Optimal Monetary Policy with Sufficient Statistics*

Regis Barnichon[a] *and* *Geert Mesters*[b]

[a] Federal Reserve Bank of San Francisco and CEPR

[b] Universitat Pompeu Fabra, Barcelona GSE and VU Amsterdam

September 16, 2019

Abstract

If monetary policy instruments are set optimally, small perturbations to the instruments should have no first-order effect on welfare. Drawing on this insight, we show that the impulse responses to monetary shocks are sufficient statistics to evaluate the optimality of monetary policy, and we derive a monetary rule that allows to quantify the "distance to optimality" of any given policy. Drawing on historical records of monetary reports to Congress going back four decades, we assess the optimality of the Fed's policy since 1980 and find that policy has been remarkably close to optimal, except during 2008-2009 when it was constrained by the zero lower-bound.

*We thank Martin Eichenbaum, Ethan Ilzetzki, Sylvain Leduc, Huiyu Li, Nicolas Petrosky-Nadeau, Adam Shapiro and seminar participants at the Bank of England, McGill, the University of Ottawa and the San Francisco Fed for helpful comments. The views expressed in this paper are the sole responsibility of the authors and do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System. Mesters acknowledge support from the Spanish Ministry of Science and Technology: Grant ECO2015-68136-P, Marie Sklodowska-Curie Actions: FP7-PEOPLE-2012-COFUND Action grant agreement no: 600387, the Spanish Ministry of Economy and Competitiveness through the Severo Ochoa Programme for Centres of Excellence in R&D (SEV-2011-0075) and the Analysis of Big Data in Economics and Empirical Applications - 2016 BBVA Foundation Grants for Scientific Research.

1 Introduction

Modern central banks face two medium term objectives: price stability and full employment.[1] Achieving these two mandates is a difficult balancing act however, because the two objectives are often incompatible due to the existence of an inflation-unemployment trade-off, as captured by the Phillips curve. In fact, the central bank relies on this trade-off to "transform" unemployment into inflation (and vice-versa) through its interest rate policy, so that satisfying the two objectives often involves some trade-off between inflation and real activity.

While simple macro models have clear implications for optimal policy (e.g., Galí, 2015), i.e., for the ways to achieve the best trade-off between the two mandates, the complex, dynamic and uncertain environment in which central banks operate makes optimal policy much more difficult to achieve in practice. Despite impressive recent progress, structural models face clear limitations and their policy prescriptions are often model specific and not robust to model mis-specification (e.g., Svensson, 2003).

In this paper, we propose a sufficient-statistics approach to assess the "distance to optimality" of a given policy. The method is valid in any linear-quadratic framework with arbitrary dynamics and only requires estimates of the impulse responses of inflation and unemployment to innovations to the monetary instruments.[2] Importantly, because we do not rely on any specific structural model, the approach is largely immune to model mis-specification issues.[3]

Intuitively, when the monetary instruments are set optimally, small perturbations to the instruments should have no first-order effect on welfare. If this were not the case, the policy instruments could be adjusted until that first-order condition is satisfied, and this "optimal discretionary adjustment" is a measure of the distance to optimality of the original policy.

[1]The full employment mandate is explicit for the Fed, but often implicit for other central banks with inflation targets.

[2]The policy instrument can be a conventional tool like the fed funds rate or a more unconventional tool, such as forward guidance or quantitative easing.

[3]As we will see, our approach does require some modeling decisions to estimate the relevant sufficient statistics, but we view these as mild compared to specifying the full set of structural equations necessary to derive optimal policy prescriptions.

Importantly, in a linear-quadratic framework, this optimal adjustment can be calculated without specifying a particular underlying model. Instead, it only requires the central bank's forecasts for inflation and unemployment (absent any discretionary adjustment) and a set of sufficient elasticities that capture the transmission of monetary policy.[4]

The optimal discretionary adjustment formula has a simple interpretation as the least squares projection coefficient of current and future inflation and unemployment on the impulse responses of inflation and unemployment to monetary shocks, where the expectations are taken with respect to the central bank's information set. Intuitively, the point of optimal monetary policy is to set the central bank's instruments in order to best undo expected future movements in inflation and unemployment. Since the impulse responses of inflation and unemployment to shocks to the policy instruments capture monetary policies' transmission to unemployment and inflation, optimal policy amounts to best fitting expected future movements in inflation and unemployment with these impulse responses. In fact, the optimal discretionary adjustment is simply the solution of a static least squares problem, and the higher the population R^2 of that problem, the better the central bank is expected to be able to stabilize its target variables.

The regression interpretation clarifies a central, but generally little emphasized, component of monetary policy making: the central bank faces a *time-varying* trade-off between inflation and unemployment. This result comes from the fact that, in a dynamic setting, the central bank not only faces a trade-off across variables (stabilizing inflation versus unemployment) but also across horizons (stabilizing targets at shorter horizons versus longer horizons). Indeed, a key constraint on the central bank's technology is the coarse nature of its instruments: (i) monetary policy propagates slowly, and (ii) the transmission dynamics can be different for inflation and unemployment. Depending on how well the transmission of monetary policy lines up with the expected developments taking place in the economy, the central bank will have different abilities to affect the future path of its target variables.

[4] As our approach does not rely on any particular underlying model but only on a set of estimable sufficient elasticities, it follows the tradition of the sufficient statistic approach (e.g. Chetty, 2009).

Thus, depending on the nature of the shocks affecting the economy at any point in time, optimal policy will put different weights on stabilizing inflation versus unemployment.

Our approach extends naturally to handle uncertain environments, notably environments where the transmission of monetary policy is uncertain. The optimal discretionary adjustment generalizes Brainard (1967)'s seminal analysis to any linear setting with arbitrary dynamics, and we show how the Brainard conservatism principle can be reinterpreted as an attenuation bias due to the presence measurement error in impulse response estimates.

We use our optimal discretionary adjustment formula to assess the optimality of the Fed policy over the past four decades. This requires forecasts for inflation and unemployment and estimates of the impulse responses to monetary policy shocks. To obtain forecast data, we draw from historical records of monetary reports to Congress from 1980 until 2018 that contain the median forecasts of FOMC participants for inflation and unemployment. Importantly, these projections are conditional on the Fed following an optimal policy, as judged by the FOMC members, which allows us to measure the distance to optimality of the Fed's actions back to 1980. To estimate the sufficient statistics capturing the effects of the Fed's instruments, we group the Fed's monetary policy instruments into two groups: the first one captures conventional monetary policy and operates through the fed funds rate; and the second one captures a broad class of unconventional monetary policies that operate through the slope of the yield curve, as in Eberly, Stock and Wright (2019). We estimate the impulse responses of interest with local projection instrumental variable methods (Jordà, 2005; Stock and Watson, 2018), using external instruments derived from changes in asset prices around the FOMC meetings (Kuttner, 2001; Gürkaynak, Sack and Swanson, 2005).

We find that the Fed monetary policy has remarkably close to optimal, except during 2008-2009 when it was constrained by the zero lower-bound. For the fed funds rate policy, the deviations from optimality are relatively minor. For slope policies, which only operate after 2007, we find that a large reduction of 1.5 percentage points would have brought the slope to optimality during the financial crisis. This suggests that unconventional monetary policy measures —LSAP or QE— could have been used more aggressively to bring the slope down in line with optimality.

This paper bridges two large macro literatures: the literature on optimal monetary policy and the literature on the identification of macroeconomic shocks and their propagation. Our optimality criterion can be seen as a general "specific targeting rule" in the sense of Svensson (2003). Specific targeting rules are conditions that the forecast paths of the target variables must satisfy in order to minimize a particular loss function. As argued by Svensson in a series of influential papers (e.g., Svensson, 2003), specific targeting rules have great appeal –simplicity, transparency and immunity to time-consistency problems (Svensson and Woodford, 2005)—, but they are model specific and depend on the specification of structural equations describing the economy; the Phillips curve, the IS curve, the Okun's type law, and so on. The sufficient-statistic monetary rule that we propose is a generalized "specific targeting rule" that shares the benefits of specific targeting rules but remains agnostic about the underlying model. This allows us to derive a general measure of the distance to optimality that can be applied to any set of policies for which impulse response estimates are available from the macro-econometric literature (see Ramey, 2016, for an overview). While the structural impulse response literature has traditionally focused on estimating impulse responses to monetary shocks in order to guide model building (e.g., Christiano, Eichenbaum and Evans, 2005), our paper provides a novel and important role for impulse response estimates: as a testbed for the optimality of monetary policy.

The remainder of this paper is organized as follows. In the next sections we derive the sufficient-statistics monetary rule and the formula for the optimal discretionary adjustment first without parameter uncertainty and then taking into account parameter uncertainty. Section 5 provides the historical study for the optimality of US monetary policy. Section 6 concludes.

2 Forecast targeting with sufficient statistics

In this section, we derive a necessary condition for optimal monetary policy that only involves estimable sufficient statistics. As we will see, the optimal policy rule that we derive belongs to the class of "specific targeting rules" as advocated by Svensson in a series of influential papers (e.g., Svensson, 2003) on the "forecast targeting" approach to optimal monetary policy.[5] Different from earlier specific targeting rules however, our approach does not rely on a particular underlying model but only on a set of estimable sufficient elasticities, in the tradition of the sufficient statistic approach literature (e.g. Chetty, 2009).

2.1 Environment

Our starting point is a policy maker acting under discretion at time t to minimize a loss function of the form

$$\mathcal{L}_t = \sum_{j=0}^{J} \mathbb{E}_t L_{t+j} \ , \qquad\qquad L_{t+j} = (\pi_{t+j} - \pi^*)^2 + \lambda \left(u_{t+j} - u^*_{t+j}\right)^2 \ , \tag{1}$$

where λ is the preference parameter between the two mandates —the deviation of inflation π_t from its target π^* and the deviation of unemployment u_t from its unobserved natural rate u^*_t— and $\mathbb{E}_t$ is the expectation with respect to the central bank's time t information set $\mathcal{F}_t$.[6] The horizon J is left arbitrary at this point and can be considered infinite and with $\beta = 1$.

Using more compact notation, we can rewrite the loss function in terms of the ℓ_2-norms

[5] Forecast targeting refers to using forecasts of the target variables effectively as intermediate target variables, and in that context a specific targeting rule is a condition for the target variables, typically conditions that the forecast paths must satisfy in order to minimize a particular loss function.

[6] See e.g., Woodford (2011) for a micro-foundation of such an objective function. The results of this paper hold for any loss function of the form: $\mathcal{L}_t = \sum_{j=0}^{J} \mathbb{E}_t L_{t+j}$ with $L_{t+j} = \sum_{i=1}^{\mathcal{M}} \lambda_i \beta^j y^2_{i,t+j}$ where $\{y_{i,t+j}\}_{i=1}^{\mathcal{M}}$ is a set of mandates, $\{\lambda_i\}_{i=1}^{\mathcal{M}}$ is a set of preference parameters, β a discount factor and $\mathcal{M}$ is the total number of mandates. For concreteness we focus on the specific and common case where there are two mandates: inflation and unemployment.

of inflation and the unemployment gap[7]

$$\mathcal{L}_t = \mathbb{E}_t \left\|\tilde{\Pi}_t\right\|^2 + \lambda \mathbb{E}_t \left\|\tilde{U}_t\right\|^2 \ , \tag{2}$$

where $\tilde{\Pi}_t = (\tilde{\pi}_t, \ldots, \tilde{\pi}_{t+J})'$, with inflation deviations $\tilde{\pi}_t = \pi_t - \pi^*$, and $\tilde{U}_t = (\tilde{u}_t, \ldots, \tilde{u}_{t+J})'$, with $\tilde{u}_t = u_t - u_t^*$ the unemployment gap. Since the horizon of the loss function J is arbitrary and not of central importance, we simplify the notation by omitting J from the vectors $\tilde{\Pi}_t$ and $\tilde{U}_t$.

Suppose that the central bank has M policy instruments, denoted by $p_t = (p_{1,t}, \ldots, p_{M,t})'$, at its disposal at time t, with $M < J$. For instance, $p_{i,t}$ may correspond to the short term nominal interest rate i_t, or alternatively $p_{i,t}$ may correspond to a forward guidance announcement. In general, the objective of the central bank is to minimize the loss function (1) using its vector of instruments p_t.

We impose the following linear assumption that clarifies the mapping between the policy instruments to inflation and unemployment.

