Oil Prices, Monetary Policy and Inflation Surges

Luca Gagliardone*   Mark Gertler†

July, 2024.

(First draft: March, 2023)

Abstract

We develop a simple quantitative New Keynesian model aimed at analyzing how the reaction of monetary policy contributed to the recent rise and fall in inflation. The model includes several shocks but features oil price shocks for two reasons: (i) energy prices have been among the central factors in discussions about the surge; (ii) we can use identified oil shocks along with monetary shocks to estimate and discipline the model. We then employ the estimated framework to recover shocks without targeting inflation. Overall the model accounts for roughly three fourths of the surge in PCE inflation. Both the oil shocks and the shocks to policy accommodation played important roles in the inflation rise. Moreover, the easing of oil prices and subsequent shift to policy tightening contributed to the decline. A nonmonetary demand shock (a composite of private demand and fiscal stimulus) also contributed to inflation starting in 2022.

JEL: E32; E52.

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

1 Introduction

To account for the recent inflation surge, the literature has proposed a variety of supply and demand shocks, including supply-chain disruptions, energy and other commodity price shocks, expansionary fiscal policy, and shifts in labor market tightness. What has received somewhat less attention is the role of monetary policy. For any of the driving forces listed above, the response of monetary policy ultimately shapes their effect on inflation.

In this paper, we develop and estimate a simple New Keynesian model designed to analyze how monetary policy may have contributed to the recent inflation surge. To be clear, the model is not meant to provide a complete description of the surge, but rather a laboratory to study the role of accommodative policy. We place particular emphasis on oil shocks as a driving force for two reasons. First, a number of authors have emphasized rising energy prices as a significant factor underlying the surge. Second, given the availability of identified exogenous oil shocks due to Känzig (2021), we can use evidence on the impact of these shocks on the economy to discipline our model. Similarly, given our interest in analyzing the impact of the response of monetary policy, we can also use evidence on the reaction of the economy to identified monetary policy shocks to provide additional discipline.

Figure 1 provides motivation for focusing on the roles of oil shocks and monetary policy. 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 that the correlation remains strong as both inflation and oil prices ease after mid 2022. Our model will sort out the causal factors underlying these correlations.

The bottom panel illustrates the monetary policy accommodation over this period: 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. It is true that markets began to anticipate the liftoff as longer-term nominal yields began to rise in advance. However, real rates not only declined substantially during the policy accommodation, they remained low until well after the tightening began, as illustrated by the behavior of the one and five-year real rates in Figure 1. Conversely, real rates rose significantly in

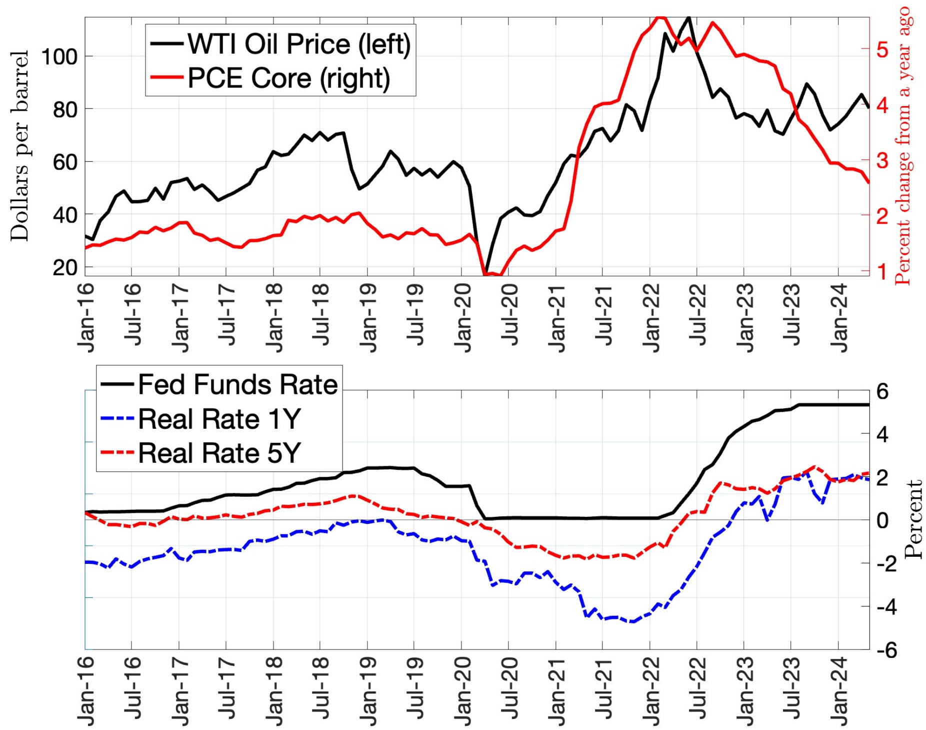

Figure 1: Time series of PCE core inflation against WTI oil prices (top panel) and Fed funds rate against the real one-year and five-year real rates (bottom panel). The one-year rate is constructed as the one-year government bond yield minus the one-year inflation expectations from the Michigan survey. The five-year rate is constructed as the five-year government bond yield minus the five-year breakeven inflation rate.

response to the subsequent tightening. Overall the figure suggests considerable policy accommodation from spring 2021 through mid 2022, followed by a clear move to a tightening stance in late 2022 to the present. The framework we develop will sort out the implications of this shifting stance of monetary policy.

Our framework also allows for other factors thought to be relevant, 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 roughly three fourths of 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 the evidence on the impact of oil shocks on inflation.[1] We also allow for unemployment via search and matching in the labor market, which enables us to consider shocks to labor market tightness as a source of inflation. In addition, we capture demand shifts by variation in the discount factor. As we will discuss, one can interpret the demand shock as a composite of private-sector demand and spending induced by fiscal policy.

For tractability, there are some factors potentially relevant to the recent inflation surge 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.

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, where the rise in demand can be due to either non-monetary factors or accommodative monetary policy. But to have a rise in marginal cost 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). We consider two types of exogenous shocks. The first is a high-frequency oil shock, identified as in Känzig (2021). The second is a high-frequency shock to monetary policy, identified as in Gertler and Karadi (2015). Each shock serves as an external instrument in the SVAR. We choose to match impulse responses from both shocks to ensure the model produces a realistic response to oil shocks along with a realistic reaction to monetary policy accommodation.

In section 5 we explore how well the model accounts for the recent inflation.

[1]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.

We quantify the role of each of the four aggregate shocks in our model, namely shocks to oil, demand, labor market tightness, and monetary policy. Here our measure of shocks to monetary policy, which is essentially our measure of policy accommodation, is the difference between the policy rate and the central bank's historical policy rule. For the monetary policy rate, we choose a measure that provides a conservative estimate of the degree of policy accommodation. A difficulty with the Federal Funds rate, which is the standard rate, is that it does not readily account for either forward guidance or balance sheet policies (i.e. QE/QT) that affect longer-term yields. Accordingly. we use as the monetary policy rate the "proxy rate" developed by Choi et al. (2022) that adjusts the Funds rate to capture these factors. Doing so reduces the overall measured policy accommodation relative to what is obtained from using the Fed funds rate.

To perform our "inflation accounting" exercise, we first use our estimated model to recover the four shocks by targeting the following four variables: unemployment, labor market tightness, the Federal funds rate, and oil prices. We leave untargeted the nominal variables, including both headline and core inflation and nominal wage inflation. We then show that the model does a good job explaining inflation, including most though not all of the recent surge. Accounting for much of the inflation surge is a combination of oil price shocks and accommodative monetary policy, even after controlling for shocks to demand and labor market tightness. Beginning at the end of 2022, however, the shift toward aggressive tightening of monetary policy and the easing of oil prices both contribute to the decline in inflation. Picking up the slack are nonmonetary demand shocks that increase inflation.

In section 6, we discuss how to think about factors left out of the model such as supply chain disruptions and fiscal policy. We provide descriptive evidence that some of the gap between the model and the data could likely be accounted for by supply chain issues. We also show that the behavior of our nonmonetary demand shock, which is a composite of private demand variation and the effects of fiscal stimulus is consistent with the evidence in Smets and Wouters (2024) that models the two phenomena separately.

