Oil Prices, Monetary Policy and Inflation Surges

Luca Gagliardone* Mark Gertler†

January, 2024.

(First draft: March, 2023)

Abstract

We develop a simple quantitative New Keynesian model aimed at accounting for the recent sudden and persistent rise in inflation. The model features oil as a complementary good for households and as a complementary input for firms, along with real wage rigidity and an explicit role for labor market tightness. We estimate the key parameters by matching model impulse responses to those from identified oil and monetary policy shocks in a structural VAR. We then show that our model does a good job explaining the recent surge, despite treating inflation as an untargeted variable in the historical decomposition. We find that a combination of oil price shocks and "accommodative" monetary policy alone can account for most of the surge, even after allowing for shocks to aggregate demand and labor market tightness. Essential for the quantitative impact of the oil price shock is a low elasticity of substitution between oil and labor, which we estimate to be the case.

JEL: E32; E52.

---

* New York University. Email: luca.gagliardone@nyu.edu; † New York University and NBER. Email: mark.gertler@nyu.edu. We thank Marco Del Negro, Jordi Gali, Diego Känzig, Jonathon Hazell, and Carolin Pflueger for helpful comments.

1 Introduction

From the mid 1990s to the summer of 2021, inflation remained relatively low and stable. This behavior reinforced conventional wisdom that, so long as long-term inflation expectations remained anchored, high inflation would remain a phenomenon of the past.[1] Seen in this context, the initial spurt in inflation in 2021 was thought by many (including the Federal Reserve!) to be transitory, the product of short-lived factors including a relative shift in spending from services to goods in conjunction with supply-chain problems. However, high inflation has persisted through 2023 despite expectations remaining reasonably anchored and supply-chain disruptions abating.[2] Recently inflation has shown signs of moderating, though it remains uncomfortably above the central bank's two percent target. These events suggest a clear need to revisit the sources of inflation.

In this paper, we develop and estimate a simple New Keynesian model designed to account for the sources of the recent inflation surge. We place particular emphasis on two factors: the post-pandemic jump in oil prices and accommodative monetary policy in the form of a delayed response by the central bank to the mounting inflation. Figure 1 provides motivation for taking this direction. The top panel shows the strong correlation between the jump in oil prices that began in early 2021 and the surge in core PCE inflation. Note also that the correlation remains strong as both inflation and oil prices ease after mid 2022. The bottom panel illustrates the initial monetary policy accommodation: it shows that the central bank kept the Federal funds rate at the zero lower bound until beginning the liftoff in the spring of 2022, roughly a year after the initial inflation surge. The framework we develop will sort out the causal factors underlying the broad correlations in Figure 1.

Our framework also allows for other factors thought to be relevant,

---

[1] For evidence on the relative importance of long-horizon inflation expectations for postwar inflation dynamics, see Hazell et al. (2022).

[2] Alcedo et al. (2022) shows that by April 2022, shortages for most categories of goods had moderated significantly since the peak two years earlier.

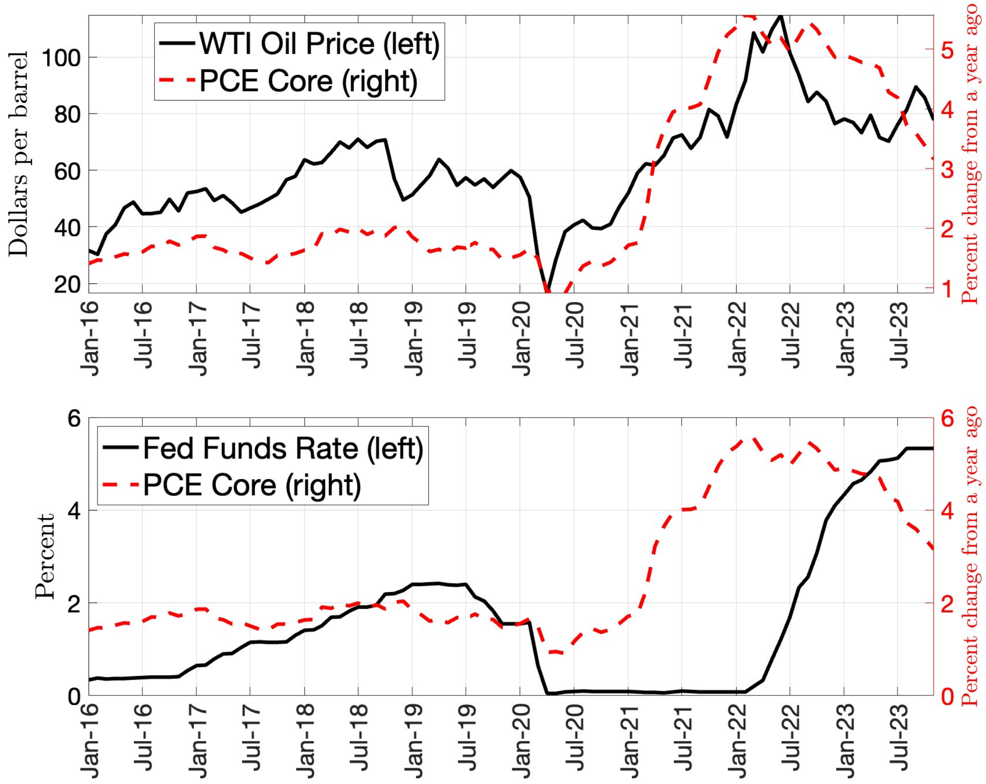

Figure 1: Time series of PCE core inflation against WTI oil prices (top panel) and Fed funds (bottom panel).

including increasing demand and shocks to labor market tightness. We show that even though we do not target inflation in our estimation, the model does a good job of explaining inflation since 2010, including the recent surge.

Section 2 presents the model, a variant of a standard New Keynesian framework with consumption goods only. We follow Blanchard and Gali (2007) by including oil as both a consumption good and an input into production. An important difference is that we allow for oil to be complementary with other consumption goods for households and a complementary input with labor for firms. As we make clear, a low elasticity of substitution between oil and labor is essential to match quantitatively the impact of oil shocks on inflation.[3] We also allow for unemployment via search and matching in the labor market,

---

[3] Bachmann et al. (2022) emphasize how oil being a strong complementary input enhances the impact of an oil shock on output. That is also true in our case though we emphasize the impact on inflation.

which enables us to consider shocks to labor market tightness as a source of inflation. For tractability, there are some factors potentially relevant to the recent inflation that we do not include, such as supply chain disruptions. However, because we do not target inflation in the historical decomposition, we allow for the possibility that these missing factors could account for the discrepancy between the model and data. In the end, though, our parsimonious model does a very good job of accounting for the recent inflation episode.

Section 3 presents the mechanism through which the model can produce inflation surges. As long as long-horizon inflation expectations remain anchored, an inflation surge requires a jump in the expected path of firms' marginal costs. Both oil price shocks and rising demand can trigger this jump. But to have a rise in marginal costs sufficient to account for the inflation surge, it is necessary to have a low elasticity of substitution between oil and labor, as we illustrate in this section.

In section 4 we estimate the key parameters of the model. We do so by matching the model-implied impulse responses to a set of impulse responses from an estimated structural vector autoregression (SVAR). Given our emphasis on the roles of oil shocks and monetary policy, we consider the following two observable exogenous shocks: a high-frequency oil shock, identified as in Känzig (2021), and a high-frequency shock to monetary policy, identified as in Gertler and Karadi (2015) and Bauer and Swanson (2023). Each shock serves as an external instrument in the SVAR. In each case, we estimate parameters without using data from the inflation surge period that we seek to explain. In section 4.5 we show that the model impulse responses match well those from the SVAR for both the oil and money shocks.

In section 5 we explore how well the model accounts for the recent inflation. We quantify the role of each of the four aggregate shocks in our model, namely shocks to oil, monetary policy, demand, and labor market tightness. We first use our estimated model to recover these shocks by targeting the following variables: unemployment, labor market tightness, the Federal funds rate, and oil prices. We leave untargeted the nominal variables, including both headline and core inflation. We then show that the model does a good job explaining inflation, including the recent surge. Accounting for the inflation surge is a combination of oil price shocks and easy money shocks, even after controlling for shocks to demand and labor market tightness.

Finally, our measure of monetary policy accommodation is based on the difference between the funds rate and the central bank's historical policy rule. A possible limitation of this measure is that it does not take into account the policy tightening that occurred in late 2021 due to forward guidance, which led to an increase in longer maturity rates in advance of the liftoff of the funds rate. Accordingly, we redo the exercise using the "proxy" funds rate developed by Choi et al. (2022) that adjusts the funds rate to factor in the role of forward guidance. It is still the case that oil shocks and monetary accommodation remain the most important factors underlying the inflation surge. However, when accounting for forward guidance, the effects of accommodative monetary policy on inflation die out by the end of 2022. In fact, monetary tightening starts to contribute to reducing inflation in 2023. Interestingly, by factoring in forward guidance, the fit of the model for inflation also improves a bit.

Concluding remarks are in section 6.

Related Literature. As suggested earlier, our theoretical framework is related to Blanchard and Gali (2007)'s model of oil shocks. It differs by making oil a complementary good and input, as well as in a number of other details. Also relevant is the literature that estimates New Keynesian DSGE models with oil, including Soto and Medina (2005), Bodenstein and Guerrieri (2011), and Bodenstein et al. (2012), among others. In the standard DSGE estimation approach, the primitive shocks are unobserved and inferred from the residuals of the model equations. Besides model details, we differ by estimating the model using observable shocks to oil and monetary policy.

Both Ball et al. (2022) and Bernanke and Blanchard (2023) also emphasize oil prices: they present estimates to suggest that oil price shocks played an important role in the inflation surge. We differ by (i) developing and estimating a structural model to quantify the importance of different forces and (ii) considering the role of accommodative monetary policy.

On the theoretical side, we build on Lorenzoni and Werning (2023), who emphasize the role of production complementarities and wage rigidity in inflation surges. We differ by presenting a quantification of this mechanism that illustrates its overall importance.

Also relevant is the rapidly growing literature that presents structural models of the recent inflation surge. A number of papers have emphasized the reallocation between goods and services and supply chain problems to explain the rise in inflation in 2021, including Guerrieri et al. (2021), Amiti et al. (2022), Di Giovanni et al. (2022), Comin et al. (2023), Ferrante et al. (2023), and Di Giovanni et al. (2023). Consistent with our findings, the latter also finds an important role for energy shocks. Finally, Benigno and Eggertsson (2023) emphasize non-linearities in the Phillips curve due to asymmetries in wage adjustment. We differ in our focus on oil prices and monetary policy. We also differ in methodology, both with the use of observable shocks in the estimation and by treating inflation as an untargeted variable in the historical decomposition, the net effect being to provide greater discipline on the ability of the model to explain the surge.

2 The Model

The starting point is a standard New Keynesian model with consumption goods only. We add oil which is a complement good for households and a complement input for firms. There are two types of firms. Competitive wholesale firms produce intermediate goods using labor and oil. These firms add workers via a search and matching process. The wholesale firms then sell their output to monopolistically competitive retailers that package the intermediate input into final goods. Retailers also set nominal prices on a staggered basis, which introduces nominal price rigidity as in the standard NK model. We also introduce several features that improve the empirical performance of the model, including habit formation and real wage rigidity.

2.1 Households

There is a representative household with a continuum of members of measure unity. The number n_t of members are currently employed. The household provides perfect consumption insurance for its members. Family members currently not employed look for a job. A search and matching process that we describe shortly determines employment n.

Each period the household consumes a composite c_t that is the following CES aggregate of final consumption goods c_{qt} and oil c_{ot}:

$$c_t = \left(\chi^{\frac{1}{\psi}} c_{ot}^{1-\frac{1}{\psi}} + (1-\chi)^{\frac{1}{\psi}} c_{qt}^{1-\frac{1}{\psi}} \right)^{\frac{1}{1-\frac{1}{\psi}}}, \tag{1}$$

where $\psi > 0$ is the elasticity of substitution between the two goods and χ determines the share of oil in consumption. As we show later, our estimates suggest that $\psi < 1$, implying that the goods are complements. Finally, c_{qt} is a composite of a continuum of differentiated retail consumption goods, but we defer a description of the demand for these differentiated goods until later.

Let β be the subjective discount factor and ε_{bt} a discount factor shock, which serves effectively as a demand shock. The household's objective depends on the utility gain from consumption, as follows:

$$E_t \sum_{i=0}^{\infty} \beta^i \varepsilon_{bt} \ln(c_{t+i} - hc_{t-1+i}), \tag{2}$$

where $h \in (0,1)$ is the degree of habit persistence. As is standard, we allow for habit formation to capture the hump-shaped dynamics in real activity as well as the delayed response to monetary policy that is present in the data.

