Web-appendix for:

A Sufficient Statistics Approach for Macro Policy Evaluation

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

[(a)] Federal Reserve Bank of San Francisco and CEPR

[(b)] Universitat Pompeu Fabra and BSE

April 27, 2022

Abstract

In this web-appendix, we provide the following additional results:

S1: Example: Optimal policy under commitment

S2: General nonlinear OPP framework

S3: General convex loss functions

S4: Estimating robust preference parameters

S5: Additional results for the empirical study

References to lemmas, equations, etc..., which start with a "S" are references to this document. References, which consist of only a number refer to the main text.

S1 Example: Optimal policy under commitment

In this section we present the details for our sufficient statistics approach when applied to the optimal policy problem under commitment in the baseline New Keynesian model (e.g. Galí, 2015, Section 5.1.2). This serves to show that our approach is not only valid under discretion but can also handle cases of commitment. In other words, regardless which policy problem the policy maker considers we will show that $\mathcal{R}'\mathbb{E}_t\mathbf{Y}_t^0 = 0$ remains a necessary and sufficient condition for optimality. Implying that with estimates for $\mathcal{R}$ and $\mathbb{E}_t\mathbf{Y}_t^0$ it is possible to verify the optimality of the proposed policy choice.

The optimal policy problem under commitment is characterized by a policy maker that aims to minimize

$$\mathcal{L}_0 = \frac{1}{2}\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t(\pi_t^2 + x_t^2) \ ,$$

with respect to $\pi_t, \pi_{t+1}, ...$ and $x_t, x_{t+1}, ...$ and subject to the constraints

$$\pi_t = \beta\mathbb{E}_t\pi_{t+1} + \kappa x_t + \xi_t \ ,$$
$$x_t = \mathbb{E}_t x_{t+1} - \frac{1}{\sigma}(i_t - \mathrm{E}_t\pi_{t+1}) \ .$$

where β is the discount factor. Note that we can view this example as case a special case of our general policy problem (14) when taking $\mathbb{E}_t$ as $\mathbb{E}_0$, i.e. starting at $t = 0$.

The optimality conditions for this problem are given by

$$x_0 = -\kappa\pi_0 \qquad \text{and} \qquad x_t = x_{t-1} - \kappa\pi_t \ , \quad \forall\ t = 1, 2, \dots \ , \tag{S1}$$

or

$$x_t = -\kappa\hat{\mathrm{p}}_t \quad \forall\ t = 0, 1, 2, \dots \ ,$$

where $\hat{\mathrm{p}}_t = \mathrm{p}_t - \mathrm{p}_{-1}$ denotes the (log) deviation between the price level and an "implicit target" given by the price level prevailing one period before the central bank chooses its optimal plan (Galí, 2015, page 135).

A possible interest rate rule that (a) implements this optimal allocation and (b) leads to a unique equilibrium is given by

$$i_t = -[\phi_p + (1-\delta)(1-\kappa\sigma)]\sum_{k=0}^{t}\delta^k\xi_{t-k} - (\phi_p/\kappa)x_t$$

for any $\phi_p > 0$ (Galí, 2015, page 138). Note that this instrument rule is a special case of the generic policy rule (17). The coefficients in the rule are given by

$$\delta \equiv \frac{1-\sqrt{1-4\beta a^2}}{2a\beta} \;, \qquad \text{with} \qquad a \equiv \frac{1}{1+\beta+\kappa^2} \;.$$

The forecasts under the optimal allocation can be written as

$$\mathbb{E}_0\pi_0 = \delta\xi_0 \qquad \mathbb{E}_0 x_0 = -\kappa\delta\xi_0 \qquad \mathbb{E}_0\pi_t = (\delta^{t+1}-\delta^t)\xi_0 \qquad \mathbb{E}_0 x_t = -\kappa\delta^{t+1}\xi_0 \tag{S2}$$

for $t \geq 1$ (Galí, 2015, page 136).

Next, we rewrite this example in our general notation. Let $\mathbf{Y}_0 = (\pi_0, x_0, \pi_1, x_1, \ldots)'$, $\mathbf{P}_0 = (i_0, i_1, \ldots)'$, $\mathbf{\Xi}_0 = (\xi_0, \xi_1, \ldots)'$ and denote by $\boldsymbol{\epsilon}_0 = (\epsilon_0, \epsilon_1, \ldots)'$ the sequence of policy shocks, which are equal to zero under the optimal rule (note that $\mathbf{W}_t$ does not exist in this application). The general model (15)-(17) becomes

$$\begin{aligned} \mathcal{A}_{yy}\mathbb{E}_0\mathbf{Y}_0 - \mathcal{A}_{yp}\mathbb{E}_0\mathbf{P}_0 &= \mathcal{B}_{y\xi}\mathbb{E}_0\mathbf{\Xi}_0 \\ \mathbb{E}_0\mathbf{P}_0 - \mathcal{A}_{py}\mathbb{E}_0\mathbf{Y}_0 &= \mathcal{B}_{p\xi}\mathbb{E}_0\mathbf{\Xi}_0 + \boldsymbol{\epsilon}_0^e \end{aligned}$$

where the coefficient maps are given by

$$\mathcal{A}_{yy} = \begin{bmatrix} 1 & -\kappa & -\beta & 0 & \ldots & \ldots \\ 0 & 1 & -1/\sigma & -1 & 0 & \ldots \\ 0 & 0 & 1 & -\kappa & -\beta & \ddots \\ 0 & 0 & 0 & 1 & -1/\sigma & \ddots \\ \vdots & \vdots & \ddots & \ddots & \ddots & \ddots \end{bmatrix} \qquad \mathcal{A}_{yp} = \begin{bmatrix} 0 & 0 & 0 & \ldots \\ 1/\sigma & 0 & 0 & \ldots \\ 0 & 0 & 0 & \ldots \\ 0 & 1/\sigma & 0 & \ldots \\ 0 & 0 & 0 & \ddots \\ 0 & 0 & 1/\sigma & \ddots \\ \vdots & \vdots & \ddots & \ddots \end{bmatrix}$$