Assumption 1 (Linearity). *Denote by $\delta_t = (\delta_{1,t}, \ldots, \delta_{M,t})'$ a vector of innovations to the M policy instruments p_t. The effects of δ_t on the inflation and unemployment gaps are given by*

$$\tilde{\Pi}_t = \mathcal{R}^\pi \delta_t + X_t^\pi \qquad \text{and} \qquad \tilde{U}_t = \mathcal{R}^u \delta_t + X_t^u \ , \tag{3}$$

where $\tilde{\Pi}_t = (\tilde{\pi}_t, \ldots, \tilde{\pi}_{t+J})'$, with $\tilde{\pi}_t = \pi_t - \pi^*$, and $\tilde{U}_t = (\tilde{u}_t, \ldots, \tilde{u}_{t+J})'$, with $\tilde{u}_t = u_t - u_t^*$, $X_t^y = (x_t^y, \ldots, x_{t+J}^y)'$ is a random vector, for $y = \pi, x$, and $\mathcal{R}^\pi = \frac{\partial \tilde{\Pi}_t}{\partial \delta_t'}$ and $\mathcal{R}^u = \frac{\partial \tilde{U}_t}{\partial \delta_t'}$ are $J \times M$ coefficient matrices capturing the impulse response functions of inflation and unemployment to the $M \times 1$ vector of innovations δ_t.

Assumption 1 in conjunction with loss function (1) restricts our results to the linear-quadratic framework, but it places no restrictions on the dynamics of the process for π_t and u_t, and it does not restrict the set of variables relevant for inflation and unemployment.

[7]We consider the usual vector norm where for any $a \in \mathbb{R}^n$ we have that $\|a\| = \sqrt{\sum_{t=1}^n a_t^2}$.

The key thing to notice is that Assumption 1 imposes a *direct* mapping between the policy perturbations and the target variables as opposed to specifying a recursive model for the targets.[8] Assumption 1 is similar to the original formulation of the optimal policy problem in Tinbergen (1952) and Theil (1957).

Assumption 1 is mild and this will provide three key benefits; (i) *generality*: the resulting optimal policy recommendations will apply for a broad class of models, (ii) *transparency*: the findings will be easy to communicate to a broad audience as all implications can be traced back to a simple linear mapping and (iii) *robustness*: when model (3) is treated as an approximation to a more general underlying model, it typically remains to produce good forecasting results (e.g. Chevillon, 2007).

To give a concrete example that satisfies Assumption 1, we can think of the central bank as setting its instruments p_t following the prescription of a Taylor rule. X_t^π and X_t^u capture the future path of inflation and unemployment conditional on that Taylor rule. and Assumption 1 posits that small perturbations δ_t to p_t have a linear effect on the future paths of inflation and unemployment.[9] δ_t can influence $\tilde{\pi}_{t+j}$ and $\tilde{u}_{t+j}$ in any arbitrary way for all j (as captured by $\mathcal{R}^\pi$ and $\mathcal{R}^u$) as long as the effect remains linear.

The linearity assumption imposes that the policy innovations δ_t do not lead agents to revise expectations about their underlying representation of the central bank's reaction function (which affects their decision making process). If the policy innovations are sufficiently small —modest interventions in the sense of Leeper and Zha (2003)—, we can reasonably consider that the economy remains in the same linear regime before and after a monetary innovation, and assumption 1 is satisfied.[10]

[8] Recursive models such as VAR models or linear state space models are more common in the optimal policy literature see for examples Chow (1972, 1973); Sack (2000); Rudebusch (2002); Orphanides (2003); Swanson (2004) or the textbook treatment of Ljungqvist and Sargent (2004). Note that any linear recursive model can be re-written to satisfy (3) simply by adjusting the definition for X_t.

[9] In this example, p_t can be seen as the systematic component of monetary policy, while δ_t is a discretionary adjustment to the systematic conduct of monetary policy.

[10] According to Leeper and Zha (2003), a *modest* policy intervention does not significantly shift agents' beliefs about policy regime and does not induce the changes in behavior that Lucas (1976) emphasizes. Leeper and Zha (2003) show that a rich class of interventions routinely considered by the Fed is modest so that their impact can be reliably forecasted by an identified linear model, i.e., satisfy Assumption 1.

2.2 A Specific Targeting Rule with Sufficient Statistics: STReSS

The following lemma states the central bank's first-order condition for optimal policy.

Lemma 1 (Specific Targeting Rule with Sufficient Statistics: STReSS)**.** *Given assumption 1 and under the assumption that $\mathcal{R}^y$, for $y = \pi, u$, is known, the necessary conditions for the minimization problem* $\min_{\delta_t \in \mathbb{R}^M} \mathcal{L}_t$ *are given by*

$$\mathcal{R}^{\pi'} \mathbb{E}_t \tilde{\Pi}_t = -\lambda \mathcal{R}^{u'} \mathbb{E}_t \tilde{U}_t \ . \tag{4}$$

Lemma 1 states that, when policy is set optimally, the *projection* of the inflation forecast on the impulse response of inflation should equal the *projection* of the unemployment forecast on the impulse response of unemployment, scaled by λ, which captures society's preference between stabilizing inflation and unemployment (λ).

We will refer to the set of necessary conditions (4) as a Specific Targeting Rule with Sufficient Statistics, or STReSS for short. It is a specific targeting rule (STR) in the sense of Svensson (2003) —a set of conditions that the forecasts of the target variables must satisfy in order to minimize the loss function—. At the same time however, STReSS can be seen as a generalization of STR rules, as it only involves a set of estimable sufficient statistics that capture the transmission of monetary policy —the impulse responses of inflation and unemployment to monetary shocks—.

2.3 Optimality discretionary adjustment and distance to optimality

An important implication of Lemma 1 is that it is possible to solve for the vector of policy innovations δ_t that would ensure that policy satisfies the necessary optimality condition (4). This is summarized in the following proposition.

Proposition 1 (Optimal discretionary adjustment)**.** *Given Assumption 1 and under the assumption that $\mathcal{R}$ is known with $\mathcal{R}'\mathcal{R} \succ 0$, the discretionary adjustment $\hat{\delta}_t$ needed to make* *a given policy optimal is given by*

$$\hat{\delta}_t = -\left(\mathcal{R}'\mathcal{R}\right)^{-1}\mathcal{R}'\mathbb{E}_t X_t \ , \tag{5}$$

where

$$X_t = \begin{bmatrix} X_t^\pi \\ \sqrt{\lambda} X_t^u \end{bmatrix} \ , \qquad \mathcal{R} = \begin{bmatrix} \mathcal{R}^\pi \\ \sqrt{\lambda}\mathcal{R}^u \end{bmatrix} .$$

The estimate $\hat{\delta}_t$ can be thought of as the discretionary adjustment needed to bring policy in line with the requirements of STReSS, i.e., to make sure that policy does satisfy the necessary condition of optimality. In other words, we can see $\hat{\delta}_t$ as a measure of the distance to optimality for the vector of policy instruments p_t at time t. This expression for optimal policy adjustment can be traced back to Theil (1957) and an insightful textbook treatment is given by Preston and Pagan (1982).

The key insight here is that the measure $\hat{\delta}_t$ of the distance to optimality does not require the estimation of a fully specified structural model. The only requirement is to have estimates of the sufficient statistics, i.e., the impulse responses to monetary shocks and the policy makers' expected paths for inflation and unemployment conditional on (what they judge to be) an appropriate policy. As we will see, the sufficient statistics can be directly estimated from the data, i.e., from empirical estimates of treatment effects following random variations in monetary policy (Ramey, 2016). This avoids specifying structural equations that relate policy changes to welfare, and contrasts with the traditional optimal control approach that involves specifying a set of structural equations; in our case a Phillips curve, linking inflation to real activity, an (IS) curve, linking real activity to the short-term policy rate, and an Okun-type law linking the output gap to the unemployment gaps. All these equations are prone to mis-specification (e.g., Blanchard, 2016; Barnichon and Mesters, 2019).

3 Interpretation and implications

The optimal discretionary adjustment formula has a simple interpretation as the least squares projection coefficient of the expected future paths of inflation and unemployment on their corresponding impulse responses to monetary shocks. This projection interpretation clarifies a central, but generally little emphasized, component of monetary policy making: because of the blunt nature of its instruments, monetary policy will have to strike different trade-offs between inflation and unemployment, depending on the nature of the shocks affecting the economy at any point in time. This result comes from the fact that, in a dynamic setting, the central bank not only faces a "static" trade-off across variables (stabilizing inflation versus unemployment) but also a "dynamic" trade-off across horizons (stabilizing targets at shorter horizons versus longer horizons).

3.1 Optimal adjustment as a least-squares projection

To get some intuition behind the optimal discretionary adjustment given by (5), it is helpful to think of $\hat{\delta}_t$ as a vector of population regression coefficients. Consider the following linear regression model

$$X_t = -\mathcal{R}\delta_t + \tilde{Y}_t \ , \tag{6}$$

where

$$X_t = \begin{bmatrix} X_t^{\pi} \\ \sqrt{\lambda} X_t^u \end{bmatrix} , \qquad \mathcal{R} = \begin{bmatrix} \mathcal{R}^{\pi} \\ \sqrt{\lambda}\mathcal{R}^u \end{bmatrix} , \qquad \tilde{Y}_t = \begin{bmatrix} \tilde{\Pi}_t \\ \sqrt{\lambda}\tilde{U}_t \end{bmatrix} .$$

Equation (6) is obviously equivalent to Assumption 1 but it also makes clear that we can see $\hat{\delta}_t$ as the vector of least squares projection coefficients of X_t on $\mathcal{R}$ (defined conditional on $\mathcal{F}_t$).[11]

The purpose of the optimal discretionary adjustment is to make the residual $\tilde{Y}_t$ as small

[11]Recall that for a standard linear regression, say $y = x'\beta + u$, we have that the least squares projection coefficients are given by $\beta^* = \mathbb{E}(xx')^{-1}\mathbb{E}(xy)$, see for instance Hayashi (2000) section 2.9. In Proposition 1 we assume that $\mathcal{R}$ is known and fixed, and hence $\mathcal{R}$ can be taken out of the expectation. In Section 4 we relax this assumption, but we note that the intuition that is discussed in this section continues to apply.

as possible (in an ℓ_2-norm sense), i.e., to generate the highest R^2 of a regression of X_t on $\mathcal{R}$. A high R^2 means that the central bank's instruments can fit well the future developments in the economy, i.e., best undo them —note the minus sign in equation (5)— and stabilize its targets with the appropriate set of policies.[12] The larger the number M of instruments (e.g., fed funds rate and slope policy), the better the central bank can span expected future developments (i.e., get a higher population R^2) and therefore fulfill its mandates, provided the effects of the different instruments are not co-linear.[13]

Importantly, note how $\tilde{Y}_t$ stacks all the variables entering the loss function — $\{\tilde{\pi}_{t+j}\}_{j=0}^{J}$ and $\{\tilde{u}_{t+j}\}_{j=0}^{J}$ — in one single vector, meaning that all these variables enter the policy problem in the same fashion. In other words, we can see the central bank as facing not just two dynamic targets (inflation and unemployment) but instead $2(J+1)$ *static* targets ($\{\tilde{\pi}_{t+j}\}_{j=0}^{J}$ and $\{\tilde{u}_{t+j}\}_{j=0}^{J}$). In this static representation of the optimal policy problem, the central bank has $M \ll J$ instruments to best hit $2(J+1)$ (equal-weight) targets.

The benefit of this static representation is that optimal control in a static setting is well understood and very transparent (e.g., Preston and Pagan, 1982). In the context of monetary policy where communication and transparency are paramount, this is an important advantage. In particular, the simplicity of the projection interpretation will allow to clearly spell out the technological constraints faced by the central bank and the trade-offs at play at any point in time.

3.2 Trade-offs between mandates and horizons

The key constraint on the central bank's optimal control problem is that the instruments of monetary policy are *blunt*: monetary policy (i) moves inflation and unemployment in opposite directions, and (ii) transmits slowly. In turn, these "technological constraints"

[12]Naturally, since X_t is not observed as it contains future values, the regression can only be considered in expectation conditional on the central bank's information set $\mathcal{F}_t$.