The household receives wage income from its employed members and unemployment insurance from the unemployed ones. Let w_{ct} denote the real wage and b_t unemployment insurance, both in units of the consumption composite. In addition, the household has the option of saving in the form of a nominal bond B_t that pays the gross nominal rate R_t^n. Let p_{ct} be the nominal price of c_t. The overall budget constraint is then given by:

$$c_t = w_{ct}n_t + b_t(1 - n_t) + R^n_{t-1}\frac{p_{ct-1}}{p_{ct}}B_{t-1} - B_t + \Pi_t, \tag{3}$$

where Π_t are total net payments to the household, which includes dividends from ownership of firms and net lump sum taxes paid to the government. Conditional on n_t, the household chooses c_t, B_t, c_{qt} and c_{ot} to maximize (2) given (3) and (1). Let $u_{ct} = \frac{1}{c_t - hc_{t-1}} - \frac{\beta h}{c_{t+1} - hc_t}$ be the marginal utility of consumption. Then, from the household's consumption/saving decision:

$$E_t\left\{\Lambda_{t,t+1}R^n_t\frac{p_{ct}}{p_{ct+1}}\right\} = 1,$$

where $R^n_t\frac{p_{ct}}{p_{ct+1}}$ is the real return on the nominal bond and $\Lambda_{t,t+1} = \beta\frac{u_{ct+1}}{u_{ct}}$ is the household's stochastic discount factor.

Next, let p_{qt} and p_{ot} be the nominal prices of c_{qt} and c_{ot}, respectively, and $s_t = p_{ot}/p_{ct}$ the relative price of oil. From cost minimization, we obtain demand functions for consumption goods and oil:

$$c_{qt} = (1 - \chi)\left(\frac{p_{qt}}{p_{ct}}\right)^{-\psi}c_t, \qquad c_{ot} = \chi s_t^{-\psi}c_t. \tag{4}$$

Combining with (1) yields a price index for p_{ct}:

$$p_{ct} = \left(\chi p_{ot}^{1-\psi} + (1 - \chi)p_{qt}^{1-\psi}\right)^{\frac{1}{1-\psi}}.$$

2.2 Unemployment, Vacancies, and Matching

As we noted earlier, production and employment take place in the wholesale sector. Following Mortensen and Pissarides (1994), at time t, each wholesale firm i employs $n_t(i)$ workers and posts $v_t(i)$ vacancies to attract new workers. To post each vacancy a firm must pay the fixed cost c. Total employment and vacancies are given by $n_t = \int_0^1 n_t(i)di$ and $v_t = \int_0^1 v_t(i)di$. All unemployed workers at t look for jobs. We assume that those unemployed who find a job go to work immediately within the period. Accordingly, normalizing the total labor force to unity implies that unemployment u_t is given by:

$$u_t = 1 - n_{t-1}.$$

The number of new hires Φ_t is governed by a matching function with constant returns to scale that is increasing in vacancies and unemployment:

$$\Phi_t = \varepsilon_{\Phi t} u_t^{\sigma} v_t^{1-\sigma}, \tag{5}$$

where the random variable $\varepsilon_{\Phi t}$ is a shock to match efficiency. The shock could also reflect shifts in the search effort by the unemployed or recruiting intensity by firms. Note that a decline in $\varepsilon_{\Phi t}$ acts like a negative shock to labor supply, as it implies that more vacancies are needed to create the same amount of matches. This leads to an outward shift in the Beveridge curve (the relation between vacancies and unemployment) and therefore an increase in labor market tightness.

Next, the probability q_t a firm fills a vacancy in period t and the probability a worker finds a job f_t are given by, respectively:

$$q_t = \frac{\Phi_t}{v_t}, \qquad f_t = \frac{\Phi_t}{u_t}. \tag{6}$$

Firms and workers take both q_t and f_t as given.

Finally, in each period an exogenous fraction of workers $1-\rho$ separate from the firm at which they were employed and become unemployed.

2.3 Wholesale Firms

Competitive wholesale firms produce and sell output to retail firms. Wholesale firm i makes output y_t using input of labor n_t and oil o_t according to the following CES production function (where we drop the firm subscript i):

$$y_t = \left(\alpha^{\frac{1}{\epsilon}} n_t^{1-\frac{1}{\epsilon}} + (1-\alpha)^{\frac{1}{\epsilon}} o_t^{1-\frac{1}{\epsilon}}\right)^{\frac{1}{1-\frac{1}{\epsilon}}}, \tag{7}$$

where ϵ is the elasticity of substitution between labor and oil. As we show, our estimates suggest a value of ϵ well below unity, implying that oil and labor are strong complementary inputs.

Employment at t is the sum of surviving workers from the previous period, ρn_{t-1} and new hires, where the latter is the product of the vacancy filling probability and total vacancies, $q_t v_t$. That is, we can write:

$$n_t = \rho n_{t-1} + q_t v_t. \tag{8}$$

The firm can thus adjust employment by posting vacancies, taking q_t as given.[4]

We next turn to the firm's objective. Let p_{wt} be the wholesale firm's relative price, $w_{qt} = w_{ct}(p_{ct}/p_{qt})$ the real product wage, and $s_{qt} = s_t(p_{ct}/p_{qt})$ the relative price of oil, all in units of final good output. The firm's objective then is to maximize the discounted stream of profits, F_t, given by:

$$F_t = p_{wt} y_t - w_{qt} n_t - c v_t - s_{qt} o_t + E_t \left\{ \Lambda^q_{t,t+1} F_{t+1} \right\}, \tag{9}$$

where $\Lambda^q_{t,t+1} = \beta \left(\frac{u_{ct+1}}{u_{ct}} \right) \left(\frac{p_{qt+1}/p_{ct+1}}{p_{qt}/p_{ct}} \right)$ is the household's stochastic discount factor in terms of final good output. Profits each period are the difference between revenues $p_{wt} y_t$ and the sum of the wage bill $w_{qt} n_t$, vacancy posting costs $c v_t$, and oil costs $s_{qt} o_t$. The optimization problem is then the following: firms choose vacancies v_t, employment n_t, and oil o_t to maximize (9) subject to (7) and (8).

Let a_{nt} be the marginal product of labor. The first-order conditions for v_t and n_t along with the envelope condition yield the following standard first-order condition for hiring:

$$\begin{aligned} \frac{c}{q_t} =& \sum_{i=0}^{\infty} \rho^i E_t \left\{ \Lambda^q_{t,t+i} (p_{wt+i} a_{nt+i} - w_{qt+i}) \right\} \\ =& p_{wt} a_{nt} - w_{qt} + \rho E_t \left\{ \Lambda^q_{t,t+i} \frac{c}{q_{t+1}} \right\}, \end{aligned} \tag{10}$$

where the marginal product of labor is given by:

$$a_{nt} = \left(\alpha \frac{y_t}{n_t} \right)^{\frac{1}{\epsilon}}. \tag{11}$$

Let a_{ot} be the marginal product of oil. The firm's demand for oil is given by the condition that the marginal value of oil equals the marginal cost:

$$p_{wt} a_{ot} = s_{qt}, \tag{12}$$

---

[4] We assume the law of large numbers applies so that $q_t v_t$ is the number of new hires.

where the marginal product of oil is:

$$a_{ot} = \left((1-\alpha)\frac{y_t}{o_t}\right)^{\frac{1}{\epsilon}}.$$

So far we have described the firm's hiring decision conditional on the path of wages. Before describing how wages are determined, it is useful to characterize the value J_t of a worker to the firm, after hiring costs have been paid. From differentiating equation (9) with respect to n_t and applying the envelope theorem, we obtain:[5]

$$\begin{aligned} J_t =& \sum_{i=0}^{\infty} \rho^i E_t \left\{\Lambda^q_{t,t+i}(p_{wt}a_{t+i} - w_{qt+i})\right\} \\ =& p_{wt}a_t - w_{qt} + \rho E_t \left\{\Lambda^q_{t,t+1} J_{t+1}\right\}. \end{aligned} \tag{13}$$

2.4 Workers

We next develop an expression for the worker's surplus from a job. Recall that $w_{ct} = w_{qt}(p_{qt}/p_{ct})$ is the real wage in units of the consumption composite. Let V_t be the value to a worker of employment at t and U_t the value of being unemployed. Then V_t and U_t are:

$$V_t = w_{ct} + E_t \left\{\Lambda_{t,t+1}\left(\rho V_{t+1} + (1-\rho)U_{t+1}\right)\right\},$$

$$U_t = b_t + E_t \left\{\Lambda_{t,t+1}\left(f_{t+1}V_{t+1} + (1-f_{t+1})U_{t+1}\right)\right\},$$

where w_{ct} and $b_t = b(p_{qt}/p_{ct})$ are the flow values of work and unemployment respectively, ρ is the job survival probability, and f_{t+1} is the probability of moving from unemployment in t to employment in $t+1$.

The job surplus H_t is then given by:

$$H_t = V_t - U_t = w_{ct} - b_t + E_t \left\{\Lambda_{t,t+1}\left((\rho - f_{t+1})H_{t+1}\right)\right\}. \tag{14}$$

---

[5] Because production is constant returns and there is a continuum of workers, the value of the marginal worker is the same as the value of the average worker.

2.5 Wage Determination

In the conventional Mortensen-Pissarides (MP) framework, wages are determined by period-by-period Nash bargaining. Absent any stickiness in wage determination, however, it is difficult to explain the large effects of oil price shocks, as wages could freely adjust to dampen the impact on the economy. Here we introduce a simple form of real wage rigidity within the MP framework: we assume that the wage depends on the gap between the value that would arise under Nash bargaining and its steady-state value. The degree of stickiness is parsimoniously characterized by a single parameter that we estimate.

2.5.1 Nash Bargaining Wage

Let us start by characterizing the product wage under Nash bargaining. In this hypothetical case the firm and its workers choose w_{qt} to maximize the joint surplus from the match, as follows:

$$\max_{w_{qt}} H_t^{\varsigma} J_t^{1-\varsigma},$$

where $\varsigma \in [0,1]$ is the relative bargaining power of workers and H_t and J_t are as in equations (13) and (14). The solution to the maximization problem then leads to the product wage that would arise under Nash Bargaining:

$$w_{qt}^{o} = \frac{\varsigma\left(p_{wt}a_{nt} + \rho E_t\left\{\frac{c}{q_{t+1}}\left(\Lambda_{t,t+1}^{q} - \Lambda_{t,t+1}\right)\right\} + E_t\left\{\Lambda_{t,t+1}c\theta_{t+1}\right\}\right) + (1-\varsigma)\frac{p_{qt}}{p_{ct}}b}{\varsigma + (1-\varsigma)\frac{p_{qt}}{p_{ct}}}.$$

As is standard, the Nash wage is a convex combination of the period surplus the worker brings to the match and the worker's outside option, where the weights depend on relative bargaining power. The term $\rho E_t\left\{\frac{c}{q_{t+1}}\left(\Lambda_{t,t+1}^{q} - \Lambda_{t,t+1}\right)\right\}$ reflects differences between the parties in the evaluation of the value of a worker due to differences in the stochastic discount factors. The presence of the relative price of goods, p_{qt}/p_{ct}, reflects how workers value the nominal wage and unemployment insurance payments differently from firms.

In what follows, we assume that the bargaining weight ς and $1-\varsigma$ equal the corresponding weights σ and $1-\sigma$ in the matching function, implying the Hosios condition holds: the equilibrium with wages determined by Nash bargaining is thus constrained efficient, in the sense that the social value of the marginal hire equals the marginal recruiting cost.

2.5.2 Real Wage Rigidity

Though the details differ, we follow Blanchard and Gali (2007) in introducing real wage rigidity. We suppose that the percent adjustment of the real wage relative to steady state is the fraction $1-\gamma$ of the percent fluctuation in the Nash wage w^o_{qt}, where $\gamma \in [0,1]$ reflects the degree of real wage rigidity and is a parameter we will estimate. In particular,

$$w_{qt} = (w^o_{qt})^{1-\gamma}(w^o_q)^{\gamma}, \tag{15}$$

where w^o_q is the steady state Nash wage. Under reasonable parametrizations, equation (15) is consistent with rational behavior: because the implied wage lies within the bargaining set, i.e. it is never above firm's reservation wage nor is it ever below worker's reservation wage. Equation (15) can be interpreted as the firm providing some insurance to workers by offering a smoother real wage than would be the case under period-by-period Nash bargaining. However, we do not motivate this argument from first principles.

2.6 Retail Firms and Core Inflation

There is a continuum of monopolistically competitive retail firms indexed by $j \in [0,1]$. Retailers buy intermediate goods from the wholesale firms described earlier. Retailers then transform intermediate goods into a differentiated final good. Households buy and consume these differentiated products. Finally, retail firms set prices on a staggered basis à la Calvo: we denote with $1-\lambda$ the probability the firm is able to change price in the current period, where the draw is i.i.d. across time and firms.

The consumption good composite for each household, c_{qt}, is given by a CES aggregate of each retail firm's output y_{jt}. From cost minimization, we obtain the household's demand for each retail good as an inverse function of the relative price, p_{jt}/p_{qt},

$$y_{jt} = \left(\frac{p_{jt}}{p_{qt}}\right)^{-\eta} c_{qt}, \tag{16}$$

where η is the elasticity of substitution across intermediate goods.

In each period, the fraction λ of retail firms that are unable to adjust price simply meet demand for their differentiated final good. They do so by buying enough input from wholesalers as long as the relative output price, $\frac{p_{jt}}{p_{qt}}$, is not less than the cost of inputs, p_{wt}.

