The Systematic Origins of Monetary Policy Shocks*

Lukas Hack Klodiana Istrefi Matthias Meier

This version: February 2025
First version: May 2024

Abstract

Conventional strategies to identify monetary policy shocks rest on the implicit assumption that systematic monetary policy is time-invariant. In an environment with time-varying systematic monetary policy, we formally show that these strategies yield shocks that are contaminated, leading to bias in estimated impulse responses. In line with our theoretical results, we empirically show that conventional monetary policy shocks are predictable by measured fluctuations in systematic monetary policy. We propose new shocks that are purged of this predictability. Our preferred new shocks show that U.S. monetary policy affects inflation and output more strongly and faster compared to the corresponding conventional shocks.

Keywords: Monetary policy shocks, systematic monetary policy, identification.
JEL Codes: E32, E43, E52, E58.

*We thank Klaus Adam, Christian Bayer, Thomas Drechsel, Luca Gambetti, Refet Gürkaynak, Carolin Pflueger, Giovanni Ricco, Fabio Stohler (discussant), Eric Swanson, and Christian Wolf, as well as numerous participants at various seminars and conferences, for insightful comments and suggestions. Lukas Hack and Matthias Meier acknowledge support from the German Research Foundation (DFG) through CRC TR 224 (Project C05). Matthias Meier acknowledges support from the UniCredit & Universities Foundation. The views expressed in this article are solely the responsibility of the authors, and should not be interpreted as reflecting the views of the Banque de France or the Eurosystem. Lukas Hack (ETH Zürich): hack@kof.ethz.ch. Klodiana Istrefi (Banque de France and CEPR): klodiana.istrefi@banque-france.fr. Matthias Meier (University of Mannheim): m.meier@uni-mannheim.de.

1 Introduction

Empirical monetary policy shock series form the backbone of a large literature in monetary economics. The estimated responses to these shocks are used to assess the effectiveness of monetary policy, construct policy counterfactuals, study the optimality of monetary policy, estimate structural macroeconomic equations, and estimate DSGE models.[1] These applications require empirical monetary policy shocks that are well identified, meaning they capture unpredictable and exogenous changes in monetary policy that are orthogonal to other macroeconomic shocks.

The central point of this paper is that fluctuations in systematic monetary policy pose a challenge to conventional strategies for identifying monetary policy shocks. The fundamental problem of conventional identification strategies is the implicit assumption that systematic monetary policy is constant across time. Any time variation in systematic monetary policy will be contained in conventional empirical monetary policy shocks, consistent with common views about these shocks:[2]

> *We do not have many good economic theories for what a structural monetary policy shock should be. Other than "random coin flipping," the most frequently discussed source of monetary policy shocks is shifts in central bank preferences, caused by changing weights on inflation vs unemployment in the loss function or by a change in the political power of individuals on the FOMC.* Ramey (2016, Handbook of Macroeconomics, Vol. 2A, p.89)

This paper makes a theoretical and empirical contribution to the identification of monetary policy shocks. Our theoretical contribution is to formally show that conventional empirical identification strategies do not isolate monetary policy shocks in an environment with time-varying systematic monetary policy. Instead, they are contaminated by other macroeconomic shocks, leading to bias in estimated impulse response functions. Empirically, we revisit monetary policy shocks (i) estimated in the seminal Romer and Romer (2004), (ii) the refinement of Aruoba and Drechsel (2022), and (iii) the shocks based on high-frequency identified monetary policy surprises in Miranda-Agrippino and Ricco (2021). The empirical contribution is threefold. First, we show that the three types of monetary policy shocks are predictable by a measure of time-varying systematic monetary policy. Second, we propose new monetary policy shock series that are orthogonal to measured fluctuations in systematic

[1] See, e.g., Romer and Romer (1989); Bernanke and Blinder (1992); Bernanke, Gertler, and Watson (1997); Christiano, Eichenbaum, and Evans (1999); Romer and Romer (2004); Christiano, Eichenbaum, and Evans (2005); Gertler and Karadi (2015); Barnichon and Mesters (2020, 2023); McKay and Wolf (2023).

[2] Similarly, in the first Handbook of Macroeconomics, Christiano et al. (1999, p.71-72) argue that an empirical monetary policy shock *[..] reflects exogenous shocks to the preferences of the monetary authority, perhaps due to stochastic shifts in the relative weight given to unemployment and inflation. These shifts could reflect shocks to the preferences of the members of the Federal Open Market Committee (FOMC), or to the weights by which their views are aggregated.*

monetary policy. Third, we find that inflation and output respond significantly more strongly and quickly to the new shocks of type (i) and (ii) compared to the corresponding original shocks. For the third type of shock, we also uncover significant differences in responses between the original and the new shock.

Our theoretical analysis builds on monetary policy following a general type of time-varying policy rule. The rule determines a policy instrument as a function of inputs to the rule, e.g., inflation and output, time-varying slope coefficients describing the policy response to macroeconomic conditions, i.e., systematic monetary policy, and a monetary policy shock.[3] In contrast, many conventional empirical identification strategies implicitly assume a policy rule with time-invariant slope coefficients. Empirical monetary policy shocks are estimated as deviations from such rule. Identification strategies following this approach include Taylor rule-type regressions (e.g., Romer and Romer, 2004) and linear monetary VAR models using exclusion restrictions (e.g., Christiano et al., 1999), sign restrictions (e.g., Uhlig, 2005), narrative restrictions (e.g., Antolín-Díaz and Rubio-Ramírez, 2018), or external instruments (e.g., Gertler and Karadi, 2015).

Against the backdrop of a time-varying Taylor rule, we show that the conventionally estimated monetary policy shock contains the (actual) monetary policy shock but also time variation in systematic monetary policy interacted with the inputs to the policy rule. To the extent that other macroeconomic shocks (present and past) affect the inputs, the empirical monetary policy shock is contaminated by these macroeconomic shocks.

We formally show that contaminated shocks do not identify the causal effects of (actual) monetary policy shocks, but lead to biased impulse responses. We characterize three sources of bias reflecting endogeneity and attenuation. The estimated impulse response function remains biased even if time variation in systematic monetary policy is exogenous, i.e., if time variation in the slope coefficients of the Taylor rule is independent of other macroeconomic shocks. The bias also remains if the inputs to the policy rule are predetermined. The bias only disappears under strong assumptions, such as assuming monetary policy is fully exogenous.

Our theoretical insights similarly apply to monetary policy shocks constructed from high-frequency monetary policy surprises (e.g., Nakamura and Steinsson, 2018). Identification rests on the implicit assumption that systematic monetary policy, as perceived by financial market participants, is constant in a time window around monetary announcements. Otherwise, the shocks are contaminated and lead to biased estimates. Bauer and Swanson (2023a) provide evidence consistent with such high-frequency belief changes. We further show that

[3] For evidence on fluctuations in the coefficients of the policy rule, see, e.g., Clarida, Gali, and Gertler (2000), Orphanides (2004), Bordo and Istrefi (2023), and Hack, Istrefi, and Meier (2023).

regressing high-frequency monetary surprise on publicly available macroeconomic forecasts (Bauer and Swanson, 2023b) or Greenbook forecasts (Miranda-Agrippino and Ricco, 2021) does generally not resolve the contamination problem.

While previous work has noted that time-varying systematic monetary policy may complicate the identification of monetary policy shocks (e.g., Coibion, 2012; Bauer and Swanson, 2023a; McMahon and Munday, 2023), our paper is the first to formally characterize (i) how time-varying systematic monetary policy leads to contamination in the monetary policy shocks obtained from a broad set of conventional empirical identification strategies, and (ii) how contamination leads to biased impulse response estimates. We further go beyond previous work by providing new empirical evidence on shock contamination and a new identification strategy that tackles this problem.

Our empirical analysis starts from the testable prediction of the theory that conventional monetary policy shocks are predictable by time variation in systematic monetary policy interacted with the inputs to the policy rule. We measure time variation in systematic U.S. monetary policy through the historical composition of hawks and doves in the Federal Reserve's Federal Open Market Committee (FOMC). This composition builds on the narrative classification of FOMC members by Istrefi (2019). Hawks are more concerned about inflation. Doves are more concerned about supporting employment and growth.[4] We consider two measures of systematic monetary policy, the Hawk-Dove balance across all voting FOMC members and the balance across the four FOMC members currently with voting rights through the annual rotation. The former is a more comprehensive measure, whereas the latter primarily reflects exogenous variation through the rotation.

We test our prediction using empirical monetary policy shocks as estimated in Romer and Romer (2004), hereafter RR.[5] We regress the RR shock on the Taylor rule inputs considered by RR, notably Greenbook forecasts for various macroeconomic variables and horizons, interacted with measured fluctuations in systematic monetary policy. We consider the original RR sample 1969-1996, the extended Wieland and Yang (2020) sample 1969-2007, and the post-Volcker disinflation sample 1983-2007. The regression explains between 10 and 54% of the variance of RR shocks depending on sample and regressors (contemporaneous or lagged). Using the regressors lagged by one FOMC meeting yields the highest R^2, ranging between 0.33 and 0.54. Overall, our evidence strongly suggests that RR shocks are contaminated by

[4] Istrefi (2019) shows that these preferences match with narratives on monetary policy, preferred interest rates, dissents, and forecasts of FOMC members. Bordo and Istrefi (2023) study the origins of these preferences, linking them to early-life experiences and education. Hack et al. (2023) use the Hawk-Dove classification to study the effects of systematic monetary policy on the propagation of macroeconomic shocks.

[5] The RR identification strategy has been applied to the U.K. (Cloyne and Hürtgen, 2016), Germany (Cloyne, Hürtgen, and Taylor, 2022), Norway (Holm, Paul, and Tischbirek, 2021), Canada (Champagne and Sekkel, 2018) and many other countries (Choi, Willems, and Yoo, 2024).

fluctuations in systematic monetary policy.

The empirical evidence motivates us to construct new series of monetary policy shocks that are not predictable by fluctuations in measured systematic monetary policy. We estimate an extension of the Taylor rule regression in RR that includes the interaction of the Hawk-Dove balance with the Taylor rule inputs. The correlation between the original RR shock and our new shock is 0.67. The sign-correlation between the two series is lower, meaning many shocks flip sign. The distribution of new shocks is less dispersed, with a standard deviation of 0.23, compared to 0.34 for the RR shock.

Finally, we compare impulse responses between our new monetary policy shock and the RR shock. We focus on the post-Volcker disinflation sample 1983-2007 because the estimated responses to many conventional monetary policy shock series appear puzzling in this sample (e.g., Ramey, 2016).[6] For comparability, we normalize the size of both shocks to the same peak increase of the FFR. The dynamic FFR response to our new shock is less persistent. In contrast, the decline in GDP and inflation is substantially larger for the new shock. The trough GDP response is about twice as large for the new shock compared to the RR shock. The differences between the responses to the two shocks are statistically significant at the five percent level for many horizons. Importantly, the RR shock seems to operate with a long lag, not affecting inflation up until two years after the shock. The GDP response is broadly insignificant. In contrast, inflation and GDP respond to our new shock with a lag of one year. Beyond the first year, the responses of inflation and GDP are significantly different from zero at the 5% level.[7] Our findings suggest that the puzzling effects of RR shocks in the 1983-2007 sample reflect contamination from time-varying systematic monetary policy.

We further revise the refined RR shocks in Aruoba and Drechsel (2022), who use textual analysis to create sentiment indicators about the Fed staff's assessment of the economy to better capture the Fed's information set about the state of the economy. Measured systematic monetary policy also has predictive power for the refined RR shocks in Aruoba and Drechsel (2022). In addition, orthogonalizing the Aruoba and Drechsel (2022) shock with respect to measured systematic monetary policy leads to similar differences in the estimated responses compared to the RR shock. Finally, we also provide evidence on contamination and bias in the impulse responses to the monetary policy shocks estimated in a proxy VAR with high-frequency monetary policy surprises as external instruments in Miranda-Agrippino and Ricco (2021). A limitation of our revision of the Miranda-Agrippino and Ricco (2021) shocks is that we cannot clean the underlying instrument, the high-frequency monetary policy

[6] Relatedly, Barakchian and Crowe (2013) show that a variety of conventional monetary policy shock series raise GDP when raising the federal funds rate in a post-1988 sample.

[7] In the 1969-2007 sample, we also find that output and inflation respond more strongly to the new shock, albeit with a sluggish inflation response.

surprises, from potential contamination arising from expectation revisions about systematic monetary policy. The reason is that we only have a low-frequency measure of time variation in systematic monetary policy but no high-frequency measure of expectation revisions around monetary announcements. Therefore, our preferred shocks in this paper are the new versions of the Romer and Romer (2004) and the Aruoba and Drechsel (2022) shocks.

Our paper highlights the importance of accounting for the time-varying nature of systematic monetary policy when identifying monetary policy shocks. An alternative approach addresses time-varying systematic monetary policy by modeling it as latent variable or time-varying coefficients, see, for example, regime-switching models (e.g., Owyang and Ramey, 2004; Sims and Zha, 2006), time-varying coefficient monetary VAR models (e.g., Primiceri, 2005), and Taylor rules with time-varying coefficients (e.g., Boivin, 2006; Coibion, 2012; Bauer, Pflueger, and Sunderam, 2022). Particularly related is Coibion (2012) who uses the latter approach to estimate a monetary policy shock series. The estimated shock is highly correlated with the RR shock and yields similar impulse responses as the RR shock. The difference between this finding and ours might reflect the challenge of time-varying coefficient models to identify genuine time variation in the parameters of interest while avoiding overfitting.

2 Identification challenge in theory

In this section, we study the identification of monetary policy shocks in an environment with time-varying systematic monetary policy. We show that a wide spectrum of conventional identification strategies yield monetary policy shocks that are contaminated by other macroeconomic shocks. The contaminated shocks lead to bias in estimated impulse response functions unless we impose strong assumptions on monetary policy.

2.1 Time-varying systematic monetary policy

We depart from the common assumption that systematic monetary policy is constant across time, and assume monetary policy follows the time-varying Taylor rule

$$i_t = \alpha + (\phi + \tilde{\phi}_t)\, x_t + w_t^m, \quad \mathsf{E}[\tilde{\phi}_t] = \mathsf{E}[x_t] = \mathsf{E}[w_t^m] = \mathsf{E}[\tilde{\phi}_t w_t^m] = 0, \tag{2.1}$$

where i_t R is a policy instrument, x_t $\mathsf{R}^{n\times 1}$ are the n inputs of the policy rule, e.g., present and lagged inflation and GDP (forecasts), $\tilde{\phi}_t$ $\mathsf{R}^{n\times 1}$ is a vector of time-varying (slope) coefficients describing fluctuations in systematic monetary policy, with ϕ $\mathsf{R}^{n\times 1}$ the average coefficient vector, and w_t^m denotes a random monetary policy (intercept) shock.

We assume the inputs in x_t are mean zero and set $\alpha = -\mathsf{E}[\tilde{\phi}_t x_t]$, which simplifies some subsequent derivations but is not critical for our results.[8]

Time variation in the coefficients of the rule $\tilde{\phi}_t$ may be driven by changes in the preferences of central bankers. Preference changes can occur for exogenous reasons, e.g., the FOMC rotation of voting rights (Hack et al., 2023), or for endogenous reasons, e.g., monetary policy may become more responsive to inflation when inflation is high (Davig and Leeper, 2008). Our main results hold even if $\tilde{\phi}_t$ fluctuates only for exogenous reasons. Finally, we assume that ϕ_t does not co-move with monetary policy shocks, $\mathsf{E}\left[\phi_t w_t^m\right] = 0$, which allows for a sharp conceptual distinction between systematic monetary policy and monetary policy shocks.

2.2 Contamination under conventional shock identification

In this section, we show that the presence of fluctuations in systematic monetary policy poses a challenge for conventional strategies to identify monetary policy shocks. Conventionally identified shocks are contaminated, they do not isolate the monetary policy shock.

We consider as conventional identification strategies (a) Taylor rule-type regressions (e.g., Romer and Romer, 2004), (b) linear structural vector-autoregressive (SVAR) models identified using *inter alia* exclusion restrictions (e.g., Christiano et al., 1999), sign restrictions (e.g., Uhlig, 2005), narrative restrictions (e.g., Antolín-Díaz and Rubio-Ramírez, 2018), or external instruments (e.g., Gertler and Karadi, 2015), and (c) using high-frequency monetary policy surprises directly as shock series (e.g., Nakamura and Steinsson, 2018).

Identification strategies of type (a) and (b) both estimate monetary policy shocks as residual from a time-invariant Taylor rule-type model

$$i_t = b\, x_t + e_t^m, \tag{2.2}$$

where the estimated residual, $\hat{e}_t^m$, is an empirical monetary policy shock. This is a broad description of a wide variety of identification strategies which differ mainly in how the coefficients in equation (2.2), and thus the residual, are estimated. In particular, SVAR models to identify monetary policy shocks contain an equation consistent with equation (2.2) irrespective of identifying assumptions and estimation method.[9]

[8] A richer formulation of (2.1) may contain time-varying target variables, e.g., $i_t = \alpha + (\phi + \tilde{\phi}_t)\,(x_t - x_t) + w_t^m$, where $x_t \quad \mathsf{R}^{n\times 1}$ is the target, e.g., the inflation target. Shocks to the target generate a third type of monetary policy shock, the effect of which is correlated with fluctuations in systematic monetary policy. Throughout this paper, we abstract from fluctuations in target variables.

[9] A SVAR model is defined by $B(L)Y_t = W_t$, where Y_t is a vector of variables, $B(L)$ a lag polynomial, and W_t a vector of structural shocks. Let Y_t include the policy instrument i_t, and W_t include a monetary policy shock, *wlog* the first element of W_t. Then, the first equation of the SVAR model is a monetary policy rule, identical with equation (2.2) given a corresponding specification of Y_t.

Against the backdrop of the time-varying monetary policy rule in (2.1), the time-invariant regression in (2.2) is misspecified, leading to contamination in the estimated monetary policy shock. The following proposition formally characterizes the estimated empirical shocks for a given estimate $\hat{b}$.

Proposition 1 (**Shock contamination**)**.** *Let monetary policy follow* (2.1)*. Given an estimate $\hat{b}$, the associated estimate of the monetary policy shock $\hat{e}_t^m$ in* (2.2) *is given by*

$$\hat{e}_t^m = w_t^m + \omega_t^{\hat{b}} + \omega_t^{\tilde{\phi}},$$

with the two wedges defined by

$$\omega_t^{\hat{b}} = (\phi - \hat{b})\, x_t, \quad \textit{and} \quad \omega_t^{\tilde{\phi}} = \tilde{\phi}_t x_t - \mathsf{E}[\tilde{\phi}_t x_t].$$

The proof is straightforward when combining (2.1) and (2.2). The proposition characterizes two wedges between the actual monetary policy shock w_t^m and the estimated shock $\hat{e}_t^m$.
The first wedge, $\omega_t^{\hat{b}}$, arises whenever the estimate $\hat{b}$ does not equal the average policy coefficient vector ϕ. This wedge may be present even in the absence of time-variation in systematic monetary policy $\tilde{\phi}_t = 0$. For example, if b is estimated via OLS, a well-known endogeneity bias arises if the monetary policy shock correlates with x_t (Cochrane, 2011; Carvalho, Nechio, and Tristão, 2021). The presence of time-varying systematic monetary policy generates an additional type of endogeneity bias. Formally, the OLS estimate $\hat{b}$ of the regression model (2.2) satisfies $\hat{b} \xrightarrow{p} \phi + \mathsf{E}\left[x_t x_t\right]^{-1} \mathsf{E}\left[x_t w_t^m\right] + \mathsf{E}\left[x_t x_t\right]^{-1} \mathsf{E}\left[x_t x_t \tilde{\phi}_t\right]$. Hence, the bias remains even if $\mathsf{E}\left[x_t w_t^m\right] = 0$.[10] Whatever the method by which (2.2) is estimated, if $\hat{b} = \phi$ the estimated monetary policy shock $\hat{e}_t^m$ will correlate with x_t.
The second wedge captures the misspecification of (2.2) in the presence of fluctuations in systematic monetary policy. Fluctuations in $\tilde{\phi}_t$ interacted with x_t are therefore captured by the regression residual. The wedge disappears if we assume away fluctuations in systematic monetary policy $\tilde{\phi}_t = 0 \quad t$. If systematic monetary policy fluctuates, the wedge is present regardless of the estimate $\hat{b}$. Even if $\hat{b} = \phi$, the estimated monetary policy shock is still contaminated by $\tilde{\phi}_t x_t$. In the next two subsections, we impose additional structure to study the nature of contamination and its implications.