$$\mathcal{A}_{py} = \begin{bmatrix} 0 & \phi_p/\kappa & 0 & 0 & \ldots & \ldots \\ 0 & 0 & 0 & \phi_p/\kappa & 0 & \ldots \\ 0 & 0 & 0 & 0 & 0 & \phi_p/\kappa \\ \vdots & \vdots & \ddots & \ddots & \ddots & \ddots \end{bmatrix}$$

and

$$\mathcal{B}_{y\xi} = \begin{bmatrix} 1 & 0 & 0 & \dots \\ 0 & 0 & 0 & \dots \\ 0 & 1 & 0 & \dots \\ 0 & 0 & 0 & \dots \\ 0 & 0 & 1 & \ddots \\ 0 & 0 & 0 & \ddots \\ \vdots & \vdots & \ddots & \ddots \end{bmatrix} \qquad \mathcal{B}_{p\xi} = \begin{bmatrix} \gamma_0 & 0 & 0 & \dots \\ \gamma_1 & \gamma_0 & 0 & \dots \\ \gamma_2 & \gamma_1 & \gamma_0 & \ddots \\ \vdots & \ddots & \ddots & \ddots \end{bmatrix}$$

where $\gamma_j = -[\phi_p + (1-\delta)(1-\kappa\sigma)]\delta^j$. It follows that $\mathcal{R}$, after some tedious manipulations, can be written —under the optimal policy rule— as

$$\mathcal{R} = \begin{bmatrix} \kappa/(\sigma v) & \kappa^2/(\sigma^2 v^2) + \kappa/\sigma v^2 + \kappa/(\sigma v) & \dots \\ 1/(\sigma v) & \kappa/(\sigma^2 v^2) + 1/(\sigma v^2) & \dots \\ 0 & \kappa/(\sigma v) & \dots \\ 0 & 1/(\sigma v) & \\ 0 & 0 & \dots \\ \vdots & \vdots & \ddots \end{bmatrix}$$

where $v = 1 - \phi_p/(\kappa/\sigma)$. Note that we only show the first two columns for ease of exposition. Given $\mathcal{R}$ and the forecasts (S2) we can verify the equivalence condition, similar as shown in equation (8) for the problem under discretion we have

$$\begin{aligned} \left.\frac{\partial \mathcal{L}_0}{\partial \boldsymbol{\epsilon}_0}\right|_{i_t} &= \mathcal{R}' \mathbb{E}_0 \mathbf{Y}_0 \\ &= \begin{bmatrix} \kappa/(\sigma v) & \kappa^2/(\sigma^2 v^2) + \kappa/\sigma v^2 + \kappa/(\sigma v) & \dots \\ 1/(\sigma v) & \kappa/(\sigma^2 v^2) + 1/(\sigma v^2) & \dots \\ 0 & \kappa/(\sigma v) & \dots \\ 0 & 1/(\sigma v) & \\ 0 & 0 & \dots \\ \vdots & \vdots & \ddots \end{bmatrix}' \begin{bmatrix} \delta\xi_0 \\ -\kappa\delta\xi_0 \\ (\delta^2 - \delta)\xi_0 \\ -\kappa\delta^2\xi_0 \\ \vdots \\ \vdots \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ \vdots \end{bmatrix}. \end{aligned}$$

This shows that $\mathcal{R}'\mathbb{E}_0\mathbf{Y}_0 = 0$ must also hold under optimality when considering problems with commitment.

In addition, note that as in the case under discretion, the impulse response matrix $\mathcal{R}$ is sufficient to characterize the optimal targeting rule. Working under perfect foresight, the condition $\mathcal{R}'\mathbf{Y}_0 = 0$ corresponds exactly to the optimal targeting rule (S1).

S2 General nonlinear OPP framework

In this section we formalize the generalization of our sufficient statistics approach for nonlinear models. First, we discuss a high level framework that expounds the exact class of models for which the OPP approach to policy evaluation can work. Second, we provide the details for the specific cases of state dependence and multiple policy regimes that were discussed in Section 7.

B1: Generic nonlinear OPP

Consider an economy that can be represented at time t by the moment equations

$$\left\{ \begin{array}{lcl} \mathbf{0} & = & \mathbb{E}_t f(\mathbf{Y}_t, \mathbf{P}_t, \mathbf{X}_{-t}, \boldsymbol{\Xi}_t, ; \phi) \\ \mathbf{0} & = & \mathbb{E}_t \phi(\mathbf{Y}_t, \mathbf{P}_t, \mathbf{X}_{-t}, \boldsymbol{\Xi}_t) \end{array} \right. , \tag{S3}$$

where $f()$ and $\phi()$ are possibly nonlinear functions. The function $f()$ describes the general economy and takes as inputs the policy variables $\mathbf{Y}_t$, the policy instruments $\mathbf{P}_t$, the initial conditions $\mathbf{X}_{-t}$ and the structural shocks $\boldsymbol{\Xi}_t$. The policy equation is characterized by ϕ, which takes similar inputs. A key difference with the set up in the main text is that we allow f to depend directly on the policy rule ϕ. This generalization directly corresponds to the generic set-up in Lucas (1976), see equations 16 and 17 in his paper, which allows the policy rule to alter the function $f()$ in an arbitrary way.

The optimal allocation, for the loss function $\mathcal{L}_t = \mathbb{E}_t \mathbf{Y}_t' \mathcal{W} \mathbf{Y}_t$, in this economy is characterized by

$$\min_{\mathbf{Y}_t, \mathbf{P}_t, \phi} \mathcal{L}_t \qquad s.t. \qquad (S3)$$

where the difference is that the optimal allocation now also depends on ϕ, the choice for the policy rule that may affect the structure in the economy. The optimal expected policy path is again denoted by $\mathbf{P}_t^{e\,\mathrm{opt}}$.