On the other hand, retail firms that are able to adjust their price within the period choose the reset price p^*_{jt} and output y_{jt} to maximize expected discounted profits, subject to the demand curve (16):

$$\max_{p^*_{jt}, y_{jt}} E_t \left\{ \sum_{i=0} \lambda^i \Lambda^q_{t,t+i} \left(\frac{p^*_{jt}}{p_{qt}} - p_{wt} \right) y_{jt+i} \right\},$$

where the probability λ^i that the firm's price remains fixed i periods into the future. Note that the relative wholesale price p_{wt} corresponds to the marginal cost of production. The first-order condition for the retailer's reset price is:

$$E_t \left\{ \sum_{i=0} \lambda^i \Lambda^q_{t,t+i} \left(\frac{p^*_{jt}}{p_{qt+i}} - (1+\mu) p_{wt+i} \right) y_{jt+i} \right\} = 0, \tag{17}$$

where $\mu = 1/(1 - 1/\eta)$ is desired net markup. When able to adjust, a firm chooses to reset the price p^*_{jt} so that, over the period in which its price is expected to remain fixed, its relative price equals a discounted average of the desired gross gross markup $(1+\mu)$ over its real marginal cost p_{wt+i}.

Finally, from cost minimization by the retailer and from using the law of large numbers, we can express the price index as:

$$p_{qt} = \left((1-\lambda)(p^*_t)^{1-\eta} + \lambda p^{1-\eta}_{t-1} \right)^{\frac{1}{1-\eta}}. \tag{18}$$

Equations (17) and (18) govern the path of goods inflation conditional on p_{wt}.

2.7 The Oil Market and Resource Constraints

We suppose that there is a representative oil producer who acts competitively. Each period the producer receives an endowment of oil equal to $S\exp(-\varepsilon_{ot})$, where ε_{ot} is a shock to the oil supply and S is the steady-state oil supply. The producer takes the price of oil as given. All profits are paid out as dividends to households. Each period the sum of the firm demand for oil o_t and the household demand c_{ot} must equal the total supply, as follows:

$$o_t + c_{ot} = S\exp(-\varepsilon_{ot}),$$

where the respective firm and household oil demand functions are given by equations (4) and (12). The relative price of oil s_t adjusts to clear the market.[6] For produced goods, the relevant resource constraint is given by the condition that consumption goods c_{qt} must equal output y_{qt} net hiring costs cv_t:

$$c_{qt} = y_{qt} - cv_t.$$

Finally, the supply of nominal bonds is zero, $B_t = 0$.

2.8 Government Policy

We suppose that the central bank adjusts the nominal interest rate according to a simple Taylor rule augmented by a persistent exogenous money shock ε_{rt}. Let ϕ_π be the feedback coefficient on inflation, ρ^R be the interest rate smoothing parameter, and $\pi_{qt} = \ln(p_{qt}/p_{qt-1})$ net core inflation. The rule is then given by:

$$R_t^n = (R^n(1+\pi_{qt})^{\phi_\pi})^{(1-\rho^R)}(R_{t-1}^n)^{\rho^R}e^{\varepsilon_{rt}}$$

We assume the central bank responds to core inflation (inflation absent oil prices) so as to avoid temporary gyrations associated with headline inflation.

---

[6] In practice, oil prices depend on both oil production and the existing stock of inventories. In particular, inventories can be used strategically to manipulate the price of oil in the short run. Despite abstracting from modeling inventories explicitly, we account for speculative behavior resulting in temporary fluctuations of oil prices in the empirical analysis, as discussed in detail in section 5.

The rule allows for persistence in the interest rate due to both interest smoothing, which is governed by ρ^R, and the persistence of the monetary shock. In the empirical exercise in Section 4, we find that it is the latter that mainly accounts for the persistence of the interest rate.[7]

The only fiscal expenditures are unemployment insurance payments. We suppose payments are financed by lump-sum taxes on households: $b_t u_t = \tau_t$.

3 Sources of Inflation Surges

We now characterize the features of our model that can help account for inflation surges. As discussed in section 2.6, inflation depends on the real marginal cost of final goods firms. In our model, the marginal cost corresponds to the relative price of wholesale goods p_{wt}. Let $\pi_{qt} = \ln(p_{qt}/p_{qt-1})$ be goods market inflation and $\widehat{p}_{wt} = \ln(p_{wt}/p_w)$ be the log deviation of the relative wholesale price from its steady state. Loglinearizing (17) around the zero-inflation steady state and using equation (18), then yields the following Phillips curve relation for π_{qt}:

$$\pi_{qt} = \kappa \widehat{p}_{wt} + E_t\{\pi_{qt+1}\},$$

where $\kappa = (1-\lambda)(1-\lambda\beta)/\lambda$ is the slope of the Phillips curve. As in the standard NK formulation, inflation depends on real marginal cost, which in this case is $\widehat{p}_{wt}$. Iterating forward implies that inflation depends on an expected discounted stream of present and future marginal costs, as follows:

$$\pi_{qt} = \kappa \sum_{i=0}^{\infty} E_t \left\{ \widehat{p}_{wt+i} \right\}.$$

In the model, a large inflation surge originates from a significant and persistent increase in the expected path of real marginal cost.

We next decompose the movement in p_{wt} into three terms: the real wage, the marginal hiring cost, and the marginal product of labor. As we show, all

---

[7] In fact, the estimate of ρ^R is not statistically different from zero. However, results are robust to calibrating it to conventional values.

three factors could play a significant role in an inflation surge. However, given the strong complementarities between labor and oil, the marginal product of labor plays a particularly important role.

From the hiring condition (10) we can derive a relation for the marginal hiring cost. Let ω_t be the net marginal cost of hiring a worker.[8] From equation (10), we can express ω_t as:

$$\omega_t = \frac{c}{q_t} - \rho E_t \left\{ \Lambda^q_{t,t+1} \frac{c}{q_{t+1}} \right\}, \tag{19}$$

which is the gross cost of adding a worker at t, c/q_t, net the expected discounted benefit that the additional worker at t will generate in the future $\rho E_t \left\{ \Lambda^q_{t,t+1} c/q_{t+1} \right\}$.[9] We can then express the marginal cost of producing a unit of output as the sum of the wage w_{qt} and net hiring costs ω_t, normalized by the marginal product of labor a_{nt}, as follows:

$$p_{wt} = \frac{w_{qt} + \omega_t}{a_{nt}}. \tag{20}$$

From loglinearizing equation (20), we can decompose marginal cost $\widehat{p}_{wt}$ into a convex combination of the real product wage $\widehat{w}_{qt}$ and net hiring costs $\widehat{\omega}_t$ minus the marginal product of labor $\widehat{a}_{nt}$, all expressed in log deviations:

$$\widehat{p}_{wt} = \zeta \widehat{w}_{qt} + (1-\zeta)\widehat{\omega}_t - \widehat{a}_{nt}, \tag{21}$$

where $\zeta = \frac{w_q}{w_q+\omega}$ is the relative weight on the real product wage.

Equation (21) highlights how inflation surges are generated in our model. First, the presence of complementarities enhances the sensitivity of the marginal product of labor (and hence marginal cost) to fluctuations in oil intensity, measured by the ratio of oil to labor input, o_t/n_t. After combining equations (7) and (11), we can obtain the following loglinear approximation for the marginal product of labor:

---

[8] Ferrante et al. (2023) also develop a role for hiring costs in affecting marginal cost in their model of sectoral reallocation with quadratic labor adjustment costs.

[9] From the hiring condition, we can infer that c/q_{t+1} equals the present value of earnings at $t+1$ and beyond generated by a worker who is with the firm at t. From the vantage of time t we take expectations and discount this value by the job survival probability ρ and the household stochastic discount factor $\Lambda^q_{t,t+1}$.

$$\widehat{a}_{nt} = \frac{1}{\epsilon}(1-\overline{\alpha})(\widehat{o}_t - \widehat{n}_t),$$

with:

$$\overline{\alpha} = \frac{\alpha}{\alpha + \alpha^{1-\frac{1}{\epsilon}}(1-\alpha)^{\frac{1}{\epsilon}}(\frac{o}{n})^{1-\frac{1}{\epsilon}}} \approx \alpha.$$

Note first that under our calibration $\overline{\alpha} \approx \alpha$ since $o/n \approx (1-\alpha)/\alpha$. The equation then makes clear how, as the elasticity of substitution ϵ declines, the sensitivity of $\widehat{a}_{nt}$ to $\widehat{o}_t - \widehat{n}_t$ increases. With sufficiently strong complementarities, even a small oil shock that reduces oil intensity can produce a sharp decline in $\widehat{a}_{nt}$ contributing to a surge in inflation via its impact on costs $\widehat{p}_{wt}$. Similarly, given that the oil supply is fixed in the short run, a positive demand shock that reduces $\widehat{o}_t - \widehat{n}_t$ by increasing labor demand also generates inflationary pressures that are stronger when the elasticity of substitution ϵ is small.

Second, wage rigidity also matters. With flexible wages, in response to an increase in oil prices, wages may drop significantly, moderating the impact of the oil shock on marginal cost. Therefore, wage rigidity dampens this offsetting adjustment and thus amplifies the transmission of supply shocks to inflation.

Third, labor market tightness $\theta_t = v_t/u_t$ affects the marginal cost. Note first that from equations (5) and (6) the vacancy filling probability q_t varies inversely with tightness, i.e. $q_t = \varepsilon_{\Phi t}\theta_t^{-\sigma}$. Replacing this in equation (19), net hiring costs become a linear function of current and expected market tightness:

$$\widehat{\omega}_t = \frac{1}{1-\rho\beta} E_t \left\{ \sigma\widehat{\theta}_t - \rho\beta\sigma\widehat{\theta}_{t+1} - \rho\beta\widehat{\Lambda}^q_{t,t+1} + \ln\varepsilon_{\Phi t} - \rho\beta\ln\varepsilon_{\Phi t+1} \right\}. \tag{22}$$

Equations (21) and (22) illustrate how, via net hiring costs $\widehat{\omega}_t$, market tightness $\widehat{\theta}_t$ affects marginal cost $\widehat{p}_{wt}$. In addition, from the equation for the Nash wage, the real wage is increasing in expected labor market tightness as the latter increases the value of unemployment. Both forces imply that a tightening of labor market conditions raises marginal cost, which thus applies upward pressure on prices.

In section 5 we show that underlying the inflation surge was indeed a surge in real marginal cost. Playing a key role, further, was the strong complementarity between labor and oil.

4 Model Estimation

We estimate the key parameters of the model by matching the model impulse responses to a set of impulse responses generated from an estimated structural vector autoregression (SVAR). We consider two types of observable shocks that serve as external instruments in our SVAR: a high-frequency oil shock, identified as in Känzig (2021); and a high-frequency shock to monetary policy, obtained as in Gertler and Karadi (2015) and Bauer and Swanson (2023).

4.1 Data

Our SVAR is monthly, estimated over the period 1973:01 to 2019:12. We use the post 2019 data for model validation in Section 5. We include seven reasonably standard macroeconomic variables: log real gross domestic output, unemployment in levels, log real oil prices, the Federal funds rate, log headline PCE, log real wages, and the Gilchrist and Zakrajšek (2012) excess bond premium. The latter we include in the SVAR to improve the precision of the impulse responses but do not target it.

Monthly real GDP is log cumulated real GDP growth constructed by Brave-Butters-Kelley. The real oil price is the log spot West Texas Intermediate crude oil price deflated by core PCE. The real wage is measured as log average hourly earnings by production and nonsupervisory employees deflated by core PCE. Unemployment is the number of unemployed as a percentage of the labor force (16 years or older).

4.2 Identification of the Effects of Shocks

We begin by estimating the reduced form of our monthly seven-variable VAR, using twelve lags of each variable. As is standard, we can represent the seven reduced-form residuals as linear combinations of seven structural shocks. Our goal is to identify how the structural shocks to oil and monetary policy affect the contemporaneous reduced form residuals. Once we have estimated these effects, we can then use the VAR to trace out the dynamic effects.

To identify exogenous variation for the oil and money shocks, we use as external instruments the surprises in futures market prices constructed around OPEC and FOMC announcements, respectively. Let s_t^i be the surprise in the log price of a futures contract for variable i at the announcement date t. The key assumption is that the news revealed within the window that leads to the surprise in the futures price can be treated as exogenous with respect to the other variables in the VAR. Let $\mathbb{E}_t(P_{t+h}^i)$ be the log expected spot price conditional on the information available after the announcement and $\mathbb{E}_{t-w}(P_{t+h}^i)$ be the log forecast of the same variable just prior to the window opening. Then assuming that the risk premium does not change within the window around the announcement, the surprise simplifies to:

$$s_t^i = \mathbb{E}_t(P_{t+h}^i) - \mathbb{E}_{t-w}(P_{t+h}^i)$$

Each surprise $s_t^i \in \{s_t^o, s_t^r\}$ is used as an instrumental variable to identify the impact of the respective structural shock on the set of contemporaneous reduced form residuals. We normalize the impact of the money shock on the Fed funds and the impact of the oil shock on the real oil price to be one standard deviation.[10]

To construct oil price surprises we follow Känzig (2021) exactly.[11] We consider the surprise in the futures price for oil on the day on which the Organization of the Petroleum Exporting Countries (OPEC) has a meeting. The relevant time window over which the surprise takes place is between the day of the announcement and the last trading day before the OPEC meeting.[12]

For monetary policy surprises, we start with Gertler and Karadi (2015) by using unexpected movements in interest rate futures around the Federal

---

[10] See footnote 4 in Gertler and Karadi (2015) for the details.