Finally, we study contamination for conventional identification strategy (c). High-frequency monetary policy surprises are constructed as changes in interest rate futures (or swaps) in a narrow window around monetary announcements. The idea is that they capture changes

[10] In a New Keynesian model with time-varying systematic monetary policy, which we study in Section 2.3, it generally holds that $\mathsf{E}\left[x_t x_t \tilde{\phi}_t\right]$ is non-zero.

in expectations about monetary policy. For simplicity, we assume that financial markets precisely captures changes in expectations around monetary announcements. Formally, the high-frequency identified monetary policy surprise then corresponds to $\hat{e}_t^m = \mathsf{E}_{t+\Delta}[i_{t+\tau}] - \mathsf{E}_{t-\Delta}[i_{t+\tau}]$, where $\mathsf{E}_{t+\Delta}[i_{t+\tau}]$ denotes the period $t+\Delta$ expectation of period $t+\tau$ interest rates as measured by the price of a future (or swap) contract. If monetary policy follows (2.1), the monetary policy surprise for $\tau = 0$ is given by

$$\hat{e}_t^m = w_t^m + \phi \left(\mathsf{E}_{t+\Delta}[x_t] - \mathsf{E}_{t-\Delta}[x_t]\right) + \left(\mathsf{E}_{t+\Delta}[\tilde{\phi}_t x_t] - \mathsf{E}_{t-\Delta}[\tilde{\phi}_t x_t]\right). \tag{2.3}$$

The first term is the actual monetary policy shock. The second term captures expectation revisions about x_t, which is the private central bank information effect (Nakamura and Steinsson, 2018; Jarociński and Karadi, 2020). The third term reflects time-varying systematic monetary policy. If monetary announcements convey no news about x_t, the third term becomes $\left(\mathsf{E}_{t+\Delta}[\tilde{\phi}_t] - \mathsf{E}_{t-\Delta}[\tilde{\phi}_t]\right) \mathsf{E}_{t-\Delta}[x_t]$. Thus, expectation revisions about systematic monetary policy in a narrow window around monetary policy announcements contaminate monetary policy surprises. This result has previously been noted by Bauer and Swanson (2023a), who also provide evidence consistent with such contamination.[11] Compared to Proposition 1, whether (perceived) systematic monetary policy varies outside the narrow announcement windows is irrelevant for strategy (c), but not for (a) and (b).

2.3 A New Keynesian model

In this section, we impose structure on the macroeconomy in the form of a stylized New Keynesian (NK) model, which allows us to think about the nature of contamination and its implications. The subsequent section considers a more general environment.

We depart from the textbook NK model of Galí (2015) and enrich it with time-varying systematic monetary policy. The model is given by the following four equations:

$$y_t = \mathsf{E}_t[y_{t+1}] - (i_t - \mathsf{E}_t[\pi_{t+1}]) + (d_t - \mathsf{E}_t[d_{t+1}]) \tag{2.4}$$
$$\pi_t = \beta \mathsf{E}_t[\pi_{t+1}] + \kappa y_t - \kappa a_t \tag{2.5}$$
$$i_t = \alpha + (\phi + \tilde{\phi}_t)\pi_t + m_t \tag{2.6}$$
$$\tilde{\phi}_t = \rho_\phi \tilde{\phi}_{t-1} + \psi^d d_t + \psi^a a_t + q_t \tag{2.7}$$

[11] Importantly, regressing high-frequency identified monetary policy shocks on $\mathsf{E}_{t-\Delta}[x_t]$, whether that is publicly available macroeconomic forecasts (Bauer and Swanson, 2023b) or Greenbook forecasts (Miranda-Agrippino and Ricco, 2021), does not resolve the contamination by time-varying systematic monetary policy. In general, the residual of the regression $\hat{e}_t^m = \gamma \mathsf{E}_{t-\Delta}[x_t] + v_t$ still contains variation in $\tilde{\phi}_t x_t$. A special case in which such variation is not contained in v_t is the absence of expectation revisions about x_t together with expectation revisions about $\tilde{\phi}_t$, which are constant in sign and magnitude across monetary announcements.

Variable y_t denotes output, π_t the inflation rate, i_t the nominal interest rate, all in (log) deviation from steady state. The first equation is the dynamic IS equation followed by the New Keynesian Phillips Curve, a time-varying Taylor rule and a law of motion for systematic monetary policy in line with (2.1). The economy fluctuates in response to exogenous movements in the discount rate (demand) d_t, technology (supply) a_t, monetary policy (intercept) m_t, and systematic monetary policy (slope) q_t. Fluctuations in systematic monetary policy are (partly) endogenous if $\psi^d = 0$ or $\psi^a = 0$. All exogenous variables are mutually independent and follow stable stable AR(1) process with with shocks w_t^d, w_t^a, w_t^m, w_t^q.
Based on Hack et al. (2023), the approximate equilibrium dynamics of z_t $\{y_t, \pi_t\}$ follow

$$z_t = \alpha_z + \delta_z^d d_t + \delta_z^a a_t + \delta_z^m m_t + \gamma_z^d d_t \tilde{\phi}_t + \gamma_z^a a_t \tilde{\phi}_t + \gamma_z^m m_t \tilde{\phi}_t + \delta_z^{\phi} \tilde{\phi}_t. \tag{2.8}$$

It is straightforward to show (given $\beta, \kappa > 0$) that the demand shock raises output and inflation, $\delta_y^d, \delta_\pi^d > 0$, and both responses are dampened by a higher $\tilde{\phi}_t$, that is, $\gamma_y^d, \gamma_\pi^d < 0$. The supply shock raises output and lowers inflation, $\delta_y^d > 0$ and $\delta_\pi^d < 0$, and $\tilde{\phi}_t$ increases both responses, $\gamma_y^d, \gamma_\pi^d > 0$.
In the context of our model, the wedges in Proposition 1 are given by

$$\omega_t^{\hat{b}} = (\phi - \hat{b})\pi_t, \quad \text{and} \quad \omega_t^{\tilde{\phi}} = \tilde{\phi}_t \pi_t - \mathsf{E}[\tilde{\phi}_t \pi_t].$$

Given that the inflation rate responds to all macroeconomic shocks, the conventionally estimated monetary policy shock $\hat{e}_t^m$ does not isolate the monetary policy shock but is contaminated by the demand and supply shocks. This holds even if $\hat{b} = \phi$ and even if we additionally assume $\tilde{\phi}_t$ fluctuates only for exogenous reasons ($\psi^d = \psi^a = 0$).

2.4 Bias in impulse response estimates

Empirical monetary policy shocks are often not the object of interest *per se*, but rather impulse response functions (IRF) estimated based on these shocks. We analytically show that the contamination of monetary policy shocks generally leads to biased IRF estimates, including relative IRF estimates. We provide results for a general data-generating process and use the New Keynesian model from Section 2.3 as illustrative example.
Suppose we are interested in the causal effects of the monetary policy shock w_t^m on some scalar outcome z_{t+h}, that is, h periods after the shock. Let z_t follow the stationary Moving

Average (MA) process

$$z_t = \gamma_z + \sum_{h=0}^{\infty} \left(\theta_z^h \, w_{t-h}^m + v_{z,t-h}^h \right), \qquad \mathsf{E}[v_{z,t-h}^h] = \mathsf{E}[w_{t-h}^m v_{z,t-j}^h] = 0 \quad h, j, \tag{2.9}$$

where θ_z^h denotes the causal effect of w_t^m on z_{t+h} and γ_z is a constant. All fluctuations in z_t not explained by $\{w_{t-h}^m\}$ are explained by the 'residual' $\{v_{z,t-h}^h\}$. The MA process describes a general data generating process in the sense that we leave the residual largely unrestricted other than assuming that $v_{z,t+h}^h$ is uncorrelated with the monetary policy shock at all lags and leads. In general, θ_z^h can be defined as the best linear prediction.

In the NK model of Section 2.3, the residual $v_{z,t}^0$ collects all terms on the right-hand side of (2.8) except $\delta_z^m w_t^m$.[12] Given our assumptions that all shocks are mutually independent and that systematic monetary policy does not respond to w_t^m, the residual in the NK model is uncorrelated with the monetary policy shock, in line with the assumption in (2.9). In the model, the causal effect of w_t^m on z_t is given by $\theta_z^0 = \delta_z^m$.

To state the next proposition, it is convenient to rewrite (2.9) as

$$z_{t+h} = \gamma_z + \theta_z^h \, w_t^m + \tilde{v}_{z,t+h}^h, \quad \tilde{v}_{z,t+h}^h = \sum_{j=0}^{\infty} v_{z,t+h-j}^h + \sum_{j=0,\ j=h}^{\infty} \theta_z^h \, w_{t+h-j}^m \tag{2.10}$$

and it follows from (2.9) that $\mathsf{E}[w_t^m \tilde{v}_{z,t+h}^h] = 0 \quad h \geq 0$.[13] Next, suppose an econometrician aims to estimate the effects of monetary policy via the local projection

$$z_{t+h} = c_z^h + d_z^h \hat{e}_t^m + u_{z,t+h}^h, \tag{2.11}$$

where $\hat{e}_t^m$ denotes the estimated monetary policy shock as described in Proposition 1.[14] If $\hat{e}_t^m = w_t^m$, the econometrician will uncover the causal effect via the OLS estimate as $\hat{d}_z^h \xrightarrow{p} \theta_z^h$. In general, however, the estimate $\hat{d}^h$ will be biased, as the following proposition shows.

Proposition 2 (**IRF bias**). *Let monetary policy follow* (2.1) *and* z_t *follow the MA process in* (2.9)*. Consider the local projection in* (2.11) *with* $\hat{e}_t^m$ *as described in Proposition 1. As* $T \to \infty$*, the OLS estimate* $\hat{d}_z^h$ *of the local projection satisfies*

$$\hat{d}_z^h \xrightarrow{p} \theta_z^h + \vartheta_z^{\hat{b}} + \vartheta_z^{\tilde{\phi}} + \vartheta_z^a$$

[12] Given the AR(1) process, $z_t = \rho^z z_{t-1} + w_t^z$, the residual also contains $\delta_z^m \rho^z z_{t-1}$.

[13] We assume $\tilde{\phi}_t$, x_t, z_t, $\tilde{v}_{z,t+h}^h$ jointly follow a stable and ergodic process with finite fourth moments.

[14] We further consider an extension of the local projection in (2.11) that includes lagged control variables. This leads to broadly similar results, as we discuss further below.

with the three bias terms given by

$$
\begin{aligned}
\vartheta_z^{\hat{b}} &= \mathsf{E}\left[(\hat{e}_t^m)^2\right]^{-1}\left(\phi - \hat{b}\right)\left(\theta_z^h \mathsf{E}\left[x_t w_t^m\right] + \mathsf{E}\left[x_t \tilde{v}_{z,t+h}^h\right]\right), \\
\vartheta_z^{\tilde{\phi}} &= \mathsf{E}\left[(\hat{e}_t^m)^2\right]^{-1}\left(\theta_z^h \mathsf{E}\left[\tilde{\phi}_t x_t w_t^m\right] + \mathsf{E}\left[\tilde{\phi}_t x_t \tilde{v}_{z,t+h}^h\right]\right), \\
\vartheta_z^{a} &= \mathsf{E}\left[(\hat{e}_t^m)^2\right]^{-1}\theta_z^h\left(\mathsf{E}\left[(w_t^m)^2\right] - \mathsf{E}\left[(\hat{e}_t^m)^2\right]\right).
\end{aligned}
$$

The proof is straightforward and omitted here.

We next discuss the three bias terms. The first bias $\vartheta_z^{\hat{b}}$ derives from the wedge $\omega_t^{\hat{b}}$. A sufficient condition for the wedge to be zero is $\hat{b} = \phi$. But, as we argued in Section 2, the presence of time-varying systematic monetary policy further complicates obtaining an unbiased estimate of ϕ. Another sufficient condition is $\mathsf{E}\left[x_t w_t^m\right] = 0$ and $\mathsf{E}[x_t \tilde{v}_{z,t+h}^h] = 0$. In the context of the NK model, the first subcondition requires the strong assumption that the inflation rate does not respond to the monetary policy shock. Beyond the model, the subcondition is satisfied if x_t only includes variables that are predetermined with respect to w_t^m, e.g., forecasts shortly before a monetary policy decision (cf. Romer and Romer, 2004). The second subcondition requires the even stronger assumption that x_t does not respond to any other macroeconomic shock (present or past), or the knife-edge case that the terms in $\mathsf{E}[x_t \tilde{v}_{z,t+h}^h]$ sum up to zero. Even the case in which x_t includes only predetermined variables does not necessarily help as long as $\tilde{v}_{z,t+h}^h$ contains past macroeconomic shocks, as is the case in the NK model.

The second bias $\vartheta_z^{\tilde{\phi}}$ derives from the wedge $\omega_t^{\tilde{\phi}}$, which captures the misspecification of the linear Taylor rule regression. The bias depends on two expectations, $\mathsf{E}[\tilde{\phi}_t x_t w_t^m]$ and $\mathsf{E}[\tilde{\phi}_t x_t \tilde{v}_{z,t+h}^h]$. Similar to the first wedge, the first expectation is not zero in our NK model but becomes zero if x_t includes only predetermined variables. The second expectation is only zero under much stronger assumptions. In particular, assuming that x_t is predetermined is not sufficient. Similarly, it is not sufficient to assume that $\tilde{\phi}_t$ is predetermined or that $\hat{b} = \phi$. Instead, we need to assume that $\tilde{\phi}_t x_t$ is exogenous, which effectively assumes that monetary policy is fully exogenous. That is, it does not respond to macro conditions.[15] Alternatively, we need to assume away time-varying systematic monetary policy $\tilde{\phi}_t = 0 \quad t$. Hence, the bias arising from $\mathsf{E}[\tilde{\phi}_t x_t \tilde{v}_{z,t+h}^h] = 0$ is unlikely zero.

The third term, ϑ^a, can be interpreted as a type of attenuation bias. If the estimated monetary policy shock satisfies $\mathsf{E}\left[(\hat{e}_t^m)^2\right] > \mathsf{E}\left[(w_t^m)^2\right]$, the estimate $\hat{d}_z^h$ will be biased toward zero relative to θ_z^h. However, given that w_t^m may correlate with the wedges $\omega_t^{\hat{b}}$ and $\omega_t^{\tilde{\phi}}$, the estimated monetary policy shock is not classical measurement error. If $\mathsf{E}\left[(\hat{e}_t^m)^2\right] < \mathsf{E}\left[(w_t^m)^2\right]$, the estimate $\hat{d}_z^h$ will be biased away from zero. Even if the first two biases are zero, the third

[15] Importantly, neither assuming exogeneity of $\tilde{\phi}_t$ or x_t separately is sufficient.

bias remains non-zero as long as $\tilde{\phi}_t = 0 \quad t$.

The local projection specified in (2.11) is highly parsimonious. It does not include any endogenous control variables. Consider instead the extended local projection $z_{t+h} = c_z^h + d_z^h \hat{e}_t^m + \Gamma(L) Y_t + u_{z,t+h}^h$, where $\Gamma(L) = \sum_{i=1}^{\infty} \Gamma_i L^i$ is a lag polynomial, and Y_t a vector of control variables. The additional control vector means we need to replace $\hat{e}_t^m$ and $\tilde{v}_{z,t+h}^h$ by projections of these variables on $\{Y_{t-1}, Y_{t-2}, ...\}$ in the bias terms in Proposition 2. Importantly, the control variables do not eliminate the bias. To see this, consider a macroeconomic shock that realizes between t and $t-1$. This shock may affect x_t even if x_t is predetermined regarding the monetary policy decision.[16] If the shock is persistent, then it may also affect z_{t+h} and, hence, $\tilde{v}_{z,t+h}^h$. Since the lagged controls are determined before the realization of the oil shock, we have that $\mathsf{E}[x_t \tilde{v}_{z,t+h}^h]$ and $\mathsf{E}[\tilde{\phi}_t x_t \tilde{v}_{z,t+h}^h]$ are generally non-zero.

In some empirical applications, the econometrician aims to identify the relative effect of monetary policy shocks rather than its absolute effect. If $\theta_{z_1}^h$ is the absolute causal effect of w_t^m on $z_{1,t+h}$, the relative causal effect is $\theta_{z_1}^h / \theta_{z_2}^h$, where $z_{2,t+h}$ denotes another outcome. For example, it is common to study the effects of monetary policy shocks that raise the nominal interest rate by 25 or 100 basis points. This requires dividing the response of some outcome variable of interest by the interest rate response. For some empirical questions, a bias in the estimated absolute effect may be acceptable as long as the bias cancels out in the estimated relative effect. The following proposition provides a condition for the relative estimate to be unbiased.

Proposition 3 (**Relative IRF bias**). *Let monetary policy follow* (2.1) *and* $z_{1,t}$ *and* $z_{2,t}$ *follow MA processes as in* (2.9). *Consider two local projections as in* (2.11) *to estimate the e ects of* $\hat{e}_t^m$*, as described in Proposition 1, on* $z_{1,t+h}$ *and* $z_{2,t+h}$. *As* $T \to \infty$*, the two OLS estimates* $\hat{d}_{z_1}^h$ *and* $\hat{d}_{z_2}^h$ *satisfy*

$$\frac{\hat{d}_{z_1}^h}{\hat{d}_{z_2}^h} \xrightarrow{p} \frac{\theta_{z_1}^h}{\theta_{z_2}^h}$$

if and only if

$$\frac{\left(\phi - \hat{b}\right) \mathsf{E}\left[x_t \tilde{v}_{z_1,t+h}^h\right] + \mathsf{E}\left[\tilde{\phi}_t x_t \tilde{v}_{z_1,t+h}^h\right]}{\theta_{z_1}^h} = \frac{\left(\phi - \hat{b}\right) \mathsf{E}\left[x_t \tilde{v}_{z_2,t+h}^h\right] + \mathsf{E}\left[\tilde{\phi}_t x_t \tilde{v}_{z_2,t+h}^h\right]}{\theta_{z_2}^h}.$$

The condition under which the relative IRF is not biased is a knife-edge condition, which is generally not satisfied. A (strong) sufficient condition is $\hat{b} = \phi$ and $\tilde{\phi}_t = 0 \quad t$, which assumes

[16] In practice, e.g., an oil shock that materializes at the beginning of the month should affect Greenbook forecasts for a monetary policy meeting taking place later in the same month.

away the identification problem by imposing time-invariant systematic monetary policy.

3 Empirical evidence on shock contamination

In this section, we show that fluctuations in U.S. systematic monetary policy, measured by the composition of hawks and doves in the FOMC, are strongly predictive for the empirical monetary policy shocks identified from Taylor rule-type regressions in Romer and Romer (2004), the refinement in Aruoba and Drechsel (2022), but also the shocks identified from high-frequency monetary policy surprises in Miranda-Agrippino and Ricco (2021).