To build up to our sufficient statistics approach consider a policy maker who can introduce exogenous surprises in the policy rule ϕ. As before we denote such policy shocks by $\boldsymbol{\epsilon}_t$ and we postulate that each element of $\mathbf{P}_t$ corresponds to a specific element in $\boldsymbol{\epsilon}_t$. The economy becomes

$$\left\{ \begin{array}{lcl} \mathbf{0} & = & \mathbb{E}_t f(\mathbf{Y}_t, \mathbf{P}_t, \mathbf{X}_{-t}, \boldsymbol{\Xi}_t, ; \phi) \\ \mathbf{0} & = & \mathbb{E}_t \phi(\mathbf{Y}_t, \mathbf{P}_t, \mathbf{X}_{-t}, \boldsymbol{\Xi}_t, \boldsymbol{\epsilon}_t) \end{array} \right. . \tag{S4}$$

In this setting a policy choice is determined by the function $\phi \in \Phi$, where Φ denotes an arbitrary function class and a sequence of policy shocks $\boldsymbol{\epsilon}_t$. Let the proposed policy choice be denoted by $(\phi^0, \boldsymbol{\epsilon}_t^0)$, which implicitly characterizes the expected policy path $\mathbf{P}_t^{e0}$ and allocation $\mathbb{E}_t \mathbf{Y}_t^0$.

The following high-level assumption exactly determines the class of models for which our sufficient statistics approach continues to work.

Assumption S1. *There exists a non-empty subset $\Phi^{\text{opt}} \subset \Phi$ such that*

1. *for all $\phi \in \Phi^{\text{opt}}$ and $\phi = \phi^0$ we have a unique and determinate equilibrium given by*

$$\begin{cases} \mathbb{E}_t \mathbf{Y}_t &= h_y(\mathbf{X}_{-t}, \mathbb{E}_t \mathbf{\Xi}_t, \boldsymbol{\epsilon}_t^e; \phi) \\ \mathbf{P}_t^e &= h_p(\mathbf{X}_{-t}, \mathbb{E}_t \mathbf{\Xi}_t, \boldsymbol{\epsilon}_t^e; \phi) \end{cases} .$$

   where $h_y()$ is continuously differentiable with respect to all $\boldsymbol{\epsilon}_t^e \in \mathcal{E}$, where $\mathcal{E}$ is an open convex subset of $\mathbb{R}^\infty$.

2. *$\mathcal{L}_t(\phi, \mathbf{0}) \leq \mathcal{L}_t(\tilde{\phi}, \tilde{\boldsymbol{\epsilon}}_t^e)$ for all $\phi \in \Phi^{\text{opt}}$, $\tilde{\phi} \in \Phi \backslash \Phi^{\text{opt}}$ and $\tilde{\boldsymbol{\epsilon}}_t^e \in \mathcal{E}$.*

The first part of the assumption imposes the existence of a unique equilibrium under the optimal policy rules ϕ^{opt} and the proposed policy rule ϕ^0. The second part defines the optimal rules as those that minimize the loss function. The key part in the assumption is that the loss function is minimized by ϕ^{opt} with $\boldsymbol{\epsilon}_t^e = 0$. That is, under the optimal rule it is not possible to further lower the loss function by introducing exogenous policy news shocks.[1]

For nonlinear models like (S3) this assumption would need to be verified on a case by case basis. To better understand Assumption S1, it is helpful to consider our baseline (linear) framework and discuss how AssumptionS1 is satisfied in this simpler case.

First, for the first part of AssumptionS1, note that Assumption 2 in the main text ensures that a unique equilibrium exists under ϕ^0. In addition, AssumptionS1 also require that a unique equilibrium exists under ϕ^{opt}. The reason is that we need to ensure that the gradient of $\mathbf{Y}_t$ with respect to $\boldsymbol{\epsilon}_t$ exists under the optimal policy. In the linear model this is not necessary as the optimal policy allocation does not depend on ϕ, see equation (16). With $f(\cdot)$ depending on ϕ however, this is no longer the case and the optimal policy does depend on ϕ. In that case, we also need a "well-behaved" optimal rule, i.e., that ϕ^{opt} implies a unique equilibrium. The second part of Assumption S1 holds in our linear framework as the optimal allocation $\mathbf{P}_t^{e\,\text{opt}}$ does not depend on $\boldsymbol{\epsilon}_t^e$ and can be attained by a rule of form (17). For instance, we may take $\mathcal{A}_{pp} = \mathcal{A}_{pw} = \mathcal{B}_{px} = \mathcal{B}_{p\xi} = 0$ and $\mathcal{A}_{py} = \mathcal{R}^{0'}\mathcal{W}$, to obtain the optimal targeting rule $\mathcal{R}^{0'}\mathcal{W}\mathbb{E}_t\mathbf{Y}_t^0 = 0$.

[1]If the the functional form of the policy rule (the function $\phi()$) is unrestricted, this assumption is mild as there is no reason for the policy maker to introduce exogenous shocks: Since the loss function is conditioned on time t information, the policy maker should be able to minimize the loss function by choosing an arbitrarily complex ϕ that optimally responds to the time-t information set. That said, if the class of rules is restricted (say because there is a limit to the complexity of the rules that policy makers can follow in practice), additional policy surprises may lower the loss function. In other words, whether assumption S1-part 2 holds depends on the interplay between the complexity of the economy and the class of policy rules that can be implemented.

