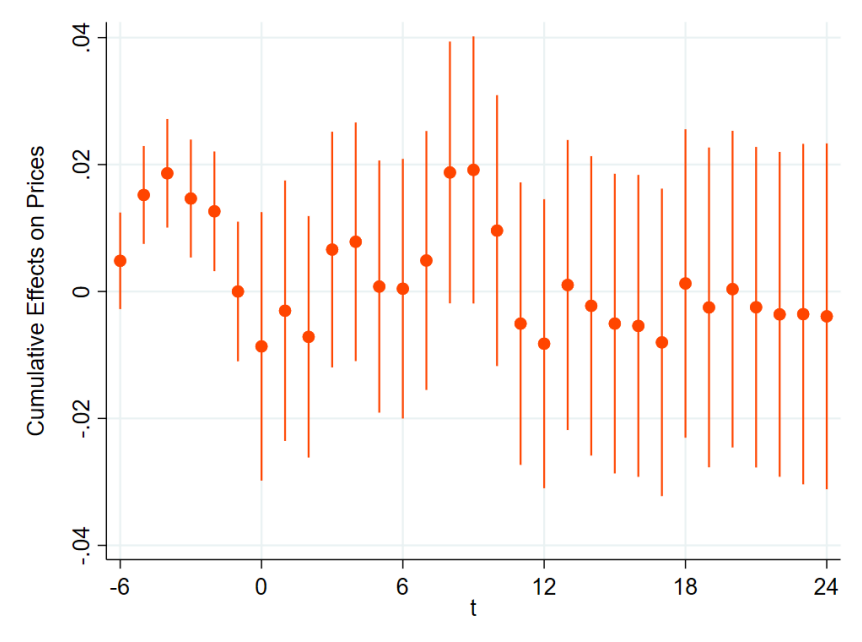
Note: Reduced-form test of whether firms respond to movements in futures prices in addition to movements in spot prices. For oil prices, we confirm that firms respond to both movements, with full pass-through in the long-run (a long-run point estimate of 1). For non-oil commodity prices, we instead find evidence that firm's respond to movements in the spot price only. There is no statistically significant evidence of pass-through differing with changes in the futures curve, holding changes in the spot price constant.

Figure G.22 – Estimation of the Intercept Required for Aggregating Industry Price Responses

Panel 1: Intercept using Regression Equation 25



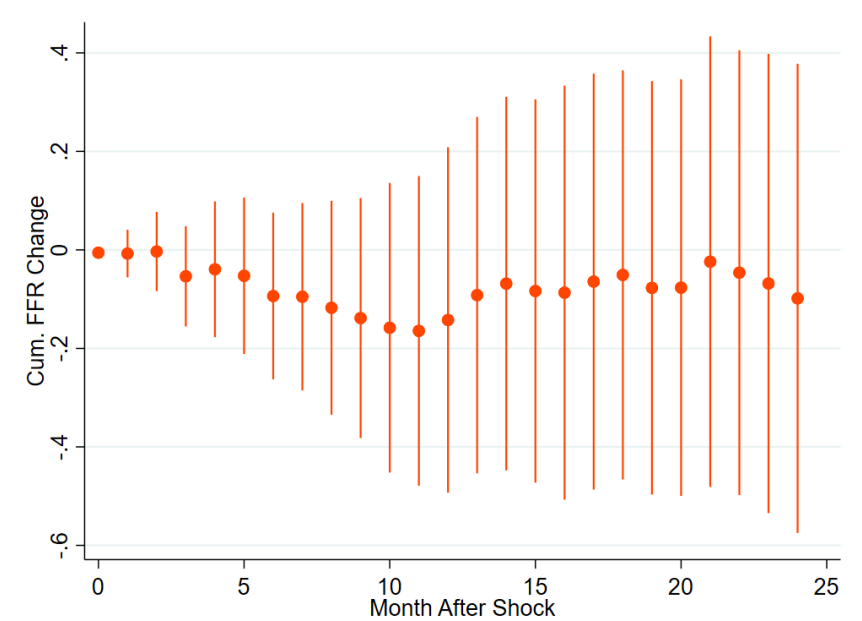
Panel 2: Intercept using Regression Equation 26