[13]If the central bank has sufficiently many instruments that satisfy the conditions of Proposition 1, it is possible to get an R^2 of one, i.e. to fully stabilize all target variables. If the central bank has less instruments than targets it can still get an R^2 of one *if* the target variables are correlated. An example of this happens in the baseline New Keynesian model under the divine coincidence (e.g. Blanchard and Galí, 2007): where there is one horizon, one instrument and two targets which are perfectly correlated.

imply that the central bank faces (i) a "static" trade-off between mandates, as well as (ii) a "dynamic" trade-off between horizons.

To see that, we can re-write the expression for the optimal discretionary adjustment as

$$\hat{\delta}_t = W\hat{\delta}_t^{\pi} + (I - W)\hat{\delta}_t^{u} \ , \tag{7}$$

with $W = (\mathcal{R}'\mathcal{R})^{-1}\mathcal{R}^{\pi\prime}\mathcal{R}^{\pi}$ a weighting matrix and $\hat{\delta}_t^{y}$ the optimal discretionary adjustment to p_t in the case of a single mandate:

$$\hat{\delta}_t^{y} = -(\mathcal{R}^{y\prime}\mathcal{R}^{y})^{-1}\mathcal{R}^{y\prime}\mathbb{E}_t X_t^{y} \ , \qquad \text{for} \quad y = \pi, u \ . \tag{8}$$

The optimal discretionary adjustment $\hat{\delta}_t$ is a *weighted average* of $\hat{\delta}_t^{\pi}$ and $\hat{\delta}_t^{u}$, the optimal adjustments needed to best stabilize each target, with the weights W capturing the average inflation-unemployment trade-off.[14]

Expression (7) highlights how the blunt nature of the monetary instruments imply the two types of trade-offs.

First, monetary policy moves both mandates at the same time and in opposite directions, so that the optimal individual adjustments $\hat{\delta}_t^{\pi}$ and $\hat{\delta}_t^{u}$ can have opposite signs, implying that the central bank must choose between stabilizing one mandate or the other. This is the well-known "static" trade-off between mandates.

Second, monetary policy propagates slowly with (possibly) different transmission dynamics for inflation and unemployment. Indeed, unless an instrument can be precisely targeted at one (and only one) of the $(J + 1)$ horizons, each instrument will affect mandates over multiple horizons and this will imply some trade-offs between horizons.[15]

[14] W is a function of society's preference between the two mandates (λ) and the relative ability of the central bank to move its targets (determined by $\mathcal{R}$). Note that $W \underset{\lambda \to 0}{\longrightarrow} 1$ (a single inflation mandate), and similarly $W \underset{\lambda \to \infty}{\longrightarrow} 0$ (a single unemployment mandate). When the central bank only has one policy instrument (e.g., the fed funds rate), the weight simplifies to $W = \frac{1}{1+\lambda/\kappa}$ and depends on the ratio of society's preference between the two mandates (λ) and the central bank's instrument "average" ability to transform unemployment into inflation ($\kappa = \frac{\|\mathcal{R}^{\pi}\|}{\|\mathcal{R}^{u}\|}$).

[15] This does not happen if the transmission dynamics of an instrument (the impulse response) is degenerate at some horizon, in the sense that it is zero everywhere except at some horizon j. This is the case in the

3.3 Illustration

The interactions between the static and dynamic trade-offs are central to optimal policy making. We now provide a number of examples to illustrate how our sufficient-statistics approach (i) can quantify the distance to optimality of any given policy, and (ii) clearly spells out the trade-offs at play at any point in time, which is helpful both for decision making within the central bank and for communication with the outside public.

Our approach relies on estimates of the effect of the policy instruments on the mandates. In this section, we assume that the central bank has one instrument (say the fed funds rate) and we posit some (empirically plausible) transmission mechanism of that instrument, as depicted in Figure 1 which shows the impulse response of inflation to an innovation to the monetary instrument. Note how monetary policy transmits slowly and how the response of inflation lags that of unemployment, in line with empirical evidence (e.g., Christiano, Eichenbaum and Evans, 1999). We consider three different scenarios faced by the central bank.

Target smoothing across horizons

First, we illustrate the constraints imposed by transmission dynamics of monetary policy. Figure 2 considers a strict inflation targeting central bank ($\lambda = 0$) facing a deflationary shock, i.e., an aggregate shock that lowers inflation.

The top panel reports the expected paths for inflation before any Optimal Discretionary Adjustment (ODA, filled circles) and after optimal adjustment (empty circles). The middle panel reports the impulse response of inflation to an innovation to the monetary instrument, as in Figure 1.

Because the transmission of monetary policy to inflation ($\mathcal{R}^{\pi}$) does not line up perfectly with the expected future path of inflation ($\mathbb{E}_t X_t^{\pi}$), the central bank cannot perfectly offset the shock. However, it can lower the loss function by smoothing the inflation deviations across

baseline New-Keynesian model (e.g., Galí, 2015) where the interest rate instrument only affects macro variables only on impact ($j = 0$) or in the double-lag model of Svensson (1997) where the interest rate instrument affects the output gap only after one period ($j = 1$) and inflation only after two periods ($j = 2$).

horizons. This can be seen in the bottom panel that shows scatter plots of the expected future paths of inflation $\mathbb{E}_t X_t^\pi$ against the impulse responses of inflation $\mathcal{R}^\pi$ both before (left panel) and after ODA (right panel). The red dashed-line depicts the best linear fit for inflation. The fit is not perfect but the more expansionary policy ($\hat{\delta}_t < 0$) allows to further reduce the loss function by trading smaller inflation deviations at shorter horizons term with larger deviations at longer horizons.

Target smoothing across mandates and horizons

Second, consider a central bank with a dual inflation-unemployment mandate. The interaction between static and dynamic trade-offs has interesting implications for the conduct of optimal policy. Depending on how well the propagation of monetary policy lines up with the expected future paths of inflation and unemployment, the central bank may have different abilities to affect the future path of one mandate versus another, and it can be optimal to stabilize *one mandate now* at the cost of destabilizing the *other mandate later*.

To illustrate this point, consider an economy dominated at time t by an adverse aggregate demand (AD) shock in the sense that the expected future path of inflation is negative while the expected path of unemployment is positive, as illustrated with dashed-lines in the top panel of Figure 3. The expected path for inflation is the same as the previous example, but now the central bank has two mandates and must take into account the expected path of unemployment. The structure of the figure 3 is the exact same as with figure 2 except that it also adds information related to unemployment (in blue).

In the case depicted in Figure 3, the optimal adjustment (dashed-line) has the central bank offset the AD shock by even more than in a strict inflation targeting case. This will allow to mitigate the increase in unemployment, although at the cost of more inflation at longer horizons. In other words, the central bank trades off lower unemployment now with higher inflation later. This smoothing across horizons *and* mandates is welfare improving, because it substitutes the large (and costly) unemployment deviations with small (and hence less costly) inflation deviations.

To better understand the trade-off at play, we can go back to our regression interpretation of the optimal adjustment $\hat{\delta}_t$, as depicted in the bottom panel. With an (AD) shock, there is no standard 'static" trade-off between inflation and unemployment, in that monetary policy works in the same direction as the (AD) shock: $\hat{\delta}_t^u$ and $\hat{\delta}_t^\pi$ have the same sign (the dashed blue and red lines both slope upwards) and the optimal adjustment calls for a more accommodating policy. A "dynamic" trade-off between inflation and unemployment exists nonetheless, because monetary policy is (in this case) more powerful at influencing the expected future path of unemployment than that of inflation. This can be seen by the fact that the regression line for unemployment (blue dashed-line) provides a better fit to the data than the regression line for inflation (red dashed-line). For that reason, the optimal adjustment (black dashed line) calls for more stabilization of unemployment at the expense of inflation. Note however how the regression R^2 is only 0.72: because of the blunt nature of its instrument, the central bank cannot offset a sizable fraction of the effect of the (AD) shock.

As a third example, suppose that the deflationary path of inflation is instead caused by some Aggregate Supply (AS) shock that lowers both inflation and unemployment (prior to any discretionary adjustment to policy). As the (AS) shock moves inflation and unemployment in the same direction, but monetary policy does not, the policy maker must strike a trade-off between stabilizing one mandate versus another. This case is depicted in Figure 4.

The standard trade-off is apparent in the fact that the regression lines for inflation and unemployment have opposite slopes ($\delta_t^\pi > 0$ but $\delta_t^\pi < 0$): it is desirable to run a more contractionary policy to lower inflation *and* at the same time run a more expansionary policy to lower unemployment, which is not possible. In fact, in this example, the trade-off is so strong that the optimal adjustment is close to zero as δ_t^π and δ_t^π almost exactly cancel out each other. In other words, because of the blunt nature of its instrument, there is nothing more the central bank can do to offset the effect of the (AS) shock, and policy is optimal.

4 Optimal monetary policy with uncertainty

So far our discussion of optimal monetary policy has ignored parameter uncertainty and we have operated under the assumption that the effects of monetary policy were known with certainty. In this section we consider adjusting the optimal policy adjustment of Proposition 1 for parameter uncertainty. Our approach for incorporating uncertainty generalizes the seminal work of Brainard (1967) to a setting with arbitrary dynamics.[16]

The following assumption complements Assumption 1 by prescribing how the central bank constructs its forecasts for the case where the impulse responses are estimated. We follow Brainard (1967) by treating the impulse responses $\mathcal{R}^\pi$ and $\mathcal{R}^u$ as random variables and their estimates are obtained by the posterior means that are determined after updating prior beliefs based on the information set $\mathcal{F}_t$. The benefit of this approach is that the initial forecasts and the impulse response estimates are treated symmetrically. However, as noted by Brainard (1967), an equivalent result can be obtained when treating the impulse responses as deterministic coefficients.

Assumption 2 (Forecast construction)**.** *The central bank forecasts for inflation deviations and the unemployment gap are determined by*

$$\mathbb{E}_t\tilde{\Pi}_t = \hat{\mathcal{R}}_t^\pi \delta_t + \mathbb{E}_t X_t^\pi \qquad \textit{and} \qquad \mathbb{E}_t\tilde{U}_t = \hat{\mathcal{R}}_t^u \delta_t + \mathbb{E}_t X_t^u \ , \tag{9}$$

where $\hat{\mathcal{R}}_t^y = \mathbb{E}_t \mathcal{R}^y$ *and we further define* $\Omega_{j,t}^y = \mathbb{V}_t(\mathcal{R}_j^y)$ *and* $C_{j,t}^y = \mathbb{C}_t(\mathcal{R}_j^y, x_{t+j}^y)$*, where* $\mathcal{R}_j^y$ *is the* j*th row of* $\mathcal{R}^y$*, for* $j = 0, \dots, J$ *and* $y = \pi, u$*.*

It is useful to think about $\mathbb{E}_t X_t^\pi$ and $\mathbb{E}_t X_t^u$ as the initial forecasts constructed by the central bank for inflation deviations and the unemployment gap. These forecasts are modified by the policy perturbations δ_t, via the estimated impulse responses $\hat{\mathcal{R}}_t^\pi$ and $\hat{\mathcal{R}}_t^u$, to obtain the final forecasts $\mathbb{E}_t\tilde{\Pi}_t$ and $\mathbb{E}_t\tilde{U}_t$. Note that the impulse response estimates depend explicitly on

[16] An important advantage of Brainard (1967)'s formulation of the effects of uncertainty on optimal policy is its transparency. This stands in contrast to the treatment of uncertainty in medium-scale DSGE models where parametrization and estimation uncertainty interact in a complex fashion, making the effects of uncertainty on optimal behavior often model specific and somewhat opaque.

time t as they are determined by using the time t information set $\mathcal{F}_t$.[17] Assumption 2 further describes the uncertainty in the impulse responses via $\Omega^y_{j,t}$ and their potential covariance with the random variables x^y_{t+j} via $C^y_{j,t}$, across the different horizons $j = 0, \ldots, J$.

The following proposition generalizes the optimal discretionary adjustment of Corollary 1 to account for parameter uncertainty.