Finally, concluding remarks are in section 7.

Related Literature. As suggested earlier, our theoretical framework is related to Blanchard and Gali (2007)'s NK model with 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 Erceg et al. (2024), 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 the 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 in inflation surges. We differ by presenting a quantification of this mechanism.

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) who also consider the response of monetary policy, Ferrante et al. (2023), and Di Giovanni et al. (2023). Consistent with our results, the latter also finds an important role for energy shocks. Next, Benigno and Eggertsson (2023) emphasize non-linearities in the Phillips curve due to asymmetries in wage adjustment. Pfäuti (2023) focuses on fluctuations in attention. Finally, see Smets and Wouters (2024) and the references therein for an analysis of the contribution of fiscal policy. We differ from this literature in our focus on oil prices and monetary policy. We also differ in methodology, both with the use of observable shocks to oil and monetary policy in the estimation and by treating inflation as an untargeted variable in the historical decomposition, placing greater discipline on the analysis.

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 and oil shocks 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:

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

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:

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

where $R_t^n \frac{p_{ct}}{p_{ct+1}}$ 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:

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

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.[2]

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:

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

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

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:

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

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:[3]

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

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:

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

[3] 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.

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

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:

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

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.

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_{qt}^o, where $\gamma \in [0,1]$ reflects the degree of real wage rigidity and is a parameter we will estimate. In particular,

$$w_{qt} = (w_{qt}^o)^{1-\gamma}(w_q^o)^{\gamma}, \tag{15}$$

where w_q^o 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 the following standard condition:

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

where $\mu = 1/(1 - 1/\eta)$ is desired net markup.

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.[4] 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 with interest rate smoothing. 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:

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

where ε_{rt} is an exogenous money shock that obeys a first-order autoregressive process. We assume the central bank responds to core inflation so as to avoid temporary gyrations associated with headline. [5]

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

[4]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.

[5]We explored having the policy rate respond to the unemployment rate as well but found the estimated feedback coefficient to not be statistically different from zero.

3 Sources of Inflation Surges

We now characterize how the model can capture inflation surges. As discussed in section 2.6, core inflation depends on the real marginal cost of final goods firms. In our model, this marginal cost corresponds to the relative price of wholesale goods p_{wt}. Let $pwt = \ln(p_{wt}/p_w)$ 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. 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 can then express the marginal cost of producing a unit of retail 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:

$$p_{wt} = \frac{w_{qt} + \omega_t}{a_{nt}}. \tag{19}$$

where from equation (10), we can express ω_t as:

$$\omega_t = \frac{c}{q_t} - \rho E_t \left\{\Lambda^q_{t,t+1} \frac{c}{q_{t+1}}\right\}, \tag{20}$$

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 $\rho E_t \left\{\Lambda^q_{t,t+1} c/q_{t+1}\right\}$.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 $\Lambda^q_{t,t+1}$.

From loglinearizing equation (19), 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 marginal product of labor plays a key role particularly due to the presence of complementarities that enhance its sensitivity 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:

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

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, marginal cost is increasing in labor market tightness $\theta_t = v_t/u_t$ through two channels. First, a rise in market tightness increases the marginal hiring cost by reducing the vacancy-filling probability $q_t = \varepsilon_{\Phi t}\theta_t^{-\sigma}$. Second, the real wage is increasing in expected 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. As we show, both demand and supply factors can drive market tightness.

In sum, in this section, we have described how, with anchored expectations an inflation surge reflects a large and persistent increase in real marginal cost. We also discussed how production complementaries can facilitate this surge. Eventually, we identify the role of each of the supply and demand factors within our model, including monetary accommodation, in driving this sharp increase in marginal cost.

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 macroeconomic variables: log real gross domestic output, unemployment in levels, log real oil prices, log headline PCE, log real wages, a measure of the monetary policy rate that we describe shortly, 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 in the estimation. 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).

In choosing a measure of the monetary policy rate we opt for an indicator

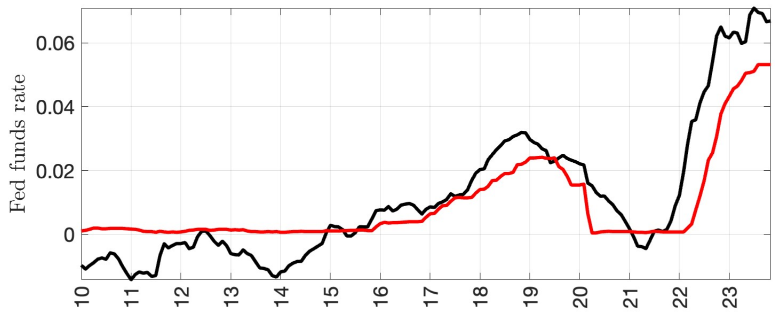

Figure 2: Comparison of the Fed funds rate (red) against the proxy rate (black).

that provides a conservative estimate of the degree of policy accommodation over the pandemic era period. The conventional indicator, the Federal Funds rate, suggests that monetary tightening did not begin until the early spring of 2022. However, as we noted in the discussion of Figure 1, longer-term nominal rates began creeping up in the latter part of 2021, consistent with tightening due to forward guidance. In addition, the central bank began "quantitative tightening" (QT) in early 2022 which likely placed further upward pressure on longer-term rates.

Accordingly, we use as the monetary policy instrument the "proxy" funds rate developed by Choi et al. (2022) that adjusts the funds rate to factor in the roles of forward guidance and balance sheet policies (i.e., QE/QT) that affect term premia. To derive the proxy rate, 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.[6] The proxy rate is then the fitted value of a regression of the effective funds rate on the index, with coefficients estimated on pre-2008 data. The identification exploits the fact that, prior to 2008, the Federal Reserve did not materially rely on either forward guidance or balance sheet 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 rates and rate spreads, which enter the index, are not correlated with the effective funds in the same way they were

[6]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 government bond yields.

pre-2008 due to the active use of forward guidance and balance sheet policies.

Figure 2 shows the proxy rate (black line) relative to the funds rate (red line) for the period 2010 to the present. Importantly for our purposes, over the recent tightening period, the proxy rate leads the runup in the Funds by several months thus providing a more conservative estimate of initial policy accommodation. In addition, the proxy rate reaches a higher peak than the Funds rate, likely the product of the accompanying QT pushing term premia up. Finally, note that for periods where the ZLB is binding, the proxy rate moves below the funds rate. During the ZLB periods, longer-term rates were lower than normal relative to the funds rate, due to forward guidance and QE.

We show in Appendix C that our results are similar if we follow convention by using the Funds rate as the policy rate, though with a somewhat larger contribution of policy accommodation to the surge. To be conserative, however,, use the proxy rate for our baseline.

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 proxy funds rate and the impact of the oil shock on the real oil price to be one standard deviation.[7]

To construct oil price surprises we follow Känzig (2021) exactly.[8] 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.[9]

For monetary policy surprises, we start with Gertler and Karadi (2015) by using unexpected movements in interest rate futures around the Federal 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.[10] 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 rate

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

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

[9]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.

[10]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.

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.[11] 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.[12]

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 3 reports the impulse responses for the identified money and oil shocks (the black lines) along with ninety-five percent confidence bands. The red line is our estimated model response which we discuss later.

The IRFs for the money shock are similar to previous estimates obtained in the literature: A monetary policy tightening of nearly 20 basis points implies

[11] 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.

[12] This is consistent with findings from Bauer and Swanson (2023), which orthogonalizes 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.