3.1 Measuring time-varying systematic monetary policy

We describe two time series of systematic monetary policy, the Hawk-Dove balance among all FOMC members, and the Hawk-Dove balance among the subset of rotating FOMC members. The FOMC decides U.S. monetary policy and consists of 12 voting members, among which four members serve one-year terms on a rotating basis. We use the narrative classification of FOMC members as hawks and doves in Istrefi (2019). Hawks are perceived to be more concerned with inflation, while doves are more concerned with employment and growth. The hawk-dove classification is a panel that tracks FOMC members over time at FOMC meeting frequency.[17] Istrefi (2019) shows that the perceived policy preferences match well with policy tendencies that are unknown in real-time to the public, as expressed by preferred interest rates, with forecasting patterns of individual FOMC members, and with dissents. In addition, Bordo and Istrefi (2023) show that the FOMC members' educational background and early life experience have predictive power for individual policy preferences.

To measure variation in systematic monetary policy over time, we aggregate the individual FOMC member preferences into a Hawk-Dove balance for each meeting (cf. Istrefi, 2019). We do so because the nature of monetary policy-making involves the aggregation of diverse individual policy preferences in a collective decision. We first map the qualitative hawk-dove classification on a numerical scale for FOMC member i at meeting τ ranging from $Hawk_{i\tau}$ =+1 for consistent hawks, +1/2 for hawks who have been doves before, 0 for unclassified member, and -1/2 (-1) for swinging (consistent) doves.[18] We then construct

[17] Among the 147 FOMC members between 1960 and 2023, 129 are classified as hawk or dove. The news coverage for the remaining 18 members is insufficient for classification. 95 classified members are consistently hawks or doves, while the others switch camps at least once. The 34 swinging members switch camps at 1.8% of member-meeting pairs.

[18] Hack et al. (2023) show that alternative aggregation schemes lead to similar empirical findings.

the aggregate Hawk-Dove balance in the FOMC by

$$Hawk_\tau^{\mathsf{F}} = \frac{1}{|\mathsf{F}_\tau|} \sum_{i \in \mathsf{F}_\tau} Hawk_{i\tau} \tag{3.1}$$

where F_τ denotes the (full) set of voting FOMC members i at meeting τ.[19] The Hawk-Dove balance may respond to the state of the economy, through swings in types or through new appointments. For example, the Federal Reserve may become more dovish in response to high unemployment or more hawkish in response to high inflation (cf. Davig and Leeper, 2008). Systematic monetary policy may also change in response to political pressure (e.g., Abrams, 2006; Bianchi, Gómez-Cram, Kind, and Kung, 2023). To address the endogeneity of the Hawk-Dove balance, we construct the Hawk-Dove balance among the set of FOMC members who currently have voting rights through the annual rotation.[20] The mechanical nature of the rotation renders it orthogonal to the state of the economy and political cycles. Formally, the Rotation Hawk-Dove balance is defined by

$$Hawk_\tau^{\mathsf{R}} = \frac{1}{|\mathsf{R}_\tau|} \sum_{i \in \mathsf{R}_\tau} Hawk_{i\tau}, \tag{3.2}$$

where R_τ denotes the set FOMC members at meeting τ that had voting right through the rotation.[21] While $Hawk_\tau^{\mathsf{F}}$ is a more comprehensive measure of systematic monetary policy, $Hawk_\tau^{\mathsf{R}}$ has the advantage of primarily reflecting exogenous variation through the rotation. We present the evolution of $Hawk_\tau^{\mathsf{F}}$ and $Hawk_\tau^{\mathsf{R}}$ from 1960 through 2023 in Figure 1. Both balances vary considerably, featuring hawkish and dovish majorities. The variation reflects the turnover of rotating FOMC members, the turnover of non-rotating FOMC members, and changes in policy preferences of incumbent FOMC members. The correlation between $Hawk_\tau^{\mathsf{F}}$ and $Hawk_\tau^{\mathsf{R}}$ is 0.60; see Table A.1 for further descriptive statistics. Fluctuations in $Hawk_\tau^{\mathsf{R}}$ are more short-lived, reflecting the annual rotation of voting rights.

The Hawk-Dove balances are informative about systematic monetary policy. First, the classification matches well with narratives of monetary policy in the U.S. (Istrefi, 2019). Second, a hawkish FOMC responds to higher inflation by raising the policy rate more aggressively (Bordo and Istrefi, 2023; Hack et al., 2023). Finally, a hawkish FOMC tightens monetary policy more aggressively in response to expansionary government spending shocks, leading to a significantly dampened GDP expansion (Hack et al., 2023).

[19] Occasionally, $|\mathsf{F}_\tau| \leq 12$ because of absent members and vacant positions.

[20] This was originally proposed in Hack et al. (2023) as instrument for systematic monetary policy. Hack et al. (2023) discuss in detail that swings and new appointments are not a concern for the rotating group.

[21] In our sample, $|\mathsf{R}_\tau| = 4$ for 625 out of 634 FOMC meetings and $|\mathsf{R}_\tau| = 3$ for the remaining meetings.

Figure 1: Hawk-Dove balance in the FOMC

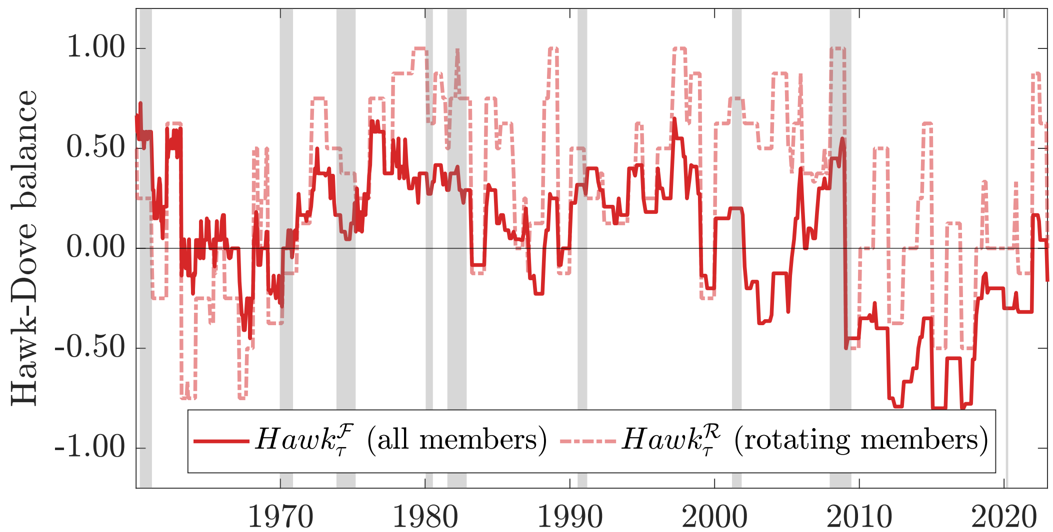

Notes: The solid red line shows the aggregate Hawk-Dove balance of the full FOMC $Hawk_\tau^{\mathcal{F}}$ at FOMC meeting frequency from 1960 through 2023. The dashed red line shows the aggregate Hawk-Dove balance of the rotation panel $Hawk_\tau^{\mathcal{R}}$. Grey bars indicate NBER dated recessions.

3.2 Predictability of Taylor rule residuals

We show that the monetary policy shocks in the seminal Romer and Romer (2004) (henceforth RR) are predictable by fluctuations in measured systematic monetary policy in a way that supports our theoretical results. This empirical finding extends to the refined RR shocks in Aruoba and Drechsel (2022) (henceforth AD).

RR shocks. The RR shocks are the estimated residuals $\hat{e}_\tau^{rr}$ of a Taylor rule-type regression

$$i_\tau = a + b\, x_\tau + e_\tau^{rr}, \tag{3.3}$$

estimated via OLS and where τ denotes FOMC meetings. RR specify i_τ as the change in the intended federal funds rate between two FOMC meetings. The right-hand side x_τ includes 18 variables: the Greenbook forecast of output growth and inflation, prepared in advance of FOMC meeting τ, respectively for the quarter preceding the FOMC meeting, the current and the two subsequent quarters; the revision of all 8 Greenbook forecasts relative to the same forecasts prepared for the preceding FOMC meeting; the Greenbook forecast of the unemployment rate in the current quarter; and the intended federal funds rate before FOMC meeting τ. We use the estimated monetary policy shocks $\hat{e}_\tau^{rr}$ and associated regressors x_τ

from Wieland and Yang (2020) who extend the RR sample 1969-1996 to 1969-2007.[22]

Predictability of RR shocks. If systematic monetary policy is time-varying, as in (2.1), the estimated residuals $\hat{e}^{rr}_\tau$ in (3.3) will contain fluctuations in systematic monetary policy multiplied with the inputs of the Taylor rule, corresponding to $\tilde{\phi}_\tau x_\tau$ in Proposition 1. Hence, a testable prediction of a time-varying Taylor rule is that fluctuations in measured systematic monetary policy interacted with x_τ (partly) explain RR shocks $\hat{e}^{rr}_\tau$. To test this prediction, we estimate the following regression

$$\begin{aligned}\hat{e}^{rr}_\tau =& \beta_0 + \beta_1\; x_{\tau-p} Hawk_{\tau-p} + \beta_2\; x_{\tau-p} \Delta Hawk_{\tau-p} \\ &+ \beta_3\; Hawk_{\tau-p} + \beta_4\; \Delta Hawk_{\tau-p} + \beta_5\; x_{\tau-p} + u_\tau, \end{aligned} \tag{3.4}$$

where τ denotes an FOMC meeting, $\hat{e}^{rr}_\tau$ is the RR shock and x_τ the RR regressors. $Hawk_\tau$ is either $Hawk^{\mathsf{F}}_\tau$ or $Hawk^{\mathsf{R}}_\tau$, and $\Delta Hawk_\tau$ is the first difference of $Hawk_\tau$.[23] We consider contemporaneous regressors ($p = 0$) or lags up to two meetings ($p = 1, 2$). Our motivation to consider lags is to capture that it may take time for FOMC members to affect policy decisions (Hack et al., 2023).[24]

Table 1 presents the R^2 for various specifications of (3.4), as well as the p-values for the null hypothesis that all regression coefficients are jointly zero. We consider the full sample, the original RR sample, as well as a post-Volcker disinflation sample. Across regression specifications, we obtain an R^2 between 0.10 and 0.54. Using a one-meeting lag ($p = 1$) yields the largest R^2 ranging from 0.33 to 0.54, consistent with the sluggish nature of decision-making in the FOMC. In other words, a sizable fraction of the variation in RR shocks can be explained by past variables, irrespective of the type of Hawk-Dove balance and the three sample specifications. For $p = 1$, we can reject the null hypothesis that all coefficient estimates are zero at the 1% significance level. The R^2 tends to be lower for $p = 0$ and $p = 2$ across sub-samples, except for the post-Volcker disinflation sample. We further find that $Hawk^{\mathsf{R}}_\tau$ tends to predict the RR shocks somewhat better than $Hawk^{\mathsf{F}}_\tau$.

[22]We follow Wieland and Yang (2020) and implement their Greenbook data corrections and adjustments. We end the sample just before the Great Recession, thus avoiding the subsequent period with interest rates close to the zero lower bound and a large expansion of unconventional monetary policy.

[23]In Hack et al. (2023), the rotation Hawk-Dove balance is proposed as an instrument to provide causal evidence on the state-dependent effects of macroeconomic shocks regarding time-variation in systematic monetary policy. In this paper, we do not use the rotation Hawk-Dove balance explicitly as an instrument because the high number of regressors in our empirical application would render an IV approach unreliable.

[24]For example, former Governor Laurence Meyer remarks: *I came to see policy decisions as often evolving over at least a couple of meetings. The seeds were sown at one meeting and harvested at the next. [...] Similarly, while in my remarks to my colleagues it sounded as if I were addressing today's concerns and today's policy decisions, in reality I was often positioning myself, and my peers, for the next meeting.* Laurence Meyer (2004), A Term at the Fed: An Insiders' View, Harper Business.

Table 1: Explaining RR shocks by systematic monetary policy

	$Hawk_\tau^F$			$Hawk_\tau^R$		
Sample	69-07	69-96	83-07	69-07	69-96	83-07
(a) Contemporaneous FOMC meeting (p=0)						
R^2	0.098	0.134	0.426	0.165	0.216	0.462
p-value	0.189	0.243	0.000	0.012	0.000	0.000
T	353	265	200	353	265	200
(b) One FOMC meeting lag (p=1)						
R^2	0.333	0.431	0.452	0.432	0.543	0.441
p-value	0.003	0.001	0.000	0.000	0.000	0.000
T	349	261	200	349	261	200
(c) Two FOMC meetings lag (p=2)						
R^2	0.241	0.310	0.369	0.278	0.359	0.423
p-value	0.000	0.000	0.000	0.000	0.000	0.000
T	347	259	200	347	259	200

Notes: The table shows results from regressions based on (3.4). The rows of the three subtables show R^2, the p-values for the null hypothesis that all coefficient estimates are jointly zero, and the number of observations T. The three left columns show results for the Hawk-Dove balance across all FOMC members, and the three right columns for the Hawk-Dove balance across all rotating FOMC members with voting rights. The three subtables differ by the specification of FOMC meeting lag p. Columns one to three differ by the sample period between 1969-2007, 1969-1996, and 1983-2007, and analogously for columns four to six.

We next investigate the contribution of subsets of regressors for explaining variation in RR shocks. We focus on the regression specification with $Hawk_{\tau-p}^R$ and $p = 1$ because it yields a comparatively large R^2 across subsamples, but we obtain similar results for other specifications. Table 2 reports the R^2 and p-value when regressing the RR shock $\hat{e}_\tau^{rr}$ separately on subsets of the regressors included in equation (3.4). Interactions between $x_{\tau-1}$ and, respectively, $Hawk_{\tau-1}^R$ and $\Delta Hawk_{\tau-1}^R$ yield high R^2. In contrast, the (non-interacted) level of $Hawk_{\tau-1}^R$ and $\Delta Hawk_{\tau-1}^R$ has practically no predictive power for the RR shock. This finding further supports the interpretation of the Hawk-Dove balance as capturing variation in systematic monetary policy, the slope of the monetary policy rule rather than its intercept. Finally, the regressor $x_{\tau-1}$ has some predictive power in explaining RR shocks, in particular for the post-Volcker disinflation sample. Overall, our results suggest that a substantial fraction of the RR shocks can be explained by variation in systematic monetary policy.

Table 2: Explaining RR shocks by subsets of regressors

	Interactions			Levels		
Sample	69-07	69-96	83-07	69-07	69-96	83-07
	(a) $Hawk^R_{\tau-1} \times x_{\tau-1}$			(b) $Hawk^R_{\tau-1}$ & $\Delta Hawk^R_{\tau-1}$		
R^2	0.111	0.136	0.121	0.007	0.011	0.002
p-value	0.087	0.050	0.032	0.325	0.286	0.826
	(c) $\Delta Hawk^R_{\tau-1} \times x_{\tau-1}$			(d) $x_{\tau-1}$		
R^2	0.251	0.290	0.066	0.091	0.133	0.256
p-value	0.000	0.000	0.000	0.029	0.005	0.000
	(e) All interactions			(f) All level terms		
R^2	0.343	0.399	0.197	0.097	0.154	0.257
p-value	0.000	0.000	0.000	0.034	0.001	0.000
T	349	261	200	349	261	200

Notes: The table shows results from regressions based on (3.4), considering different subsets of the regressors. The rows of the three subtables show R^2 and the p-values for the null hypothesis that all coefficient estimates are jointly zero, and the number of observations T. The three left columns show results for the interactions between the Hawk-Dove balance and $x_{\tau-1}$, and the three right columns show the results for the non-interacted (level) regressors. Columns one to three differ by the sample period between 1969-2007, 1969-1996, and 1983-2007, and analogously for columns four to six.

A potential concern with our results is that the large set of regressors might lead to overfitting. We may mechanically absorb variation, although there is no systematic relationship in the data. We address this concern with two exercises. First, we present the adjusted R^2 in Table B.1 in Appendix B. A positive adjusted R^2 means the regressors have explanatory power beyond the power obtained from adding unrelated random regressors in finite samples. The specification with $Hawk^R_{\tau-p}$ and $p = 1$ yields the adjusted R^2 ranging between 0.22 and 0.42.[25] As a second exercise, we present a Lasso estimation. The Lasso minimizes the sum of squared residuals (as OLS) but additionally penalizes the number of estimated parameters to keep the set of included regressors small. We choose the penalization parameter to gradually

[25]Moreover, in Section 4.1, we compare the original RR regression with a version augmented by the regressors from (3.4). We find that the adjusted R^2 more than doubles in the augmented regression, see Table 4.

increment the number of regressors from one to four. We present the results in Table B.2, in Appendix B. We find that four (scalar) regressors are sufficient to yield an R^2 of 0.15 in the full sample. All four regressors are interactions with $\Delta Hawk^{R}_{\tau-1}$ and have signs consistent with our interpretation of the hawk-dove balance. In particular, interactions between the hawk-dove balance and inflation are positive, meaning hawkish policymakers respond more strongly to higher inflation. We conclude that the predictability of RR shocks is unlikely driven by overfitting.

Predictability of AD shocks. In related work, Aruoba and Drechsel (2022) refine the RR shock by using a large vector x_t with the goal of better capturing the Fed's information set about the state of the economy. They use textual analysis to create sentiment indicators about the Fed staff's assessment of the economy before FOMC meetings. The sentiment indicators are used as additional regressors in a Taylor rule-type Ridge regression. The estimated residual is the AD monetary policy shock. To assess whether the AD shock is predictable by systematic monetary policy, we estimate the regression in (3.4) but with the AD shock as left-hand side variable. We provide the results in Table B.3 in Appendix B. The full sample runs from 1983 until 2007 because this is the period for which the AD shock is available. In this sample, we find an R^2 between 0.26 and 0.36 depending on lag order ($p = 0, 1, 2$) and the type of Hawk-Dove balance. Thus, even their refined shock is predictable and, hence, may be contaminated by time variation in systematic monetary policy.

3.3 High-frequency identified monetary policy shocks

We show that monetary policy shocks identified from high-frequency monetary policy surprises in Miranda-Agrippino and Ricco (2021) (henceforth MAR) are also predictable by time-variation in systematic monetary policy.

MAR shocks. MAR identify monetary policy shocks via a proxy VAR with high-frequency monetary policy surprises as an external instrument (proxy) for a monthly sample from 1980M1 through 2014M12. Similar proxy VAR approaches are used in the literature (e.g., Gertler and Karadi, 2015; Jarociński and Karadi, 2020; Bauer and Swanson, 2023b). In the proxy VAR, predictability may arise from the misspecification of the linear VAR or from contaminated high-frequency monetary policy surprises. Using the identified shock allows us to capture both sources of predictability. We consider the monetary policy shock series associated with the maximum likelihood estimation of their six-variate proxy VAR model.