With assumption S1 in hand we can proceed with our policy evaluation, similar as in the main text we are interested in testing $H_0 : \mathbf{P}_t^{e0} = \mathbf{P}_t^{e\text{opt}}$. As a test statistic we rely on the generalized OPP statistic given by

$$\boldsymbol{\delta}_t^{g*} = -(\mathfrak{R}^{0\prime}\mathcal{W}\mathfrak{R}^0)^{-1}\mathfrak{R}^{0\prime}\mathcal{W}\mathbb{E}_t Y_t^0 \ ,$$

where

$$\mathfrak{R}^0 = \left.\frac{\partial h_y(\mathbf{X}_{-t}, \mathbb{E}_t\boldsymbol{\Xi}_t, \boldsymbol{\epsilon}_t^e; \phi)}{\partial \boldsymbol{\epsilon}_t^e}\right|_{\phi=\phi^0, \boldsymbol{\epsilon}_t=\boldsymbol{\epsilon}_t^0} .$$

Two comments are in order. First, note that for the linear model considered in the main text we have $h_y(\mathbf{X}_{-t}, \boldsymbol{\Xi}_t, \boldsymbol{\epsilon}_t^e; \phi^0) = \mathcal{R}^0\boldsymbol{\epsilon}_t^e + \boldsymbol{\Upsilon}_t$, see Proposition 1, such that we immediately have $\mathfrak{R}^0 = \mathcal{R}^0$. Second, we do not claim that $\mathfrak{R}^0$ can generally be estimated using conventional econometric methods. Clearly, when the function $h_y()$ is unknown this is complicated as non-parametric methods will need to be used, which given the typically limited time series observations available may yield uninformative causal effects. The point here is simply to illustrate the theoretical limit of our approach.

We have the following key result.

Theorem S1. *Given model* (S4) *under assumption S1 if* $\boldsymbol{\delta}_t^{g*}$ *exists we have that*

$$\boldsymbol{\delta}_t^{g*} \neq 0 \qquad \Rightarrow \qquad \mathbf{P}_t^{e0} \neq \mathbf{P}_t^{e\text{opt}} \ .$$

Proof. By Assumption S1 part 1, the loss function $\mathcal{L}_t$ is continuously differentiable on $\mathcal{E}$, thus by Lemma 4.3.1 in Dennis and Schnabel (1996) and Assumption S1 the optimal policy $\mathbf{P}_t^{e\text{opt}} = h_p(\mathbf{X}_{-t}, \mathbb{E}_t\boldsymbol{\Xi}_t, \mathbf{0}; \phi^{\text{opt}})$ for any $\phi^{\text{opt}} \in \Phi^{\text{opt}}$ satisfies the gradient condition $\nabla_{\boldsymbol{\epsilon}_t}\mathcal{L}_t|_{\phi=\phi^{\text{opt}}, \boldsymbol{\epsilon}_t^e=0} = \mathbf{0}$. Hence, if $\mathbf{P}_t^{e0}$ is optimal we must have that $\nabla_{\boldsymbol{\epsilon}_t}\mathcal{L}_t|_{\phi=\phi^0, \boldsymbol{\epsilon}_t^e=\boldsymbol{\epsilon}_t^{e0}} = \mathfrak{R}^{0\prime}\mathcal{W}\mathbb{E}_t\mathbf{Y}_t^0 = \mathbf{0}$, with $\mathbb{E}_t\mathbf{Y}_t^0 = h_y(\mathbf{X}_{-t}, \mathbb{E}_t\boldsymbol{\Xi}_t, \boldsymbol{\epsilon}_t^{e0}; \phi^0)$, which if $\boldsymbol{\delta}_t^{g*}$ exists implies that $\boldsymbol{\delta}_t^{g*}$ must satisfy $\boldsymbol{\delta}_t^{g*} \neq 0$ if $\mathbf{P}_t^{e0} \neq \mathbf{P}_t^{e\text{opt}}$. The existence of $\boldsymbol{\delta}_t^{g*}$ depends on whether the inverse of $\mathfrak{R}^{0\prime}\mathcal{W}\mathfrak{R}^0$ exists, which can be verified for any given model. □

The key purpose for stating Theorem S1 is twofold. First, the key underlying Assumption (S1) characterizes the class of models for which our sufficient statistics approach can be used for policy evaluation. Second, Theorem S1 is useful for characterizing optimal policy in complicated possibly nonlinear models. Indeed while the traditional approach to optimal policy — as documented in Section 2 for the baseline NK model — may become complex for nonlinear models, the gradient based approach underlying Theorem S1 may provide a workable necessary condition for optimal policy. That is determining $\boldsymbol{\delta}_t^{g*} \neq 0$ may be easier then solving the model and computing the optimal policy.

Next, we discuss the details for the specific examples — state dependence and multiple regimes — as mentioned in the main text.

B1: State dependence

Under state dependence with two states as in Auerbach and Gorodnichenko (2013) the generic model for the economy remains conditionally linear, i.e.

$$\begin{cases} \mathcal{A}_{yy}(z_t)\mathbb{E}_t\mathbf{Y}_t - \mathcal{A}_{yw}(z_t)\mathbb{E}_t\mathbf{W}_t - \mathcal{A}_{yp}(z_t)\mathbf{P}_t^e & = & \mathcal{B}_{yx}(z_t)\mathbf{X}_{-t} + \mathcal{B}_{y\xi}(z_t)\mathbb{E}_t\Xi_t \\ \mathcal{A}_{ww}(z_t)\mathbb{E}_t\mathbf{W}_t - \mathcal{A}_{wy}(z_t)\mathbb{E}_t\mathbf{Y}_t - \mathcal{A}_{wp}(z_t)\mathbf{P}_t^e & = & \mathcal{B}_{wx}(z_t)\mathbf{X}_{-t} + \mathcal{B}_{w\xi}(z_t)\mathbb{E}_t\Xi_t \end{cases} , \quad \text{(S5)}$$

where z_t is some predetermined time-t measurable variable and for $\mathcal{D} = \mathcal{A}, \mathcal{B}$

$$\mathcal{D}_{..}(z_t) = F(z_t)\mathcal{D}_{..(1)} + (1 - F(z_t))\mathcal{D}_{..(2)}$$

where $F(z_t)$ can be interpreted as a measure of probability of being in state 1 at time t based on some time t predetermined variable z_t.[2]