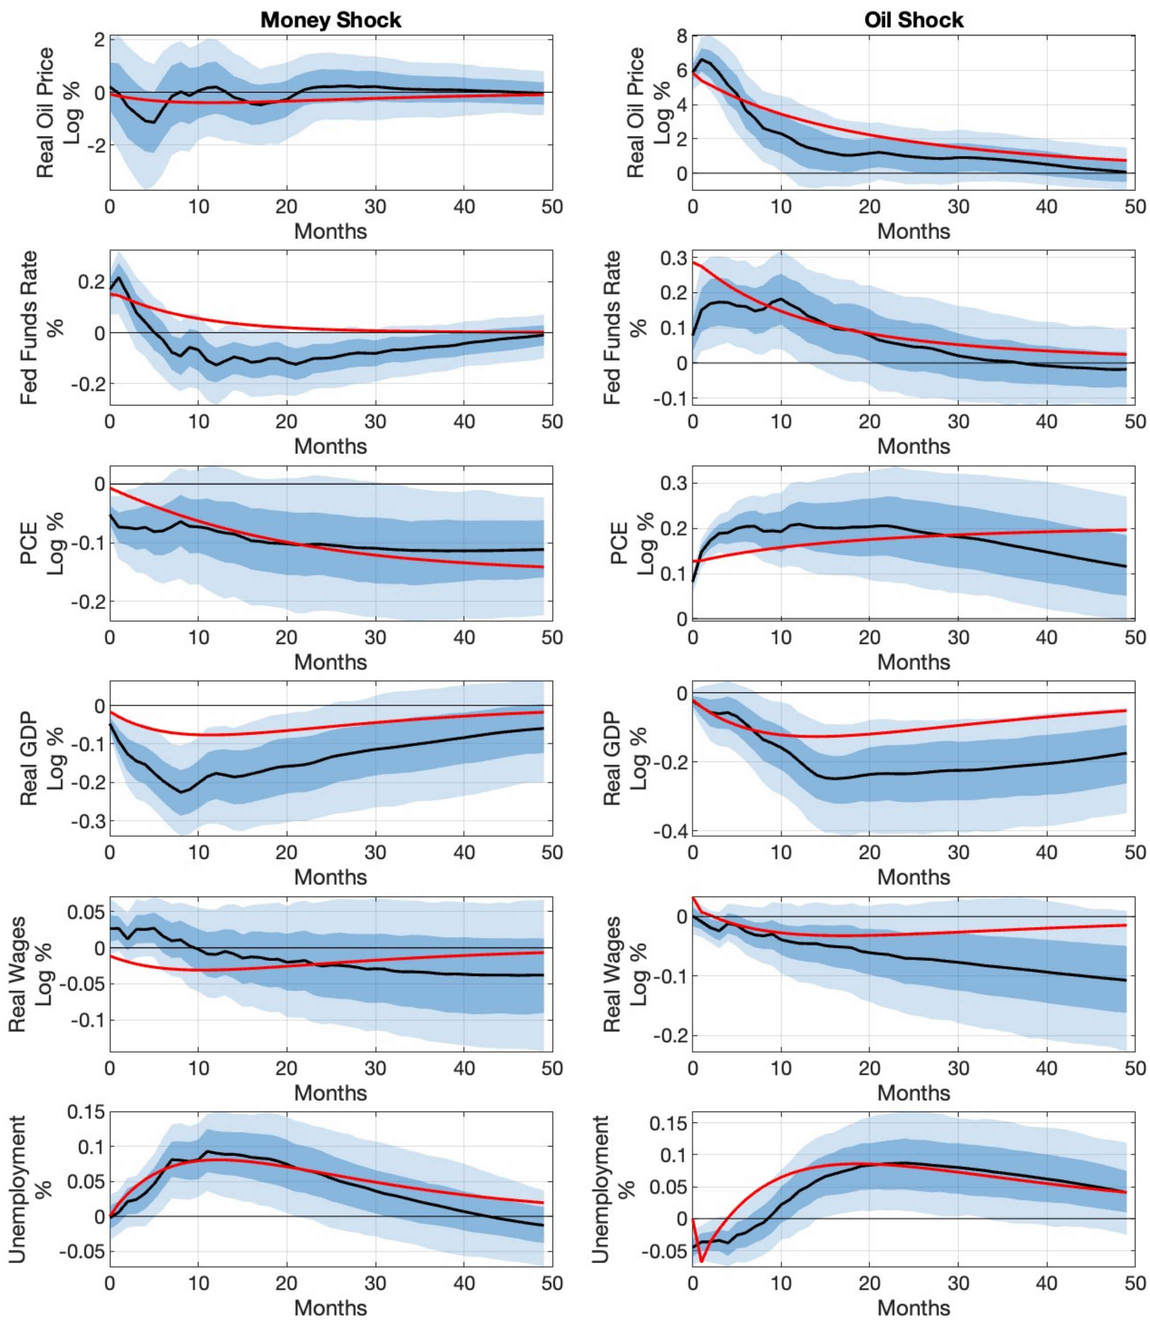

Figure 3: SVAR-based impulse responses for identified money and oil shocks vs model-based impulse responses (in red). The solid line is the point estimate and the dark and light-shaded areas are 68 and 95 percent confidence bands, respectively, computed using the wild bootstrap.

a decline in GDP of about 20 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 ten basis points. Real wages also decline slightly after ten months, 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 25 basis points and increases the unemployment rate by 10 points, while it increases the price level by about 20 points.[13] The proxy funds rate increases about 10-15 basis points on impact and persists above zero for 30 months, suggesting that the central bank reacts 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 conditionally 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 3. 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 Appendix A.

[13] 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.

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.[14] The steady-state ratio of firm to household expenditures on oil pins down the share of oil in households' expenditures χ at 2 percent.[15]

[14] 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).

[15] 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

Parameter	Θ_1	Value	Parameter	Θ_2	Estimate	S.E.
Discount factor	β	.998	F complementarity	ϵ	.370	.138
Demand curvature	η	4	H complementarity	ψ	.020	.301
Job survival	ρ	.96	Habit persistence†	h	.906 (.74)	.034
Matching elasticity	σ	.5	Price stickiness†	λ	.946 (.83)	.010
Bargaining power	ς	.5	Wage stickiness	γ	.705	.120
Outside option	b	.7	Taylor coefficient	ϕ_π	2.16	.754
SS oil/consumption	o/c_o	1.5	Interest smoothing	ρ^R	.063	.214
SS oil/output	o/y	.03	Money persistence	ρ^m	.946	.012
SS unemployment	u	.05	Oil persistence	ρ^o	.964	.012
			Normalization	σ^m	.023	.006
			Normalization	σ^o	.062	.027

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.

4.4.2 Estimated Parameters

Conditionally on the calibrated parameters, we then estimate eleven parameters that govern: complementarities with oil in production and consumption (ϵ and ψ), habit persistence (h), price rigidity (λ), wage rigidity (γ), the feedback coefficient on inflation in the Taylor rule (ϕ_π), the interest-rate smoothing parameter (ρ^R), and the persistences and volatilities of the money and oil shocks (ρ^m, ρ^o, σ^m, σ^o).

Table 1 presents the results. The estimates of $\epsilon = 0.37$ 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 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

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.

impulse responses. The point estimates and standard errors in Table 1 suggest we can soundly reject the null hypothesis of no complementarities.

The habit formation parameter h, which has a quarterly value of 0.74 (monthly 0.9) 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).[16] 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 thirty percent.

The feedback coefficient on the Taylor rule ϕ_π is estimated to be 2.16, which is in the range of estimates of the literature (see e.g. Carvalho et al. 2021).[17] The results also suggest that the persistence in the policy rate is not due to interest rate smoothing but rather to the persistence in the money shock. In particular, we estimate a low interest-rate smoothing parameter ($\rho^R = 0.06$) that is not statistically different from zero. The reason we find an insignificant smoothing parameter is that our estimation approach targets the sharp immediate response of the policy rate to the oil shock. With a high smoothing parameter, the policy rate would not jump sufficiently in response to the oil shock. Conversely, we estimate a high degree of persistence of the exogenous monetary shock ($\rho^m = 0.94$). As suggested by Bauer and Swanson (2021), the high persistence of the money shock likely reflects the occurrence of sustained departures from the normal Taylor rule, such as occurred in the wake of the pandemic, as Figure 1 illustrates. For robustness in section 6, we re-estimate the policy rule using an alternative approach that identifies a significant degree of interesting rate smoothing and shows that our main findings hold. Finally, the persistence parameter for oil ($\rho^o = 0.96$) matches the persistent effect of an oil shock on the oil price dynamics.

[16] 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. The estimates are also consistent with those in Hazell et al. (2022).

[17] As standard errors for this parameter are quite large, we performed substantial robustness regarding this value. In particular, results are robust to calibrating this parameter to lower numbers, which would lead to a larger contribution of monetary policy in the inflation surge. See section 6 for the results estimating the Taylor rule over the sample of the decomposition.