[11] For classic approaches to identifying oil shocks, see Hamilton (1983) and Kilian (2009).

[12] Unfortunately, intraday oil futures are not available until the latter part of the sample. As discussed by Känzig (2021), markets react to OPEC announcements slower compared to FOMC announcements, and this gives further justification for using a daily window rather than a tighter one.

Open Market Committee (FOMC) dates. We then follow Bauer and Swanson (2023) by also measuring surprises around non-FOMC dates where the Federal Reserve revealed information.[13] To measure the futures market surprise we use the unexpected movement in the first principal component of the first four quarterly Eurodollar future contracts. Given data availability, we are able to use a very tight window of thirty minutes: the money shock surprise is thus the log difference between the realized value twenty minutes after the announcement and the forecast ten minutes prior to the meeting. To identify contemporaneous effects of interest rate surprises, we begin in 1988:01 given that interest rates futures data are not available until then. Note that we still use the whole sample to estimate the reduced form coefficients in the VAR.

One challenge we need to address is that oil prices have predictability for interest rate surprises: an increase in the growth of oil prices prior to the FOMC meeting predicts an increase in the surprise, which appears to violate our maintained hypothesis that the surprises are exogenous. A likely explanation involves endogeneity: monetary policy tends to ease when oil prices fall and vice versa when they rise.[14] Accordingly, we purge our measure of the monetary surprise from the information contained in oil prices, as follows: we run the regression of money surprises on the log change in oil spot prices calculated between the day before the meeting and the previous month Δp_{ot}:

$$s_t^r = \underset{(.038)}{+.073} \cdot \Delta p_{ot} + \xi_t$$

We find that monetary policy surprises can be predicted by oil prices.[15]

---

[13] We also do not include the measured money shock during the month of the Lehman Brothers collapse. Because the markets were expecting a larger easing, our measure shows an unanticipated tightening. At the same time, there was a huge drop in GDP and industrial production due to the financial collapse. Because factors beyond monetary policy were relevant to the drop in real activity, we thought it was prudent to drop this observation. Including would slightly reduce the impact of a surprise tightening on real GDP.

[14] As discussed in Bauer and Swanson (2021), one might argue that the effects of oil prices prior to FOMC dates on interest rates should be captured in futures markets. A reason why this might not be the case is uncertainty regarding the central bank's reaction function, leading financial markets to underestimate feedback effects from oil prices.

[15] This is consistent with findings from Bauer and Swanson (2023), that orthogonalizes

We then use the residuals of this regression, $\hat{\xi}_t$, as the monetary policy surprises, giving us an instrument that is orthogonal to oil prices. We note that without this adjustment, our SVAR would predict that a surprise monetary tightening would increase oil prices, an outcome that is clearly the product of not properly controlling for the endogeneity of monetary policy.

4.3 Impulse Responses to Money and Oil Shocks

Figure 2 reports the impulse responses for the identified money and oil shocks along with ninety-five percent confidence bands.[16]

The IRFs for the money shock are similar to previous estimates obtained in the literature: A monetary policy tightening of 15 basis points implies a decline in GDP of about 10 basis points after ten months along with a decline in the price level of about 10 basis points. Associated with the decline in output is a rise in unemployment of roughly half a percentage point. Real wages also decline slightly, though the estimate is not statistically different from zero. After forty to fifty months all the real variables have reverted to their initial values. The real oil price declines moderately but is not statistically different from zero, in line with previous evidence (e.g. Soriano and Torró 2022) as well as high-frequency evidence (e.g. Rosa 2014).

The IRFs for the oil shock behave similarly to those in Känzig (2021), though with some differences due to the variables in the VAR not being identical. The oil shock has a stagflationary effect: a shock that generates a 6 percent increase in the real price of oil reduces GDP by roughly 20 to 30 basis points and increases the price level by about 20 basis points.[17] Interestingly, we find that the Fed funds rate increases about 20 basis points on impact and persists above zero for twenty months, suggesting that the central bank reacts

---

the money shock with respect to additional observables. When using their measure of money surprises, we still find a positive and significant impact of money shocks on oil prices.

[16] Confidence bands are computed using the wild bootstrap.

[17] When we replace headline PCE with core in the VAR, the oil shock also leads to a (statistically significant) increase in the core PCE price level of roughly 10 basis points after 10 months. As one would expect, the rise is delayed compared to the increase in headline.

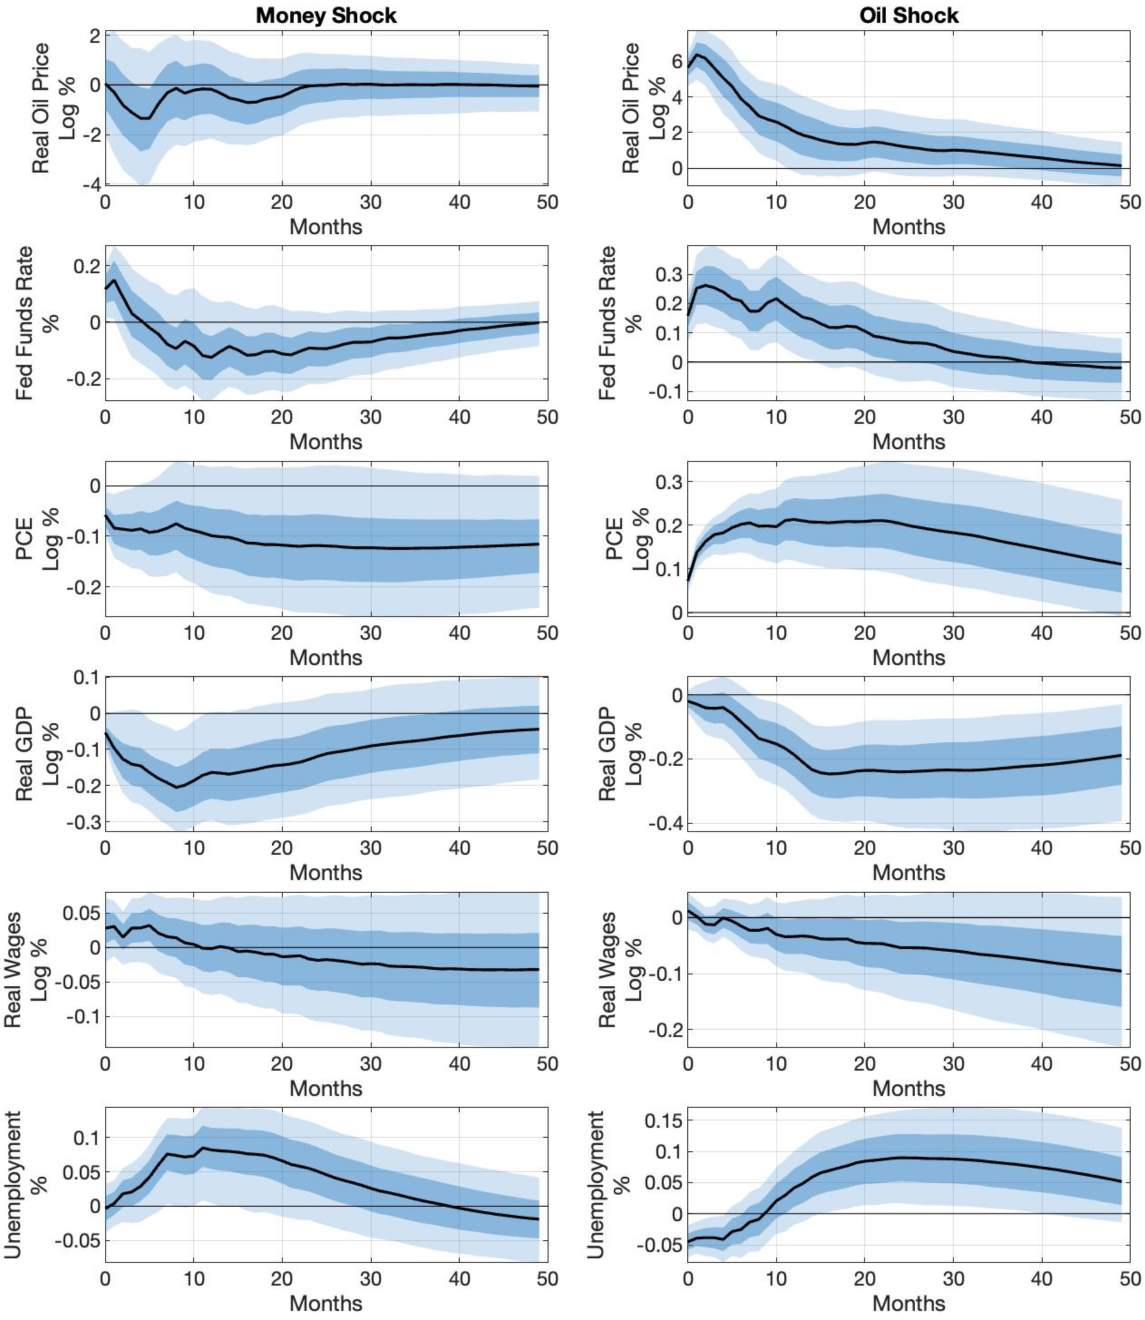

Figure 2: SVAR-based impulse responses for identified money and oil shocks. The solid line is the point estimate and the dark and light-shaded areas are 68 and 95 percent confidence bands, respectively, computed using the wild bootstrap.

to the increase in inflation with a monetary policy tightening. Real wages decline persistently by about 5 to 10 basis points, mainly due to nominal wages increasing by less than core inflation.

4.4 Parameter Estimation

We first calibrate a set Θ_1 of parameters and then estimate the remaining parameters in the set Θ_2 conditional on the calibrated parameters. Parameters are estimated using the simulated method of moments to match the model impulse response functions with those from the SVAR with identified money and oil shocks, as portrayed in Figure 2. Impulse responses are weighted using the estimated precision. Confidence intervals for the parameters are derived using the delta method. We describe the details of the estimation procedure in Online Appendix A.

4.4.1 Calibrated Parameters

We begin with the parameters in Θ_1 which we calibrate to standard values. We start with conventional parameters. We choose the discount factor β to generate a steady-state annual real interest of two percent. We pick the elasticity of substitution between the differentiated consumption goods η to generate a steady-state gross markup of 1.3.

We next turn to the labor market parameters. We set the job survival rate ρ to a monthly value of 0.96, implying an average employment duration of two and a half years, consistent with the evidence. As noted earlier, we also choose worker's bargaining power ς and the match elasticity σ to each equal 0.5, so that the Hosios condition is satisfied, implying that when wages are perfectly flexible and there is Nash bargaining, job creation is efficient. Next, we choose the worker's flow outside option b so that the ratio to the steady-state contribution of the worker to the match is 0.72, consistent with Hall (2009) and implying a value of b of 0.7. Finally, we set the steady-state unemployment rate equal to the sample mean of 5 percent. The results are robust to calibrating the steady state to a range of values between 3 and 6 percent. We can then use the steady-state level of unemployment to pin down the cost of posting a vacancy c at 0.09.

Finally, we turn to the oil sector. Using data on energy expenditures from the U.S. Information Energy Administration, we set the steady-state ratio of oil used in production to output o/y to 3 percent, and the steady-state ratio of firm to household expenditures on oil o/c_o to 1.5. The steady-state ratio of oil to output pins down the share of labor in production α at 0.97.[18] The steady-state ratio of firm to household expenditures on oil pins down the share of oil in households' expenditures χ at 2 percent.[19]

4.4.2 Estimated Parameters

Conditional on the calibrated parameters, we then estimate ten parameters that govern: complementarities with oil in production and consumption (ϵ and ψ),; wage rigidity (γ), the feedback coefficient on inflation in the Taylor rule (ϕ_π), habit persistence (h), price rigidity (λ) and the persistences and volatilities of the money and oil shocks (ρ^r, ρ^o, σ^r, σ^o).

Table 1 presents the results. The estimates of $\epsilon = 0.39$ and $\psi = 0.02$ imply strong complementarities with oil in both production and consumption. What gives the high degree of complementarity in production is the simultaneous drop in output and increase in unemployment in response

---

[18] We calculate the 3 percent share of oil in production as follows: first, as in Bodenstein et al. (2012), we include natural gas along with petroleum in the measure of the oil. According to the US Energy Information Administration, petroleum, and natural gas expenditures average 4.5% as a share of domestic GDP over the period 2010-2020. Finally, oil inputs in production account for about 2/3 of total oil usage, giving an estimate of the production share of 3.1% (see the next footnote).

[19] In 2021, according to the U.S. Information Energy Administration, 67.2% of petroleum consumption is accounted for by transportation, 26.9% by industrial use, 2.8% by residential, 2.5% by commercial, and 0.5% by electricity production. Transportation includes usage that can be partially attributed to the household sector and partially to the production sector. In particular, 63% of it is motor gasoline (including transportation for commercial purposes), 23% is distillate fuel oil and 10% is jet fuel and aviation gasoline. Splitting transportation usage in half between households and firms gives a division of total oil usage in 2/3 for production and 1/3 for final consumption.