Table 3: Explaining HF identified shocks by systematic monetary policy

	$Hawk_t^F$			$Hawk_t^R$		
Sample	80-14	80-96	83-07	80-14	80-96	83-07
	(a) Contemporaneous FOMC meeting (p=0)					
R^2	0.274	0.548	0.265	0.244	0.478	0.299
p-value	0.000	0.000	0.000	0.000	0.000	0.000
T	283	139	200	283	139	200
	(b) One FOMC meeting lag (p=1)					
R^2	0.265	0.511	0.260	0.279	0.550	0.267
p-value	0.000	0.000	0.000	0.000	0.000	0.000
T	282	138	200	282	138	200
	(c) Two FOMC meetings lag (p=2)					
R^2	0.267	0.514	0.302	0.264	0.494	0.263
p-value	0.000	0.000	0.000	0.000	0.000	0.000
T	283	138	201	283	138	201

Notes: The table shows results from regressions based on (3.4), but the left-hand side is the shock from Miranda-Agrippino and Ricco (2021) at monthly frequency, which is identified via their six-variate proxy VAR and their high-frequency instrument. The rows of the three subtables show R^2, the p-values for the null hypothesis that all coefficient estimates are jointly zero, and the number of observations T. The three left columns show results for the Hawk-Dove balance across all FOMC members, and the three right columns for the Hawk-Dove balance across all rotating FOMC members with voting rights. The three subtables differ by the specification of FOMC meeting lag p. Columns one to three differ by the sample period between 1980-2014, 1980-1996, and 1983-2007, and analogously for columns four to six.

Predictability of MAR shocks. To test whether the MAR shock is predictable by systematic monetary policy, we estimate the regression in (3.4) but with the MAR shock as left-hand side variable. The regression results are provided in Table 3. The R^2 is always above 0.24 and gets as high as 0.53. Thus, the MAR shock also seems predictable by time variation in systematic monetary policy. For two out of the three subsamples, we find that the specification with $Hawk_{\tau-p}^R$ and $p = 1$ yields the largest R^2, comparable to the RR shock. The subsample from 1983 until 2007 is suitable for a quantitative comparison with the RR results in Table 1 since we have an identical subsample. The R^2 is generally lower for MAR shocks. However, the R^2 for the MAR shock is at least 60 percent of the R^2 for RR, across all specifications. Thus, we still find considerable predictability relative to RR shocks.

Overall, we conclude that the predictability of monetary policy shocks is supported by empirical evidence and not limited to shocks computed as Taylor rule-type residuals. This confirms the empirical relevance of our theory for various types of monetary policy shocks, including those identified from high-frequency changes in asset prices.

4 New monetary policy shocks

The results in Section 3 motivate us to construct new monetary policy shocks that are no longer predictable by measured systematic monetary policy. We estimate new RR shocks that are orthogonal to interactions between Greenbook forecasts and measured time variation in systematic monetary policy. We find that our new shock affects output and inflation with a substantially shorter delay, more strongly, and at higher statistical significance compared to the RR shock, in particular for a post-Volcker disinflation sample. We further document that impulse responses change significantly when estimated via new AD and new MAR shocks.

4.1 New (RR) shock

We estimate a new monetary policy shock series via the augmented RR regression

$$\begin{aligned} i_\tau = &\beta_0 + \beta_1 x_\tau + \beta_2 x_{\tau-1} + \beta_3\ x_{\tau-1} Hawk_{\tau-1} + \beta_4\ x_{\tau-1}\Delta Hawk_{\tau-1} \\ &+ \beta_5\ Hawk_{\tau-1} + \beta_6\ \Delta Hawk_{\tau-1} + e_\tau^{new}, \end{aligned} \tag{4.1}$$

where the policy instrument i_τ and the Greenbook forecast x_τ are specified as in Section 3, and $Hawk_\tau$ is the Rotation Hawk-Dove balance. Our new monetary policy shock is the estimated residual $\hat{e}_\tau^{new}$ when estimating (4.1) via OLS.[26] The specification nests the original RR regression if we restrict $\beta_j = 0 \quad j > 1$, in which case we denote the estimated residual by $\hat{e}_\tau^{rr}$. Our baseline sample to identify the shock is the full sample from 1969 through 2007. We discuss the sensitivity of our results to alternative samples towards the end of Section 4.2. Table 4 provides descriptive statistics comparing the RR shock and our new (RR) shock.[27] The regression R^2 increases from 0.28 in the RR regression to 0.68 in our augmented regression. The new shock displays no serial correlation, and the correlation between new and RR

[26] While we follow RR in using OLS, this leads to endogeneity bias as discussed in Section 2. If the monetary policy shocks explain a sufficiently small fraction of aggregate fluctuations, the endogeneity bias may be quantitatively negligible (Carvalho et al., 2021). Finally, note that estimating (4.1) via IV is practically infeasible because it would require a large number of instruments.

[27] For five FOMC meetings, x_τ is missing because not all Greenbook forecasts are available. The regression for the new shock (4.1) includes $x_{\tau-1}$ creating five additional missing observations relative to the RR shock.

Table 4: Descriptive statistics of monetary policy shocks

	R^2	Adj. R^2	SD	Autocorr	Corr	Sign-corr	Min	Max	N
RR shock	0.28	0.24	0.34	0.12	-	-	-3.25	1.86	353
New (RR) shock	0.68	0.59	0.23	-0.08	0.67	0.44	-1.02	1.04	348

Notes: The table shows descriptive statistics for the new shock ($\hat{e}_{\tau}^{new}$) and the RR shock ($\hat{e}_{\tau}^{rr}$) at FOMC meeting frequency from 1969 through 2007. R^2 and adjusted R^2 refer to the regression used to estimate the shocks in equation (4.1). Autocorr refers to the meeting-over-meeting autocorrelation. Corr refers to the correlation between new and RR shock. Sign-corr refers to the correlation of the sign of both shock series.

shock is 0.67. The correlation between the sign of both shocks is 0.44, meaning both shocks frequently have the opposite sign.

The new (RR) shock series is substantially less dispersed than the RR shock, with the standard deviation falling from 0.34 to 0.23. The RR shock includes a few exceptionally large shocks, as can be seen from the shock time series in Figure C.1 in Appendix C. These shocks are concentrated during the Volcker disinflation between 1979 and 1982. RR argue that their shocks in this period reflect changes in the Federal Reserve's operating procedures and an increased distaste for inflation. In fact, we do observe a relatively hawkish FOMC, in particular among rotating FOMC members, see Figure 1. Hence, one reason for our new shocks being smaller is that accounting for variation in systematic monetary policy better explains monetary policy decisions during this episode.

The key advantage of our new (RR) shock is that it is orthogonal to measured time-variation in systematic monetary policy. This is valuable because we have demonstrated in Sections 2 and 3 that many conventional shocks, including high-frequency identified shocks, may suffer from contamination by systematic monetary policy, leading to biased impulse responses. Even abstracting from contamination and bias, our new shock has some advantages compared to high-frequency identified shocks. First, our new shock series features substantially larger fluctuations. That is an advantage because it may reduce the need to extrapolate from small local effects to construct typical policy scenarios. Second, we span a long time series for which high-frequency identified shocks are largely unavailable.

4.2 Responses to the new (RR) shock

We study the impulse responses to our new and, for comparison, to the original RR shock. We find that the effects of monetary policy are stronger, monetary transmission is faster, and the effects are more precisely estimated when using the new shock.

Econometric framework. We estimate impulse responses using the local projections

$$z_{t+h} - z_{t-1} = \alpha_z^h + \beta_z^h \; \hat{e}_t^i + \Gamma \; Y_t + v_{t+h}^h, \quad h = 0, \ldots, H, \tag{4.2}$$

where z_t is an outcome variable of interest. The main outcomes of our analysis are the effective federal funds rate, the inflation rate, and the natural logarithm of real GDP. The monetary policy shock $\hat{e}_t^i$ is either the new (RR) shock $\hat{e}_t^{new}$ or the RR shock $\hat{e}_t^{rr}$. The control vector Y_t includes twelve lags of the federal funds rate, the inflation rate, the log of real GDP, and a linear time trend. A period t is a month, which limits the need to aggregate the monetary policy shocks.[28] Monthly log real GDP and the monthly GDP deflator inflation rate are obtained by interpolation using the procedure of Chow and Lin (1971).[29] The baseline sample of our analysis is 1983M1 through 2007M12, so post-Volcker disinflation and pre-Great Recession. We consider this sample particularly interesting because the estimated responses to many conventional monetary policy shock series appear to be implausible in post-Volcker disinflation samples (e.g., Ramey, 2016). This sample further avoids potential structural breaks around the Great Inflation episode. Toward the end of this section, we present evidence on the sensitivity of our results with respect to various modeling choices, including control variables, interpolation, and the sample.

Responses of main outcomes. Figure 2 presents the estimated responses of our main outcome variables, the federal funds rate (FFR), the inflation rate, and the log of real GDP, to the new shock and the RR monetary policy shock. The left column shows the 68% and 95% confidence bands for the new shock, and the right column shows the corresponding confidence bands for the RR shock. All confidence bands are based on standard errors robust to serial correlation and heteroskedasticity. Both shocks are normalized to a peak FFR increase of 100 basis points to facilitate comparability.

Panel (a) shows the estimated responses of the FFR. The FFR significantly increases in response to both the RR and the new shock. The response to the new shock is moderately more transitory. It becomes statistically insignificant at the five percent level after 9 months, and the point estimate is virtually zero after 12 months. In contrast, the response to the RR shock remains significant at 5% for 17 months, and the point estimate reaches zero only after 28 months. Given that the new monetary policy shock leads to a more transitory

[28]Only 4 months (all between 1969 through 1971) contain more than one FOMC meeting with a monetary policy shock $\hat{e}_\tau^i$, while a large fraction of quarters across the entire sample contain multiple $\hat{e}_\tau^i$. In months in which we observe at least one $\hat{e}_\tau^i$, we construct $\hat{e}_t^i$ as the sum of $\hat{e}_\tau^i$ contained in t. Otherwise, we set $\hat{e}_t^i = 0$.

[29]The related monthly series we use for interpolating GDP and the GDP deflator are CPI, industrial production, one-year treasury yield, and excess bond premium. For similar monthly interpolations based on Chow and Lin (1971), see, e.g., Bernanke et al. (1997) and Uhlig (2005).

Figure 2: Responses of main outcomes to monetary policy shocks

(a) Federal funds rate

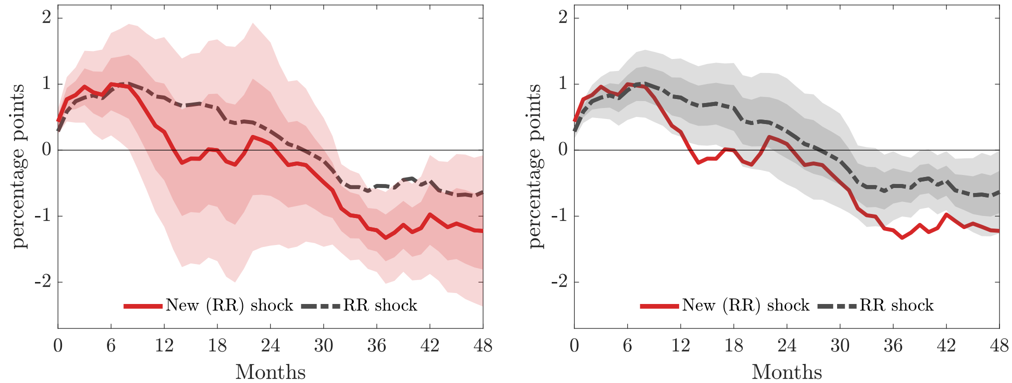

(b) Real GDP

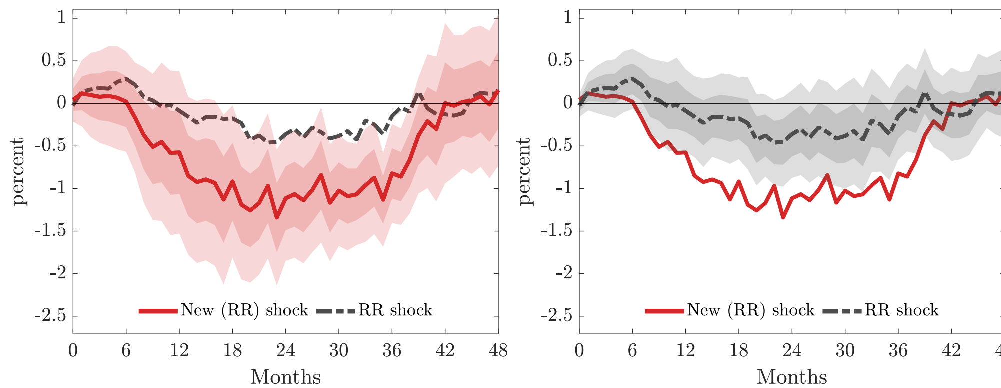

(c) Inflation

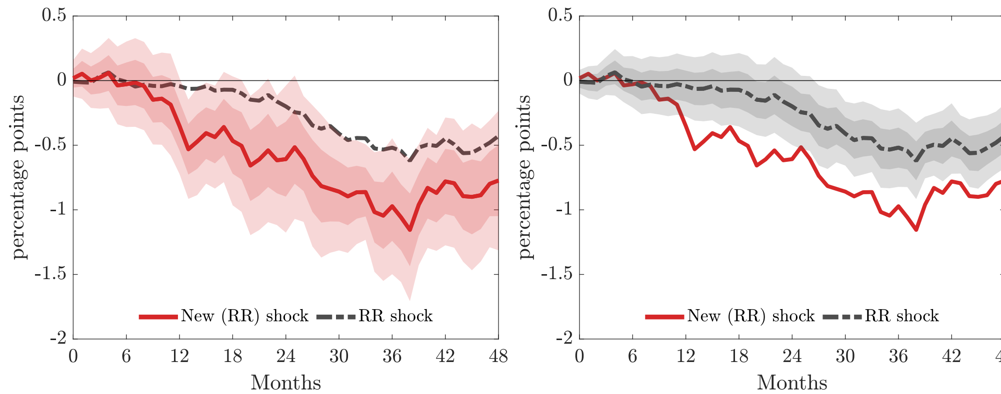

Notes: The figure shows responses of the federal funds rate, log real GDP, and the inflation rate to a monetary policy shock based on the local projection as specified along with (4.2). The new (RR) monetary policy shock is identified as the residual from the Taylor rule regression in (4.1) whereas the RR shock is based on the same regression when $b_j = 0$ for $j > 1$, as in Romer and Romer (2004). The shaded areas indicate 68% and 95% confidence bands using standard errors robust to serial correlation and heteroskedasticity.

dynamic federal funds rate response, we might expect a larger demand contraction from the RR shock, at least if both shocks are well-identified.

Panel (b) shows, however, that the estimated response of real GDP is substantially stronger for the new (RR) shock. The new shock leads to a contraction of GDP that is significant at the 5% level between 17 and 37 months after the shock. In contrast, the GDP response to the RR shock is not significantly different from zero at the 5% level for all horizons that we consider. If we use the (much) lower 32% significance standard, the new (RR) shock leads to a significant GDP contraction starting 9 months after the shock, while it takes 20 months for the RR shock. In addition, the RR shock generates a short-lived expansion around 6 months after the shock (an output puzzle). Panel (b) of Figure 3 displays the difference between the two GDP responses, along with confidence bands.[30] The difference is statistically significant at the 5% level for most horizons between 13 and 37 months after the shock. The shocks further differ strongly in the magnitude of the GDP response. The trough response is -1.34% for the new (RR) shock and -0.46% for the RR shock.

Figure 3: Response to new (RR) shock “minus” response to RR shock

(a) FFR (b) Real GDP (c) Inflation

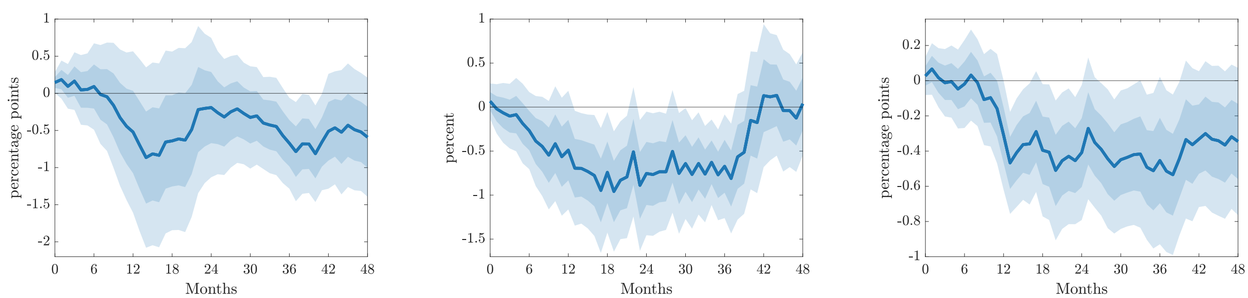

Notes: The figure shows the differences across impulse responses for the federal funds rate, log real GDP, and the inflation rate to a monetary policy shock based on the local projection as specified along with (4.2). The difference is computed as the response to the new shock minus the response to the old shock for each outcome, respectively. The new (RR) monetary policy shock is identified as the residual from the Taylor rule regression in (4.1) whereas the RR shock is based on the same regression when $b_j = 0$ for $j > 1$, as in Romer and Romer (2004). The shaded areas indicate 68% and 95% confidence bands using standard errors robust to serial correlation and heteroskedasticity.

Finally, panel (c) shows the estimated responses of inflation. Arguably, the most striking finding of Figure 2 is the difference in the lag of monetary policy affecting inflation. The inflation response becomes significant at the 5% level only after 27 months for the RR shock but after 13 months for the new (RR) shock. Thus, our new shock shows that monetary policy shocks affect inflation at substantially shorter lags compared to the RR shock. The difference between inflation responses is particularly significant between 13 and 39 months, see panel (c) of Figure 3. Quantitatively, the trough response is -0.62 percentage points for

[30]The standard errors for the difference across impulse responses are constructed by estimating both local projections as seemingly unrelated regressions and estimating the joint covariance matrix via Driscoll-Kraay.

the RR shock but sizable -1.15 percentage points for the new (RR) shock.

Overall, our results suggest that accounting for time variation in systematic monetary policy is critically important when identifying monetary policy shocks. Disregarding variation in systematic monetary policy may lead to strongly biased impulse response estimates and an inaccurate assessment of the effectiveness of monetary policy. It may further bias analyses that use the estimated impulse responses for structural estimation (e.g., Barnichon and Mesters, 2020), to assess the optimality of monetary policy (e.g., Barnichon and Mesters, 2023), or to construct policy counterfactuals (e.g., McKay and Wolf, 2023).

Impulse response decomposition. To further understand the differences across the responses to the RR and the new shock, we propose an exact impulse response decomposition that can be implemented via the local projection from (4.2). We define the difference between RR and the new shock by $\hat{\varepsilon}_t^{sys} = \hat{\varepsilon}_t^{rr} - \hat{\varepsilon}_t^{new}$, where $\hat{\varepsilon}_t^{sys}$ captures interest rate fluctuations due to variation in systematic monetary policy that the RR shock treats as if it was an exogenous shock to the intercept of the Taylor rule. We then decompose the impulse response for any outcome z as

$$\beta_{z,rr}^h = \beta_{z,new}^h \, \omega_{new}^h + \beta_{z,sys}^h \, \omega_{sys}^h, \tag{4.3}$$

where $\beta_{z,i}^h$ denotes the response of z_{t+h} to $\hat{\varepsilon}_t^i$, for $i \quad \{rr, new, sys\}$, estimated from the local projection in (4.2). The weight $\omega_i^h = \mathsf{E}[((\hat{\varepsilon}_t^i) \quad)^2]/\mathsf{E}[((\hat{\varepsilon}_t^{rr}) \quad)^2]$ measures the variance contribution to the RR shock, and $(\hat{\varepsilon}_t^i) \quad$ denotes the shocks residualized with respect to the control vector. Two remarks are in order. First, the decomposition relies only on the properties of linear projections, notably the Frisch-Waugh-Lovell Theorem, and, therefore, holds exactly even in finite samples. Second, the decomposition does not rely on the particular relation of $\hat{\varepsilon}_t^{rr}$ and $\hat{\varepsilon}_t^{new}$. Instead, the decomposition can be applied to any two arbitrary shocks or regressors of interest.