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, with emphasis on the contributions of oil shocks and accommodative monetary policy. 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 oil shock ε_{st}, the demand shock ε_{bt}, the shock to match efficiency $\varepsilon_{\Phi t}$, and the monetary policy shock ε_{rt}. The latter shock captures the impact of monetary policy accommodation.

We 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.[18] Priors are set to standard values. Results are reported in the 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

[18] For robustness, in section 6 we re-estimate the monetary policy rule over this sample.

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:[19]

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

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 high-frequency noise in oil prices only has a minor effect on the results.[20]

Since it is an untargeted variable, we can judge how well the model captures inflation by using the four shocks to construct model-implied series for year-over-year headline PCE inflation and core PCE inflation.

5.1 Historical Shock Decompositions

5.1.1 Targeted Variables

Figure 4 presents a historical decomposition for the four targeted variables over the sample 2010:01-2024:05.[21] Overall, the results are very sensible. The (non-monetary) demand shock accounts for most of the variation in unemployment. In this vein, the model treats the sharp rise in unemployment

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

[20]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.

[21]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.

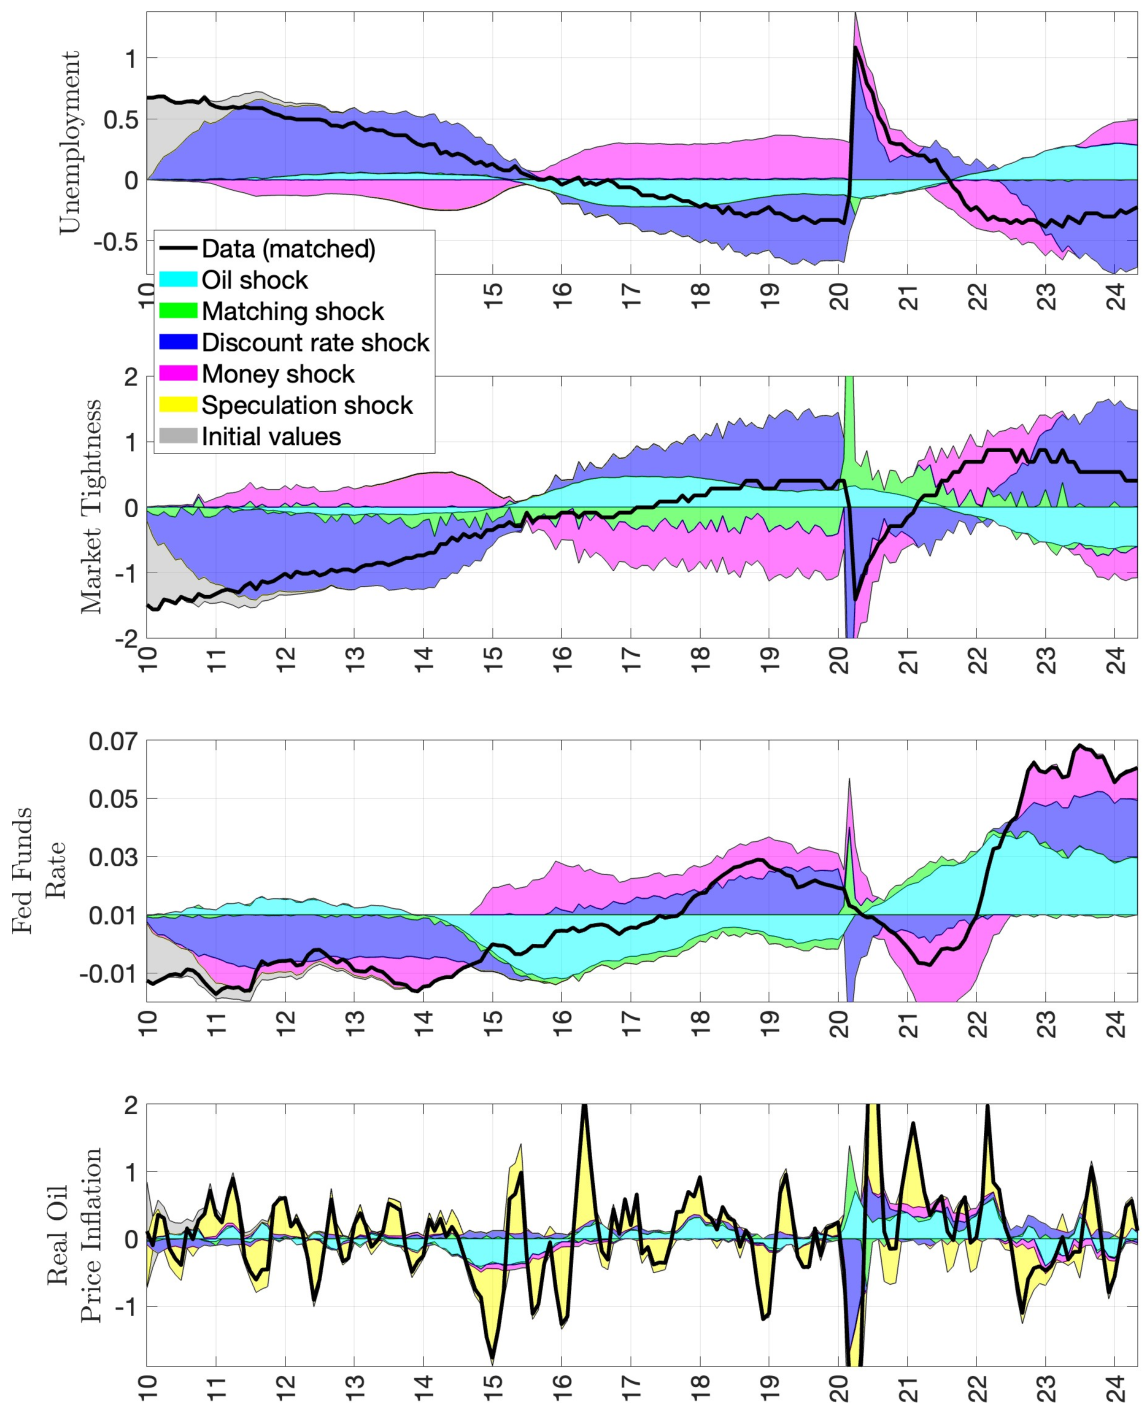

Figure 4: Historical shock decomposition of the targeted variables. Unemployment and labor market tightness are in log-deviations from the steady-state value for the model and log-deviations from the sample mean for the data. The decomposition for Fed funds is computed in deviations from steady state/sample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

during the pandemic as largely the product of a sharp drop in demand.[22] Unemployment then drops to steady state as demand improves. Interestingly, the drop in unemployment that continues in 2021-22 is in part the product of accommodative monetary policy. However, starting in early 2022 rapidly improving non-monetary demand helps push unemployment lower. Conversely, from mid 2020 onward, oil shocks contribute a roughly two percentage point increase in unemployment. Nonetheless, over this period, demand forces, including both monetary and non-monetary, more than offset the contractionary effect of oil supply shocks. At the very end of the sample, though, the post-pandemic tightening of monetary policy starts putting upward pressure on unemployment, partially offsetting the impact of rising demand.

Labor market tightness mirrors the behavior of unemployment: it is highly sensitive to the demand shock throughout the pandemic era. From mid 2021 to the end of 2022 accommodative monetary policy stimulates tightness, while oil shocks do the reverse. Interestingly, the matching shock is nontrivial during and after the pandemic but is not the leading driver of market tightness. Overall, the sharp increase in tightness over the pandemic is mainly due to demand (non-monetary and monetary) and supply (oil) factors. Note also that tightness does not contribute to unemployment variation over the sample.

For most of the sample, the demand shock plays an important role in driving the variation of the policy rate via its impact on inflation. During the surge, the oil shock became important, reflecting its contribution to the jump in inflationary pressure (which under the standard policy rule should produce an increase in the policy rate.) Because the central bank initially accommodated the oil shock, large deviations of the policy rate from the standard feedback rule emerged, inducing a period of "easy money" shocks. Beginning in early 2022, the demand shock becomes a significant driver of the Funds rate. Finally, because the monetary tightening at the end of the sample is greater than the policy rule would suggest, "tight money" shocks partly account for roughly 150 basis points of the policy rate in 2023 through early 2024.