Parameter	Θ_1	Value	Parameter	Θ_2	Estimate	S.E.
Discount factor	β	.998	F complementarity	ϵ	.374	.160
Demand curvature	η	4	H complementarity	ψ	.020	.337
Job survival	ρ	.96	Habit persistence†	h	.914 (.75)	.036
Matching elasticity	σ	.5	Price stickiness†	λ	.945 (.83)	.011
Bargaining power	ς	.5	Wage stickiness	γ	.697	.145
Outside option	b	.7	Taylor coefficient	ϕ_π	2.29	0.70
SS oil/consumption	o/c_o	1.5	Money persistence	ρ^r	.952	.011
SS oil/output	o/y	.03	Oil persistence	ρ^o	.967	.013
SS unemployment	u	.05	Normalization	σ^r	.019	.006
			Normalization	σ^o	.060	.025

Table 1: Values for the monthly calibration of the model parameters and steady-state targets. The first three columns report the calibrated parameters in Θ_1, the last four columns report the estimated parameters in Θ_2 with their point estimates and standard errors. † Quarterly calibration in parenthesis.

to the oil shock. With oil substitutability instead, the impact of the oil price shock on output would be muted as firms switch to labor and/or households substitute towards consumption goods, making it impossible for the model to replicate the empirical impulse responses. The point estimates and standard errors in Table 1 suggest we can soundly reject the null hypothesis of no complementarities, both for firms and households.

The habit formation parameter h, which has a quarterly value of 0.75 (monthly 0.91) is within the range of estimates using macro data (see Havranek et al. 2017). The estimate of the monthly degree of price rigidity λ implies a value for the monthly slope of the Phillips curve ($\kappa = 0.003$), consistent with the estimates provided by Gagliardone et al. (2023).[20] The estimate of the rigidity parameter γ of 0.7 implies that a unitary change in the Nash wage from the steady state leads to a change in the wage of roughly one third.

The feedback coefficient on the Taylor rule ϕ_π is estimated to be 2.3,

---

[20] While our estimate of the slope of the Phillips curve is similar, we obtain a higher estimate of the degree of price rigidity. The reason for the difference is that Gagliardone et al. (2023) includes strategic complementarities in price setting, which is here equivalent to an increase in price stickiness.

which is in the range of estimates of the literature (see e.g. Carvalho et al. 2021).[21] We estimate a high degree of persistence of the exogenous monetary shock ($\rho^r = 0.95$). On the other hand, we find the interest rate smoothing parameter ρ^R to not be statistically different from zero (0.27, standard error 0.19). Therefore, we set ρ^R to zero and re-estimate the model to improve the precision of the remaining parameters.[22]

Finally, the persistence parameter for oil ($\rho^o = 0.96$) is precisely estimated to match the persistence of oil prices in the impulse responses.

4.5 Results: Model versus Data

Figure 3 portrays the impulse response functions from the model versus those generated by the data. The left column portrays the effect of the money shock while the right side does the same for the oil shock. In each case, the black line is the data along with ninety-five percent confidence intervals, while the red line is the model. Overall the fit is good: the model always stays within the confidence intervals. While the model response of output to each shock is below the point estimate from the data, the response of unemployment is on target, as are the responses of the other variables (to a reasonable degree).

5 Accounting for Inflation

We now explore the extent to which the model can account for the recent inflation surge. To do so, we use the estimated model to perform a historical shock decomposition. Specifically, we identify the contribution of each of the four aggregate shocks in our model, namely the demand shock ε_{bt}, the monetary policy shock ε_{rt}, the oil shock ε_{st} and the shock to match efficiency $\varepsilon_{\Phi t}$. As discussed earlier, these shocks are among the popular explanations for

---

[21] Our results are robust to calibrating values of ϕ_π as low as 1.2. Notice that lowering the coefficient would increase the contribution of money shocks in the inflation surge.

[22] To be clear, our results in Section 5 are robust to calibrating ρ^R to other values used in the literature (e.g. $0.6 - 0.8$).

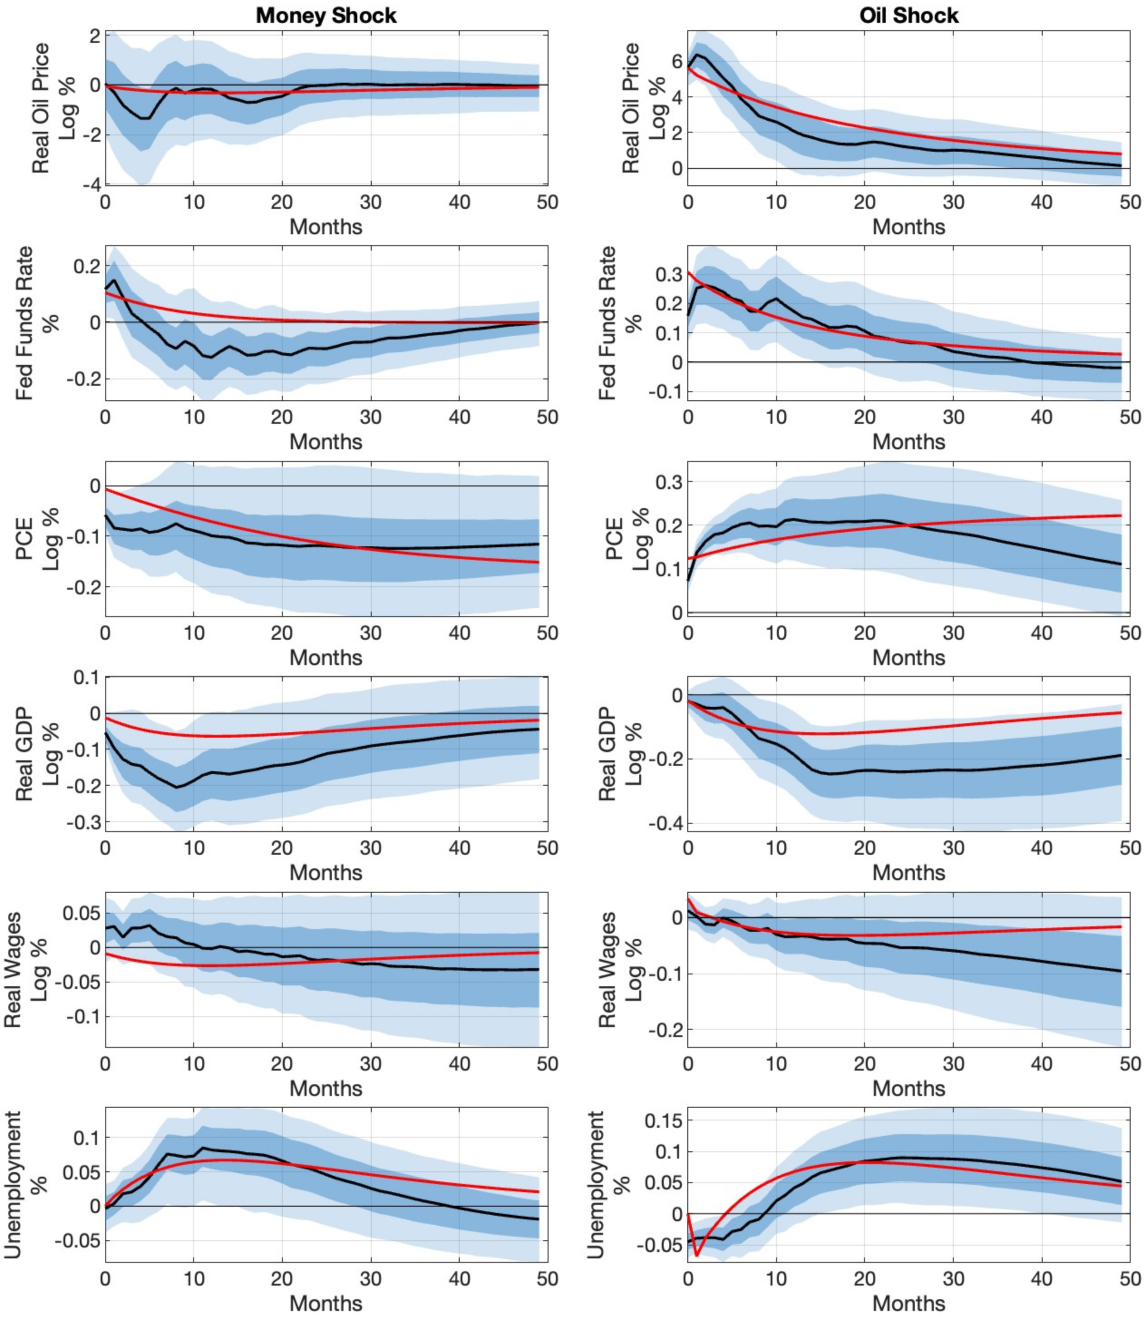

Figure 3: SVAR-based impulse responses for identified money and oil shocks vs model-based impulse responses (in red). The solid line is the point estimate and the dark and light-shaded areas are 68 and 95 percent confidence bands, respectively, computed using the wild bootstrap.

the persistent increase in inflation.

We then proceed as follows. Using the estimated model from section 4.4, we recover the shocks by targeting variables other than inflation. We then feed the estimated shocks into the model to identify how well the framework explains inflation along with the contribution of each shock. Using standard Bayesian methods, we then estimate the standard deviations of all shocks and also the persistence of the demand and matching shocks only, since we obtained the persistences of the money and oil shocks from the earlier estimation.[23] Priors are set to standard values. Results are reported in Online Appendix B.

To identify the four shocks, we target four variables: the unemployment rate, real oil price inflation (in terms of PCE core), the Federal funds rate, and labor market tightness. Real oil price inflation is the quarter-to-quarter annualized percent change in the real oil price; market tightness is obtained from JOLTS as the ratio between job openings and unemployed persons. From the four targeted variables, we obtain the smoothed series for the shocks using the Kalman smoother. We can then construct historical decompositions.

One complication in doing this exercise is that the oil price series displays considerable high-frequency volatility, possibly due to speculation in financial markets. Some of these high-frequency gyrations do not appear to immediately translate into prices that households and firms face, as a comparison of wholesale oil prices with the PCE price index for energy would suggest. Accordingly, we assume that oil price inflation ($\pi_{ot} = \ln(p_{ot}/p_{ot-1})$) is the sum of a persistent component ($\bar{\pi}_{ot} = \ln(\bar{p}_{ot}/\bar{p}_{ot-1})$), which translates into retail oil prices, and an i.i.d. component ε_{mt}, which reflects speculative noise:[24]

$$\pi_{ot} = \bar{\pi}_{ot} + \varepsilon_{mt}$$

The volatility of ε_{mt}, σ^m, is residually identified from the persistence of the oil shock that we previously estimated. We note however that cleaning off the

---

[23] Notice that matching impulse responses pins down the standard deviation of the two shocks up to a normalization, therefore we need to re-estimate these parameters.

[24] The price index becomes $p_{ct} = (\chi \bar{p}_{ot}^{1-\psi} + (1-\chi) p_{qt}^{1-\psi})^{\frac{1}{1-\psi}}$.

high-frequency noise in oil prices only has a minor effect on the results.[25]

Since all the nominal variables are untargeted, we can judge how well the model captures inflation by constructing model-implied series for the following year-over-year annualized variables: PCE inflation, core PCE inflation, nominal wage inflation, and real product wage inflation. Variables are demeaned using the sample mean. As our sample overlaps with the slow recovery from the 2008 recession, we take SS unemployment to be 5 percent.[26]

5.1 Historical Shock Decompositions

5.1.1 Targeted Variables

Figure 4 presents a historical decomposition for the four targeted variables over the sample 2010:01-2023:11. Overall, the results are very sensible. The demand shock accounts for most of the variation in unemployment. In this vein, the model treats the sharp rise in unemployment during the pandemic as largely the product of a sharp drop in demand.[27] Unemployment then drops to steady state as demand improves. Interestingly, however, the drop in unemployment that continues in 2021-22 is largely the product of easy monetary policy. Indeed, over this period, monetary policy shocks more than offset the contractionary effect of oil shocks: from mid 2020 onward, oil shocks contribute a roughly two percentage point increase in unemployment.

Labor market tightness mirrors the behavior of unemployment: it is highly sensitive to the demand shock until mid 2021. After that, accommodative monetary policy stimulates tightness, while oil shocks do the reverse. Interestingly, the matching shock is nontrivial during and after the

---

[25] In particular, there are two data points with unusually large oil price shocks that quickly revert. Without cleaning off the noise, the model would predict that these shocks would generate counterfactually large changes in the real economy in those two months.

[26] The sample mean over the period 2010-2022 is 6 percent. We choose to demean using 5 percent for consistency with the model calibration as well as the sample mean over the full sample. Results are robust to demeaning with 6 percent instead.

[27] As we show shortly, both headline and core PCE declined during the pandemic recession, consistent with the interpretation that the demand shock is a key driving force.