Figure 4 presents the results for all three components of the decomposition. The responses to the new shock, $\hat{\varepsilon}_t^{new}$, and to systematic monetary policy (Sys MP), $\hat{\varepsilon}_t^{sys}$, are scaled by the respective variance weight so that they add up to the response to $\hat{\varepsilon}_t^{rr}$. Focusing on the impulse response to the systematic monetary policy term, $\hat{\varepsilon}_t^{sys}$, one can see that this component of the RR shock leads to a significant increase in the FFR, as expected. However, it also leads to a significant increase in real GDP and inflation. This suggests that $\hat{\varepsilon}_t^{sys}$ captures the systematic monetary policy response to demand shocks and not a shock to the intercept of the Taylor rule. This rationalizes that the responses to the RR shock are considerably attenuated compared to the responses to our new shock.

Figure 4: Decomposition of the responses to the RR shock

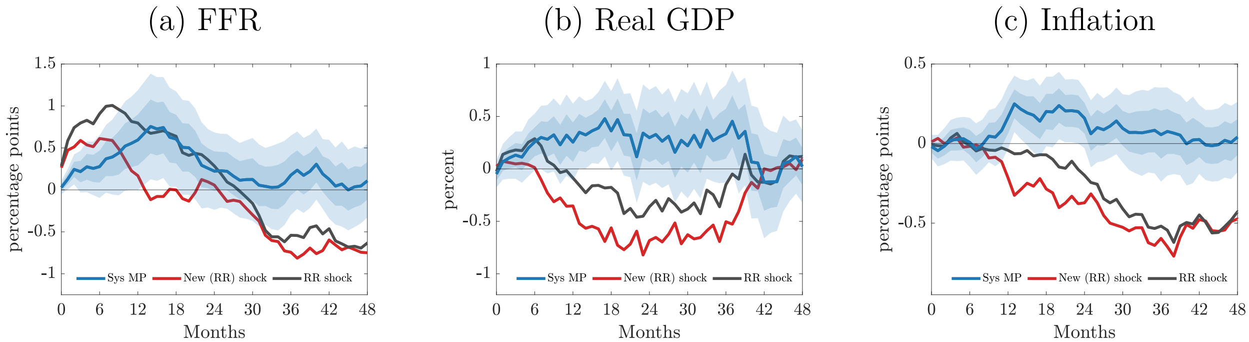

Notes: The figure shows an exact decomposition of the impulse responses to the RR shock into the responses to the new shock and the responses to the systematic monetary policy component (sys MP) that is contained in the RR shock, based on (4.3), and based on the local projection as specified along with (4.2). The outcomes are the federal funds rate, log real GDP, and the inflation rate. The new (RR) monetary policy shock is identified as the residual from the Taylor rule regression in (4.1) whereas the RR shock is based on the same regression when $b_j = 0$ for $j > 1$, as in Romer and Romer (2004). The shaded areas indicate 68% and 95% confidence bands using standard errors robust to serial correlation and heteroskedasticity.

Further outcome variables. In Figure 5, we extend the analysis to additional outcome variables, notably capacity utilization, unemployment, hours worked (in manufacturing), consumption, inventories, and a corporate credit spread. The local projection remains specified as for the main outcome variables presented above. The additional outcomes are informative about the transmission mechanism of monetary policy. The estimated responses further underscore that accounting for systematic monetary policy matters.

In response to the new (RR) shock, we find a decrease in capacity utilization, an increase in the unemployment rate, and a decrease in hours worked, all significant at the 5% level. All measures suggest an increase of slack in the economy. The responses to the RR shock are broadly similar. However, they suggest (again) a substantially longer lag of monetary policy, and the responses are less precisely estimated. The differences between RR and the new shock are significant at the 5% level for all variables and at many response horizons, see Figure C.2 in Appendix C.

The response of consumption expenditures to the new shock is much quicker and occurs within the first six months. Beyond the short-run, however, the response of consumption is similar across the new shock and the RR shock, suggesting that investment, government spending, or net exports respond quite differently to the two shocks. Business inventories initially increase, consistent with a surprise reduction in demand, and then fall. The reduction in inventories is significantly more pronounced for the new shock consistent with the more rapid decline in capacity utilization. Finally, the yield spread between BAA- and AAA-rated corporate bonds responds more strongly and significantly to the new shock.

Overall, the difference between the responses to the RR and new shock are similar as for the main outcome variables. Monetary policy transmission appears to be stronger and faster,

Figure 5: Response of further outcomes to monetary policy shocks

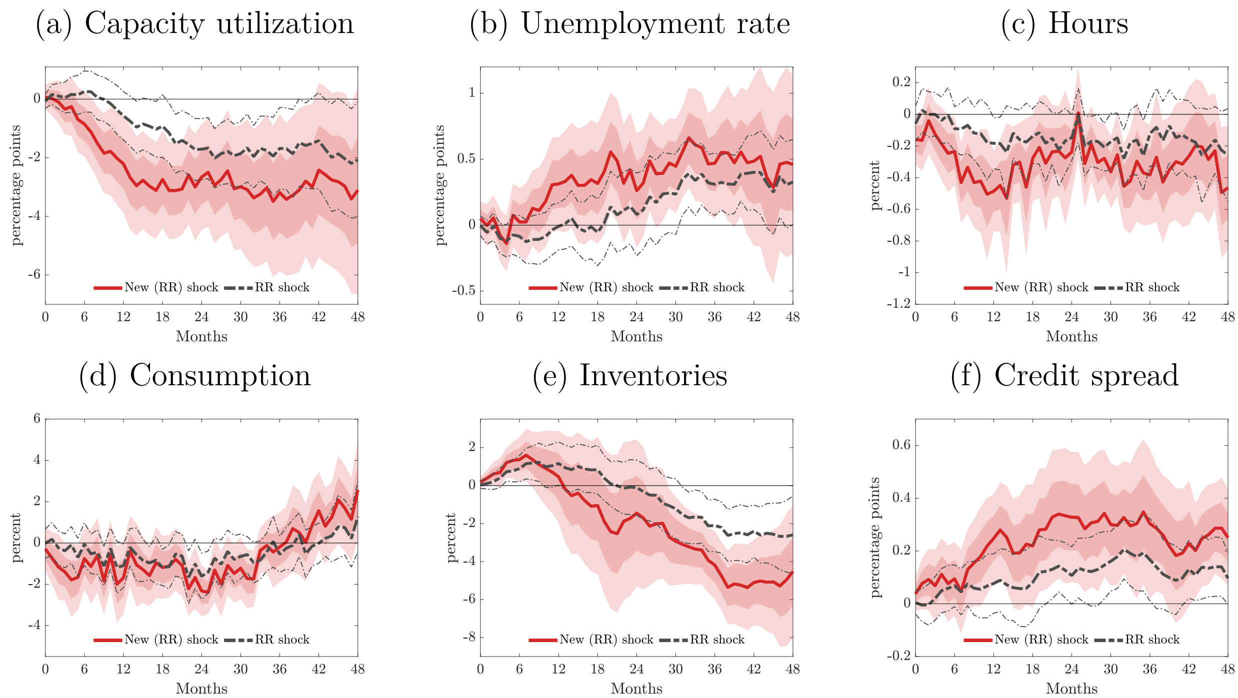

Notes: The figure shows responses of capacity utilization, the unemployment rate, log consumption expenditures, log business inventories, log hours (in manufacturing), and credit spreads (BAA- minus AAA-rated corporate bond yield) to a monetary policy shock based on the local projection as specified along with (4.2). The new (RR) monetary policy shock is identified as the residual from the Taylor rule regression in (4.1) whereas the RR shock is based on the same regression when $b_j = 0$ for $j > 1$, as in Romer and Romer (2004). The shaded areas indicate 68% and 95% confidence bands for the new shock, and the dotted lines indicate the 95% confidence band for the conventional shock using standard errors robust to serial correlation and heteroskedasticity for all bands.

with more significant responses, when using the new shock. This underscores the value of our new (RR) shock relative to the original RR shock.

Responses to new (AD) shocks. Section 3 provides evidence suggesting that the AD shock constructed by Aruoba and Drechsel (2022) is also contaminated by systematic monetary policy. Thus, we compare impulse responses to the AD shock with responses to a new (AD) shock, which we estimate as the residual in (4.1), when using the AD shock as left-hand side variable. Figure C.3 shows that the new (AD) shock leads to a more short-lived response of the FFR, a stronger decline of real GDP, and a substantially shorter lag in the inflation response, when compared to the original AD shock. The differences are sizeable and statistically significant at the 5% level for multiple horizons. Overall, impulse responses change in similar ways as for the RR shock.

Alternative shock identification. The additional regressors by which we augment the original RR regression are partly motivated by the evidence from Section 3. In particular, we choose to use the rotation hawk-dove balance since it performs best in predicting the existing shocks. We next investigate how the results change when using the overall hawk-dove balance, $Hawk^{A}_{\tau-1}$ instead. We provide the impulse responses to this alternative new shock in Figure C.4. We find that our baseline results are quite robust, although the responses to this version of the new shock become slightly less significant.

The regression to identify our new (RR) shock augments the original specification from RR by additional regressors involving $Hawk^{R}_{\tau-1}$ but also an additional lag of the RR regressors $x_{\tau-1}$ in levels. We do so to separately capture the level effects of the RR regressors and their dependence on systematic monetary policy. However, one may be concerned that the properties of our new shock are primarily driven by the inclusion of $x_{\tau-1}$. To address this concern, we run a regression where we drop all regressors involving $Hawk^{R}_{\tau-1}$ but keep $x_{\tau-1}$. In Figure C.5, we present the responses to this alternative shock and compare it with the original RR shock. We find that responses hardly differ, confirming that the inclusion of $Hawk^{R}_{\tau-1}$ in the regression is responsible for the change in responses with the new shock.

Alternative sample periods. Our new (RR) shock is estimated on the full sample of Greenbook forecasts from 1969 through 2007, but the impulse responses presented above are estimated on the post-Volcker disinflation subsample. We analyze whether our estimated responses differ if the shock identification regressions (4.1) for both the RR shock and our new shock are estimated on the post-Volcker disinflation subsample. Figure C.6 shows that the inflation response to the RR shock features a similarly long lag as in the baseline. The GDP response to the RR shock is insignificant but rather expansionary. In contrast, the response of inflation to our new shock remains similar to the baseline. The GDP response remains negative and significant at the 5% level for a few horizons.

We further estimate impulse responses on the full sample (1969-2007) and report the results in Figure C.7. Similar to the baseline, we find that the new (RR) shock delivers a significantly stronger contraction in real GDP. Interestingly, the inflation response is similar across both shocks for around two years and features a price puzzle.[31] At longer horizons, however, the new shock leads to a stronger inflation decline.

Additional control variables. Romer and Romer (2004) and Coibion (2012) impose a recursiveness assumption by including contemporaneous real GDP and inflation as control variables. In effect, these variables cannot contemporaneously respond to the monetary policy shock. Figure C.8 shows that our results are highly similar to the baseline imposing the recursiveness assumption. Parts of the related literature control for lags of the log

[31] Including twelve lags of the log commodity price index (or its growth rate) resolves the price puzzle.

S&P 500 and the Gilchrist and Zakrajšek (2012) excess bond premium (e.g., Jarociński and Karadi, 2020). Figure C.9 shows that our estimated responses are similar to the baseline when adding twelve lags of the two control variables. Finally, some of the related literature controls for lags of the RR shocks (see, e.g., Ramey, 2016). Figure C.10 shows that our results hardly change when adding twelve lags of the shock under consideration to the baseline set of control variables.

Alternative outcome variables. Our baseline results use interpolated real GDP and the GDP deflator to measure economic activity and prices at monthly frequency as similarly done in Jarociński and Karadi (2020) and Aruoba and Drechsel (2022). An alternative is to use industrial production (IP) and CPI inflation, which are readily available at monthly frequency (e.g. Gertler and Karadi, 2015; Bauer and Swanson, 2023b). Figure C.11 shows the responses of IP and CPI. The differences between the new and the RR shock remain similar to the baseline. However, the IP response is less precisely estimated (compared with the GDP response) for the new shock. If we further control for twelve lags of the Gilchrist and Zakrajšek (2012) excess bond premium and the log S&P 500, then the IP response becomes statistically significant at the 5% level for many response horizons, see Figure C.12.

4.3 High-frequency identified monetary policy shocks

In Section 3.3, we demonstrated that even the high-frequency identified monetary policy shocks from Miranda-Agrippino and Ricco (2021) (MAR) are predictable by fluctuations in systematic monetary policy. Thus, we hypothesize that responses to their shock may be biased, similar to the RR shock. Our empirical evidence is consistent with this hypothesis.

A new (MAR) shock. As in Section 3.3, our investigation rests on MAR's identified monetary policy shock from their six-variate proxy VAR model. We propose a new version of their shock, the new (MAR) shock, that is given by the regression residual from the predictive regression from Section 3.3. Specifically, we use the specification with $Hawk^R_{t-p}$ and $p = 1$, consistent with (4.1) that yields the new RR shock. The monthly sample runs from 1980 until 2014. In the scope of this paper, a downside of the Miranda-Agrippino and Ricco (2021) shocks is that we cannot clean the underlying instrument, the high-frequency monetary policy surprises, from contamination arising from expectation revisions about systematic monetary policy. The reason is that we only have a low-frequency measure of time variation in systematic monetary policy, but no high-frequency measure of expectation revisions around monetary announcements.

Figure 6: Responses of main outcomes to high-frequency identified shocks

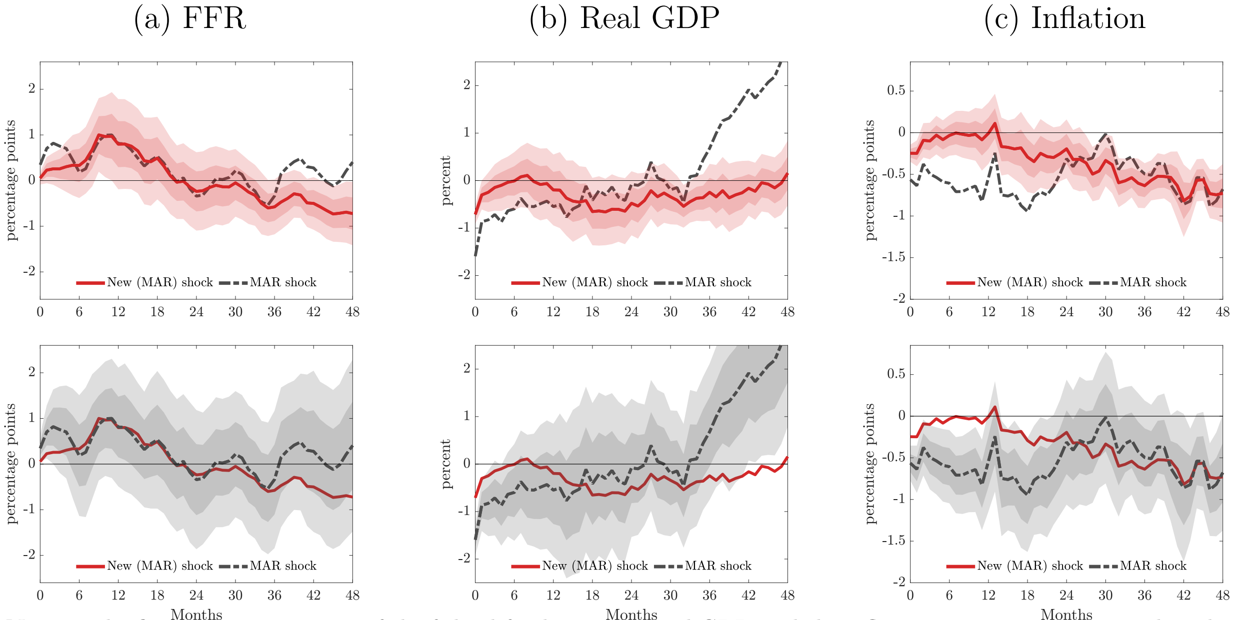

Notes: The figure shows responses of the federal funds rate, log real GDP, and the inflation rate to a monetary policy shock based on the local projection as specified along with (4.2). The new (MAR) shock is identified as the residual from the predictive regression in (3.4) where we put the MAR shock on the left-hand side. The MAR shock is taken directly from the six-variate VAR in Miranda-Agrippino and Ricco (2021). The shaded areas indicate 68% and 95% confidence bands using standard errors robust to serial correlation and heteroskedasticity.

Responses to the new (MAR) shock. To analyze the differences in impulse responses between the original and our new (MAR) shock, we run the local projection from (4.2) as specified in Section 4.2 with the estimation sample from 1983M1 to 2007M12. This choice facilitates comparability with our baseline results.

Figure 6 presents the estimated responses. The FFR responds persistently to the MAR shock, but the estimates are largely insignificant at the 5% level. Instead, the FFR response to the new (MAR) shock is slightly more significant for some horizons. The difference between both responses is small. The MAR shock leads to a significant fall in GDP and inflation in the month of the impact. The GDP response is very transitory. Instead, the GDP response to the new (MAR) shock reaches its trough only after roughly two years and is considerably smaller. For inflation, we find that the MAR shock delivers a more persistent decline that fades away after two years. Instead, the inflation response to the new (MAR) shock is zero for around one year before inflation starts to fall. This pattern somewhat resembles the inflation response to the new (RR) shock. However, the inflation response to the new (MAR) shock becomes significant with a substantially longer delay.

Overall, there are two key results. First, the high-frequency identified shock from MAR performs relatively well compared with the original RR shock, as the responses have signs consistent with theory, and are statistically significant. Second, the impulse responses to the

new (MAR) shock differ meaningfully from the responses to the original MAR shock.
Finally, we acknowledge that the responses to the original MAR shock differ from the responses reported in their paper. This is not surprising because our setup necessarily yields an imperfect replication of the original results by MAR since the sample, the choice of control variables, and the estimation method differ. That said, our local projection approach is more immune to misspecification (Olea, Plagborg-Møller, Qian, and Wolf, 2024). Further, when the MAR shock is well identified, then the results should be robust across subsamples, and the choice of lagged control variables is irrelevant for causal identification in large samples.

5 Conclusion

This paper revisits conventional empirical strategies to estimate monetary policy shock series. We show theoretically that fluctuations in systematic monetary policy lead to contaminated shocks and bias in the estimated impulse responses for a broad set of identification strategies. These strategies include Taylor rule-type regressions (e.g., Romer and Romer, 2004) and linear monetary VAR models using exclusion restrictions (e.g., Christiano et al., 1999), sign restrictions (e.g., Uhlig, 2005), narrative restrictions (e.g., Antolín-Díaz and Rubio-Ramírez, 2018), or external instruments (e.g., Gertler and Karadi, 2015). The problematic assumption common among these approaches is that systematic monetary policy ought to be constant over time. Similar problems arise for high-frequency identified monetary policy surprises, which also impose a time-constancy assumption regarding beliefs about systematic monetary policy around monetary policy announcements.
Our theory predicts that contaminated monetary policy shocks should be predictable by measured time-variation in systematic monetary policy. We provide empirical evidence suggesting that the monetary policy shocks from Romer and Romer (2004), Aruoba and Drechsel (2022), and Miranda-Agrippino and Ricco (2021) are indeed predictable and contaminated. To address this problem, we construct respective new shock series that are orthogonal to systematic monetary policy and assess their effects on the U.S. economy. Importantly, evidence from our new shocks suggests that monetary policy has shorter lags and stronger effects on inflation and output compared to the corresponding evidence for the shocks in Romer and Romer (2004) and Aruoba and Drechsel (2022).