[22] 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.

After filtering out background noise with the speculation shock (as described earlier), over the pandemic era period the oil shock mainly drives the behavior of the oil price.[23] One exception is that the demand shock that pushed the economy into the pandemic recession placed significant downward pressure on oil prices. Also, the nonmonetary demand shock and the accommodative monetary policy shock account jointly for roughly twenty-five to thirty percent of the rise in oil prices from mid 2021 through early 2022. Finally, we note that our smoothed oil price shock is reasonably correlated with the raw measure of oil price inflation over the pandemic era period.

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 core PCE and headline PCE inflation over the entire sample reasonably well. In particular, it generates a rise and fall in each series over the pandemic era that is broadly consistent with the data. In each case, as in the data, the model-predicted inflation surge begins in the spring of 2021, peaks early in 2022, and then moderates to roughly 3 percent by the spring of 2024.

The model, though, does not account for the entire rise in inflation. It explains roughly three quarters of the rise in core inflation. It undershoots this series by roughly a percentage point on average from spring 2021 through the end of 2022. For headline, the model is close except during 2022 when it falls short by nearly 2 percentage points. We note that the absence of food inflation in the model is likely an important factor in the undershooting of headline. As we discuss in section 6, some key factors left out of the model, including supply chain problems and fiscal policy are likely to have been relevant to this divergence between model and data.

We next use the model to perform a historical decomposition. For the runup in inflation, both the oil price shock and the monetary policy

[23] 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.

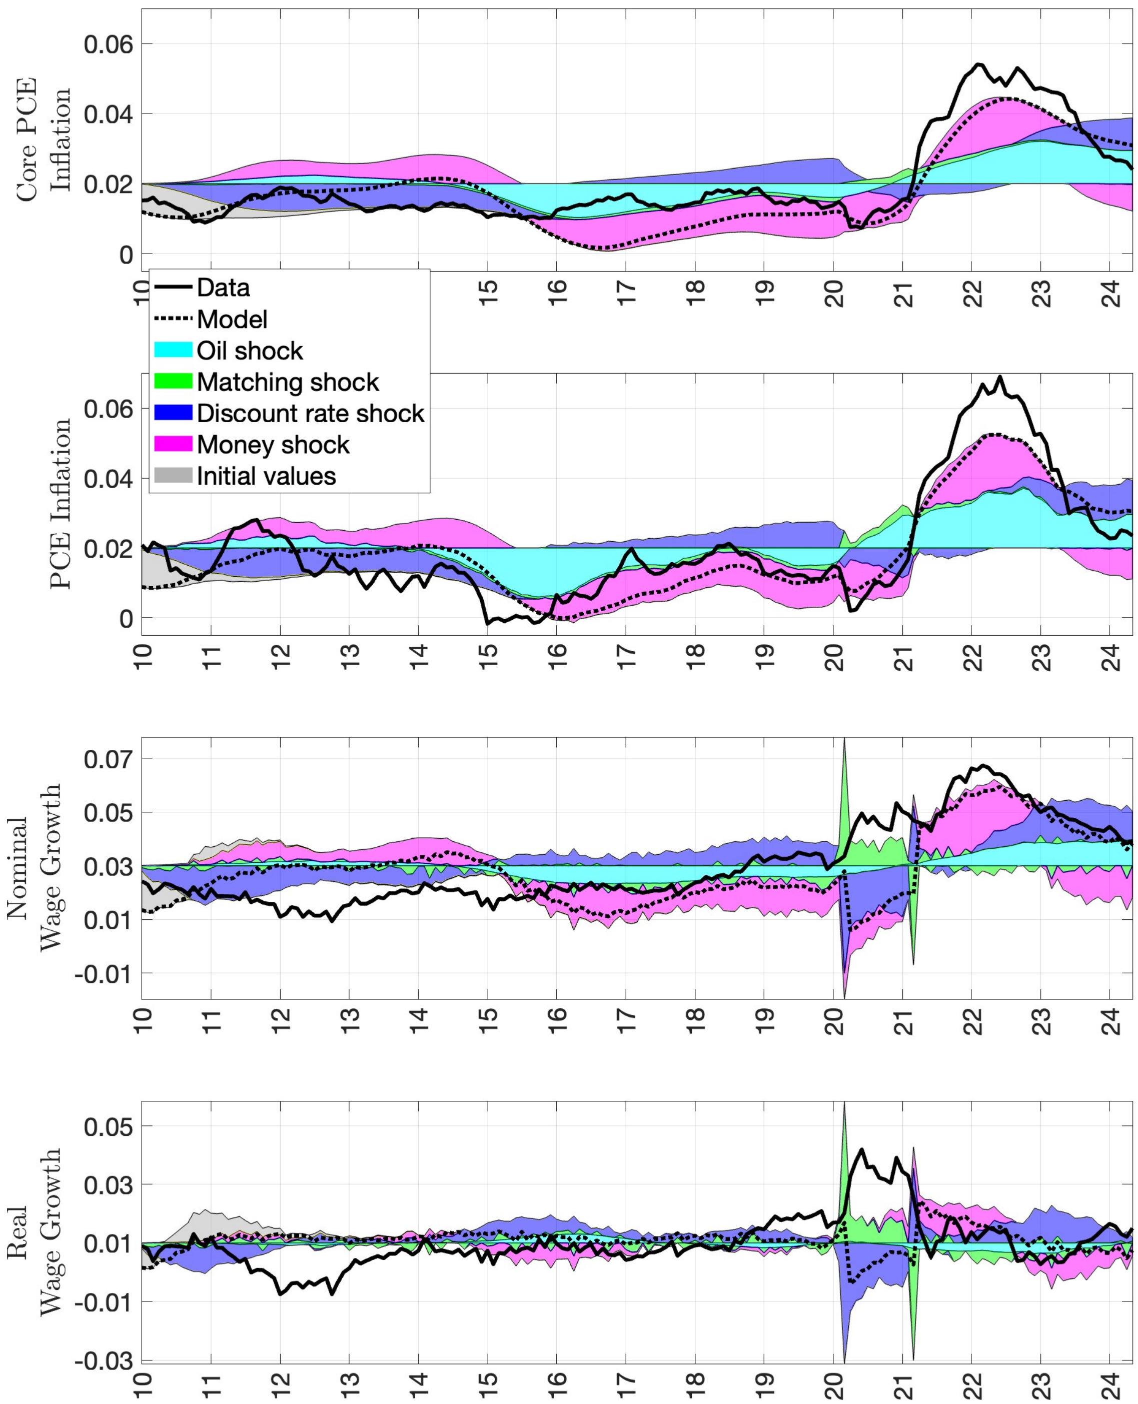

Figure 5: Historical shock decomposition of untargeted variables. The decomposition is computed in deviations from steady state/sample mean and then rescaled up by the sample mean. All the variables are annualized.

accommodation shock are important. Up to the peak in mid 2022, the shocks jointly account for roughly two and a quarter percentage points of the increase in core inflation, with each factor accounting for a roughly similar amount. The results are largely the same for headline, except the overall increase is a slightly greater amount, roughly two and a half percentage points. After this period the reversion in oil prices and the shift to an aggressively tight monetary policy contribute to the decline in inflation. Indeed from mid 2022 on, the tightening of monetary policy (relative to what a standard policy rule would predict) induces a roughly one-and-a-half percentage point decline in core inflation and a nearly two-percent decline in headline. These results on the role of oil and money shocks are consistent with the data portrayed in Figure 1. The significant and prolonged departure of the policy rate from the prepandemic policy rule is consistent with the large role of policy accommodation shocks portrayed in Figure 5.[24]

The (non-monetary) demand shock, which is a composite of demand influenced by private factors (e.g. the virus) and the component induced by fiscal stimulus, also contributes to inflation. Beginning in late 2022, it induced a roughly percentage point rise in inflation through the end of the sample, both for core and headline. That this shock did not affect inflation during 2021 does not imply that fiscal policy was irrelevant then, as we discuss in the next section. The remaining model disturbance, namely the shock to labor market tightness, had no effect on inflation. Except around the peak of the pandemic, virtually all the variation in tightness is endogenous.[25]

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

[24] Appendix C shows the case where the Federal Funds rate is the policy rate. The results are similar to the baseline except the contribution of policy accommodation to the inflation surge increases by roughly twenty-five percent, consistent with the fact that the Funds rate does not account for either forward guidance or QE/QT while the proxy rate does.