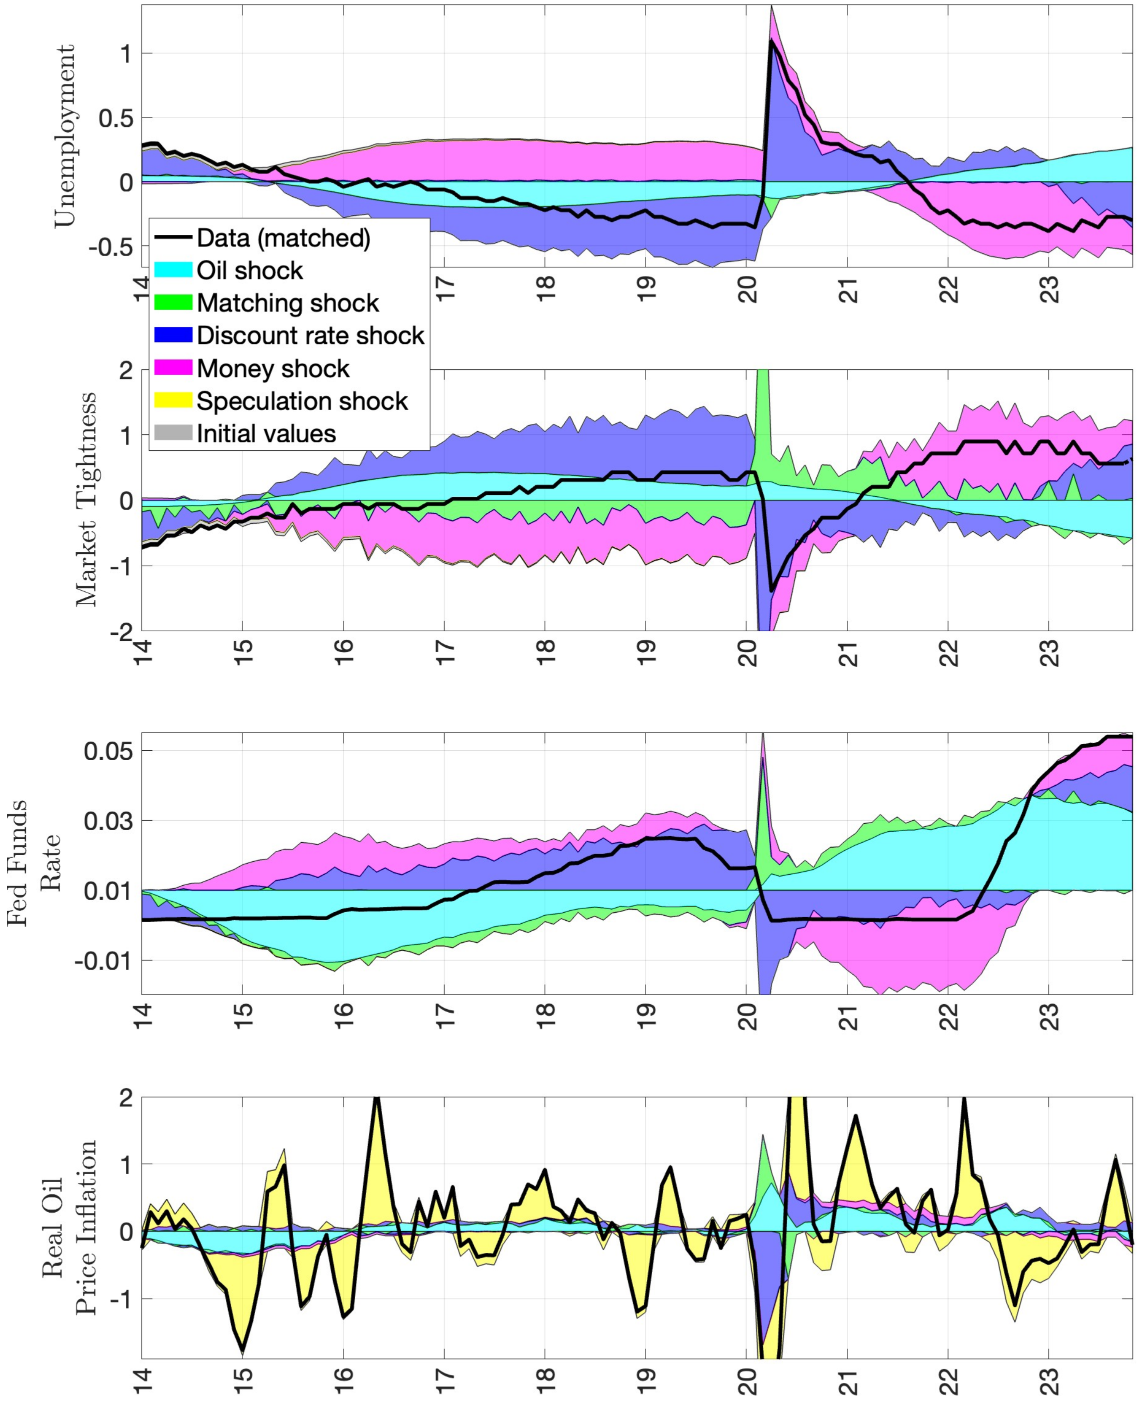

Figure 4: Historical shock decomposition of the targeted variables. Unemployment and labor market tightness are in log-deviations from the steady-state value for the model and log-deviations from the sample mean for the data. The decomposition for Fed funds is computed in deviations from steady state/sample mean and then rescaled up by the sample mean. Fed funds and oil inflation are annualized.

pandemic but is not the leading driver of market tightness. It also does not materially contribute to unemployment variation over the sample.

The Federal funds rate was fixed at the zero lower bound for much of the sample. The rise in the funds rate just before the pandemic recession and the decline just after was largely in response to the rise and fall of the demand shock. From the end of 2020 to early 2022 the easy money shock keeps the funds rate low. A rapid tightening then follows.

After filtering out background noise with the speculation shock (as described earlier), the oil shock mainly drives the behavior of the oil price. The one exception is that the demand shock that pushed the economy into the pandemic recession placed significant downward pressure on oil prices.[28]

5.1.2 Untargeted Variables: Inflation and Wages

Figure 5 reports the shock decomposition for year-over-year headline PCE inflation, core PCE inflation, nominal wage growth, and real wage growth.

The model tracks both PCE and core PCE inflation over the entire sample reasonably well. In particular, the model captures well the rise in core PCE inflation in 2021, the acceleration in early 2022, and the decline beginning in 2023. It understates the rise for a brief period in 2021, which can potentially be explained by the absence of supply-chain factors in the model.[29] Similarly, the model also does a good job at explaining the surge in headline inflation, though it misses part of the decline at the end of the sample, which is largely due to the drop in oil and commodity prices.

The figure shows that the two main driving forces underlying the inflation surge were oil shocks and shocks to monetary policy. That both these shocks remain important throughout 2022 despite both oil prices and the Fed funds

---

[28] Recent data on high-frequency oil shocks that extend the analysis by Känzig (2021) are consistent with the increase in oil prices being driven by a reduction in the OPEC oil supply in both 2021 and 2022. We thank Diego Känzig for sharing the updated series.

[29] Our interpretation of this evidence is consistent with Di Giovanni et al. (2022), which shows that supply chain factors accounted for about a third of the inflation runup in the second half of 2021, but gave way to other forces in 2022.

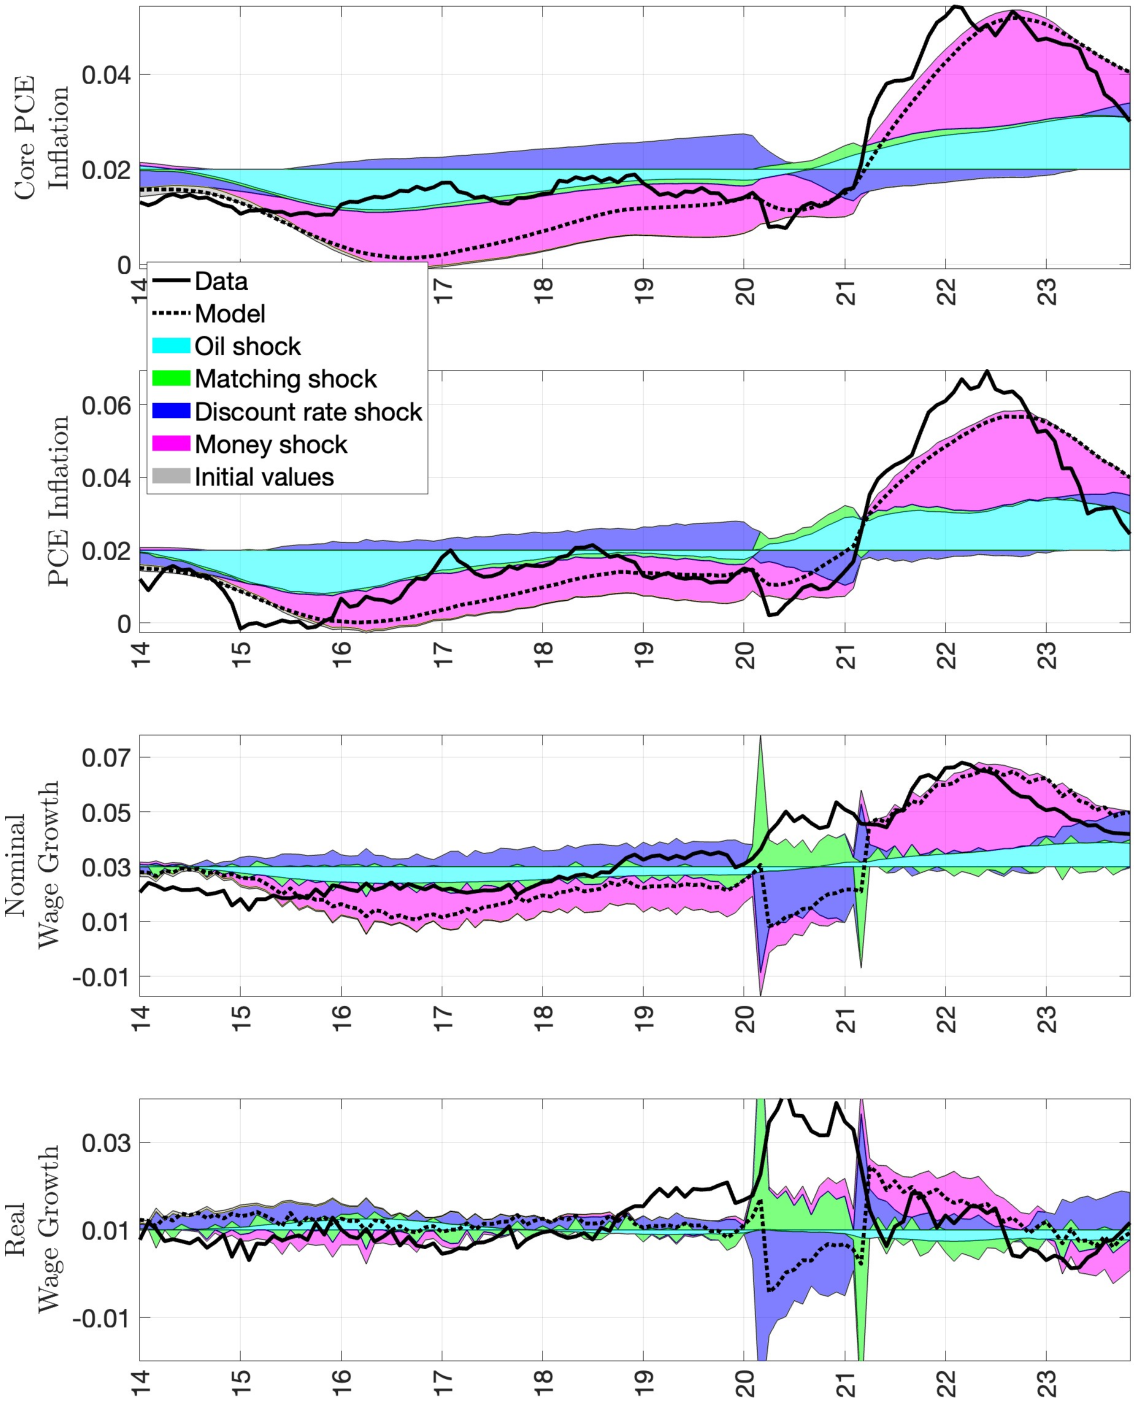

Figure 5: Historical shock decomposition of untargeted variables. The decomposition is computed in deviations from steady state/sample mean, and then rescaled up by the sample mean. All the variables are annualized.