References

Abrams, B. (2006): "How Richard Nixon Pressured Arthur Burns: Evidence from the Nixon Tapes," *Journal of Economic Perspectives*, 20, 177–188.

Antolín-Díaz, J. and J. F. Rubio-Ramírez (2018): "Narrative Sign Restrictions for SVARs," *American Economic Review*, 108, 2802–29.

Aruoba, B. and T. Drechsel (2022): "Identifying Monetary Policy Shocks: A Natural Language Approach," CEPR Discussion Papers 17133, C.E.P.R. Discussion Papers.

Barakchian, S. M. and C. Crowe (2013): "Monetary Policy Matters: Evidence from New Shocks Data," *Journal of Monetary Economics*, 60, 950–966.

Barnichon, R. and G. Mesters (2020): "Identifying Modern Macro Equations with Old Shocks," *Quarterly Journal of Economics*, 135, 2255–2298.

——— (2023): "A Sufficient Statistics Approach for Macro Policy," *American Economic Review*, 113, 2809–2845.

Bauer, M. D., C. Pflueger, and A. Sunderam (2022): "Perceptions about Monetary Policy," NBER Working Paper 30480.

Bauer, M. D. and E. T. Swanson (2023a): "An Alternative Explanation for the "Fed Information Effect"," *American Economic Review*, 113, 664–700.

——— (2023b): "A Reassessment of Monetary Policy Surprises and High-Frequency Identification," *NBER Macroeconomics Annual*, 37, 87–155.

Bernanke, B. S. and A. S. Blinder (1992): "The Federal Funds Rate and the Channels of Monetary Transmission," *American Economic Review*, 82, 901–921.

Bernanke, B. S., M. Gertler, and M. Watson (1997): "Systematic Monetary Policy and the Effects of Oil Price Shocks," *Brookings Papers on Economic Activity*, 1997, 91–157.

Bianchi, F., R. Gómez-Cram, T. Kind, and H. Kung (2023): "Threats to Central Bank Independence: High-Frequency Identification with Twitter," *Journal of Monetary Economics*, 135, 37–54.

Boivin, J. (2006): "Has U.S. Monetary Policy Changed? Evidence from Drifting Coefficients and Real-Time Data," *Journal of Money, Credit and Banking*, 38, 1149–1173.

Bordo, M. D. and K. Istrefi (2023): "Perceived FOMC: The Making of Hawks, Doves and Swingers," *Journal of Monetary Economics*, forthcoming.

Carvalho, C., F. Nechio, and T. Tristão (2021): "Taylor Rule Estimation by OLS," *Journal of Monetary Economics*, 124, 140–154.

Champagne, J. and R. Sekkel (2018): "Changes in Monetary Regimes and the Identification of Monetary Policy Shocks: Narrative Evidence from Canada," *Journal of Monetary Economics*, 99, 72–87.

Choi, S., T. Willems, and S. Y. Yoo (2024): "Revisiting the Monetary Transmission Mechanism Through an Industry-Level Differential Approach," *Journal of Monetary Economics*, 103556.

Chow, G. C. and A. Lin (1971): "Best Linear Unbiased Interpolation, Distribution, and

Extrapolation of Time Series by Related Series," *Review of Economics and Statistics*, 53, 372–375.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (1999): "Monetary Policy Shocks: What Have We Learned and to What End?" *Handbook of Macroeconomics*, 1, 65–148.

——— (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113, 1–45.

Clarida, R., J. Gali, and M. Gertler (2000): "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics*, 115, 147–180.

Cloyne, J. and P. Hürtgen (2016): "The Macroeconomic Effects of Monetary Policy: A New Measure for the United Kingdom," *American Economic Journal: Macroeconomics*, 8, 75–102.

Cloyne, J., P. Hürtgen, and A. M. Taylor (2022): "Global Monetary and Financial Spillovers: Evidence from a New Measure of Bundesbank Policy Shocks," Tech. rep., National Bureau of Economic Research.

Cochrane, J. H. (2011): "Determinacy and Identification with Taylor Rules," *Journal of Political Economy*, 119, 565–615.

Coibion, O. (2012): "Are the Effects of Monetary Policy Shocks Big or Small?" *American Economic Journal: Macroeconomics*, 4, 1–32.

Davig, T. and E. M. Leeper (2008): "Endogenous Monetary Policy Regime Change," in *NBER International Seminar on Macroeconomics 2006*, National Bureau of Economic Research, NBER Chapters, 345–391.

Galí, J. (2015): *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*, Princeton University Press.

Gertler, M. and P. Karadi (2015): "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, 7, 44–76.

Gilchrist, S. and E. Zakrajšek (2012): "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, 102, 1692–1720.

Hack, L., K. Istrefi, and M. Meier (2023): "Identification of Systematic Monetary Policy," CEPR Discussion Paper 17999.

Holm, M. B., P. Paul, and A. Tischbirek (2021): "The Transmission of Monetary Policy under the Microscope," *Journal of Political Economy*, 129, 2861–2904.

Istrefi, K. (2019): "In Fed Watchers' Eyes: Hawks, Doves and Monetary Policy," Banque de France Working Paper.

Jarociński, M. and P. Karadi (2020): "Deconstructing Monetary Policy Surprises—the Role of Information Shocks," *American Economic Journal: Macroeconomics*, 12, 1–43.

McKay, A. and C. K. Wolf (2023): "What Can Time-Series Regressions Tell Us About Policy Counterfactuals?" *Econometrica*, 91, 1695–1725.

McMahon, M. and T. Munday (2023): "Element(s) of Surprise," *Essays on Monetary Policy*, Chapter 4 of Tim Munday's dissertation.

Miranda-Agrippino, S. and G. Ricco (2021): "The Transmission of Monetary Policy Shocks," *American Economic Journal: Macroeconomics*, 13, 74–107.

Nakamura, E. and J. Steinsson (2018): "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect," *Quarterly Journal of Economics*, 133, 1283–1330.

Olea, J. L. M., M. Plagborg-Møller, E. Qian, and C. K. Wolf (2024): "Double Robustness of Local Projections and Some Unpleasant VARithmetic," Tech. rep., National Bureau of Economic Research.

Orphanides, A. (2004): "Monetary Policy Rules, Macroeconomic Stability, and Inflation: A View from the Trenches," *Journal of Money, Credit and Banking*, 36, 151–175.

Owyang, M. and G. Ramey (2004): "Regime switching and monetary policy measurement," *Journal of Monetary Economics*, 51, 1577–1597.

Primiceri, G. E. (2005): "Time Varying Structural Vector Autoregressions and Monetary Policy," *Review of Economic Studies*, 72, 821–852.

Ramey, V. (2016): "Macroeconomic Shocks and Their Propagation," *Handbook of Macroeconomics*, 2, 71 – 162.

Romer, C. and D. Romer (1989): "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz," in *NBER Macroeconomics Annual 1989, Volume 4*, National Bureau of Economic Research, Inc, 121–184.

Romer, C. D. and D. H. Romer (2004): "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 94, 1055–1084.

Sims, C. A. and T. Zha (2006): "Were There Regime Switches in US Monetary Policy?" *American Economic Review*, 96, 54–81.

Uhlig, H. (2005): "What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure," *Journal of Monetary Economics*, 52, 381–419.

Wieland, J. F. and M.-J. Yang (2020): "Financial Dampening," *Journal of Money, Credit and Banking*, 52, 79–113.

Appendix A Data

Table A.1: Descriptive statistics of the Hawk-Dove balances

	Mean	Median	SD	Autocorr	Corr	Min	Max	T
$Hawk_\tau^F$	0.06	0.10	0.34	0.95	-	-0.80	0.73	630
$Hawk_\tau^R$	0.24	0.25	0.47	0.91	0.60	-0.75	1.00	630

Notes: This table shows descriptive statistics for the time series at FOMC meeting frequency from 1960 through 2023. $Hawk_\tau^F$ is the average Hawk-Dove balance of the FOMC. $Hawk_\tau^R$ is the FOMC rotation instrument. "Autocorr" refers to the meeting-over-meeting autocorrelation. "Corr" refers to the correlation between both series.

Appendix B Additional results for Section 3

Table B.1: Explaining RR shocks by systematic monetary policy, adjusted R^2

	$Hawk_{\tau}^{F}$			$Hawk_{\tau}^{R}$		
Sample	69-07	69-96	83-07	69-07	69-96	83-07
	(a) Contemporaneous FOMC meeting (p=0)					
adjusted R^2	-0.072	-0.099	0.201	0.007	0.005	0.251
p-value	0.189	0.243	0.000	0.012	0.000	0.000
T	353	265	200	353	265	200
	(b) One FOMC meeting lag (p=1)					
adjusted R^2	0.205	0.275	0.238	0.323	0.417	0.222
p-value	0.003	0.001	0.000	0.000	0.000	0.000
T	349	261	200	349	261	200
	(c) Two FOMC meetings lag (p=2)					
adjusted R^2	0.095	0.119	0.123	0.139	0.181	0.197
p-value	0.000	0.000	0.000	0.000	0.000	0.000
T	347	259	200	347	259	200

Notes: The table shows results from regressions based on (3.4). The rows of the three subtables show the adjusted R^2, the p-values for the null hypothesis that all coefficient estimates are jointly zero, and the number of observations T. The three left columns show results for the Hawk-Dove balance across all FOMC members, and the three right columns for the Hawk-Dove balance across all rotating FOMC members with voting rights. The three subtables differ by the specification of FOMC meeting lag p. Columns one to three differ by the sample period between 1969-2007, 1969-1996, and 1983-2007, and analogously for columns four to six.

Table B.2: Lasso estimation to explain RR shocks

	(1)	(2)	(3)	(4)
$\Delta Hawk^{R}_{\tau-1} \times \Delta\pi_{\tau-1,0}$	0.212	0.159	0.266	0.234
	(0.153)	(0.106)	(0.123)	(0.118)
$\Delta Hawk^{R}_{\tau-1} \times y_{\tau-1,2}$		-0.146	-0.111	-0.098
		(0.163)	(0.108)	(0.099)
$\Delta Hawk^{R}_{\tau-1} \times \Delta\pi_{\tau-1,1}$			0.239	0.232
			(0.116)	(0.106)
$\Delta Hawk^{R}_{\tau-1} \times \pi_{\tau-1,2}$				0.096
				(0.084)
Constant	-0.004	0.001	0.006	0.006
	(0.020)	(0.017)	(0.017)	(0.017)
Observations	349	349	349	349
R^2	0.046	0.069	0.139	0.149

Standard errors in parentheses

Notes: The table shows Lasso regression results based on (3.4). The Lasso shrinkage parameter is chosen to increment the number of regressors from one to four, and the associated results are presented in columns one to five, respectively. The time sample runs from 1969 through 2007, and standard errors robust to serial correlation and heteroskedasticity are in parentheses.

Table B.3: Explaining AD shocks by systematic monetary policy

	$Hawk_{\tau}^{F}$		$Hawk_{\tau}^{R}$	
Sample	83-07	83-96	83-07	83-96
(a) Contemp. FOMC meeting (p=0)				
R^2	0.315	0.649	0.328	0.616
p-value	0.000	0.000	0.000	0.000
T	192	104	192	104
(b) One FOMC meeting lag (p=1)				
R^2	0.291	0.600	0.263	0.629
p-value	0.000	0.000	0.000	0.000
T	192	104	192	104
(c) Two FOMC meetings lag (p=2)				
R^2	0.330	0.515	0.362	0.668
p-value	0.000	0.000	0.000	0.000
T	192	104	192	104

Notes: The table shows results from regressions based on (3.4), but the left-hand side is the shock from Aruoba and Drechsel (2022), which is identified via their ridge regression. The rows of the three subtables show R^2, the p-values for the null hypothesis that all coefficient estimates are jointly zero, and the number of observations T. The two left columns show results for the Hawk-Dove balance across all FOMC members, and the two right columns for the Hawk-Dove balance across all rotating FOMC members with voting rights. The three subtables differ by the specification of FOMC meeting lag p. Columns one to two differ by the sample period between 1983-2007 and 1983-1996, and analogously for columns three to four.

Appendix C Additional results for Section 4

Figure C.1: Time series of monetary policy shocks

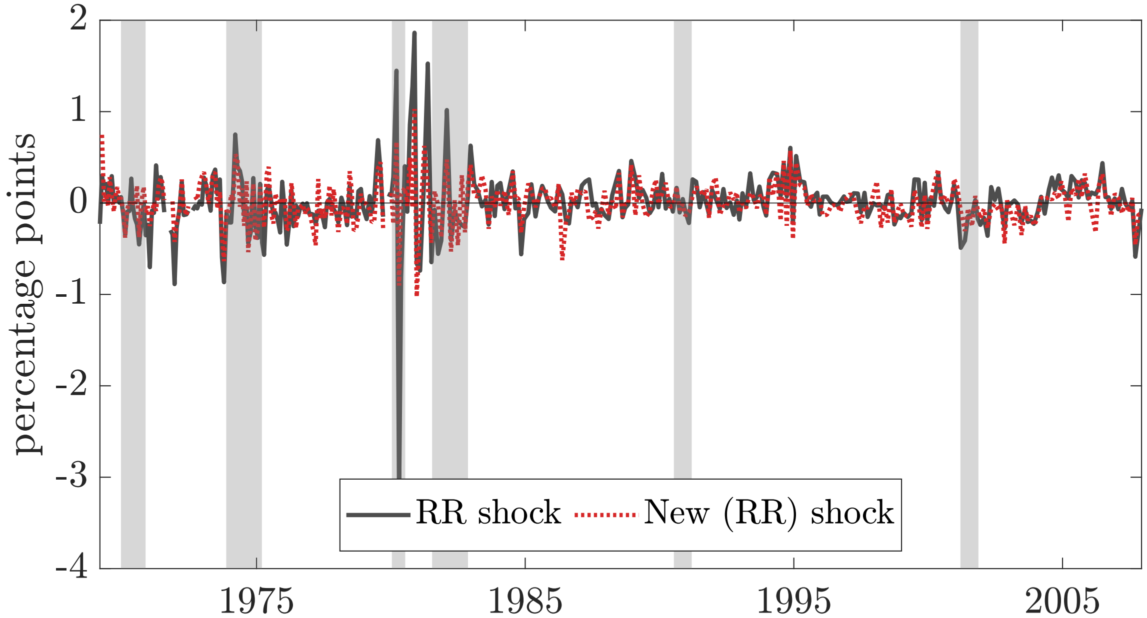

Notes: The solid black line shows the RR shock $\hat{e}_{\tau}^{rr}$ based on the regression in (4.1) when restricting $\beta_j = 0 \quad j > 1$. The dotted red line shows the new shock $\hat{e}_{\tau}^{new}$ based on the regression in (4.1). The sample period is 1969 through 2007. Grey bars indicate NBER recession.

Figure C.2: Response of further outcomes to new shock “minus” response to RR shock

Notes: The figure shows the differences across impulse responses for capacity utilization, the unemployment rate, log consumption expenditures, log business inventories, log hours (in manufacturing), and credit spreads (BAA- minus AAA-rated corporate bond yield) to a monetary policy shock based on the local projection as specified along with (4.2). The difference is computed as the response to the new shock minus the response to the old shock for each outcome, respectively. The new (RR) monetary policy shock is identified as the residual from the Taylor rule regression in (4.1) whereas the RR shock is based on the same regression when $b_j = 0$ for $j > 1$, as in Romer and Romer (2004). The shaded areas indicate 68% and 95% confidence bands using standard errors robust to serial correlation and heteroskedasticity.

Figure C.3: Comparison of main responses with Aruoba and Drechsel (2022) shock

(a) FFR (b) Real GDP (c) Inflation

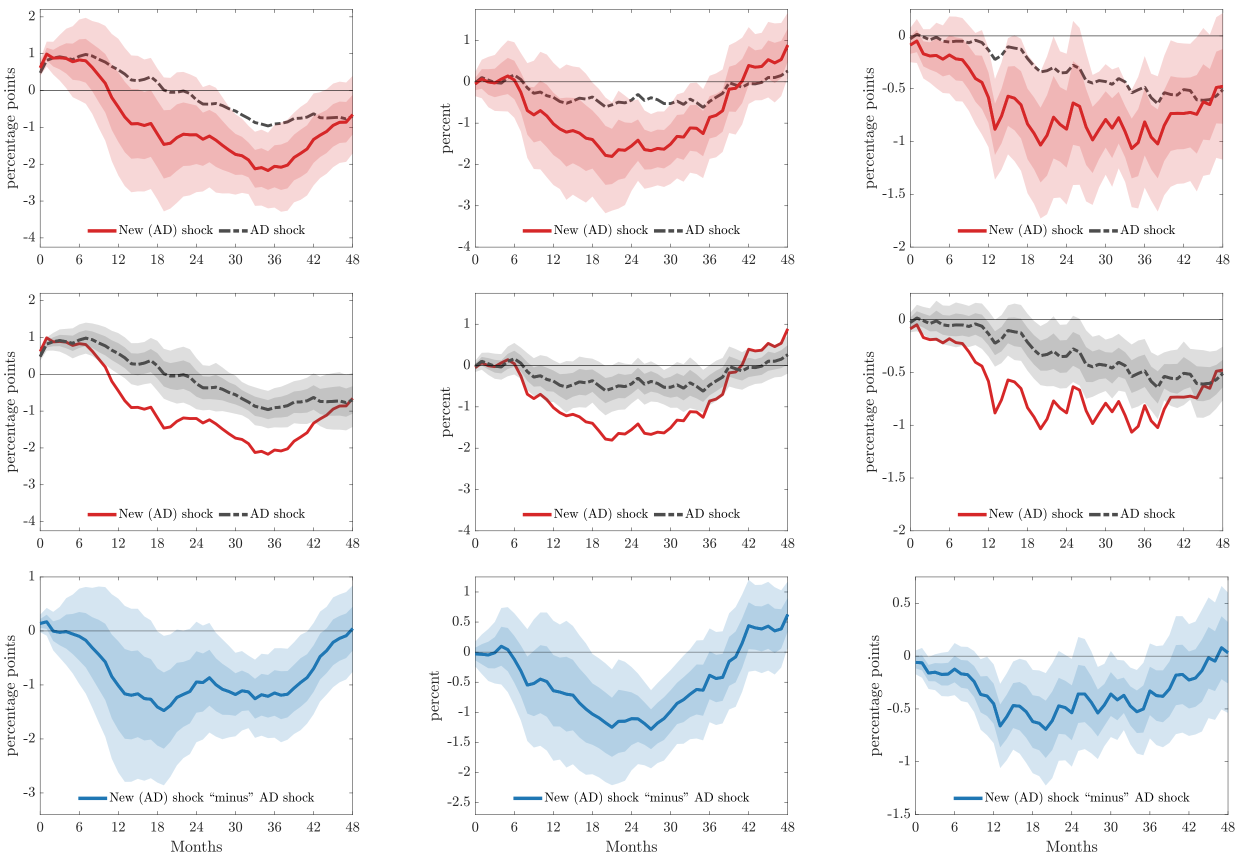

Notes: The figure shows responses of the federal funds rate, log real GDP, and the inflation rate to a monetary policy shock based on the local projection as specified along with (4.2). The new (AD) monetary policy shock is identified as the residual from the Taylor rule regression in (4.1) where we put the AD shock on the left-hand side. The AD shock is taken directly from the ridge regression in Aruoba and Drechsel (2022). The shaded areas indicate 68% and 95% confidence bands using standard errors robust to serial correlation and heteroskedasticity.