[25] We checked that our results are robust to increasing 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.

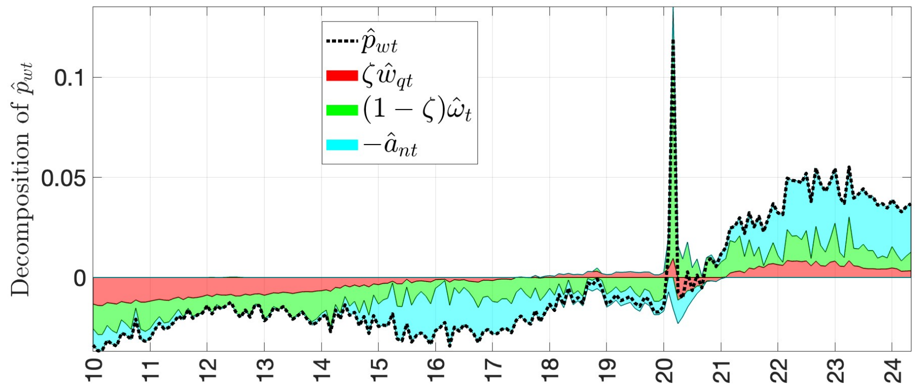

Figure 6: Historical decomposition of marginal cost into the main components from equation (21). Marginal cost is $\hat{p}_{wt}$, real wage is $\hat{w}_{qt}$ (multiplied by ζ), hiring cost is $\hat{\omega}_t$ (multiplied by $(1-\zeta)$), marginal product of labor is $\hat{a}_{nt}$ (multiplied by -1).

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. 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.

6 Robustness

As we noted earlier, we focus on the role of oil shocks and monetary policy accommodation because we have evidence of the effects of these shocks that we can use to discipline our model. In this section, we discuss how to think about the role of factors left outside the model, including supply chain issues and fiscal policy. We also discuss the robustness of our results to an alternative approach to the estimation of the monetary policy rule.

Supply Chain Disruptions. A number of authors have stressed the role of shortages induced by pandemic-related supply chain disruptions in the runup of

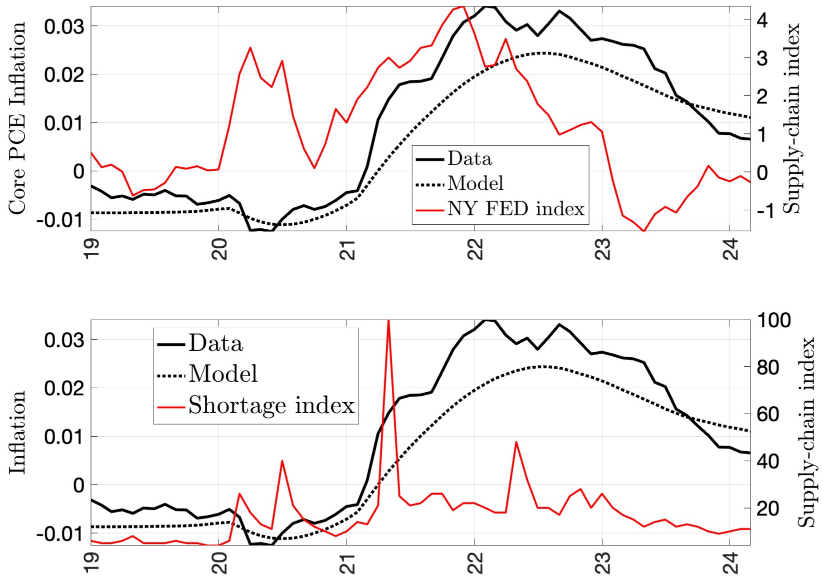

Figure 7: Inflation and the NY Fed supply-chain index.

inflation in 2021. Here we discuss the extent to which allowing for these factors may reconcile the discrepancy between our model and the data over this period.

In Figure 7, we plot two different indicators of supply chain disruptions against core PCE inflation for both the data and the model. The top panel presents the Global Supply Chain Index constructed by the Federal Reserve Bank of New York, while the bottom panel shows the shortage index used by Bernanke and Blanchard (2023), based on Google searches for the word shortage. Both indicators give a similar message, namely that supply chain issues were likely most relevant in the initial inflation runup between early and late 2021, a period where the discrepancy between our model and the data is greatest. In particular, over this sample, the model under-predicted inflation by roughly a percent. The implication is that supply factors could help reconcile the difference. Indeed the figure suggests that in percentage terms the peak effect of these supply factors would be in mid 2021, consistent with the evidence in Bernanke and Blanchard (2023). Under this interpretation, supply chain disruptions along with oil shocks and accommodative monetary policy would each account for roughly a third of the runup in inflation from early 2021 through early 2022.

Fiscal policy. Many authors have argued that fiscal policy has been a major determinant of the pandemic-era inflation. Our analysis does not rule out an important role for fiscal policy. Our non-monetary demand shock is a composite of private sector demand shocks induced by non-fiscal factors and shocks to fiscal stimulus. Thus, we cannot identify the "gross" contribution of fiscal policy, but only the net effect of our non-monetary demand shock.

However, as we noted earlier, from late 2022 to the end of the sample the non-monetary demand shock contributes roughly a percentage point to inflation. The impact of the demand shock on wage inflation is of similar magnitude though it begins earlier, in the spring of 2022. It is highly plausible that expansionary fiscal policy underlies the impact of the demand shock on inflation. Why did this effect not show up earlier, given that the lion's share of the fiscal stimulus was legislated in late 2021? Here it is useful to note that unemployment did not return to its pre-pandemic level until the beginning of 2022. The implication is prior to 2022, dampened demand due to the pandemic was offsetting the impact of the fiscal stimulus on inflation. Figure 5 suggests that the decline in unemployment following the peak in 2020 was mainly due to an increase in demand. It is reasonable to believe that fiscal stimulus played a large role. Indeed, our results are consistent with Smets and Wouters (2024), who model private sector demand and fiscal policy separately: they show that during 2021, the inflationary effect of fiscal policy exactly offset the deflationary effect from the drop in private sector demand, leading to a small net effect on inflation.

Monetary Policy Rule. In our baseline model we identified the monetary policy rule by fitting the coefficients of the rule to match the evidence on identified oil and monetary policy shocks. A virtue of this approach is that in evaluating the effects of these shocks, our model produces dynamic behavior that is consistent with the evidence. However, this estimation approach suggests an absence of interest rate smoothing for reasons we discussed in section 4. As a result, the persistence of the interest rate comes from exogenous shocks that exhibit a high degree of serial correlation. One might then wonder whether our results are robust to a rule that exhibits interest rate smoothing.

Previously we recovered the four macro shocks by using the model estimates based on the responses to the identified money and oil shocks. We now repeat the exercise except that in recovering the shocks we re-estimate the

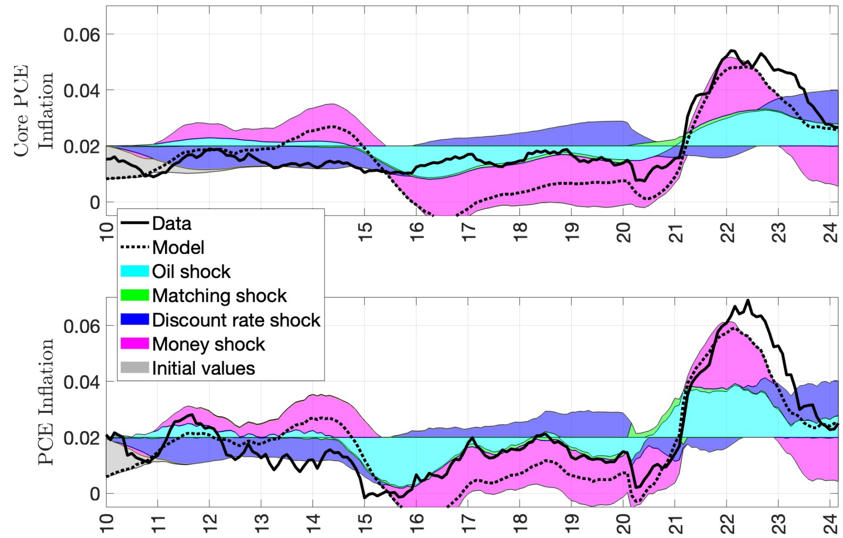

Figure 8: Historical shock decomposition of inflation under the estimated rule.

policy rule, leaving the estimates of the other equations in the model unchanged.