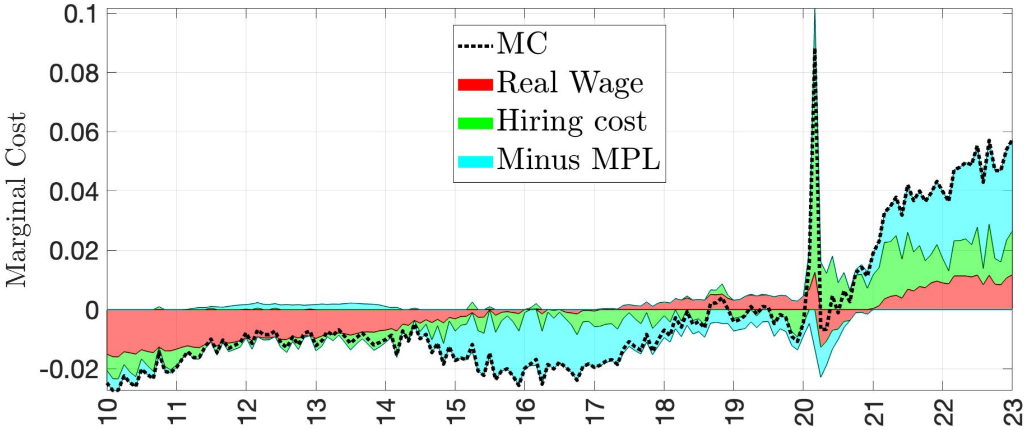

Figure 6: Historical decomposition of marginal cost into the main components from equation (21). Marginal cost is $\hat{p}_{wt}$, real wage is $\hat{w}_{qt}$ (multiplied by ζ), hiring cost is $\hat{\omega}_t$ (multiplied by $(1-\zeta)$), marginal product of labor is $\hat{a}_{nt}$ (multiplied by -1).

rate reversing course, is due to the persistent impact of each shock on the economy, as is clear from the IRFs portrayed in Figure 3. However, the effects of both shocks on inflation begin to dissipate in 2023, as monetary policy tightens and oil prices decline, contributing to a moderation of inflation. Interestingly, over this period demand emerges as an inflationary force. Finally, the matching shock does not contribute significantly to the inflation surge, consistent with its minimal impact on unemployment.[30]

The model also tracks nominal and real product wage inflation reasonably well over the whole sample. There is one caveat due to a data issue, having to do with a large spike in wage inflation at the height of the pandemic recession in mid 2020 followed by a large reversal in the subsequent quarter. The likely cause of this spike was a compositional effect arising because employment losses were concentrated among low-wage workers. Our model of course cannot capture this kind of compositional effect.

We finally illustrate the mechanics of the inflation surge. As discussed in section 3, given that long-term inflation expectations are anchored, the only way the model can produce an inflation surge is to have a large persistent

---

[30] We checked that our results are robust to increase substantially the persistence of the matching shock to 0.9 and 0.95. The matching shock picks up differences between the unemployment rate and labor market tightness, which differed substantially during the pandemic but realigned soon after.

increase in the expected path of marginal cost. Indeed our model suggests just that: it implies that at the heart of the inflation surge in 2021 was a sharp increase in marginal cost. Figure 6 shows the increase in marginal cost over this period and decomposes it into its three components: real wages, net hiring costs, and the marginal product of labor. As the figure shows, all three components play a role. However, the decline in the marginal product of labor accounts for more than half the increase. Given its importance in the dynamics of this variable (see section 3), the strong complementarity between oil and labor plays an important role in the runup of marginal cost, and hence in the runup of inflation.

5.2 Accounting for forward guidance

Our approach thus far measures monetary policy accommodation by the deviation of the Federal funds rate from what a conventional Taylor would predict. A possible criticism is that this measure does not factor in the central bank's manipulation of longer-term rates via forward guidance. Indeed government bond rates across the yield curve began to creep up in late 2021 prior to the liftoff, suggesting that monetary policy tightening was in motion prior to the liftoff of the Funds rate in early 2022. The measure of policy accommodation in our baseline case does not capture this kind of tightening. Accordingly, we will now show that our main results are robust to explicitly taking into account the tightening due to forward guidance that occurred prior to the funds rate liftoff, though with a few interesting differences.

To implement this robustness exercise, we use the "proxy" funds rate developed by Choi et al. (2022) that adjusts the funds rate to factor in the role of forward guidance and unconventional monetary policies. The authors first construct a financial conditions index based on the principal components of a set of long and short-term interest rates and interest rate spreads.[31]

---

[31] The interest rates include government bond rates on maturities ranging from two to ten years along with rates on private securities such as mortgages and corporate bonds. The spreads include both mortgage and corporate bond rates relative to similar-maturity

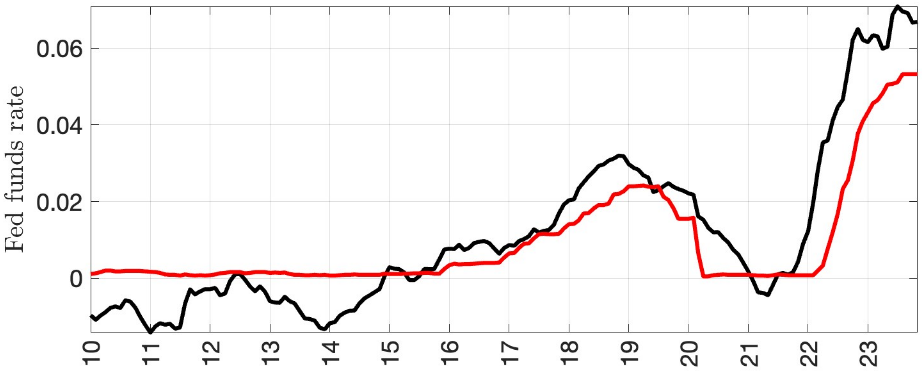

Figure 7: Comparison of the Fed funds rate (red) against the proxy rate (black).

The proxy rate is then constructed as the fitted value of a regression of the effective funds rate on the index, with coefficients estimated on pre-2008 data. The identification of forward guidance exploits the fact that, prior to 2008, the Federal Reserve did not materially rely on either forward guidance or unconventional policies, so one can use this period to identify the correlation between the funds rate and the index absent these policies. Therefore, by construction, the proxy rate aligns closely with the funds rate before 2008. It differs after that to the extent that the long-term rates, which enter the index, are not aligned with the effective funds in the same way they were pre-2008 due, for example, to the active use of forward guidance.

Figure 7 shows the proxy rate (black line) relative to the funds rate (red line) for the period 2010 to the present. For periods where the ZLB is binding, the proxy rate moves below the funds rate. This reflects the presence of longer maturity interest rates in the financial conditions index: during the ZLB periods, longer-term rates were lower than normal relative to the funds rate, due to forward guidance. Importantly for our purposes, over the recent tightening period, where long rates were high relative to the funds rate, the proxy rate led the runup in funds rate by several months.

Accordingly, to account for the role of forward guidance in the recent tightening, we repeat the same accounting exercise as in section 5.1, but replace the funds rate with the proxy rate. Figure 8 reports the historical

---

government bond yields.

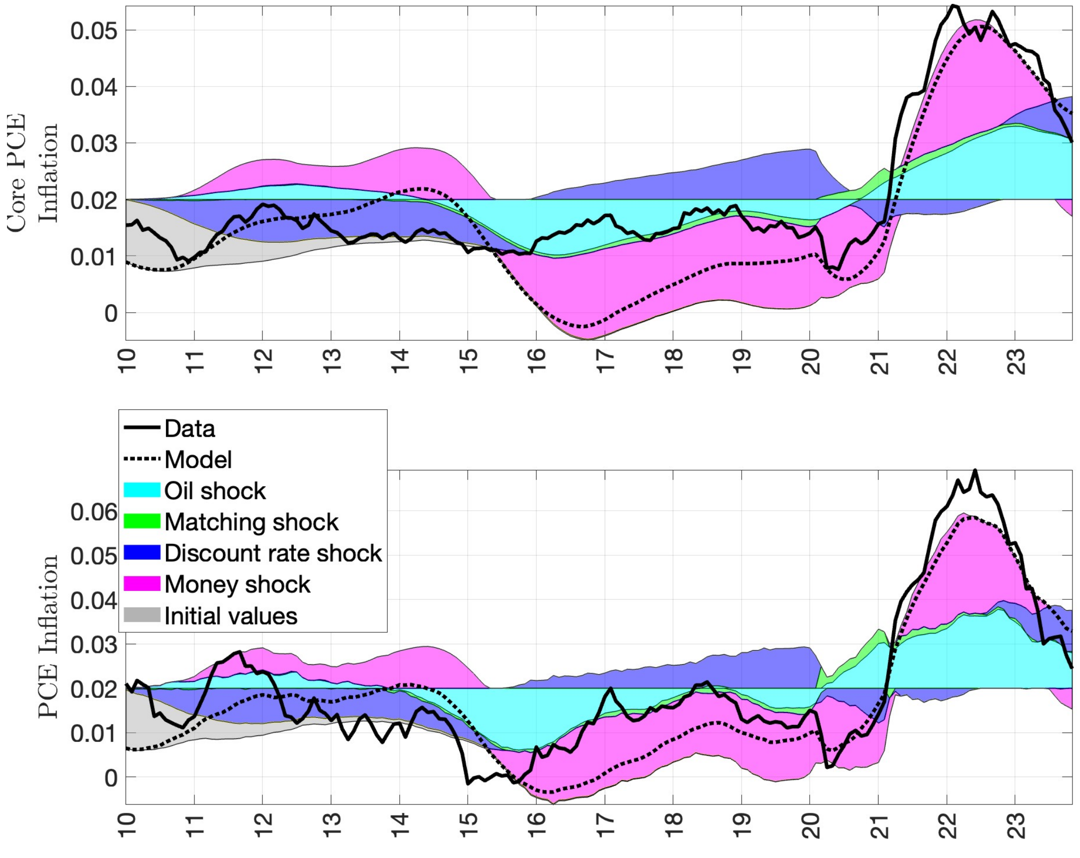

Figure 8: Historical decomposition when matching the proxy rate.

decomposition of core and headline inflation for the case with the proxy rate.[32] Interestingly, the ability of the model to explain the surge in both core and headline inflation rates improves a bit with respect to our baseline result. Beyond this, as before, oil prices and monetary accommodation are the main contributing factors to the surge prior to 2023. However, in contrast to the baseline, the impact of monetary accommodation on inflation dies out at the end of 2022, likely reflecting the impact of forward guidance. For both core and headline, the demand shock then rises in importance in 2023 while the oil shock begins to recede.

---

[32] In Online Appendix C, we report the decompositions for the complete set of variables.

6 Concluding Remarks

We have developed and estimated a simple New Keynesian model designed to account for the recent inflation surge. The model features the roles of oil shocks and accommodative monetary policy, but allows for shocks to demand and labor market tightness, factors thought to be relevant as well. Two aspects of the model important for the quantitative performance are (i) oil as a complementary good and input and (ii) real wage rigidity.

On the methodological side, we pin down parameters in a way that avoids fitting the model to the inflation surge we are attempting to explain. First, we estimate parameters by matching model impulse responses to those from identified shocks to both oil and monetary policy in a structural VAR, using pre-pandemic data. This identification via observable shocks poses tight data restrictions on the two central driving forces of our model. The estimated parameter values are sensible and consistent with the literature.

We next use the estimated framework to recover the model shocks without targeting inflation. Using these shocks, we find that the model-implied prediction tracks inflation well, including the recent surge. This provides an additional validation of our model, which we then use to perform a historical decomposition. In line with the suggestive evidence discussed in the introduction, mainly accounting for both the rise and subsequent moderation of inflation is a combination of oil and monetary policy shocks. Interestingly, the model also provides a rationalization for why during the period 2010-2019 inflation was low despite low unemployment. At work was a series of oil price declines and tight money shocks, the exact opposite of what occurred during the recent inflation surge.

We also show that our main results are robust to accounting for the policy tightening due to forward guidance that occurred in late 2021. However, factoring in forward guidance suggests that policy accommodation stopped contributing to inflation by the end of 2022. Instead, the policy tightening along with the easing of oil prices began contributing to the moderation of inflation in 2023.

Finally, we have presumed that long-run inflation expectations have remained anchored at the target as is consistent with the recent evidence. But how long we can count on inflation expectations remaining anchored as inflation persists above the target is an open question for the future.

References

Joel Alcedo, Alberto Cavallo, Bricklin Dwyer, Prachi Mishra, and Antonio Spilimbergo. E-commerce during covid: Stylized facts from 47 economies. Technical report, National Bureau of Economic Research, 2022.

Mary Amiti, Sebastian Heise, Fatih Karahan, Aysegul Sahin, et al. Pass-through of wages and import prices has increased in the post-covid period. Technical report, Federal Reserve Bank of New York, 2022.

Ruediger Bachmann, David Baqaee, Christian Bayer, Moritz Kuhn, Andreas Löschel, Benjamin Moll, Andreas Peichl, Karen Pittel, and Moritz Schularick. What if? the economic effects for germany of a stop of energy imports from russia. *Wirtschaftsdienst*, 102:251–255, 2022.

Laurence M Ball, Daniel Leigh, and Prachi Mishra. Understanding us inflation during the covid era. Technical report, National Bureau of Economic Research, 2022.

Michael Bauer and Eric T Swanson. An alternative explanation for the 'fed information effect'. 2021.

Michael D Bauer and Eric T Swanson. A reassessment of monetary policy surprises and high-frequency identification. *NBER Macroeconomics Annual*, 37(1):87–155, 2023.

Pierpaolo Benigno and Gauti B Eggertsson. It's baaack: The surge in inflation in the 2020s and the return of the non-linear phillips curve. Technical report, National Bureau of Economic Research, 2023.

Ben Bernanke and Olivier Blanchard. What caused the us pandemic-era inflation? *Hutchins Center Working Papers*, 2023.

Olivier J Blanchard and Jordi Gali. The macroeconomic effects of oil shocks: Why are the 2000s so different from the 1970s?, 2007.