The optimal state dependent policy $\mathbf{P}_t^{e\text{opt}}(z_t)$ can be defined as the solution for $\mathbf{P}_t$ to

$$\min_{\mathbf{Y}_t, \mathbf{W}_t, \mathbf{P}_t} \mathcal{L}_t \qquad \text{s.t.} \qquad \text{(S5)} \; . \quad \text{(S6)}$$

The generic policy rule is given by

$$\mathcal{A}_{pp}(z_t)\mathbf{P}_t^e - \mathcal{A}_{py}(z_t)\mathbb{E}_t\mathbf{Y}_t - \mathcal{A}_{pw}(z_t)\mathbb{E}_t\mathbf{W}_t = \mathcal{B}_{p\xi}(z_t)\mathbb{E}_t\Xi_t + \mathcal{B}_{px}(z_t)\mathbf{X}_{-t} \; ,$$

where the definition for the maps $\mathcal{A}_{..}(z_t)$ and $\mathcal{B}_{..}(z_t)$ is the same a above. We collect all the coefficients of the rule in $\phi(z_t)$. The state dependent OPP under a given rule $\phi^0(z_t)$ is given by

$$\boldsymbol{\delta}_t^*(z_t) = -(\mathcal{R}(z_t)^{0\prime}\mathcal{W}\mathcal{R}(z_t)^0)^{-1}\mathcal{R}(z_t)^{0\prime}\mathcal{W}\mathbb{E}_t\mathbf{Y}_t^0 \; ,$$

where $\mathcal{R}^0(z_t) = F(z_t)\mathcal{R}^0_{(1)} + (1 - F(z_t))\mathcal{R}^0_{(2)}$ captures the causal state dependent effect of $\boldsymbol{\epsilon}_t$ on $\mathbf{Y}_t$. The following corollary summarizes the properties of the state dependent OPP.

Corollary 1. *Given model* (S5). *Under the assumptions that (1) the optimal policy* $\mathbf{P}_t^{e\text{opt}}(z_t)$ *is unique, and (2) the rule* $\phi^0(z_t)$ *underlying the proposed policy path* $\mathbf{P}_t^{e0}(z_t)$ *leads to a unique and determinate equilibrium, we have that*

1. $\boldsymbol{\delta}_t^*(z_t) = 0 \qquad \Longleftrightarrow \qquad \mathbf{P}_t^{e0}(z_t) = \mathbf{P}_t^{e\text{opt}}(z_t)$.
2. $\mathbf{P}_t^{e0} + \boldsymbol{\delta}_t^*(z_t) = \mathbf{P}_t^{e\text{opt}}(z_t)$.

Proof of Corollary 1. The proof is identical to the proof of Proposition 1 after changing the maps $\mathcal{A}_{..}$ and $\mathcal{B}_{..}$ to $\mathcal{A}_{..}(z_t)$ and $\mathcal{B}_{..}(z_t)$, respectively. □

[2] A popular functional form for $F(.)$ is $F(z_t) = \exp(-\gamma z_t)/[1 + \exp(-\gamma z_t)]$ with γ a tuning parameter.

B2: Multiple policy regimes

Next, we consider the case where there can be a feedback between the non-policy block and the policy block, in the form of multiple multiple regimes. An important example is in the context of monetary policy is whether long-run inflation expectations are "anchored" —fixed at some value— or "unanchored" —inflation expectations depend on the state of the economy— (e.g., Bernanke, 2007). The anchoring of inflation expectations matters for the dynamics of inflation, and likely depends on the central bank's policy rule or objective function.

To capture this idea, consider a generalization of (15) with

$$\left\{\begin{array}{rcl} \mathcal{A}_{yy}(\vartheta)\mathbb{E}_t\mathbf{Y}_t - \mathcal{A}_{yw}(\vartheta)\mathbb{E}_t\mathbf{W}_t - \mathcal{A}_{yp}(\vartheta)\mathbf{P}_t^e & = & \mathcal{B}_{yx}(\vartheta)\mathbf{X}_{-t} + \mathcal{B}_{y\xi}(\vartheta)\mathbb{E}_t\Xi_t \\ \mathcal{A}_{ww}(\vartheta)\mathbb{E}_t\mathbf{W}_t - \mathcal{A}_{wy}(\vartheta)\mathbb{E}_t\mathbf{Y}_t - \mathcal{A}_{wp}(\vartheta)\mathbf{P}_t^e & = & \mathcal{B}_{wx}(\vartheta)\mathbf{X}_{-t} + \mathcal{B}_{w\xi}(\vartheta)\mathbb{E}_t\Xi_t \end{array}\right. , \tag{S7}$$

where $\mathcal{A}_{..}(\vartheta)$ and $\mathcal{B}_{..}(\vartheta)$ capture describe the economy under the policy regime ϑ.

As in the baseline model, the generic policy rule is

$$\mathcal{A}_{pp}\mathbf{P}_t^e - \mathcal{A}_{py}\mathbb{E}_t\mathbf{Y}_t - \mathcal{A}_{pw}\mathbb{E}_t\mathbf{W}_t = \mathcal{B}_{p\xi}\mathbb{E}_t\Xi_t + \mathcal{B}_{px}\mathbf{X}_{-t}, \tag{S8}$$

and we collect all the rule parameters in $\phi = \{\mathcal{A}_{pp}, \mathcal{A}_{py}, \mathcal{A}_{pw}, \mathcal{B}_{px}, \mathcal{B}_{p\xi}\}$.