Proposition 2 (Optimal discretionary adjustment with uncertainty)**.** *Given assumptions 1-2, if we have* $\mathbb{E}_t(\mathcal{R}'\mathcal{R}) \succ 0$*, then*

$$\hat{\delta}_t = -\mathbb{E}_t(\mathcal{R}'\mathcal{R})^{-1}\mathbb{E}_t(\mathcal{R}'X_t) \tag{10}$$

where we again use the compact notation as in (6)*.*

The OLS interpretation of Section 2.3 that was based on equation (6) continues to apply in the case of uncertainty. Indeed the optimal adjustment remains equivalent to the least squares projection coefficients (conditional on $\mathcal{F}_t$) of X_t on $\mathcal{R}$. Compared to the case of Proposition 1 where $\mathcal{R}$ is known, additional terms arise because $\mathcal{R}$ is unknown and needs to be estimated. We get

$$\begin{aligned}\hat{\delta}_t &= -\mathbb{E}_t(\mathcal{R}'\mathcal{R})^{-1}\mathbb{E}_t(\mathcal{R}'X_t)\\ &= -\left(\bar{\Omega}_t + \hat{\mathcal{R}}_t'\hat{\mathcal{R}}_t\right)^{-1}\left(\bar{C}_t + \hat{\mathcal{R}}_t'\mathbb{E}_t(X_t)\right),\end{aligned}$$

where $\bar{\Omega}_t = \bar{\Omega}^\pi_t + \lambda\bar{\Omega}^u_t$ and $\bar{C}^y_t = \bar{C}^\pi_t + \lambda\bar{C}^u_t$, with $\bar{\Omega}^y_t = \sum_{j=0}^H \Omega^y_{j,t}$ and $\bar{C}^y_t = \sum_{j=0}^H C^y_{j,t}$, for $y = \pi, u$.

With uncertainty in the policy multipliers —the impulse responses of π and u to policy shocks— the central bank refrains from fulling minimizing the loss function, similarly to Brainard (1967). In fact, the optimal discretionary adjustment spelled out in equation (10) generalizes Brainard's analysis to any linear setting with arbitrary dynamics.[18]

Interestingly, the projection interpretation highlights how the Brainard conservatism

[17] In the previous sections we purposely refrained from indexing $\mathcal{R}$ by t to keep the notation simple. However, the general lessons of those sections continue to hold when $\mathcal{R}$ is considered time-varying.

[18] In the static model envisioned by Brainard (1967) where the impulse responses are degenerate and the

principle can be reinterpreted as an attenuation bias. Because of measurement error in impulse response estimates, the least-squares estimate $\hat{\delta}_t$ is estimated with a downward bias with the extent of the bias depending on the noise-to-signal ratio $\bar{\Omega}_t / \hat{\mathcal{R}}_t' \hat{\mathcal{R}}_t$. Reducing the measurement error in our estimates of the effects of monetary policy is thus of direct relevance for the conduct of monetary policy.

When the central bank has multiple instruments, say the fed funds rate and slope policy, optimal policy calls for relying less on instruments with more uncertain effects. In this context, since the macro effects of forward-guidance or unconventional policies are arguably much less well understood than the effects of changing the current policy rate, ceteris paribus optimal policy calls for relying less on forward-guidance and unconventional policies and more on adjusting the current policy rate whenever possible (i.e., away from the zero lower bound).

Finally, uncertainty about the state of the economy enters only indirectly through its interaction with the uncertainty in the transmission of monetary policy and via the information set $\mathcal{F}_t$ which determines the forecasts and impulse response estimates. The interaction occurs if uncertainty in the transmission of monetary policy is correlated with uncertainty in the transmission of other shocks. To give an example, imagine that $\bar{C}_t^{\pi} \succ 0$, so that the possibility that monetary policy has less effect on inflation is positively correlated with the possibility that inflation becomes less reactive to *all types* of shocks. In that case, the central bank should be even more conservative, because should monetary policy be less powerful, then inflation would also be less reactive in general (e.g., better anchored).

---

effect of policy only takes place on impact, our optimal adjustment in the case of only one mandate y becomes

$$\hat{\delta}_t = -\frac{\hat{\mathcal{R}}_{0,t}^{y} \mathbb{E}_t y_t + C_{0,t}^{y}}{\left(\hat{\mathcal{R}}_{0,t}^{y}\right)^2 + \Omega_{0,t}^{y}} ,$$

which is the adjustment derived by Brainard (1967) (see equation (5) page 414).

5 A retrospective analysis of the Fed's monetary policy

In this section, we show how STReSS can be used to assess the optimality of the Fed's monetary policy, and we study the optimality of the Fed policy over 1980-2018.

5.1 Using STReSS to assess the distance to optimality

We rely on the optimal discretionary adjustment to monetary policy $\hat{\delta}_t$ (with parameter uncertainty) as derived in Proposition 2 to asses the distance to optimality. We restate it for convenience

$$\hat{\delta}_t = -\left(\bar{\Omega}_t + \hat{\mathcal{R}}_t'\hat{\mathcal{R}}_t\right)^{-1}\left(\bar{C}_t + \hat{\mathcal{R}}_t'\mathbb{E}_t\left(X_t\right)\right) \ . \tag{11}$$

Using the initial forecasts of the Fed and the impulse responses of the mandates to monetary policy shocks as inputs we compute $\hat{\delta}_t$ for each time period t. This is the adjustment that the Fed should have made in order to satisfy the optimal policy condition of Proposition 2. For instance, suppose that $p_{i,t}$ corresponds to the fed funds rate, then $\hat{\delta}_{i,t}$ tells us, in percentage points, the change in the fed funds rate that would have made the Fed's interest policy optimal at time t.

In general, we can translate the distance to optimality of the policies to the distance to optimality in units of the loss function $\mathcal{L}_t$, i.e., in units of the average variance of $\tilde{\pi}$ and $\sqrt{\lambda}\tilde{u}$. Indeed, we may compute

$$\Delta\mathcal{L}_t(\hat{\delta}_t) = \mathcal{L}_t(\hat{\delta}_t) - \mathcal{L}_t(0) \ , \qquad \text{with} \quad \mathcal{L}_t(\delta_t) = \mathbb{E}_t(X_t + \mathcal{R}\delta_t)'(X_t + \mathcal{R}\delta_t) \tag{12}$$

where $\mathcal{L}_t(\hat{\delta}_t)$ is the value of the loss function under the optimal policy (i.e., after the actual policy has been corrected with the optimal discretionary adjustment $\hat{\delta}_t$ and where $\mathcal{L}_t(0) = \mathbb{E}_t(X_t'X_t)$ is the value of the loss function under the policy actually implemented at time t.

5.2 Sufficient statistics and economic projections

The optimal discretionary adjustment to monetary policy $\hat{\delta}_t$ makes clear that two pieces of information are needed to assess the optimality of the Fed policy over time: (i) the Fed's projections for inflation and unemployment conditional on policy makers' desired policy path and (ii) the impulse responses of inflation and unemployment to monetary shocks — the sufficient statistics—.

A new dataset of FOMC projections going back to 1980

First, to evaluate the distance to optimality of Fed policy over time, we need data on policy makers' projections for inflation and unemployment *conditional* on the appropriate policy path, as judged by FOMC members.[19] For that purpose, we will draw on historical records of monetary reports to Congress going back four decades

Since the passage of the Full Employment and Balanced Growth Act of 1978 (also known as Humphrey-Hawkins), federal law requires the Federal Reserve Board to submit written reports to Congress containing discussions of "the conduct of monetary policy and economic developments and prospects for the future". As part of this report, the Fed provides a summary of Federal Open Market Committee (FOMC) participants' projections for the unemployment rate and inflation (among other variables): the Survey of Economic Projections (SEP).

The data were manually extracted from digital records from the archives of the House Financial Services Committee. We obtained bi-annual SEP data for the median forecasts of FOMC members for inflation and unemployment at one- and two-year ahead horizons over the period 1980-2006.[20] After 2006, SEP data are published four times a year and additionally include the median forecasts at a three-year ahead horizon. In addition, we

[19] This in contrast with, say, the Board's Greenbook/Tealbook forecasts which are constructed conditional on the policy path following a mechanical Taylor-type policy rule, which need not represent the FOMC view at a given point in time.

[20] The price index underlying the inflation measure has changed over time, ranging from the GNP deflator, CPI to PCE in the more recent period. Using a linear model of the form $\pi_t^{pce} = \alpha + \beta \pi_t^x + \varepsilon_t$ with x denoting the underlying price index, we adjusted the different inflation measures to make them consistent with a PCE-based measure.

complement these forecasts with the median FOMC estimate of the "long-run" projections for inflation and unemployment. We set the horizon for the "long-run" FOMC projections to equal 5 years.[21] Since the SEP only reports the median FOMC estimate of the long-run values of inflation and unemployment after 2007, we rely on real time estimates before 2007. Specifically, we use real-time estimates of the natural rate of unemployment constructed by (Orphanides and Williams, 2012) and we use the long-run inflation level from the Federal Reserve Board "PTR" variable, which is a measure a long-run inflation expectations derived from the Survey of Professional Forecasters (SPF).

Since the SEP projections are annual, we linearly interpolate them in order to project them on the quarterly impulse responses to monetary shocks, as required to estimate $\hat{\delta}_t$ using equation (11).

Estimation of sufficient statistics

Next, we discuss the estimation of the sufficient statistics. In principal, many different policies can be considered within our framework provided that the impulse responses are estimable. In our application we follow the recent works of Lakdawala (2019) and, more closely Eberly, Stock and Wright (2019), and group monetary policy actions into two categories: traditional monetary policy, in which the FOMC sets the Fed funds rate, and slope policy, which explicitly aims to affect the slope of the Treasury yield curve.

Broadly speaking, slope policy is a summary measure of different actions that the Fed has undertaken in the more recent years since 2007. Examples of slope policies include forward guidance, SEPs, LSAPs, and maturity management. These policies are grouped together as they all affect the slope of the safe asset term structure by either providing information about future policies or by directly changing the values of the assets. In principal, it might be possible to study these measures separately, however for the purpose of our retrospective analysis combining them leads to no loss of insights. We measure the joint effect of the policies by the spread between the yield on 10-year Treasuries and the Fed funds rate.

[21] We obtain very similar results using instead a convergence time of 10 years.

Correspondingly, for $y = \pi, u$ we define the sufficient statistics $\mathcal{R}^y_{i,j}$ as follows.

$$\mathcal{R}^y_{i,j} = \mathbb{E}(y_{t+j}|\varepsilon^i_t = 1, w_t) - \mathbb{E}(y_{t+j}|\varepsilon^i_t = 0, w_t),$$

where ε^i_t corresponds to a structural monetary policy shock of type $i = r^0, s$, with r^0 indicating the fed funds rate shock and s the slope shock, and w_t is a vector of control variables. Notice that the definition of the impulse responses corresponds with Assumption 1 in the sense that the effects of policy adjustments $\delta_{i,t}$ are measurable by the changes due to monetary policy shocks.

To estimate the sufficient statistics we use local projections with external instruments (Jordà, 2005; Stock and Watson, 2018). We aim to identify two structural monetary policy shocks (a standard monetary policy shock and a slope shock) which requires at least two instruments. For the Fed funds shock, the instrument is the difference between the target decision and the expectation implied by current-month Federal funds futures contracts, constructed as described by Kuttner (2001). For the slope shock, the instrument identifies policy induced changes in the slope of the Interbank/Treasury term structure, holding constant changes in the Fed funds rate. To this end, the slope instrument is the residual from a regression of announcement-window changes in the ten-year on-the-run Treasury yield onto the Kuttner shock. This residual is similar in spirit to the path surprise of Gürkaynak, Sack and Swanson (2005), but using a much longer maturity concept of the slope.

The local projections, for $y = \pi, u$ and $i = r^0, s$, are given by

$$y_{t+j} = x^i_t \mathcal{R}^y_{i,j} + w'_t \gamma^y_j + \varepsilon^{i,y}_{t+j} \ , \qquad j = 0, \ldots, J \ , \tag{13}$$

where $x^{r^0}_t$ is the fed funds rate and x^s_t is the slope of the term structure. The impulse responses $\mathcal{R}^y_{i,j}$ are estimated using the instruments z^i_t, where $z^{r^0}_t$ is the Kuttner shock and z^s_t is the slope instrument, that identify the monetary policy shocks ε^i_t.

Figure 5 plots the estimated impulse responses. For the effect of shocks to the fed funds rate, the IRs are consistent with previous findings in the literature (e.g. Barnichon and

Matthes, 2018): (i) monetary policy affects unemployment and inflation with a substantial delay (over a year), (ii) the response of inflation lags that of unemployment by about a year (in particular, the peak response of inflation lags the peak response of unemployment by about four quarters), and (iii) the magnitudes of the peak responses of inflation is about twice that of unemployment while the persistence of the IRFs are similar.