The results for the feedback coefficient on inflation, the interest rate smoothing, and the serial correlation of the money shock are given, respectively, by $\phi_\pi = 2.027$ (0.090), $\rho^R = 0.725$ (0.029), and $\rho^m = 0.952$ (0.004). Estimates of the feedback coefficient on inflation and the serial correlation of the money shock are nearly identical to what we had earlier. The smoothing parameter is now 0.725, as opposed to nearly zero. That the high serial correlation of the exogenous money shock remains present suggests that our finding of persistent departures from the policy rule holds even when allowing for interest smoothing.

Figure 8 shows the historical decomposition of inflation for this case. The model characterization of the inflation surge is broadly similar to the baseline presented in Figure 5. One difference is that, relative to the baseline, the model predicts slightly higher inflation from early 2021 through early 2022 and a bit lower inflation after that. A likely explanation is that the interest rate smoothing generates a stronger effect of money shocks on inflation, due to the enhanced persistent response of the policy. Hence, the accommodative monetary shocks during 2021 generate a larger rise in inflation relative to the baseline, and the tight money shocks after 2022 do the reverse. Finally, we note that the post-pandemic era was a period of sharp adjustments in the policy rate. Thus, it is reasonable to think that the central bank may have temporarily abandoned interest rate smoothing, which provides support for our baseline framework.

7 Concluding Remarks

Underlying the recent inflation surge was a perfect storm of supply and demand shocks. The response of monetary policy, however, was critical to shaping the effect of these shocks. We illustrate by developing and estimating a simple New Keynesian model with oil. We then show how oil and other candidate shocks contributed to the surge. In doing so, identify how policy accommodation influenced the jump in inflation along with how the subsequent shift to tightening impacted the decline.

On the methodological side, we pin down model 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. Doing so also makes our model well suited to both analyze the impact of oil shocks and the response of monetary policy. We then use the estimated framework to recover the model shocks without targeting inflation. Using these shocks, we find that the model-implied prediction tracks inflation reasonably well.

In line with the suggestive evidence discussed in the introduction, we find that both oil shocks and shocks to monetary policy accommodation were important contributors to the surge. Further, both the subsequent tightening of monetary policy and easing of oil prices contributed to the decline in inflation.

As for the other model shocks: the (non-monetary) demand shock, which is a composite of demand influenced by private factors (e.g. the virus) and of the component induced by fiscal stimulus, also contributes to inflation beginning in late 2022. This shock did not affect inflation during 2021, likely reflecting that expansionary fiscal policy during this period offset the drop in private sector demand due to the virus, consistent with Smets and Wouters (2024). The remaining model disturbance, the shock to labor market tightness, had no effect on inflation: except around the peak of the pandemic, virtually all the variation in tightness is endogenous. Finally, because we did not target inflation in the estimation, we leave room for the possibility that other shocks left outside the model, such as supply chain disruptions, could account for the discrepancy between the model and data.

While we have explored how shifts in monetary policy accommodation contributed to the rise and fall of inflation, we did not consider what the optimal policy response should be. Because the model exhibits a short-run tradeoff between unemployment and inflation due to the presence of real wage rigidity, some policy accommodation is likely appropriate. Understanding what the optimal policy response should be is on the agenda for future research.

References

Mary Amiti, Sebastian Heise, Fatih Karahan, Aysegul Sahin, et al. Pass-through of wages and import prices has increased in the post-covid period. 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. 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. 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. 2023.

Julian Di Giovanni, Sebnem Kalemli-Özcan, Alvaro Silva, and Muhammed A Yildirim. Global supply chain pressures, international trade, and inflation. 2022.

Julian Di Giovanni, Sebnem Kalemli-Özcan, Alvaro Silva, and Muhammed A Yildirim. Pandemic-era inflation drivers and global spillovers. 2023.

Christopher Erceg, Marcin Kolasa, Jesper Lindé, and Andreas Pescatori. Can energy subsidies help slay inflation? 2024.

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. 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.

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.

Oliver Pfäuti. The inflation attention threshold and inflation surges. *arXiv preprint arXiv:2308.09480*, 2023.

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

Frank Smets and Raf Wouters. Fiscal backing, inflation and us business cycles, 2024.

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*, (353):1, 2005.

Online Appendix

A Estimation by Matching Impulse Responses

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 Results

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	.298	.231	.352
ρ^Φ	beta	.6	.1	.515	.410	.616
σ^b	invg	.15	.15	.0577	.052	.061
σ^Φ	invg	.15	.15	.157	.143	.173
σ^m	invg	.15	.15	.232	.205	.254
σ^o	invg	.15	.15	.052	.045	.061
σ^m	invg	.15	.15	.042	.035	.046

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 ($\rho^\Phi = .51$ at the posterior mean) then the discount shock ($\rho^b = .29$ at the posterior mean). The estimates of the standard deviations are sensible, with the posterior means of the standard deviation for oil $\sigma^o = .05$ and money shock $\sigma^m = .04$ 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 $\sigma^m = .23$ 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 $\sigma^b = .057$ are larger than both oil and money (because of the lower persistence), but of the same order of magnitude.