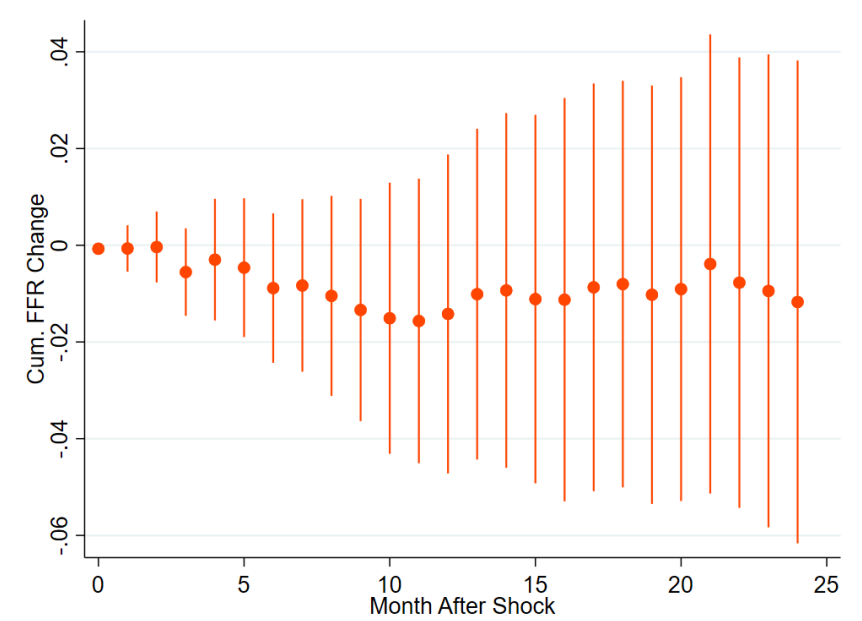
Note: We plot the intercept required for aggregation in our Application section using the 2SLS procedure described in appendix F.2. We see that, with the exception of a few periods, we cannot reject that the required intercept is 0. This suggests that sectors with no network exposure to oil largely do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