The policy regime ϑ can depend on the policy rule ϕ, creating a feedback from the policy rule to the coefficients of the non-policy block. Assume that there exists only a finite number of regimes. Here we consider the case with two regimes for clarity of exposition and following our example on anchored/unanchored inflation expectations. The maps $\mathcal{D}_{..}(\vartheta)$, for $\mathcal{D} = \mathcal{A}, \mathcal{B}$, are of the form

$$\mathcal{D}_{..}(\vartheta) = \left\{\begin{array}{lll} \mathcal{D}_{..}(\vartheta_1) & \text{if} & \phi \in \Phi_1 \\ \mathcal{D}_{..}(\vartheta_2) & \text{if} & \phi \in \Phi_2 \end{array}\right. , \tag{S9}$$

where $\Phi_1 \cup \Phi_2 = \Phi$ and $\Phi_i \cap \Phi_j = \emptyset$.

The characterization of the optimal policy choice requires more care. Here we define the optimal policy as the policy rule (and associated policy regime) that ensures the lowest loss. We consider

$$\min_{\mathbf{Y}_t, \mathbf{W}_t, \mathbf{P}_t, \phi} \mathcal{L}_t \qquad \text{s.t.} \qquad \text{(S7)-(S8)} , \tag{S10}$$

where we note that since the coefficients of the policy equation directly affect the coefficients that describe the economy, we need to take into account the policy equation when defining the optimal policy. We denote the solution for $\mathbf{P}_t$ to this minimization problem by $\mathbf{P}_t^{e\,\text{opt}}(\vartheta^{\text{opt}})$ where ϑ^{opt} corresponds to ϕ^{opt} which is the minimizing ϕ. We will assume that ϕ^{opt} lies in the interior of some Φ_i, which rules out boundary solutions.

Given a policy proposal $\mathbf{P}_t^{e0}$, implied by choices ϕ^0 and $\boldsymbol{\epsilon}^{e0}$, where ϕ^0 implies the regime

$\vartheta^0 \in \{\vartheta_1, \vartheta_2\}$, the regime specific OPP statistic is

$$\boldsymbol{\delta}_t^*(\vartheta^0) = -(\mathcal{R}^0(\vartheta^0)'\mathcal{W}\mathcal{R}^0(\vartheta^0))^{-1}\mathcal{R}^0(\vartheta^0)'\mathcal{W}\mathbb{E}_t\mathbf{Y}_t^0 \ ,$$

where $\mathbb{E}_t\mathbf{Y}_t^0$ is the expected allocation under $\mathbf{P}_t^{e0}$. Note that computing the regime specific OPP only requires $\mathcal{R}^0(\vartheta^0)$ the causal effects under the proposed policy policy and does not require knowledge of the causal effects in any of the other regimes.

The following corollary establishes the main property of the regime specific

Corollary 2. *Given model* (S7)-(S9)*, under the assumptions that (1) the optimal policy* $\mathbf{P}_t^{e\text{opt}}(\vartheta^{\text{opt}})$ *is unique and the underlying rule* ϕ^{opt} *leads to a unique and determinate equilibrium, and (2) the rule* ϕ^0 *underlying the proposed policy path* $\mathbf{P}_t^{e0}(z_t)$ *leads to a unique and determinate equilibrium, we have that*

$$\boldsymbol{\delta}_t^*(\vartheta^0) \neq 0 \qquad \Rightarrow \qquad \mathbf{P}_t^{e0}(\vartheta^0) \neq \mathbf{P}_t^{e\text{opt}}(\vartheta^{\text{opt}}) \ .$$

The corollary implies that if the regime specific OPP is non-zero we have that the policy choice is non optimal.

Proof of Corollary 2. The Lagrangian for the optimal policy problem (S10) is given by

$$\begin{aligned} \mathsf{L}_t =& \mathbb{E}_t \Big\{ \frac{1}{2}\mathbf{Y}_t'\mathcal{W}\mathbf{Y}_t + \boldsymbol{\mu}_1'(\mathcal{A}_{yy}(\vartheta)\mathbf{Y}_t - A_{yw}(\vartheta)\mathbf{W}_t - \mathcal{A}_{yp}(\vartheta)\mathbf{P}_t - \mathcal{B}_{yx}(\vartheta)\mathbf{X}_{-t} - \mathcal{B}_{y\xi}(\vartheta)\Xi_t) \\ & +\boldsymbol{\mu}_2'(\mathcal{A}_{ww}(\vartheta)\mathbf{W}_t - \mathcal{A}_{wy}(\vartheta)\mathbf{Y}_t - \mathcal{A}_{wp}(\vartheta)\mathbf{P}_t - \mathcal{B}_{wx}(\vartheta)\mathbf{X}_{-t} - \mathcal{B}_{w\xi}(\vartheta)\Xi_t) \\ & +\boldsymbol{\mu}_3'(\mathcal{A}_{pp}\mathbf{P}_t - \mathcal{A}_{py}\mathbf{Y}_t - \mathcal{A}_{pw}\mathbf{W}_t - \mathcal{B}_{px}\mathbf{X}_{-t} - \mathcal{B}_{p\xi}\Xi_t - \boldsymbol{\epsilon}_t)\Big\} \ , \end{aligned}$$

where $\mathcal{A}_{..}(\vartheta)$ and $\mathcal{B}_{..}(\vartheta)$ capture describe the economy under the policy regime ϑ determined by the policy rule $\phi = \{\mathcal{A}_{pp}, \mathcal{A}_{py}, \mathcal{A}_{pw}, \mathcal{B}_{px}, \mathcal{B}_{p\xi}\}$