Figure C.4: Responses when using $Hawk^{A}_{\tau-1}$ in augmented RR regression

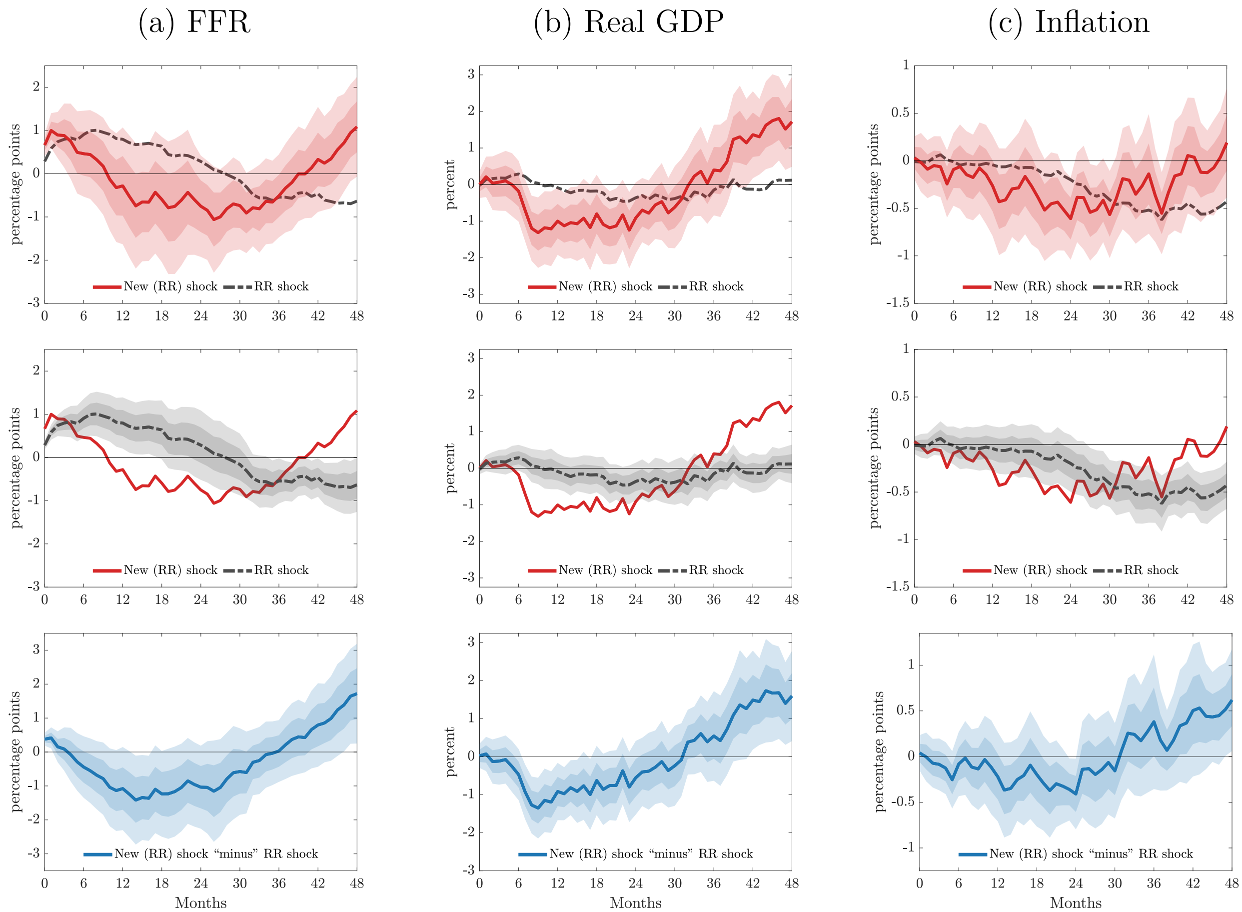

Notes: The figure shows responses of the federal funds rate, log real GDP, and the inflation rate to a monetary policy shock based on the local projection as specified along with (4.2). We use $Hawk^{A}_{\tau-1}$ instead of $Hawk^{R}_{\tau-1}$ in the augmented RR regression in (4.1). The new (RR) monetary policy shock is identified as the residual from the Taylor rule regression in (4.1) whereas the RR shock is based on the same regression when $b_j = 0$ for $j > 1$, as in Romer and Romer (2004). Columns 1 and 2 display the response to the new shock and conventional shock, respectively. Column 3 display the response to the new shock minus the response to the conventional shock. The shaded areas indicate 68% and 95% confidence bands using standard errors robust to serial correlation and heteroskedasticity.

Figure C.5: Responses when dropping the hawk-dove balance from augmented RR regression

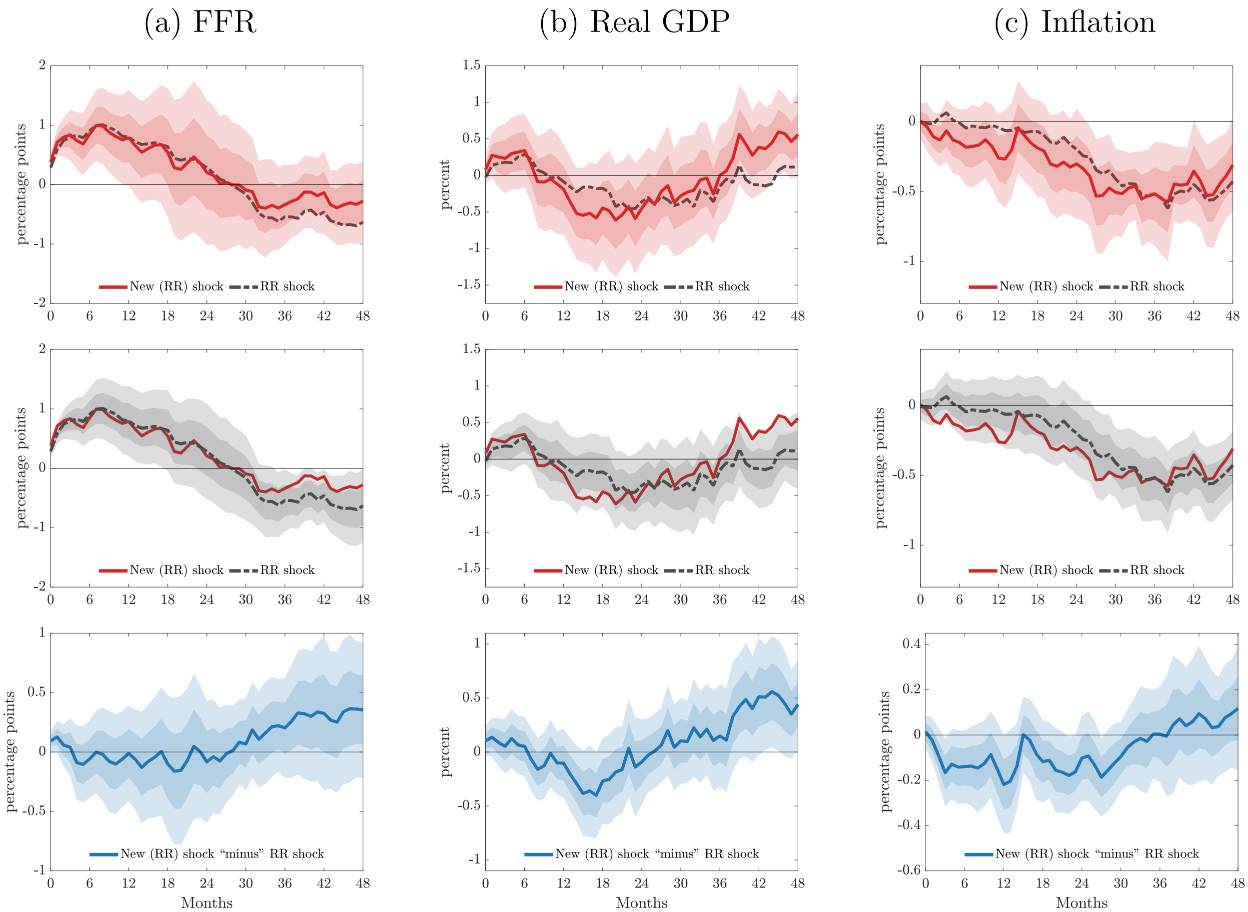

Notes: The figure shows responses of the federal funds rate, log real GDP, and the inflation rate to a monetary policy shock based on the local projection as specified along with (4.2). We drop all regressors involving the hawk-dove balance from the augmented RR regression in (4.1), i.e., we impose $\beta_i = 0, \quad i > 2$. The new (RR) monetary policy shock is identified as the residual from the Taylor rule regression in (4.1) whereas the RR shock is based on the same regression when $b_j = 0$ for $j > 1$, as in Romer and Romer (2004). Columns 1 and 2 display the response to the new shock and conventional shock, respectively. Column 3 display the response to the new shock minus the response to the conventional shock. The shaded areas indicate 68% and 95% confidence bands using standard errors robust to serial correlation and heteroskedasticity.

Figure C.6: Responses for identification sample 1983-2007

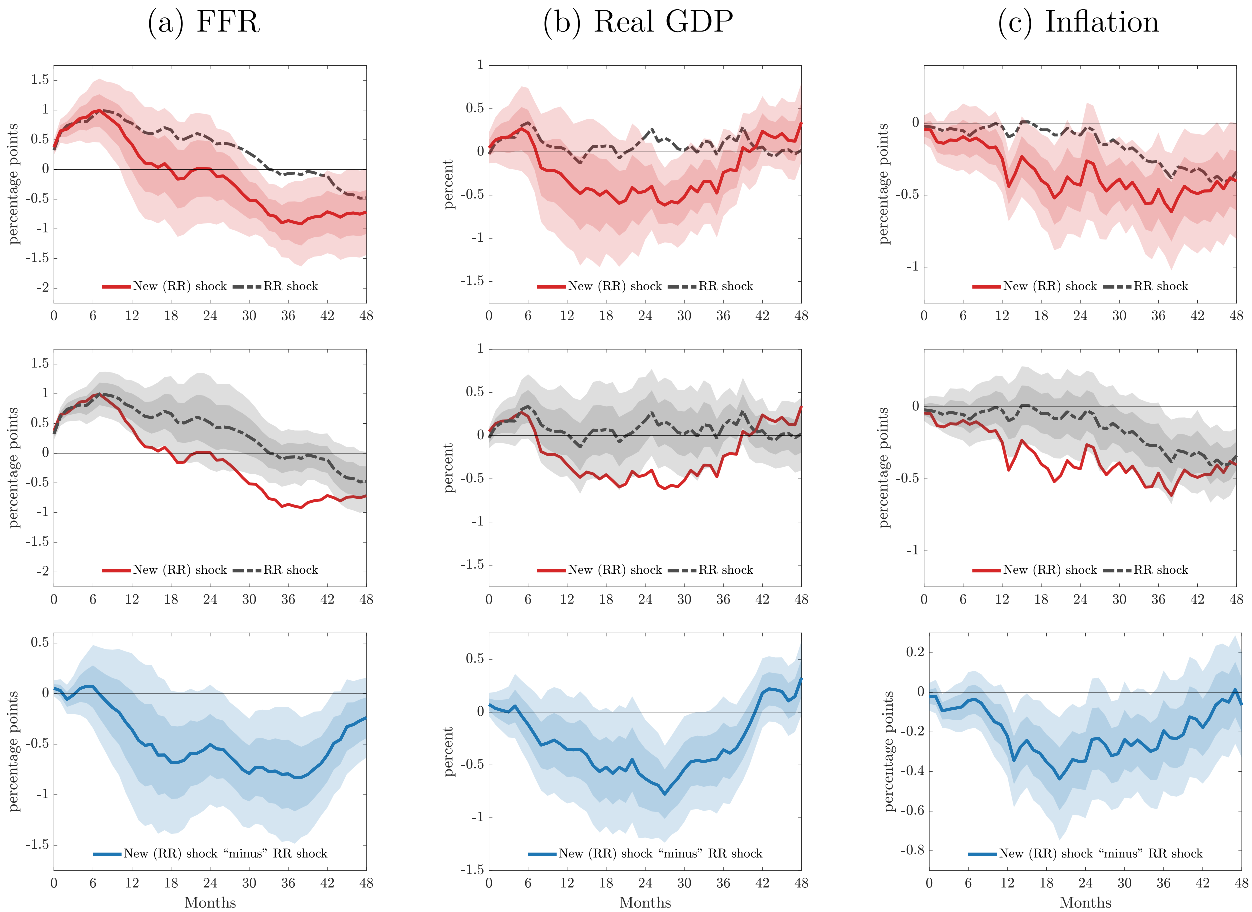

Notes: The figure shows responses of the federal funds rate, log real GDP, and the inflation rate to a monetary policy shock based on the local projection as specified along with (4.2). The new (RR) monetary policy shock is identified as the residual from the Taylor rule regression in (4.1) whereas the RR shock is based on the same regression when $b_j = 0$ for $j > 1$, as in Romer and Romer (2004). The estimation sample for shock identification coincides with the impulse response estimation sample, running from 1983 until 2007. Columns 1 and 2 display the response to the new shock and conventional shock, respectively. Column 3 display the response to the new shock minus the response to the conventional shock. The shaded areas indicate 68% and 95% confidence bands using standard errors robust to serial correlation and heteroskedasticity.

Figure C.7: Responses for estimation sample 1969-2007

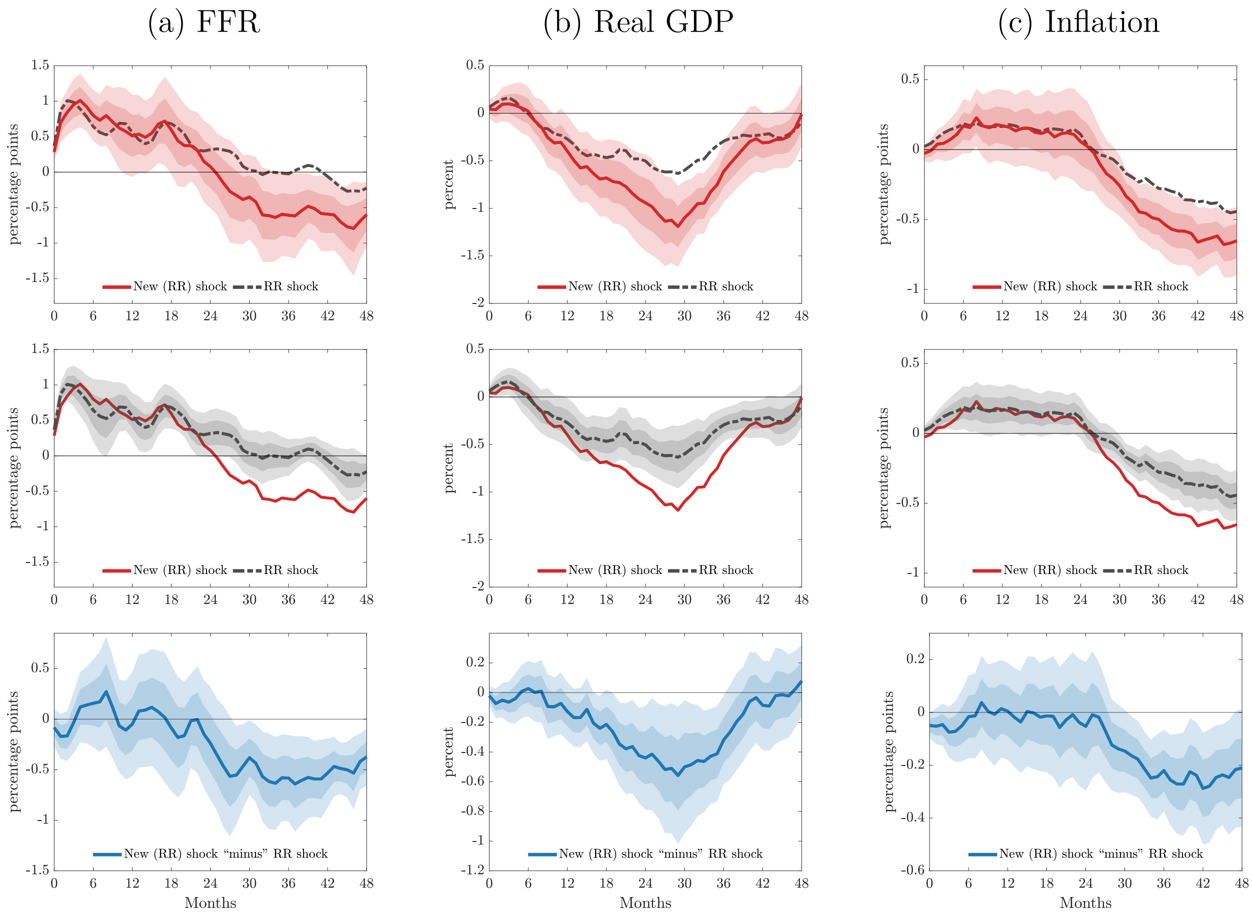

Notes: The figure shows responses of the federal funds rate, log real GDP, and the inflation rate to a monetary policy shock based on the local projection as specified along with (4.2). The new (RR) monetary policy shock is identified as the residual from the Taylor rule regression in (4.1) whereas the RR shock is based on the same regression when $b_j = 0$ for $j > 1$, as in Romer and Romer (2004). The results correspond to the full sample, running from 1969 until 2007. Columns 1 and 2 display the response to the new shock and conventional shock, respectively. Column 3 display the response to the new shock minus the response to the conventional shock. The shaded areas indicate 68% and 95% confidence bands using standard errors robust to serial correlation and heteroskedasticity.

Figure C.8: Responses when imposing a recursiveness assumption

(a) FFR (b) Real GDP (c) Inflation

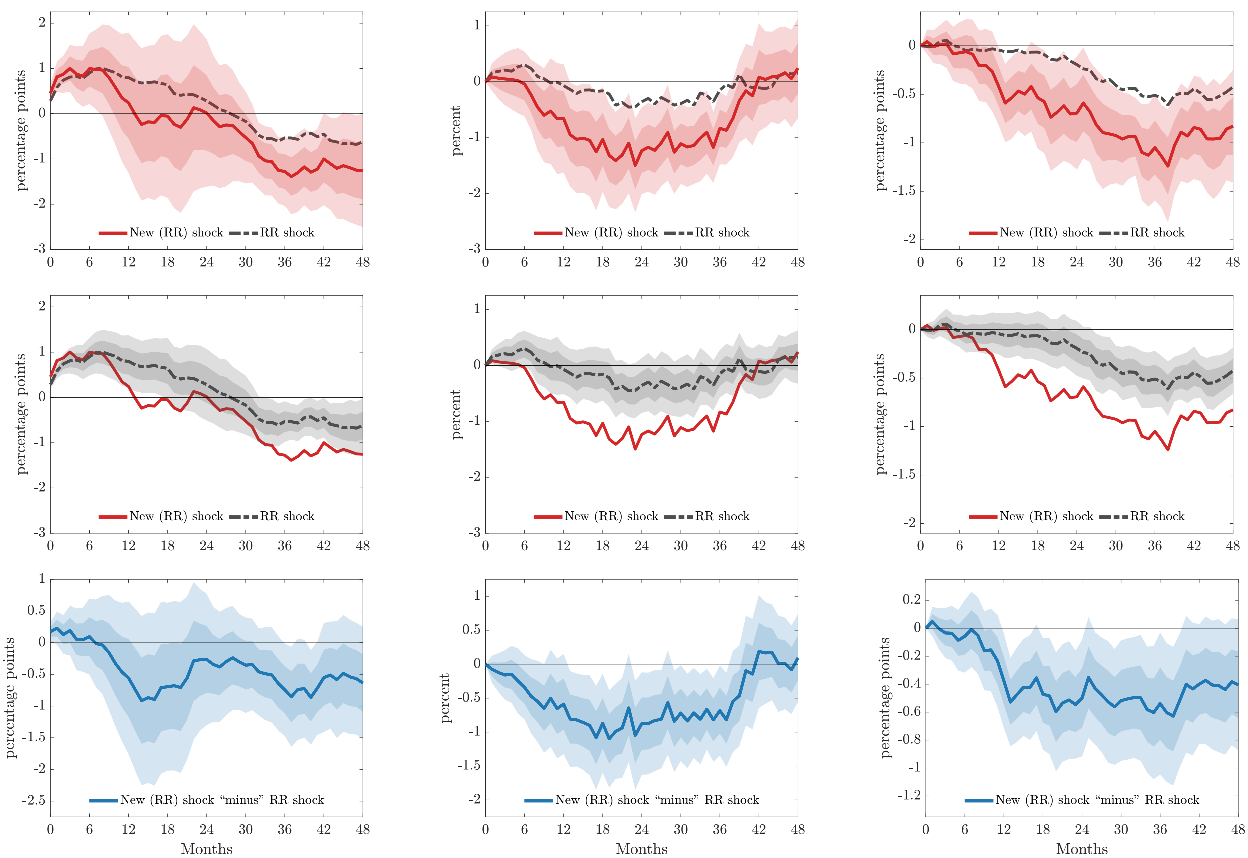

Notes: The figure shows responses of the federal funds rate, log real GDP, and the inflation rate to a monetary policy shock based on the local projection as specified along with (4.2). Additionally, we control for contemporaneous log real GDP and inflation imposing the recursiveness assumption that monetary policy shocks affect these variables only with a one-month lag. The new (RR) monetary policy shock is identified as the residual from the Taylor rule regression in (4.1) whereas the RR shock is based on the same regression when $b_j = 0$ for $j > 1$, as in Romer and Romer (2004). Columns 1 and 2 display the response to the new shock and conventional shock, respectively. Column 3 display the response to the new shock minus the response to the conventional shock. The shaded areas indicate 68% and 95% confidence bands using standard errors robust to serial correlation and heteroskedasticity.