Regarding the effects of shocks to the slope of the yield curve, the IRs of are similar the IRs to a level shock although the effect of a slope shock is less persistent. Quantitatively, a steepening of the yield curve of one percentage point raises unemployment for about two years with a peak effect of almost 2 percentage points (consistent with the results of Eberly, Stock and Wright (2019)) and lowers inflation within one year by about one percentage point.

To compare the inflation-unemployment trade-offs implied by the two instruments, we can compute their corresponding average ability to transform unemployment into inflation, defined by $\kappa = \frac{\|\mathcal{R}^{\pi}\|}{\|\mathcal{R}^{u}\|}$. The impulse responses imply $\kappa_{r^0} = .27$ and $\kappa_s = .45$ (using $J = 20$), which implies that the slope instrument has a smaller average MRT than the fed funds rate: the unemployment cost of moving inflation is smaller with the slope instrument. In other words, the slope instrument is a more powerful tool (on average) for moving inflation. That being said, the difference between κ_{r^0} and κ_s is not statistically significant.

5.3 Results

Figure 6 (top row) displays the optimal discretionary adjustment at time t as implied by equation (11) over 1980-2018 for both conventional monetary policy (through the fed funds rate) and unconventional slope policies, while the middle row displays the corresponding adjusted fed funds rate and the adjusted slope of the yield curve.

While the fed funds rate was not set exactly at its optimal level, the optimal adjustment (in absolute value) is overall relatively small averaging only 10 basis points over the full sample. The only relatively large misses are in the early 80s and perhaps a single miss during the great recession, when the policy rate could have been brought down more rapidly to zero lower-bound.[22]

The slope policies, that started in 2007, could have been used more aggressively during the financial crises and its aftermath, see also Eberly, Stock and Wright (2019). The optimal adjustment drops rapidly to -1.5 percentage points and only slowly revert back to zero indicating that for the period 2008-2012 more extended use of slope policies would have brought the slope of the term structure closer to optimal. Importantly, we note that our approach does not highlight which specific policies would have been able to induce such shift in the slope.

To summarize the distance to optimality in terms of welfare loss we show in the bottom row of Figure 6 the expected loss function under actual policy and under optimal policy, expressed in units of standard-deviation of inflation and unemployment, see equation (12). Similar as above we find two periods in which the Fed has deviated from optimality: the early 80s and during the financial crisis and its aftermath. The maximum loss, between actual and optimal, is found during the latter period and amounts to a roughly 40 percent decline in the average standard deviation of inflation and unemployment (taking $\lambda = 1$).

To understand how the Fed could have performed better, we can contrast the projected paths of inflation and unemployment with the impulse responses to monetary shocks. Figure 7 plots the median SEP forecasts $\mathbb{E}_t\Pi_t$ and $\mathbb{E}_t\Pi_t$ along with our estimated impulse response functions to slope policy shocks. Note that during that period, the fed funds rate is stuck at zero, so that the Fed has only one instrument: its slope policy. Throughout the 2009-2014 period, unemployment is considerably above target and only expected to revert slowly to its long-run value. Similarly, inflation is substantially and persistently below target. Looking at the transmission mechanism of slope policies (bottom row) makes clear how lowering the slope of the yield curve further could have brought unemployment and inflation back to target faster.

Figure 8 illustrates how fed funds rate policies could have been used to further stabilize

[22] For the early 1980s, this could be explained by the need to anchor inflation expectations. Recall that our sufficient statistics are based on the post-1990 sample where inflation expectations have been well anchored.

inflation and unemployment over 2015-2018. Throughout that period, inflation is below target but expected to revert to target relatively quickly, within about a year. Unemployment in contrast has been expected to revert to target in more than two years. This situation is the exact opposite of the transmission mechanism of fed funds rate policy: monetary policy acts faster on unemployment than on inflation. This means that by adjusting the interest rate, the Fed could engineer a faster convergence of unemployment to its target at the (smaller) cost of slower convergence of inflation to target.[23] In other words, by raising unemployment faster to target in exchange for larger inflation deviations later, the Fed could have lowered its quadratic loss function by smoothing the u and π deviations over the horizons $j = 0, ..., J$ of its mandate. This example illustrates one key lesson of this paper: taking into account the transmission dynamics of monetary policy is an essential, but generally little discussed, element of monetary policy making, and the trade-offs involved are not between static targets but instead between targets *and* horizons.

6 Conclusion

If a monetary policy path is chosen optimally, any perturbation to that path should have no first-order effect on welfare. Drawing on this insight, we show that the impulse responses to monetary shocks are sufficient statistics to evaluate the optimality of a given policy. We propose a measure of the distance to optimality of any given policy, define as the discretionary adjustment needed to bring that policy in line with the optimality first-order condition.

Drawing on historical records of monetary reports to Congress going back four decades, we assess the optimality of the Fed's policy over 1980-2018 and find that policy has been remarkably close to optimal, except during 2008-2012 when it was constrained by the zero lower-bound and slope policies could have been used more aggressively.

[23] Focusing on the more immediate situation, this trade-off is exactly the one currently facing the Fed: with an unemployment rate substantially below u_t^* but inflation at target, our optimal discretionary adjustment calls for raising the fed funds rate.

References

Barnichon, Regis, and Christian Matthes. 2018. "Functional Approximation of Impulse Responses." *Journal of Monetary Economics.*

Barnichon, Regis, and Geert Mesters. 2019. "Identifying modern macro equations with old shocks." Working Paper.

Blanchard, Olivier. 2016. "PB 16-11 Do DSGE Models Have a Future?" *Policy Brief PB*, 16–11.

Blanchard, Olivier, and Jordi Galí. 2007. "Real wage rigidities and the New Keynesian model." *Journal of money, credit and banking*, 39(s1): 35–65.

Brainard, William C. 1967. "Uncertainty and the Effectiveness of Policy." *The American Economic Review*, 57(2): 411–425.

Chetty, Raj. 2009. "Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods." *Annu. Rev. Econ.*, 1(1): 451–488.

Chevillon, Guillaume. 2007. "Direct multi-step estimation and forecasting." *Journal of Economic Surveys*, 21(4): 746–785.

Chow, Gregory C. 1972. "Optimal Control of Linear Econometric Systems with Finite Time Horizon." *International Economic Review*, 13(1): 16–25.

Chow, Gregory C. 1973. "Effect of Uncertainty on Optimal Control Policies." *International Economic Review*, 14(3): 632–645.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 1999. "Chapter 2 Monetary policy shocks: What have we learned and to what end?" In . Vol. 1 of *Handbook of Macroeconomics*, 65 – 148. Elsevier.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy*, 113(1): 1–45.

Eberly, Janice C., James H. Stock, and Jonathan H. Wright. 2019. "The Federal Reserve?s Current Framework for Monetary Policy: A Review and Assessment." Working paper.

Galí, Jordi. 2015. *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications.* Princeton University Press.

Gürkaynak, Refet, S., Brian Sack, and Eric Swanson. 2005. "The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models." *American Economic Review*, 95: 425–436.

Hayashi, F. 2000. *Econometrics.* Princeton:Princeton University Press.

Jordà, Oscar. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *The American Economic Review*, 95: 161–182.

Kuttner, Kenneth N. 2001. "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market." *Journal of Monetary Economics*, 47(3): 523–544.

Lakdawala, Aeimit. 2019. "Decomposing the Effects of Monetary Policy Using an External Instruments SVAR." *Journal of Applied Econometrics*. forthcoming.

Leeper, Eric M, and Tao Zha. 2003. "Modest policy interventions." *Journal of Monetary Economics*, 50(8): 1673–1700.

Ljungqvist, L., and T. J. Sargent. 2004. *Recursive Macroeconomic Theory.* Cambridge, Massachusetts:MIT Press.

Lucas, Robert Jr. 1976. "Econometric policy evaluation: A critique." *Carnegie-Rochester Conference Series on Public Policy*, 1(1): 19–46.

Orphanides, Athanasios. 2003. "Monetary policy evaluation with noisy information." *Journal of Monetary Economics*, 50(3): 605 – 631.

Orphanides, Athanasios, and John C Williams. 2012. "Monetary Policy Mistakes and the Evolution of Inflation Expectations." In *The Great Inflation: The Rebirth of Modern Central Banking*. 255–288. University of Chicago Press.

Preston, Alan John, and Adrian R Pagan. 1982. *The theory of economic policy: Statics and dynamics.* Cambridge University Press Cambridge.

Ramey, Valerie. 2016. "Macroeconomic Shocks and Their Propagation." In *Handbook of Macroeconomics*. , ed. J. B. Taylor and H. Uhlig. Amsterdam, North Holland:Elsevier.

Rudebusch, Glenn D. 2002. "Assessing Nominal Income Rules for Monetary Policy with Model and Data Uncertainty." *The Economic Journal*, 112(479): 402–432.

Sack, Brian. 2000. "Does the fed act gradually? A VAR analysis." *Journal of Monetary Economics*, 46(1): 229 – 256.

Stock, James H., and Mark W. Watson. 2018. "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments." *The Economic Journal*, 128(610): 917–948.

Svensson, Lars E. O. 1997. "Inflation forecast targeting: Implementing and monitoring inflation targets." *European Economic Review*, 41(6): 1111–1146.

Svensson, Lars E. O. 2003. "What Is Wrong with Taylor Rules? Using Judgment in Monetary Policy through Targeting Rules." *Journal of Economic Literature*, 41(2): 426–477.

Svensson, Lars E.O., and Michael Woodford. 2005. "Implementing Optimal Policy through Inflation-Forecast Targeting." In *The Inflation-Targeting Debate. NBER Chapters*, 19–92. National Bureau of Economic Research, Inc.

Swanson, Eric T. 2004. "Signal Extraction and Non-Certainty-Equivalence in Optimal Monetary Policy Rules." *Macroeconomic Dynamics*, 8: 27–50.

Theil, H. 1957. "A Note on Certainty Equivalence in Dynamic Planning." *Econometrica*, 25(2): 346–349.

Tinbergen, Jan. 1952. *On the Theory of Economic Policy.* North-Holland Publishing Company, Amsterdam.

Woodford, Michael. 2011. *Interest and prices: Foundations of a theory of monetary policy.* princeton university press.

Appendix

The following lemma generalizes STReSS of Lemma 1 to account for parameter uncertainty.

Lemma 2 (Specific targeting rule with sufficient statistics under uncertainty)**.** *Given assumptions 1-2, the central bank's first order conditions for the minimization problem* $\min_{\delta_t \in \mathbb{R}^M} \mathcal{L}_t$ *are given by*

$$\left(\hat{\mathcal{R}}_t^{\pi'} \mathbb{E}_t(\tilde{\Pi}_t) + \bar{\Omega}_t^{\pi} \delta_t + \bar{C}_t^{\pi}\right) = -\lambda \left(\hat{\mathcal{R}}_t^{u'} \mathbb{E}_t(\tilde{U}_t) + \bar{\Omega}_t^{u} \delta_t + \bar{C}_t^{u}\right) \tag{14}$$

where $\bar{\Omega}_t^y = \sum_{j=0}^J \Omega_{j,t}^y$ *and* $\bar{C}_t^y = \sum_{j=0}^J C_{j,t}^y$*, for* $y = \pi, u$*.*

In other words, the necessary condition is identical to the certainty equivalence case derived in Lemma 1, bar two caveats: first the average uncertainty $\bar{\Omega}_t^y$ scaled by δ_t enters the first order conditions and second the average covariance $\bar{C}_t^y$ between the impulse responses and the random variables x_{t+j}^y is included.

Given that Lemma 1 and Proposition 1 are special cases of Lemma 2 and Proposition 2, we only provide the proofs for the latter.