The first order conditions for $\mathbf{Y}_t, \mathbf{W}_t, \mathbf{P}_t$ are given by

$$\begin{aligned} \mathbf{0} &= \mathcal{W}\mathbb{E}_t\mathbf{Y}_t + \mathcal{A}_{yy}'(\vartheta^{\text{opt}})\boldsymbol{\mu}_1 - \mathcal{A}_{wy}'(\vartheta^{\text{opt}})\boldsymbol{\mu}_2 - \mathcal{A}_{py}'\boldsymbol{\mu}_3 \\ \mathbf{0} &= -\mathcal{A}_{yw}'(\vartheta^{\text{opt}})\boldsymbol{\mu}_1 + \mathcal{A}_{ww}'(\vartheta^{\text{opt}})\boldsymbol{\mu}_2 - \mathcal{A}_{pw}'\boldsymbol{\mu}_3 \\ \mathbf{0} &= -\mathcal{A}_{yp}'(\vartheta^{\text{opt}})\boldsymbol{\mu}_1 - \mathcal{A}_{wp}'(\vartheta^{\text{opt}})\boldsymbol{\mu}_2 + \mathcal{A}_{pp}'\boldsymbol{\mu}_3 \end{aligned}$$

Importantly, with a finite number of regimes if ϕ^{opt} lies in the interior of some Φ_i, all the derivatives of the maps $\mathcal{A}_{..}(\vartheta)$ and $\mathcal{B}_{..}(\vartheta)$ with respect to the elements of ϕ are zero. Intuitively, an infinitely small change in a rule coefficient does not trigger a regime change. In that case, unless the economy is already perfectly stabilized,[3] the first order conditions

[3] Technically, we exclude that all shocks until time t and all initial conditions are zero, a trivial case we can discard.

with respect to the elements of ϕ imply that $\boldsymbol{\mu}_3 = 0$. To see that, note that optimization with respect to, for instance, the jth element of the first row of $\mathcal{B}_{p\xi}$ gives

$$\mu_{3,j}\mathbb{E}\xi_t = 0$$

where $\mu_{3,j}$ is the corresponding element of the $\boldsymbol{\mu}_3$ vector. Unless $\mathbb{E}_t\xi_t = 0$,[4] this implies $\mu_{3,j} = 0$. We can proceed similarly with the other coefficients of ϕ to show $\boldsymbol{\mu}_3 = 0$.

With $\boldsymbol{\mu}_3 = 0$, note that the optimization problem (S10) has the same first order conditions as the following fictitious problem: given some optimal policy rule ϕ^{opt} the policy maker can choose $\mathbf{Y}_t$, $\mathbf{W}_t$, $\mathbf{P}_t$ *and* a sequence of policy shocks $\boldsymbol{\epsilon}_t$ to minimize $\mathcal{L}_t$. Indeed, for that problem the Lagrangian writes

$$\begin{aligned}\mathsf{L}_t^f =&\mathbb{E}_t \Big\{ \frac{1}{2}\mathbf{Y}_t'\mathcal{W}\mathbf{Y}_t + \boldsymbol{\mu}_1'(\mathcal{A}_{yy}(\vartheta^{\text{opt}})\mathbf{Y}_t - A_{yw}(\vartheta^{\text{opt}})\mathbf{W}_t - \mathcal{A}_{yp}(\vartheta^{\text{opt}})\mathbf{P}_t - \mathcal{B}_{yx}(\vartheta^{\text{opt}})\mathbf{X}_{-t} - \mathcal{B}_{y\xi}(\vartheta^{\text{opt}})\boldsymbol{\Xi}_t) \\ &+\boldsymbol{\mu}_2'(\mathcal{A}_{ww}(\vartheta^{\text{opt}})\mathbf{W}_t - \mathcal{A}_{wy}(\vartheta^{\text{opt}})\mathbf{Y}_t - \mathcal{A}_{wp}(\vartheta^{\text{opt}})\mathbf{P}_t - \mathcal{B}_{wx}(\vartheta^{\text{opt}})\mathbf{X}_{-t} - \mathcal{B}_{w\xi}(\vartheta^{\text{opt}})\boldsymbol{\Xi}_t) \\ &+\boldsymbol{\mu}_3'(\mathcal{A}_{pp}^{\text{opt}}\mathbf{P}_t - \mathcal{A}_{py}^{\text{opt}}\mathbf{Y}_t - \mathcal{A}_{pw}^{\text{opt}}\mathbf{W}_t - \mathcal{B}_{px}^{\text{opt}}\mathbf{X}_{-t} - \mathcal{B}_{p\xi}^{\text{opt}}\boldsymbol{\Xi}_t - \boldsymbol{\epsilon}_t)\Big\} \ .\end{aligned}$$

The first-order conditions with respect to $\boldsymbol{\epsilon}_t$ yield $\boldsymbol{\mu}_3 = 0$, which immediately establishes that the first-order conditions of the fictitious problem are the same as those of problem (S10). As in the proof of Proposition 1, under the stated assumption that ϕ^{opt} implies a unique equilibrium, the fictitious problem can also be stated as

$$\min_{\boldsymbol{\epsilon}_t} \mathcal{L}_t \qquad \text{s.t.} \qquad \mathbb{E}_t\mathbf{Y}_t = \mathcal{R}^{\text{opt}}\boldsymbol{\epsilon}_t^e + \mathcal{C}_x^{\text{opt}}\mathbf{X}_{-t} + \mathcal{C}_\xi^{\text{opt}}\boldsymbol{\Xi}_t \ .$$

which leads to the first order condition

$$\mathcal{R}^{\text{opt}'}\mathcal{W}\mathbb{E}_t\mathbf{Y}_t = \mathbf{0} \ .$$

This establishes that $\mathcal{R}^{\text{opt}'}\mathcal{W}\mathbb{E}_t\mathbf{Y}_t = \mathbf{0}$ is a necessary condition for the original optimization problem (S10).