C Results with the Fed Funds Rate

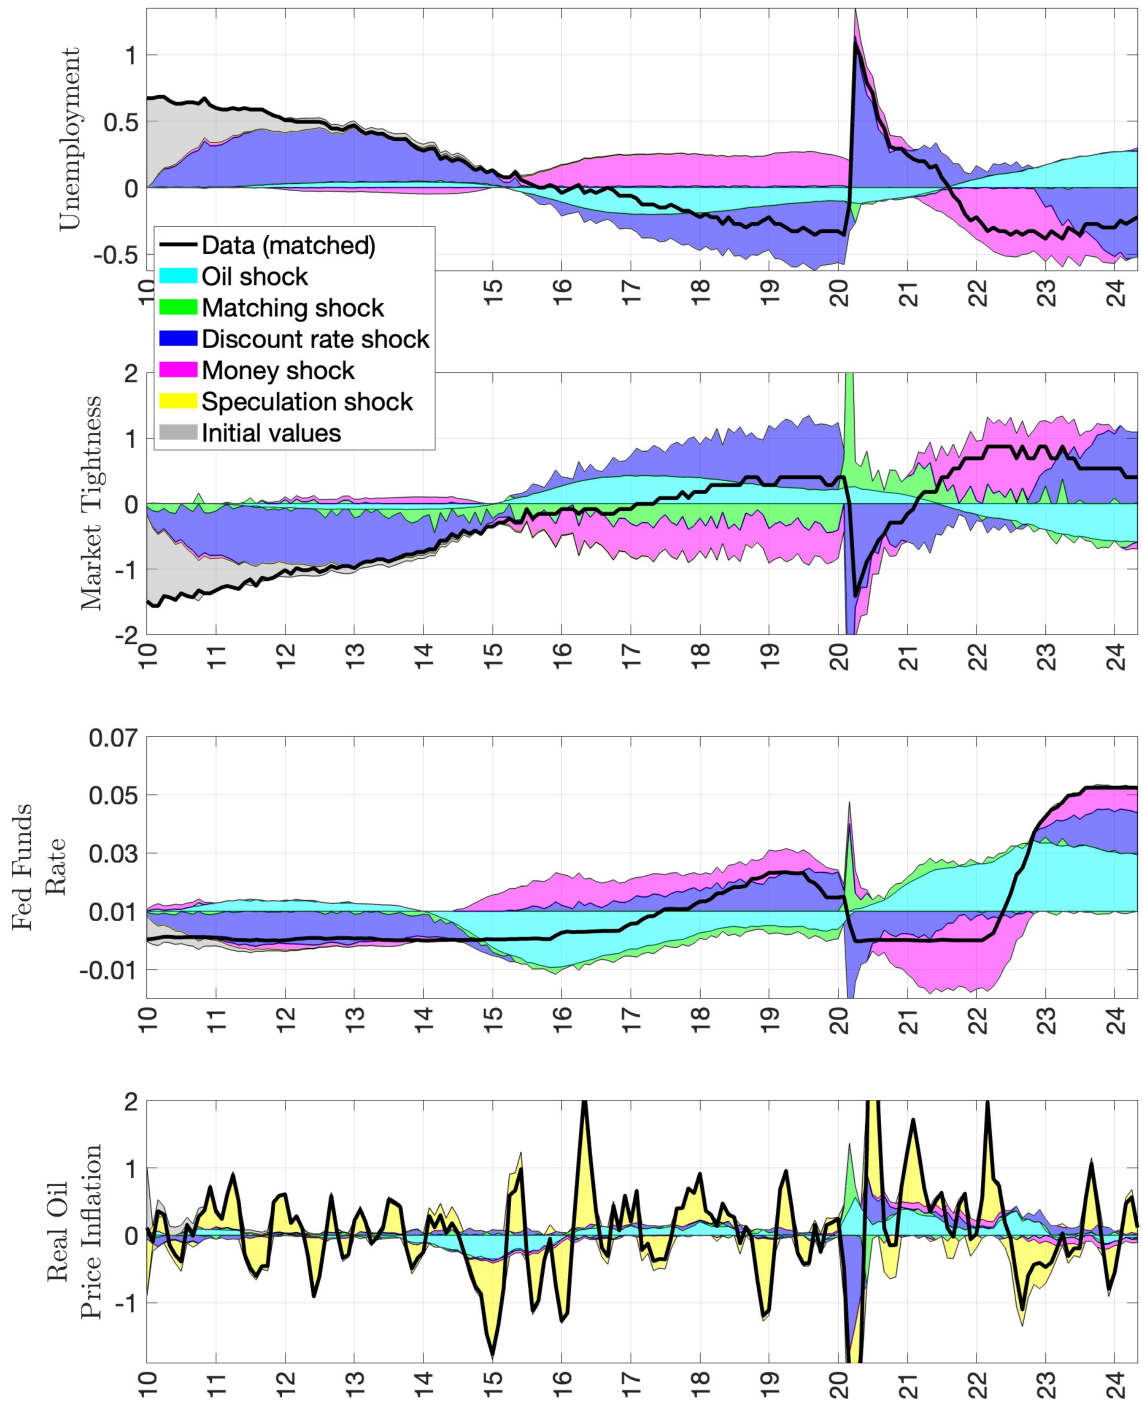

Figure 9: Historical shock decomposition of the targeted variables using the Fed Funds rate.

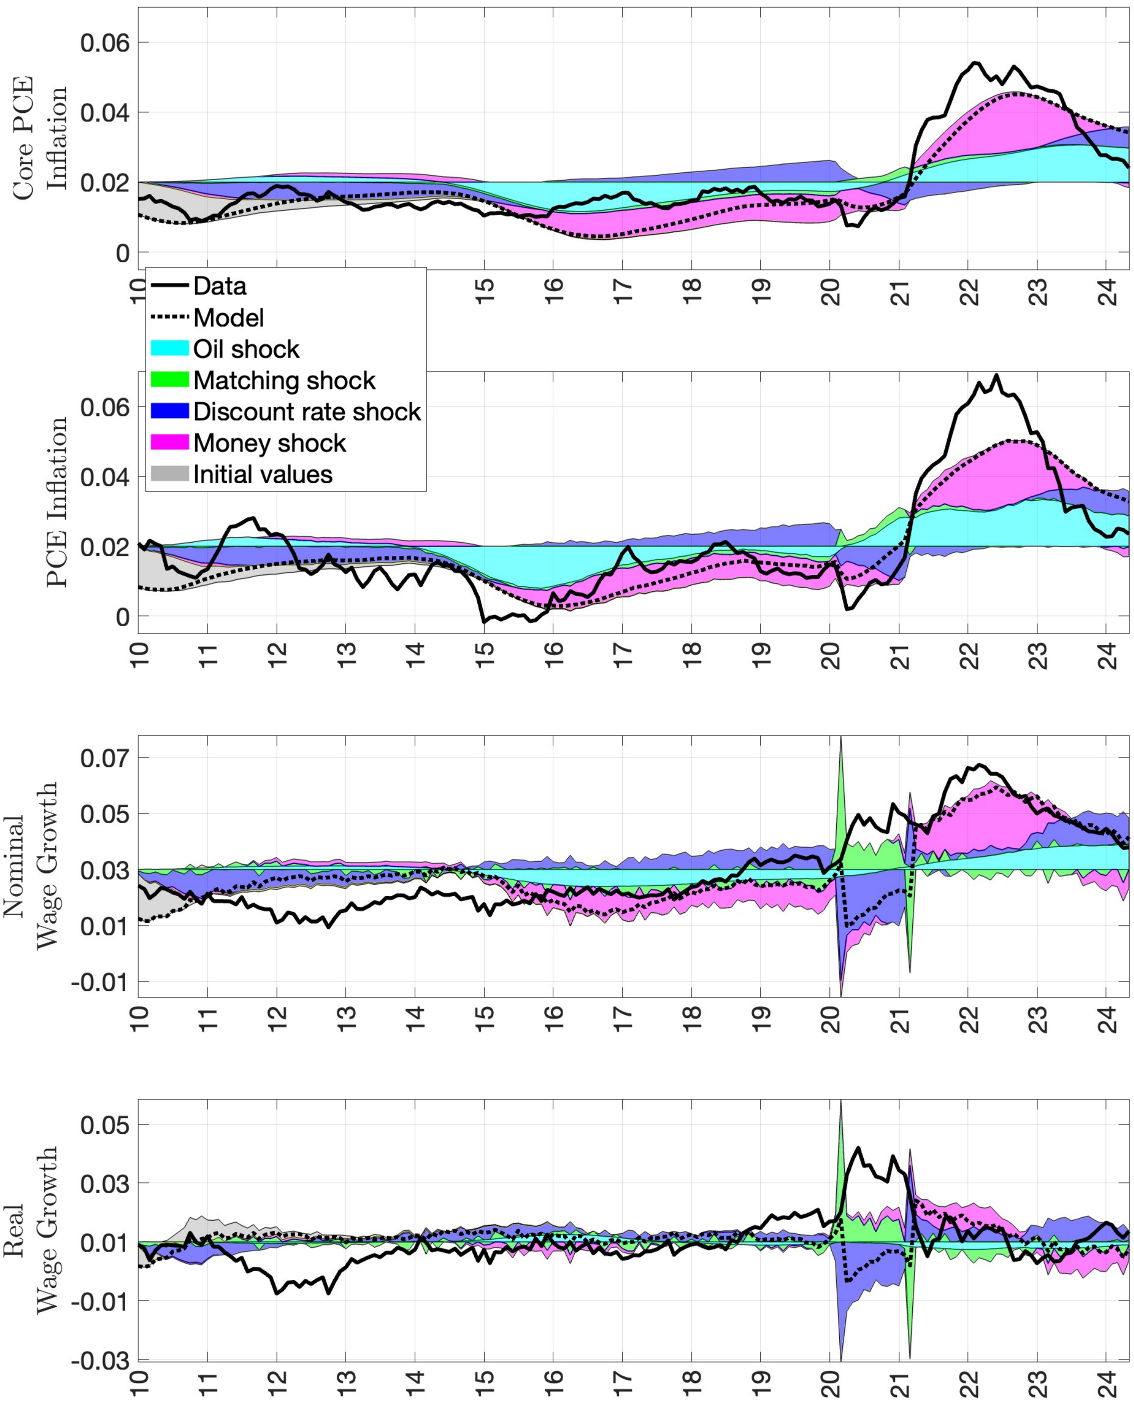

Figure 10: Historical shock decomposition of the untargeted variables using the Fed Funds rate.

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.