Proof of lemma 2. Note that we can write the loss function as

$$\begin{aligned}
\mathcal{L}_t &= \mathbb{E}_t \left\|\tilde{\Pi}_t\right\|^2 + \lambda \mathbb{E}_t \left\|\tilde{U}_t\right\|^2 \\
&= \mathbb{E}_t \left(\left\|\tilde{\Pi}_t\right\|^2 - \|\mathbb{E}_t \tilde{\Pi}_t\|^2\right) + \|\mathbb{E}_t \tilde{\Pi}_t\|^2 + \lambda \mathbb{E}_t \left(\left\|\tilde{U}_t\right\|^2 - \|\mathbb{E}_t \tilde{U}_t\|^2\right) + \lambda \|\mathbb{E}_t \tilde{U}_t\|^2 \ .
\end{aligned}$$

Now we have from assumption 1 for $y = \pi, u$, that

$$\left\|\tilde{Y}_t\right\|^2 = \delta_t' \mathcal{R}^{y'} \mathcal{R}^y \delta_t + 2\delta_t' \mathcal{R}^{y'} X_t^y + X_t^{y'} X_t^y$$

and from assumption 2

$$\|\mathbb{E}_t \tilde{Y}_t\|^2 = \delta_t' \hat{\mathcal{R}}_t^{y'} \hat{\mathcal{R}}_t^y \delta_t + 2\delta_t' \hat{\mathcal{R}}_t^{y'} \mathbb{E}_t X_t^y + \mathbb{E}_t X_t^{y'} \mathbb{E}_t X_t^y \ .$$

Combining we have that

$$\begin{aligned}
\mathbb{E}_t \left(\left\|\tilde{Y}_t\right\|^2 - \|\mathbb{E}_t \tilde{Y}_t\|^2\right) &= \delta_t' \mathbb{E}_t \left(\mathcal{R}^{y'} \mathcal{R}^y - \hat{\mathcal{R}}_t^{y'} \hat{\mathcal{R}}_t^y\right) \delta_t \\
&\quad + 2\delta_t' \mathbb{E}_t \left(\mathcal{R}^{y'} X_t^y - \hat{\mathcal{R}}_t^{y'} \mathbb{E}_t X_t^y\right) \\
&\quad + \mathbb{E}_t \left(X_t^{y'} X_t^y - \mathbb{E}_t X_t^{y'} \mathbb{E}_t X_t^y\right) \ .
\end{aligned}$$

Next, we take the derivatives

$$\frac{\partial}{\partial \delta_t} \|\mathbb{E}_t \tilde{Y}_t\|^2 = 2\hat{\mathcal{R}}_t^{y'} \hat{\mathcal{R}}_t^y \delta_t + 2\hat{\mathcal{R}}_t^{y'} \mathbb{E}_t X_t^y = 2\hat{\mathcal{R}}_t^{y'} \mathbb{E}_t \tilde{Y}_t$$

and

$$\frac{\partial}{\partial \delta_t}\mathbb{E}_t\left(\left\|\tilde{Y}_t\right\|^2 - \|\mathbb{E}_t\tilde{Y}_t\|^2\right) = 2\mathbb{E}_t\left(\mathcal{R}^{y\prime}\mathcal{R}^y - \hat{\mathcal{R}}_t^{y\prime}\hat{\mathcal{R}}_t^y\right)\delta_t + 2\mathbb{E}_t\left(\mathcal{R}^{y\prime}X_t^y - \hat{\mathcal{R}}_t^{y\prime}\mathbb{E}_tX_t^y\right)$$

We note that

$$\mathbb{E}_t\left(\mathcal{R}^{y\prime}\mathcal{R}^y - \hat{\mathcal{R}}_t^{y\prime}\hat{\mathcal{R}}_t^y\right) = \sum_{j=0}^{J}\mathbb{E}_t\left(\mathcal{R}_j^y\mathcal{R}_j^{y\prime} - \hat{\mathcal{R}}_{j,t}^y\hat{\mathcal{R}}_{j,t}^{y\prime}\right) = \sum_{j=0}^{J}\mathbb{V}_t(\mathcal{R}_j^y) = \sum_{j=0}^{J}\Omega_{j,t}^y = \bar{\Omega}_t^y$$

where $\mathcal{R}_j^y$ is the jth row of $\mathcal{R}^y$ and $\Omega_{j,t}^y = \mathbb{V}_t(\mathcal{R}_j^y)$. Further,

$$\mathbb{E}_t\left(\mathcal{R}^{y\prime}X_t^y - \hat{\mathcal{R}}_t^{y\prime}\mathbb{E}_tX_t^y\right) = \sum_{j=0}^{J}\mathbb{E}_t\left(\mathcal{R}_j^y x_{t+j}^y - \hat{\mathcal{R}}_{j,t}^y\mathbb{E}_t x_{t+j}^y\right) = \sum_{j=0}^{J}C_{j,t}^y = \bar{C}_t^y$$

where $C_{j,t}^y = \mathbb{C}_t(\mathcal{R}_j^{y\prime}, x_{t+j}^y)$. Finally, when we combine the pieces we get

$$\frac{\partial}{\partial \delta_t}\mathcal{L}_t = 2\bar{\Omega}_t^\pi\delta_t + 2\bar{C}_t^\pi + 2\hat{\mathcal{R}}_t^{\pi\prime}\mathbb{E}_t\tilde{\Pi}_t + 2\lambda\bar{\Omega}_t^u\delta_t + 2\lambda\bar{C}_t^u + 2\lambda\hat{\mathcal{R}}_t^{u\prime}\mathbb{E}_t\tilde{U}_t$$

Which when set to zero gives the result. □

Proof of proposition 2. Starting from proposition 2 we have

$$\left(\hat{\mathcal{R}}_t^{\pi\prime}\mathbb{E}_t(\tilde{\Pi}_t) + \bar{\Omega}_t^\pi\delta_t + \bar{C}_t^\pi\right) = -\lambda\left(\hat{\mathcal{R}}_t^{u\prime}\mathbb{E}_t(\tilde{U}_t) + \bar{\Omega}_t^u\delta_t + \bar{C}_t^u\right)$$

Plugging in Assumption 2 for $\mathbb{E}_t(\tilde{\Pi}_t)$ and $\mathbb{E}_t(\tilde{U}_t)$, after reordering gives

$$\left(\hat{\mathcal{R}}_t^{\pi\prime}\hat{\mathcal{R}}_t^\pi + \lambda\hat{\mathcal{R}}_t^{u\prime}\hat{\mathcal{R}}_t^u + \bar{\Omega}_t^\pi + \lambda\bar{\Omega}_t^u\right)\delta_t = -\hat{\mathcal{R}}_t^{\pi\prime}\mathbb{E}_t(X_t^\pi) - \bar{C}_t^\pi - \lambda\hat{\mathcal{R}}_t^{u\prime}\mathbb{E}_t(X_t^u) - \lambda\bar{C}_t^u$$

Since all matrices in the brackets are positive definite by assumption and $\lambda \geq 0$ we have that

$$\hat{\delta}_t = -\left(\hat{\mathcal{R}}_t^{\pi\prime}\hat{\mathcal{R}}_t^\pi + \lambda\hat{\mathcal{R}}_t^{u\prime}\hat{\mathcal{R}}_t^u + \bar{\Omega}_t^\pi + \lambda\bar{\Omega}_t^u\right)^{-1}\left(\hat{\mathcal{R}}_t^{\pi\prime}\mathbb{E}_t(X_t^\pi) + \bar{C}_t^\pi + \lambda\hat{\mathcal{R}}_t^{u\prime}\mathbb{E}_t(X_t^u) + \lambda\bar{C}_t^u\right)\ ,$$

which is the expression in the proposition. □

Figure 1: The transmission of monetary policy

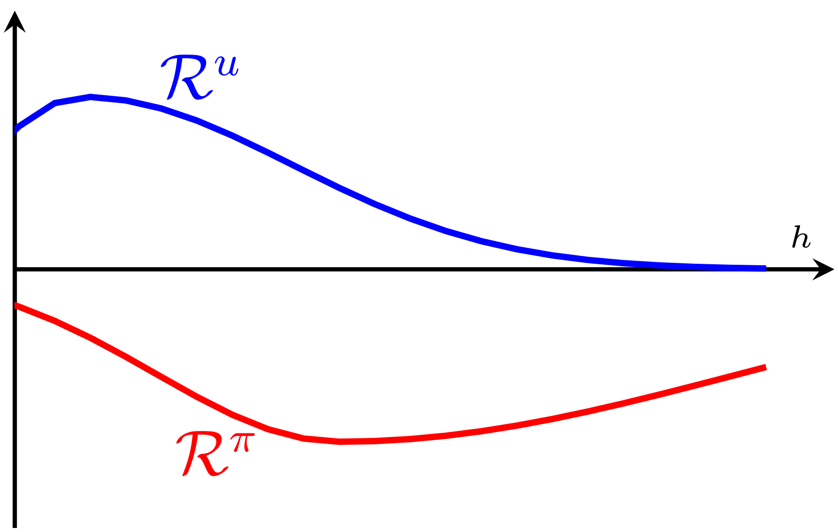

Notes: Examples of impulse responses for inflation ($\mathcal{R}^\pi_\tau$) and unemployment ($\mathcal{R}^u_\tau$) to a monetary shock.

Figure 2: ODA to a deflationary shock for a strict inflation targeter ($\lambda = 0$)

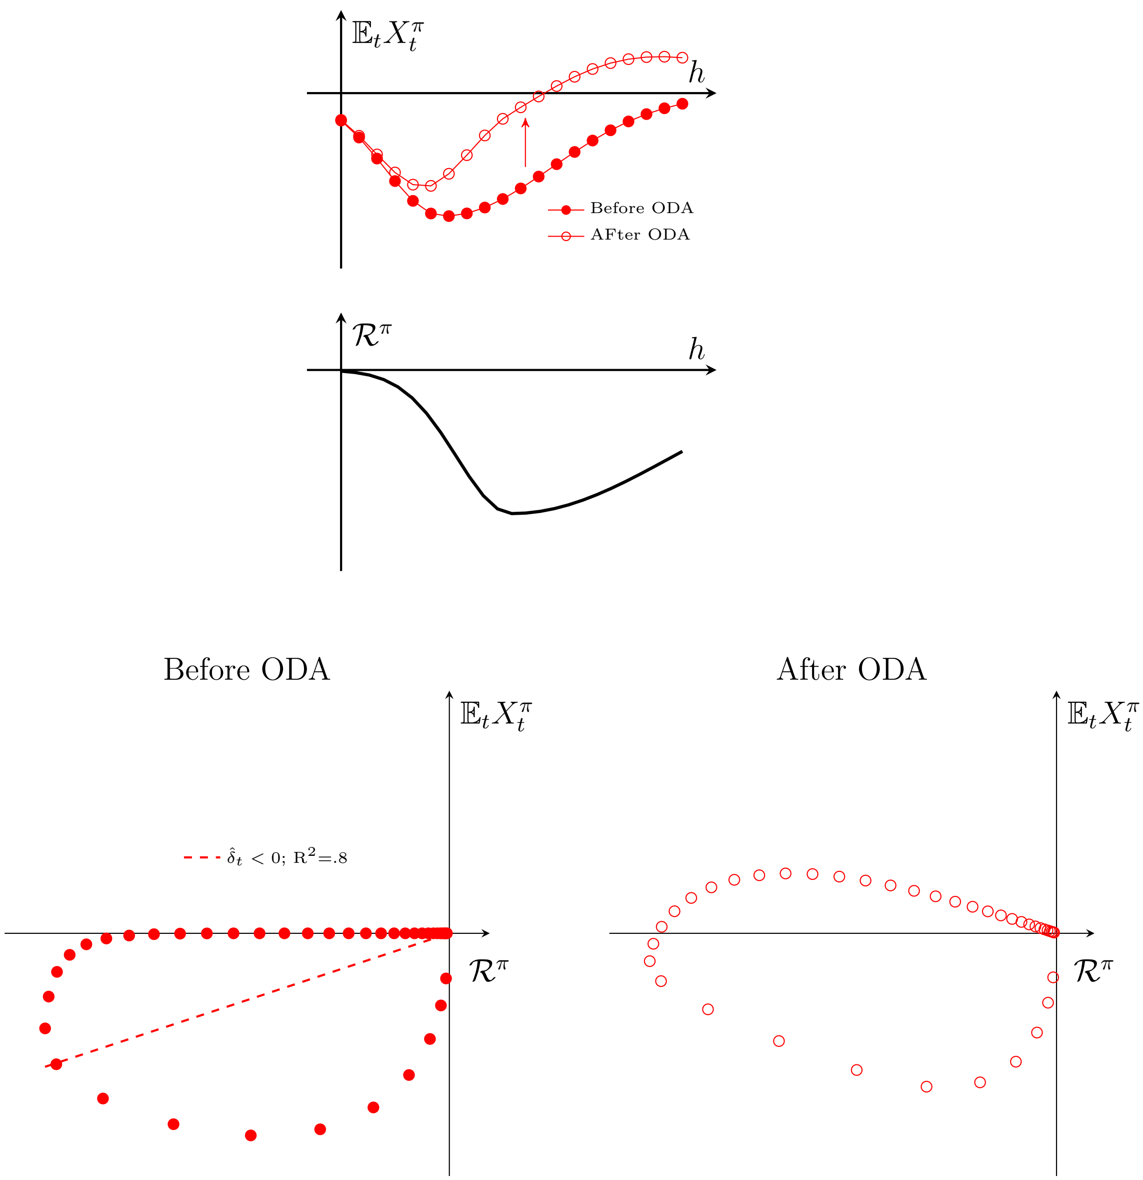

Notes. Top panel: Expected paths for inflation deviations without Optimal Discretionary Adjustment ("Before ODA", filled circles) and with Optimal Discretionary Adjustment ("After ODA", empty circles), i.e., under optimal policy under $\lambda = 0$. Middle panel: impulse responses of inflation to a monetary shock. Bottom panel: scatter plot of expected future paths of inflation against the impulse responses of inflation (in red). The red dashed-line is the best linear fit.