Figure C.9: Responses when controlling for S&P 500 and EBP

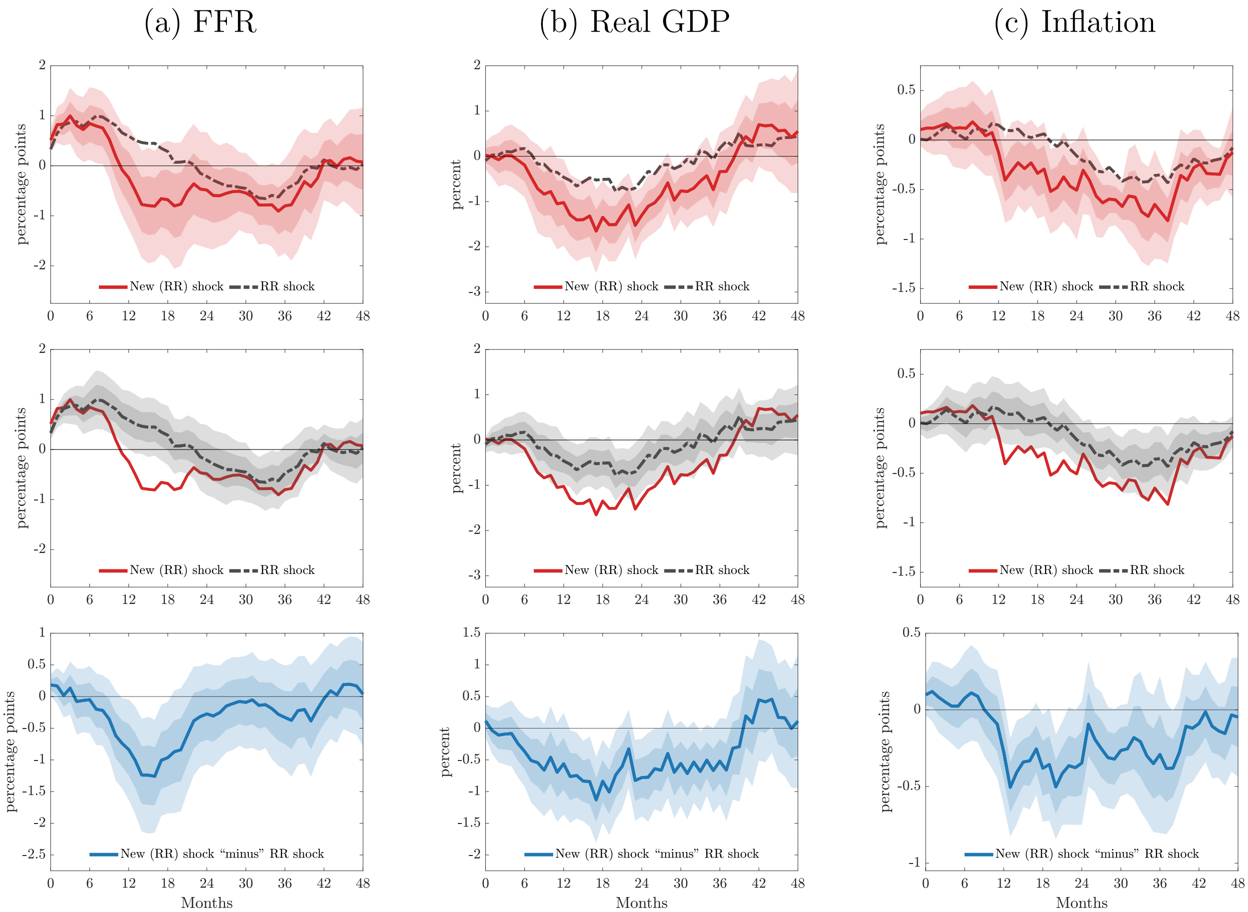

Notes: The figure shows responses of the federal funds rate, log real GDP, and the inflation rate to a monetary policy shock based on the local projection as specified along with (4.2). Additionally, we control for 12 lags of both, the S&P 500 and the excess bond premium from Gilchrist and Zakrajšek (2012). The new (RR) monetary policy shock is identified as the residual from the Taylor rule regression in (4.1) whereas the RR shock is based on the same regression when $b_j = 0$ for $j > 1$, as in Romer and Romer (2004). Columns 1 and 2 display the response to the new shock and conventional shock, respectively. Column 3 display the response to the new shock minus the response to the conventional shock. The shaded areas indicate 68% and 95% confidence bands using standard errors robust to serial correlation and heteroskedasticity.

Figure C.10: Responses when controlling for lagged shocks

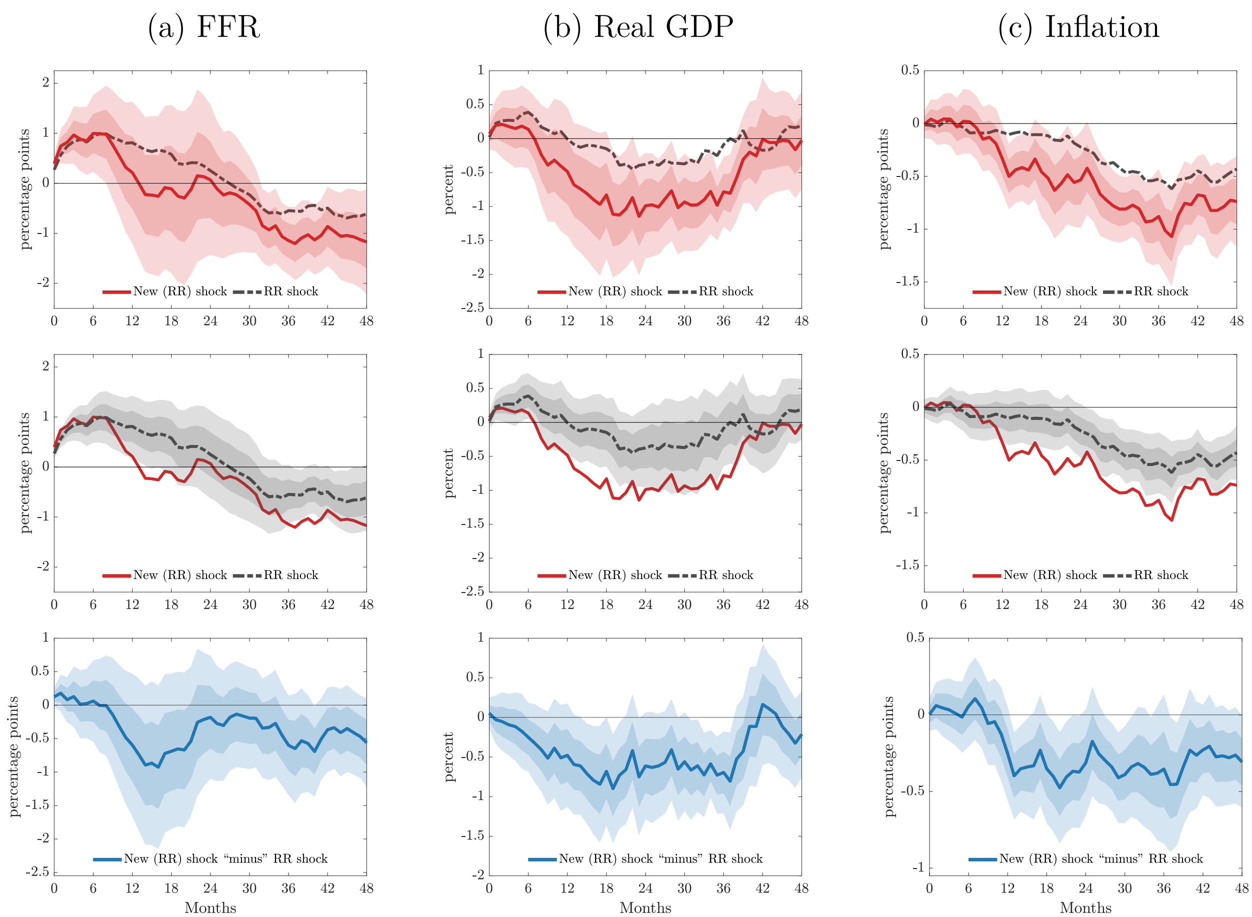

Notes: The figure shows responses of the federal funds rate, log real GDP and the inflation rate to a monetary policy shock based on the local projection as specified along with (4.2). Additionally, we control for 12 lags of monetary policy shock under consideration. The new (RR) monetary policy shock is identified as the residual from the Taylor rule regression in (4.1) whereas the RR shock is based on the same regression when $b_j = 0$ for $j > 1$, as in Romer and Romer (2004). Columns 1 and 2 display the response to the new shock and conventional shock, respectively. Column 3 display the response to the new shock minus the response to the conventional shock. The shaded areas indicate 68% and 95% confidence bands using standard errors robust to serial correlation and heteroskedasticity.

Figure C.11: Responses of IP and CPI

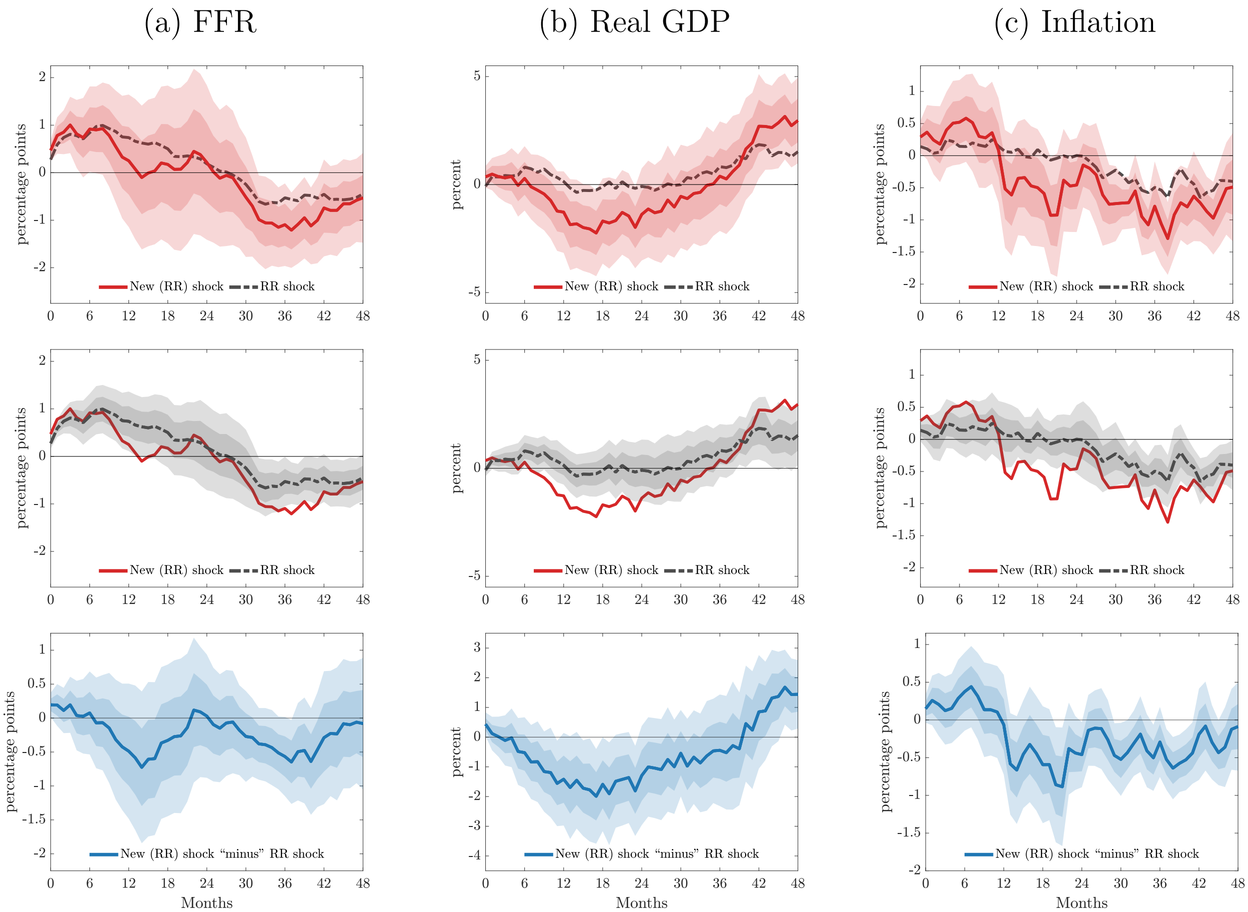

Notes: The figure shows responses of the federal funds rate, log industrial product, and the CPI inflation rate to a monetary policy shock based on the local projection as specified along with (4.2). We control for 12 lags of both, the log of industrial production and CPI inflation instead of real GDP and inflation based on the GDP deflator. The new (RR) monetary policy shock is identified as the residual from the Taylor rule regression in (4.1) whereas the RR shock is based on the same regression when $b_j = 0$ for $j > 1$, as in Romer and Romer (2004). Columns 1 and 2 display the response to the new shock and conventional shock, respectively. Column 3 display the response to the new shock minus the response to the conventional shock. The shaded areas indicate 68% and 95% confidence bands using standard errors robust to serial correlation and heteroskedasticity.

Figure C.12: Responses of IP and CPI when controlling for S&P 500 and EBP

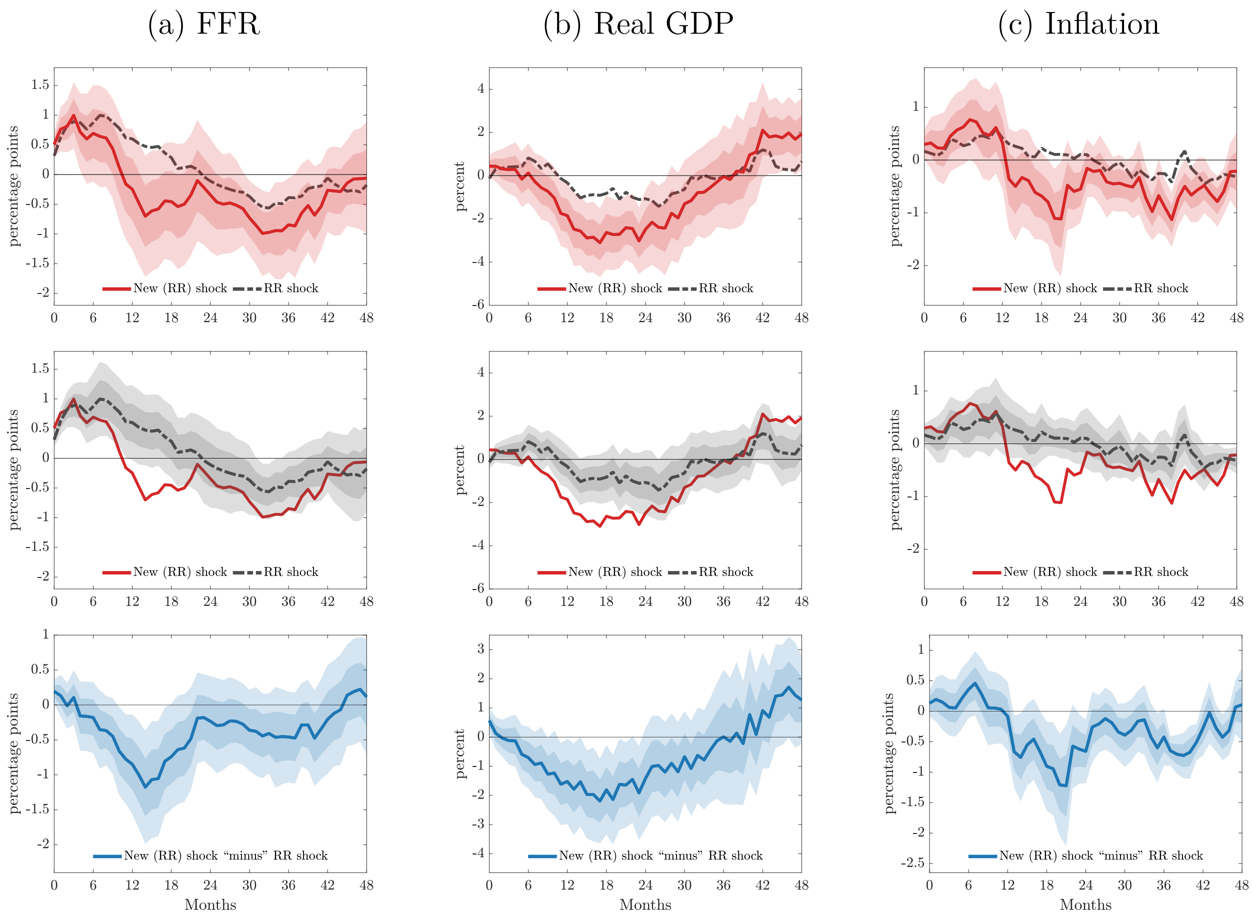

Notes: The figure shows responses of the federal funds rate, log industrial product, and the CPI inflation rate to a monetary policy shock based on the local projection as specified along with (4.2). We control for 12 lags of both, the log of industrial production and CPI inflation instead of real GDP and inflation based on the GDP deflator. Additionally, we control for 12 lags of both, the S&P 500 and the excess bond premium from Gilchrist and Zakrajšek (2012). The new (RR) monetary policy shock is identified as the residual from the Taylor rule regression in (4.1) whereas the RR shock is based on the same regression when $b_j = 0$ for $j > 1$, as in Romer and Romer (2004). Columns 1 and 2 display the response to the new shock and conventional shock, respectively. Column 3 display the response to the new shock minus the response to the conventional shock. The shaded areas indicate 68% and 95% confidence bands using standard errors robust to serial correlation and heteroskedasticity.