Martin Bodenstein and Luca Guerrieri. Oil efficiency, demand, and prices: a tale of ups and downs. 2011.

Martin Bodenstein, Luca Guerrieri, and Lutz Kilian. Monetary policy responses to oil price fluctuations. *IMF Economic Review*, 60(4):470–504, 2012.

Carlos Carvalho, Fernanda Nechio, and Tiago Tristao. Taylor rule estimation by ols. *Journal of Monetary Economics*, 124:140–154, 2021.

Jason Choi, Taeyoung Doh, Andrew Foerster, Zinnia Martinez, et al. Monetary policy stance is tighter than federal funds rate. *FRBSF Economic Letter*, 30, 2022.

Diego A Comin, Robert C Johnson, and Callum J Jones. Supply chain constraints and inflation. Technical report, National Bureau of Economic Research, 2023.

Julian Di Giovanni, Sebnem Kalemli-Özcan, Alvaro Silva, and Muhammed A Yildirim. Global supply chain pressures, international trade, and inflation. Technical report, National Bureau of Economic Research, 2022.

Julian Di Giovanni, Sebnem Kalemli-Özcan, Alvaro Silva, and Muhammed A Yildirim. Pandemic-era inflation drivers and global spillovers. Technical report, National Bureau of Economic Research, 2023.

Francesco Ferrante, Sebastian Graves, and Matteo Iacoviello. The inflationary effects of sectoral reallocation. 2023.

Luca Gagliardone, Mark Gertler, Simone Lenzu, and Joris Tielens. Anatomy of the phillips curve: Micro evidence and macro implications. Technical report, National Bureau of Economic Research, 2023.

Mark Gertler and Peter Karadi. Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, 7(1):44–76, 2015.

Simon Gilchrist and Egon Zakrajšek. Credit spreads and business cycle fluctuations. *American Economic Review*, 102(4):1692–1720, 2012.

Veronica Guerrieri, Guido Lorenzoni, Ludwig Straub, and Iván Werning. Monetary policy in times of structural reallocation. *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2021-111), 2021.

Robert E Hall. Reconciling cyclical movements in the marginal value of time and the marginal product of labor. *Journal of Political Economy*, 117(2):281–323, 2009.

James D Hamilton. Oil and the macroeconomy since world war ii. *Journal of political economy*, 91(2):228–248, 1983.

Tomas Havranek, Marek Rusnak, and Anna Sokolova. Habit formation in consumption: A meta-analysis. *European Economic Review*, 95:142–167, 2017.

Jonathon Hazell, Juan Herreno, Emi Nakamura, and Jón Steinsson. The slope of the phillips curve: evidence from us states. *The Quarterly Journal of Economics*, 137(3):1299–1344, 2022.

Diego R Känzig. The macroeconomic effects of oil supply news: Evidence from opec announcements. *American Economic Review*, 111(4):1092–1125, 2021.

Lutz Kilian. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–1069, 2009.

Guido Lorenzoni and Iván Werning. Wage price spirals. 2023.

Dale T Mortensen and Christopher A Pissarides. Job creation and job destruction in the theory of unemployment. *The Review of Economic Studies*, 61(3):397–415, 1994.

Carlo Rosa. The high-frequency response of energy prices to us monetary policy: Understanding the empirical evidence. *Energy Economics*, 45:295–303, 2014.

Pilar Soriano and Hipòlit Torró. The response of brent crude oil to the european central bank monetary policy. *Finance Research Letters*, 46:102353, 2022.

Claudio Soto and Juan Pablo Medina. Oil shocks and monetary policy in an estimated dsge model for a small open economy. *Documentos de Trabajo (Banco Central de Chile)*, (353):1, 2005.

Online Appendix

A Estimation by Matching Impulse Response

In this section, we explain how the parameters are estimated and the confidence intervals are derived. In particular, we follow Hall et al. (2012) and Mertens and Ravn (2011) who propose an estimator based on the simulated method of moments and with inference based on the delta method. Specifically, let Λ^d be the $T \cdot N \cdot S$ vector of stacked impulse responses estimated in the data, where $T = 50$ is the forecast horizon in months, $N = 6$ the number of variables that are targeted, and $S = 2$ the number of shocks considered. Also, let $\Lambda^m(\Theta_2|\Theta_1)$ be the $T \cdot N \cdot S$ vector of stacked impulse responses obtained from model simulations, where Θ_2 is the set of parameters to be estimated conditional on the calibrated parameters Θ_1. Finally, let Σ_d^{-1}be a weighting matrix. The estimator of Θ_2 is given by:

$$\hat{\Theta}_2 = \arg\min_{\Theta_2} \left[\left(\Lambda^d - \Lambda^m(\Theta_2|\Theta_1)\right)' \Sigma_d^{-1} \left(\Lambda^d - \Lambda^m(\Theta_2|\Theta_1)\right) \right]$$

For the weighting matrix Σ_d^{-1}, we follow the standard approach to use the precision of the IRFs estimated from the VAR along the main diagonal, so that estimates with a smaller variance are assigned a larger weight in the minimization. We make an exception for the contemporaneous impact of the money shock on Fed funds and the contemporaneous impact of the oil shock on the oil price, which we assign a larger weight to ensure these own impact moments are estimated more precisely.

The standard errors of $\hat{\Theta}_2$ are computed using an estimate of the asymptotic covariance matrix derived with the delta method:

$$\Sigma_{\Theta_2} = \Lambda_{\Theta_2} \frac{\partial \Lambda^m(\Theta_2|\Theta_1)'}{\partial \Theta_2} \Sigma_d^{-1} \Sigma_S \Sigma_d^{-1} \frac{\partial \Lambda^m(\Theta_2|\Theta_1)}{\partial \Theta_2} \Lambda_{\Theta_2}$$

where

$$\Lambda_{\Theta_2} = \left[\frac{\partial \Lambda^m(\Theta_2|\Theta_1)'}{\partial \Theta_2} \Sigma_d^{-1} \frac{\partial \Lambda^m(\Theta_2|\Theta_1)}{\partial \Theta_2} \right]^{-1}$$

$$\Sigma_S = \Sigma + \Sigma_s$$

and Σ denotes the covariance matrix of the estimated SVAR-based IRFs and Σ_s is the covariance matrix of the model-based impulse responses.

B Bayesian Estimation Result

We report in Table 2 the results of the Bayesian estimation of the shocks over the sample 2010-2022.

Parameter	Prior	Prior Mean	Prior stdev	Post. Mean	5%	95%
ρ^b	beta	.6	.1	.227	.177	.281
ρ^Φ	beta	.6	.1	.534	.432	.638
σ^b	invg	.15	.15	.061	.055	.066
σ^Φ	invg	.15	.15	.161	.145	.174
σ^m	invg	.15	.15	.248	.220	.281
σ^o	invg	.15	.15	.042	.035	.048
σ^r	invg	.15	.15	.038	.033	.043

Table 2: Bayesian estimation of the parameters over the sample 2010-2022.

Prior means and standard deviations are standard as in Primiceri et al. (2006). The prior standard deviations are sufficiently large to not impose serious restrictions on the parameters. The estimates imply that both the matching shock and the discount factor shock are not very persistent, with the matching shock more persistent ($\rho^\Phi = .53$ at the posterior mean) then the discount shock ($\rho^b = .23$ at the posterior mean). The estimates of the standard deviations are sensible, with the posterior means of the standard deviation for oil $\sigma^s = .04$ and money shock $\sigma^r = .04$ that are of the same order of magnitude as those estimated for the IRFs matching exercise (which were normalized to match one standard deviation of oil prices and Fed funds respectively). The mean of the standard deviation for the speculation shock $\sigma^m = .25$ is substantially larger than that of the oil shock, confirming the intuition that the speculation shock captures temporary volatility in oil prices that does not translate into a persistent effect on real variables. Finally, the posterior means for the matching shock $\sigma^\Phi = .16$ and discount factor shock $\sigma^b = .06$ are larger than both oil and money (because of the lower persistence), but of the same order of magnitude.

C Decomposition of matched variables with the proxy rate

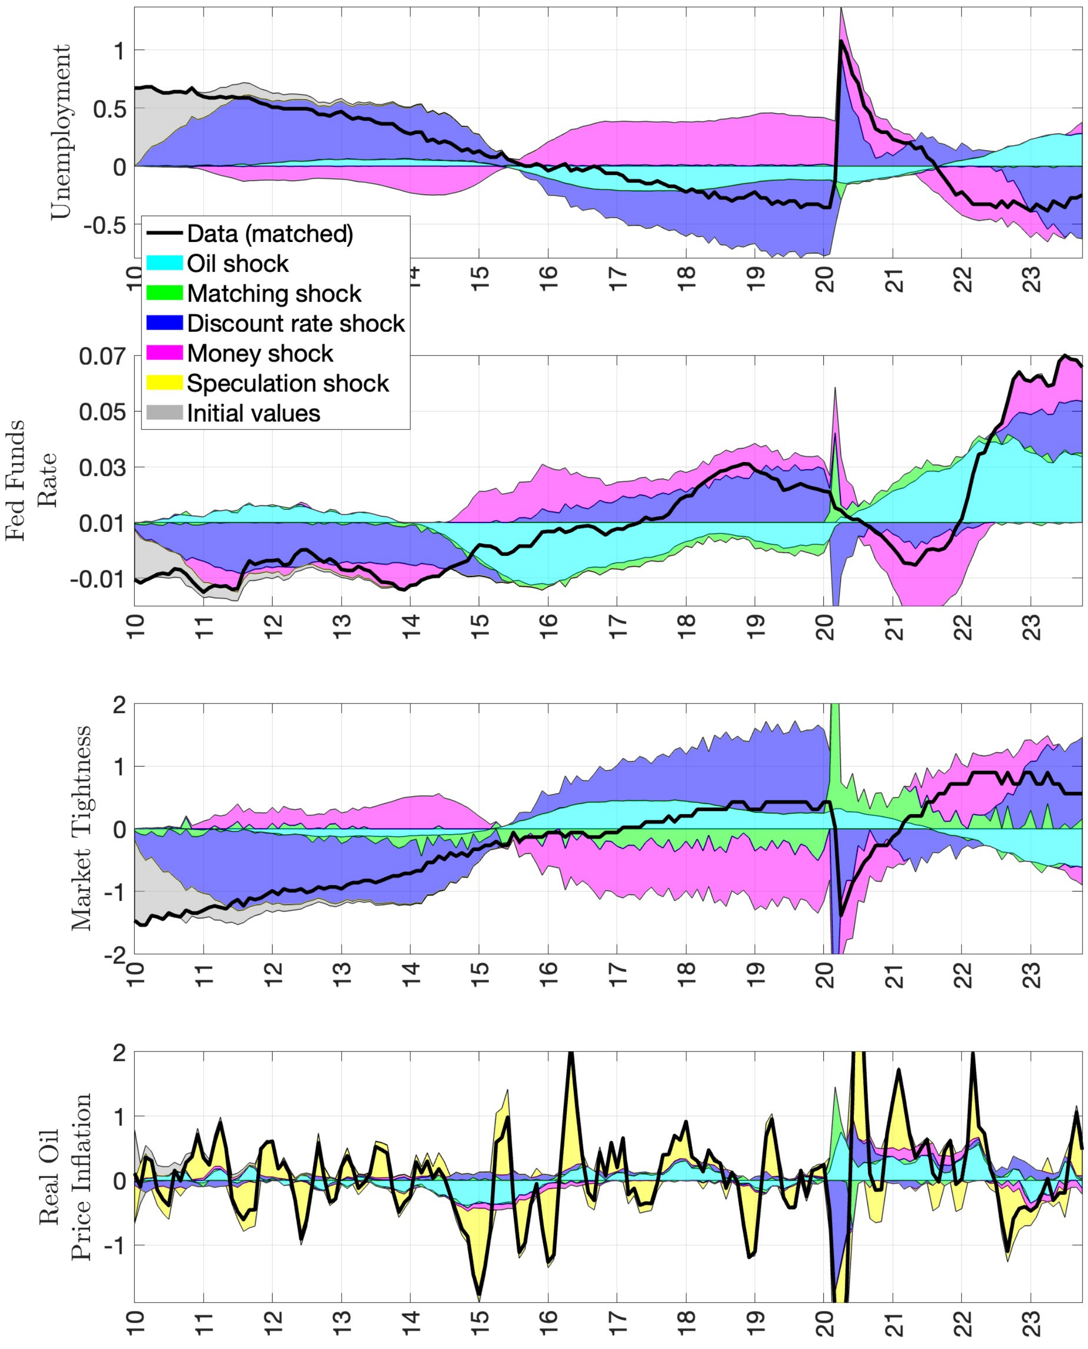

References

Alastair R Hall, Atsushi Inoue, James M Nason, and Barbara Rossi. Information criteria for impulse response function matching estimation of dsge models. *Journal of Econometrics*, 170(2):499–518, 2012.

Karel Mertens and Morten O Ravn. Understanding the aggregate effects of anticipated and unanticipated tax policy shocks. *Review of Economic dynamics*, 14(1):27–54, 2011.

Giorgio Primiceri, Ernst Schaumburg, and Andrea Tambalotti. Intertemporal disturbances, 2006.