Figure 3: ODA to an adverse (AD) shock

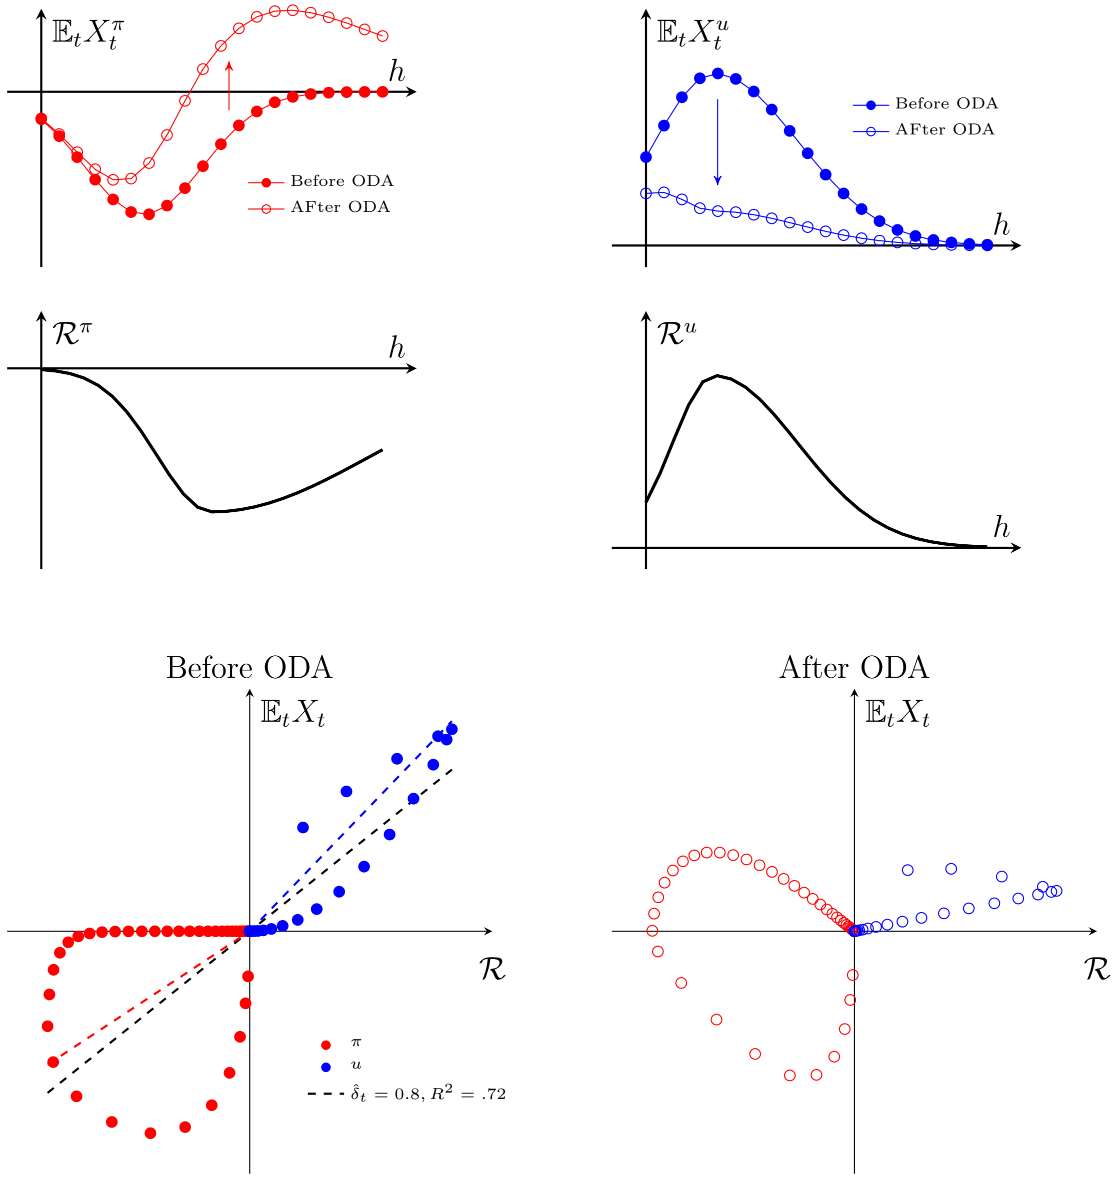

Notes. Top panel: Expected paths for inflation and unemployment deviations without Optimal Discretionary Adjustment ("Before ODA", filled circles) and with Optimal Discretionary Adjustment ("After ODA", empty circles), i.e., under optimal policy under $\lambda = 1$ (dashed lines). Middle panel: impulse responses of inflation and unemployment to a monetary shock. Bottom panel: scatter plot of expected future paths of inflation against the impulse responses of inflation (in red) and similarly for unemployment (in blue). The blue dashed-line is the best linear fit for unemployment, the red dashed-line is the best linear fit for inflation, and the dashed black line is the best-linear fit after including all points.

Figure 4: ODA to a deflationary (AS) shock

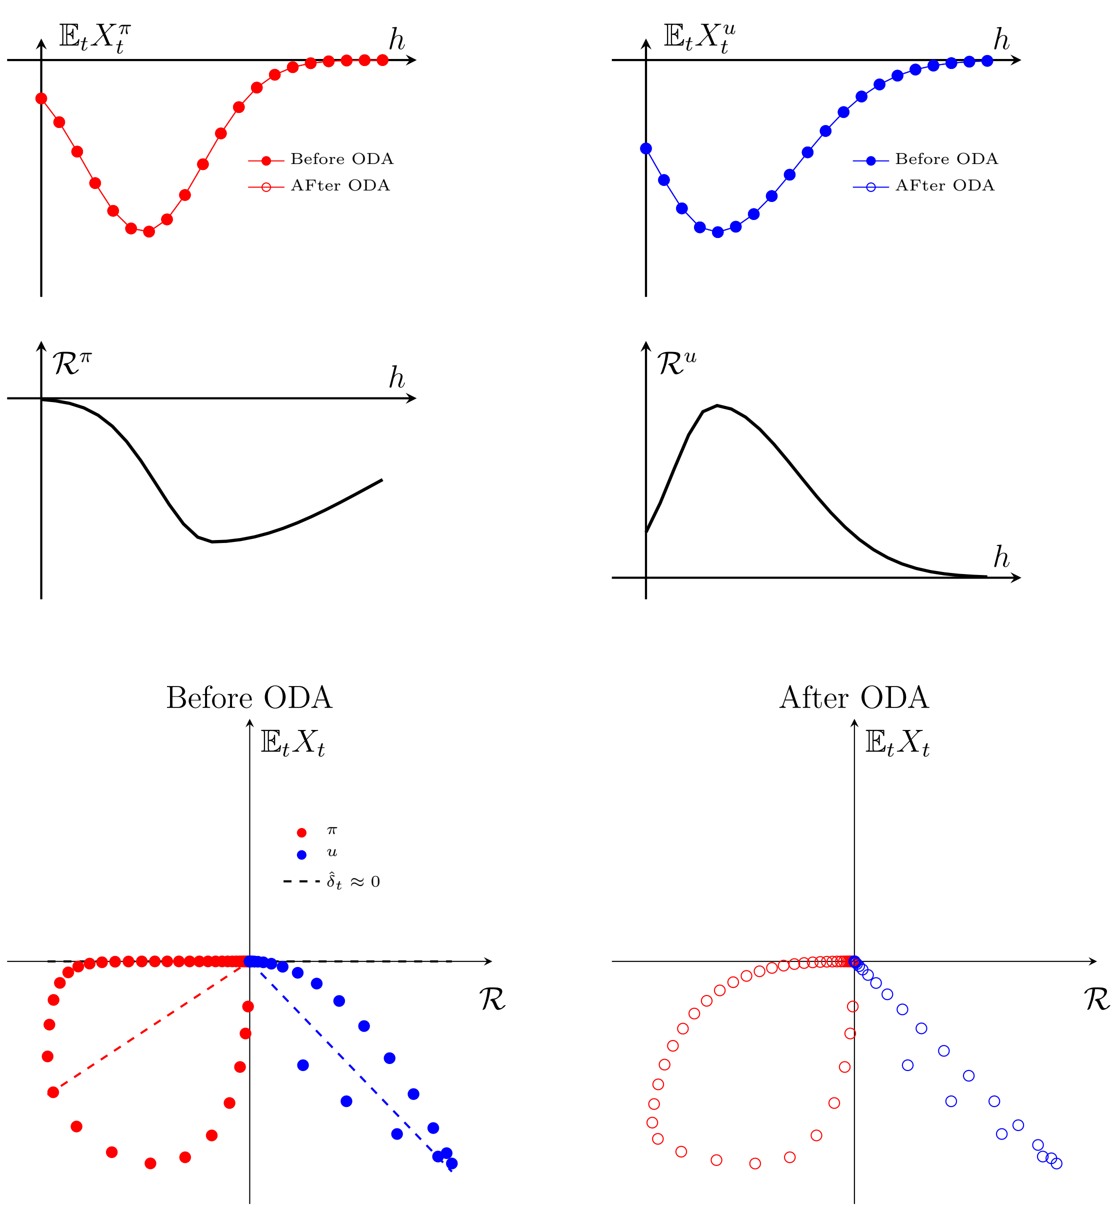

Notes. Top panel: Expected paths for inflation and unemployment deviations without Optimal Discretionary Adjustment ("Before ODA", filled circles) and with Optimal Discretionary Adjustment ("After ODA", empty circles), i.e., under optimal policy under $\lambda = 1$ (dashed lines). Middle panel: impulse responses of inflation and unemployment to a monetary shock. Bottom panel: scatter plot of expected future paths of inflation against the impulse responses of inflation (in red) and similarly for unemployment (in blue). The blue dashed-line is the best linear fit for unemployment, the red dashed-line is the best linear fit for inflation, and the dashed black line is the best-linear fit after including all points.