Figure G.23 – FFR Effects of Oil Price Movements

Panel 1: Reduced Form using Kanzig Shock

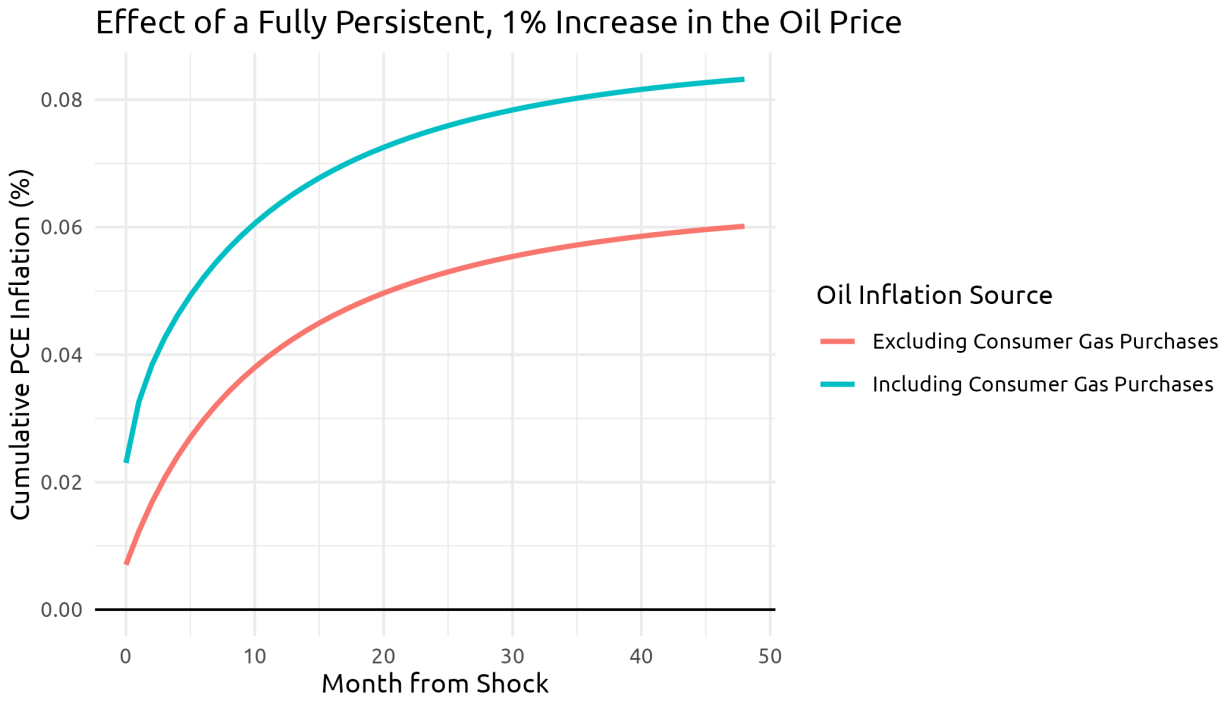


Panel 2: IV using Kanzig Shock



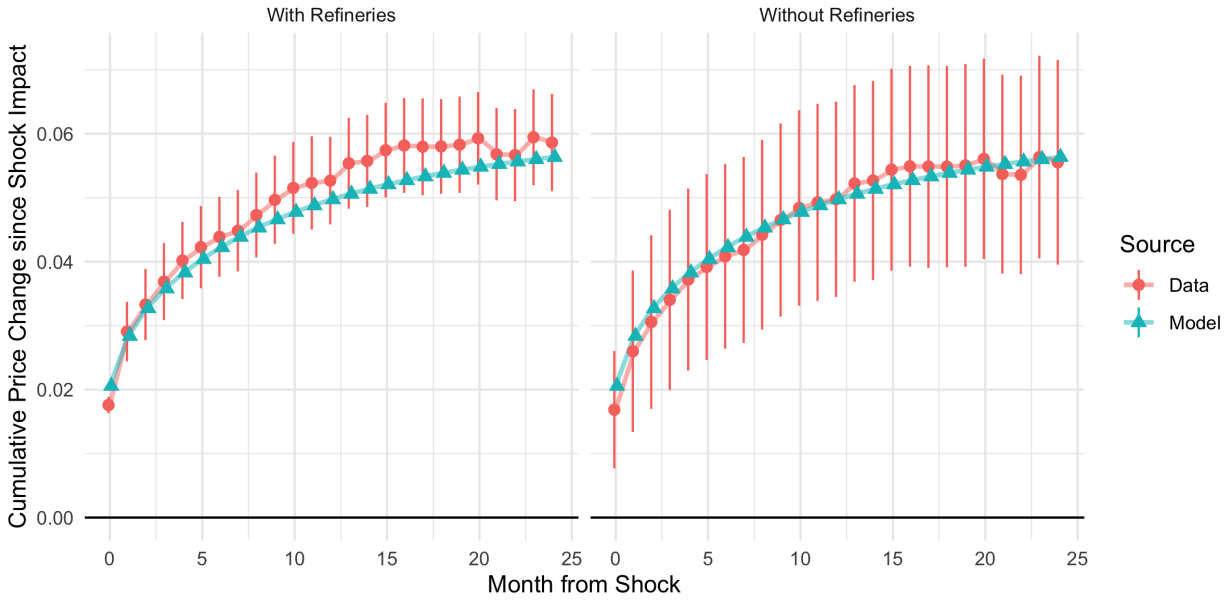
Note: We plot the effect of oil price shocks on the federal funds rate, estimated using a distributed lag model. We find no evidence of statistically distinguishable effects, and the point estimates are negative rather than the expected positive. For our oil price shocks, we use the full Kanzig shock series from 1975 - 2019. The reduced form effects are an order of magnitude larger because the Kanzig shocks are standardized to be associated with a 10% oil price increase, while the IV approach automatically re-scales effects to be associated with a 1% oil price change.

Figure G.24 – Inflationary Effects of Oil, Closed Model



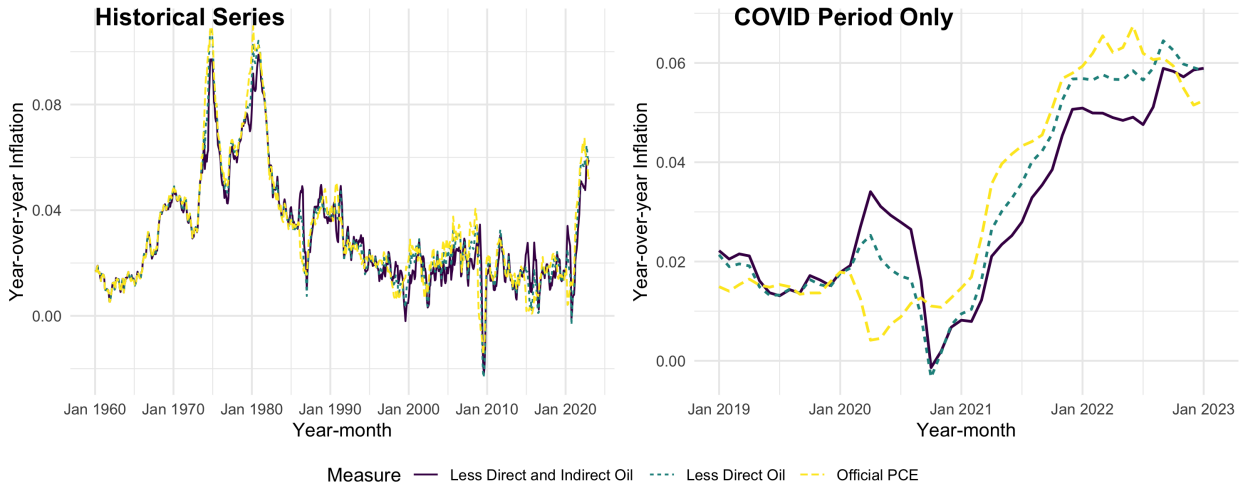
Note: The effect of a 1% oil price increase in our MD-optimal, closed model that includes a household and Taylor rule. The effects are even larger than we saw in Figure 13 because movements in the aggregate wage more than offset the contractionary monetary policy resulting from the Taylor rule.

Figure G.25 – 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.

Figure G.26 – Contribution of Oil to Official Year-over-year PCE Inflation



Note: Official PCE inflation is 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 is reduced to a peak of 5% in early 2022, down from a peak above 6.5%, when removing all the elements of oil inflation from official PCE inflation. But it also rises to nearly 6% in Jan 2023 due to the large oil price decline in mid-late 2022; this result is particularly driven by removing oil's indirect contributions, since removing just direct contributions still reveals declining inflation in late 2022. The results also suggest there was inflation during early 2020 that was hidden by the large decline in oil prices in the early part of the pandemic; this finding is amplified by removing oil's indirect contributions to aggregate inflation.