Figure 5: Impulse responses of inflation and unemployment to monetary shocks

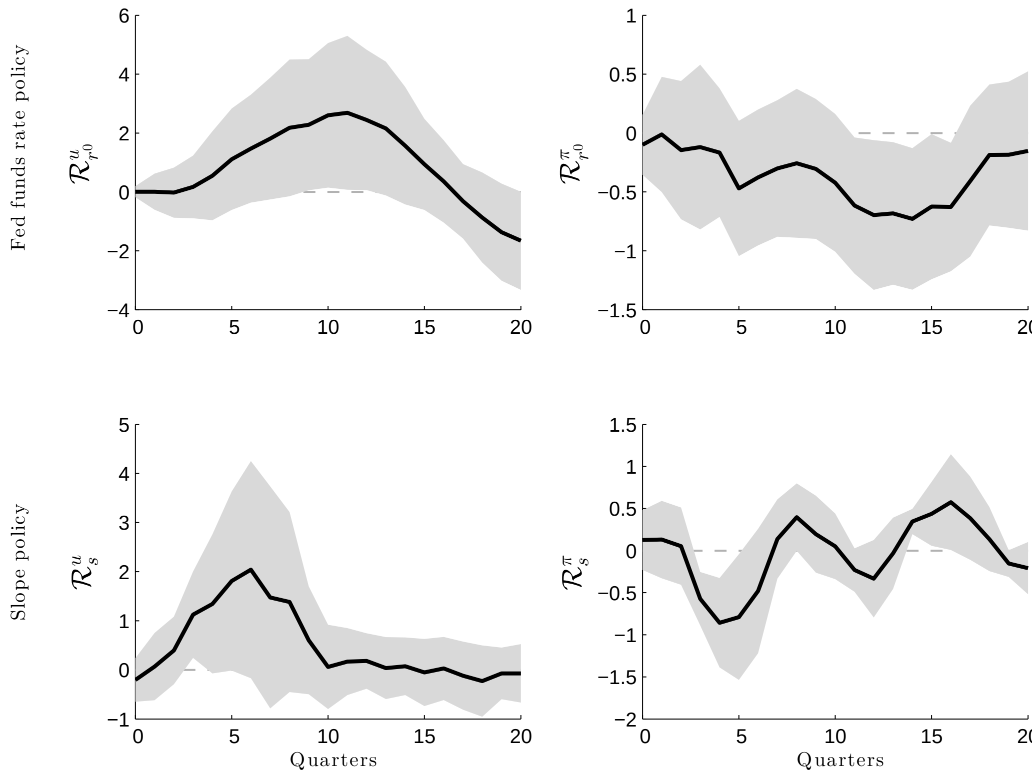

Notes: Top row: Impulse responses of unemployment and inflation to a fed funds rate shock computed using High-Frequency Identified (HFI) monetary surprises over 1990q1–2017q4. Bottom row: Impulse responses of unemployment and inflation to a shock to the slope of the yield curve (the spread between the 10-Year treasury constant maturity and the fed funds rate) computed using HFI monetary surprises over 2007q1–2017q4. 95% confidence intervals for the local projection estimates are displayed.

Figure 6: Distance to optimality for the Fed instruments (1980-2018)

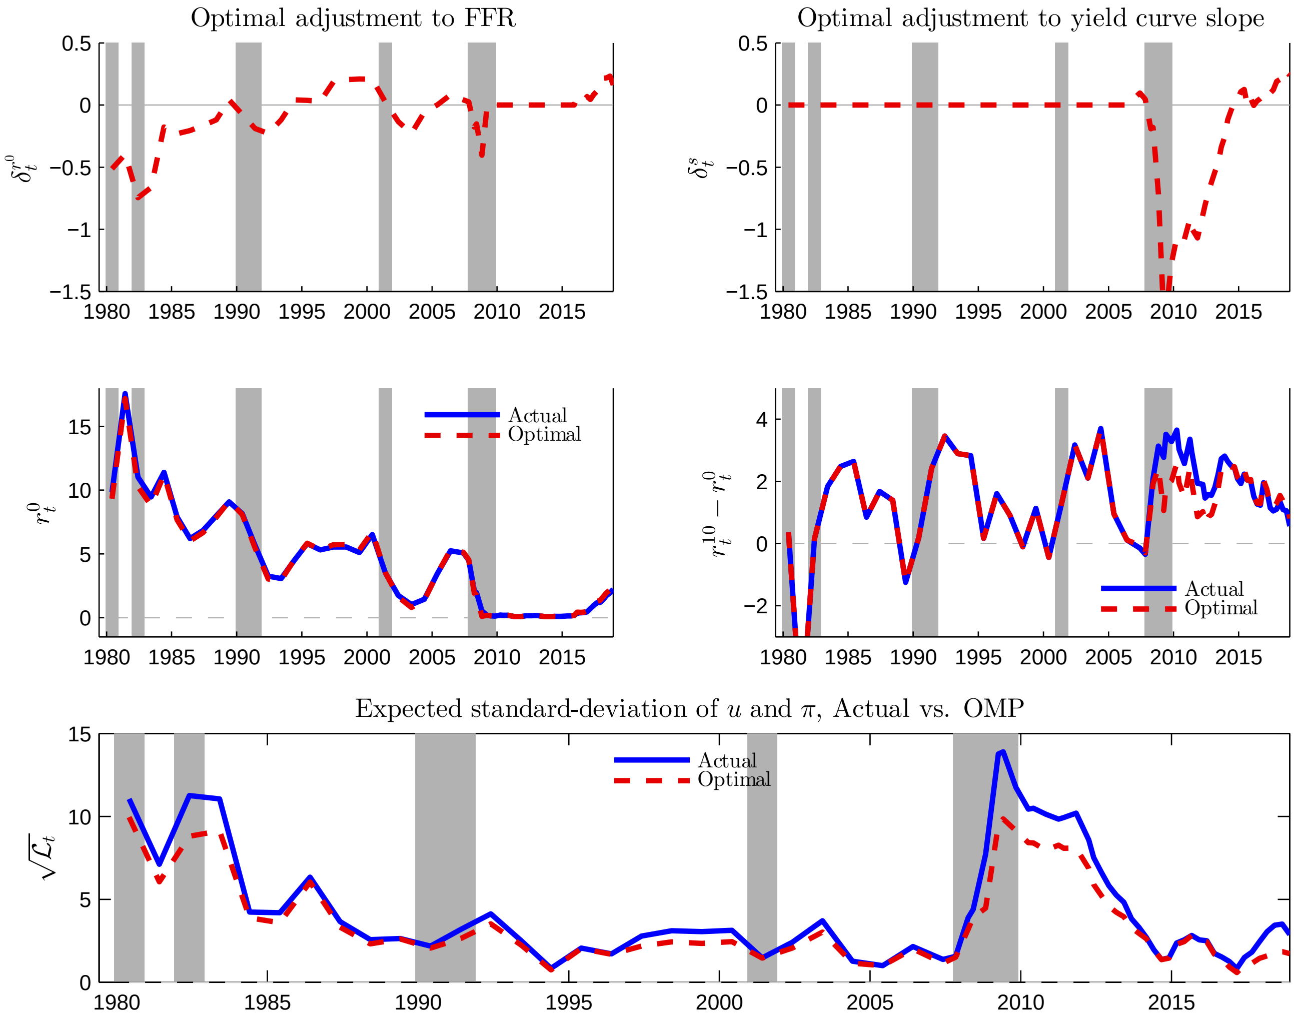

Notes: Left panel: Optimal discretionary adjustment $\delta_t^{r^0}$ to the fed funds rate (r_t^0) at time t. Right panel: Optimal discretionary adjustment δ_t^s to the slope of the yield curve ($r_t^{10} - r_t^0$, the spread between 10-year treasury bond and fed funds rate) at time t.

Figure 7: Using slope policies during the Great Recession (2009-2014)

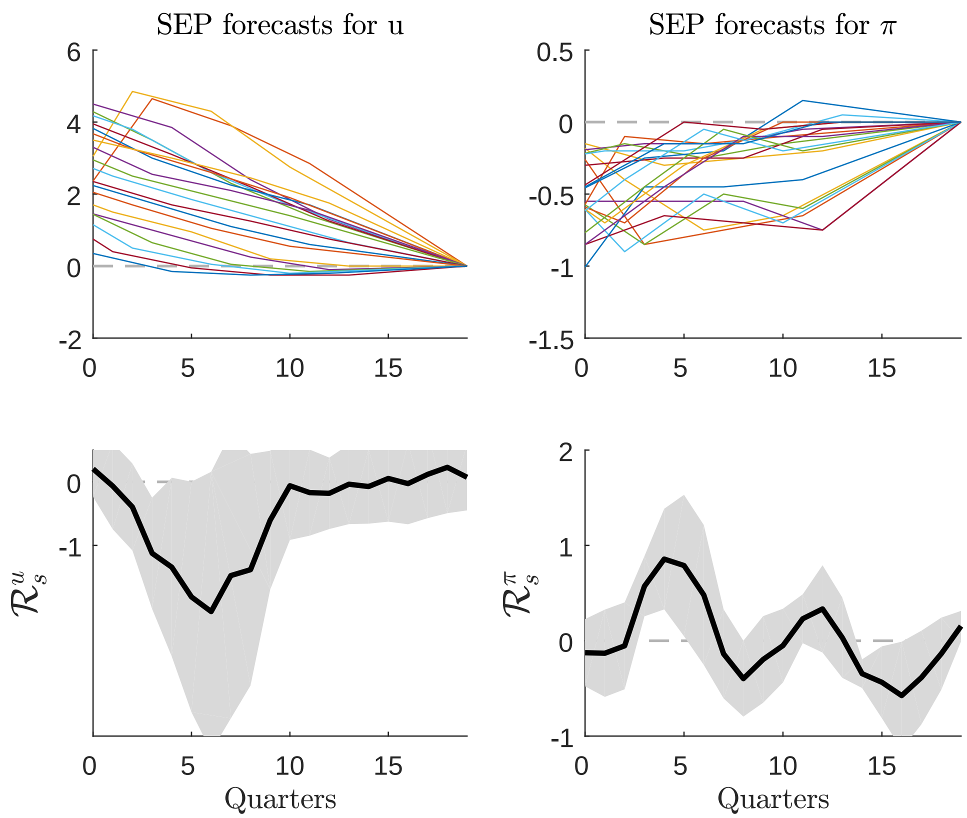

Notes: Top row: SEP median forecasts for $u_t - u_t^*$ and $\pi_t - \pi_t^*$ over 2009-2014. Bottom row: Impulse responses of inflation and unemployment to a negative slope shock computed using High-Frequency Identified (HFI) monetary surprises.

Figure 8: Using fed funds rate policies (2015-2018)

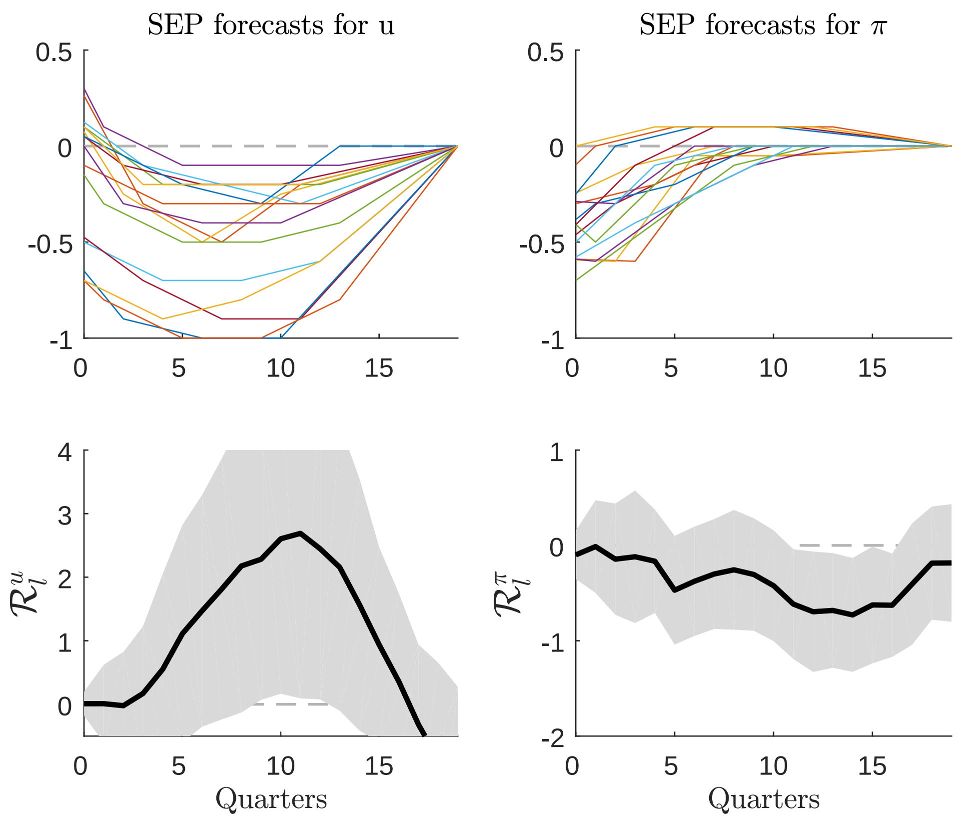

Notes: Top row: SEP median forecasts for $u_t - u_t^*$ and $\pi_t - \pi_t^*$ over 2015-2018. Bottom row: Impulse responses of inflation and unemployment to a fed funds rate shock computed using High-Frequency Identified (HFI) monetary surprises.