Next, consider the second fictitious problem: given the proposed policy rule ϕ^0 the policy maker can choose $\mathbf{Y}_t$, $\mathbf{W}_t$, $\mathbf{P}_t$ *and* a sequence of policy shocks $\boldsymbol{\epsilon}_t$ to minimize $\mathcal{L}_t$. Using the exactly the same steps as above, and noting that ϕ^0 leads to a unique equilibrium, it follows that $\mathcal{R}^{0'}\mathcal{W}\mathbb{E}_t\mathbf{Y}_t = \mathbf{0}$ is a necessary condition for optimality under this rule. This implies that if we find that $\mathcal{R}^{0'}\mathcal{W}\mathbb{E}_t\mathbf{Y}_t \neq \mathbf{0}$, then $\mathbf{P}_t^{e0} \neq \mathbf{P}_t^{e\text{opt}}$ as there can only be two cases. First, if ϕ^0 implies the same regime as ϕ^{opt} it follows immediately as $\mathcal{R}^{0'}\mathcal{W}\mathbb{E}_t\mathbf{Y}_t = \mathcal{R}^{\text{opt}'}\mathcal{W}\mathbb{E}_t\mathbf{Y}_t = \mathbf{0}$ is a necessary condition. Second, if we find that ϕ^0 does not imply the same regime as ϕ^{opt}

[4] Technically, the condition is much weaker: we only have to exclude that the economy is perfectly stabilized when $\mathbf{P}_t^e = 0$, i.e., that all shocks and all initial conditions are zero. A trivial case we can discard.

the result also follows immediately as ϕ^0 then does not minimize the loss function.

Finally, note that since $\boldsymbol{\delta}_t^*(\vartheta^0)$ is just a rescaling of $\mathcal{R}^{0'}\mathcal{W}\mathbb{E}_t\mathbf{Y}_t$ the same conclusion carries over, and $\boldsymbol{\delta}_t^*(\vartheta^0) = 0$ is a necessary condition for optimality. Importantly however, unlike in the linear case, that condition is not sufficient to characterize the optimal policy. With of a feedback from ϕ to the maps $\mathcal{A}_{..}(\vartheta)$ and $\mathcal{B}_{..}(\vartheta)$, a policy satisfying $\mathcal{R}^{0'}\mathcal{W}\mathbb{E}_t\mathbf{Y}_t = \mathbf{0}$ could be a *local* minimum, when the regime ϑ^0 associated with rule ϕ^0 is not the regime ϑ^{opt} associated with the optimal rule ϕ^{opt}.

□

S3 General convex loss functions

In the main text we restricted ourselves to quadratic loss functions when testing the optimality of a given policy choice. In this section we show that the main idea – exploiting the gradient of the loss function to evaluate optimality– continues to apply for essentially any convex loss function. The only difference is that the evaluation of the gradient will require the full forecast densities instead of only the mean oracle forecasts.

To show this, let $\mathcal{L}_t(\mathbf{Y}_t;\theta)$ denote a loss function which is convex and differentiable with respect to $\mathbf{Y}_t$ and may depend on preference parameters denoted by θ. The quadratic loss function (14) in the main text is a special case where θ includes preference parameters λ and discount factors β. Using the same generic model (15), we can summarize the policy maker's problem by

$$\min_{\mathbf{Y}_t,\mathbf{W}_t,\mathbf{P}_t} \mathbb{E}_t\mathcal{L}_t(\mathbf{Y}_t;\theta) \qquad \text{s.t} \qquad (15)$$

To evaluate whether a proposed policy choice $\mathbf{P}_t^{e0}$ solves this problem we can follow the same steps as in the main text. Specifically, using the proof of Proposition 1 it follows immediately that the equivalence $\mathbf{P}_t^{e0} = \mathbf{P}_t^{e\,\text{opt}} \iff \nabla_{\boldsymbol{\epsilon}_t}\mathbb{E}_t\mathcal{L}_t(\mathbf{Y}_t;\theta)|_{\mathbf{P}_t^{e0}} = 0$ continues to hold. However, now the gradient with respect to $\boldsymbol{\epsilon}_t$ evaluated at $\mathbf{P}_t^{e0}$ is given by

$$\nabla_{\boldsymbol{\epsilon}_t}\mathbb{E}_t\mathcal{L}_t(\mathbf{Y}_t;\theta)|_{\mathbf{P}_t^{e0}} = \mathcal{R}^{0'} \times \nabla_{\mathbf{Y}_t}\mathbb{E}_t\mathcal{L}_t(\mathbf{Y}_t;\theta)|_{\mathbf{P}_t^{e0}} \ .$$

Given that $\mathcal{L}_t(\mathbf{Y}_t;\theta)$ is convex with respect to $\mathbf{Y}_t$ and $\mathbf{Y}_t$ is an affine function of $\boldsymbol{\epsilon}_t$ in equilibrium we have that if $\nabla_{\boldsymbol{\epsilon}_t}\mathbb{E}_t\mathcal{L}_t(\mathbf{Y}_t;\theta)|_{\mathbf{P}_t^{e0}} \neq 0$ the policy choice $\mathbf{P}_t^{e0}$ is not optimal.

To evaluate the gradient we need to compute the derivative $\nabla_{\mathbf{Y}_t}\mathbb{E}_t\mathcal{L}_t(\mathbf{Y}_t;\theta)|_{\mathbf{P}_t^{e0}}$. Under a quadratic loss $\frac{1}{2}\mathbf{Y}_t'\mathcal{W}\mathbf{Y}_t$, this expression simplifies to $\nabla_{\mathbf{Y}_t}\mathbb{E}_t\mathcal{L}_t(\mathbf{Y}_t;\theta)|_{\mathbf{P}_t^{e0}} = \mathcal{W}\mathbb{E}_t\mathbf{Y}_t^0$ as in the main text, but for a general convex loss we have

$$\nabla_{\mathbf{Y}_t}\mathbb{E}_t\mathcal{L}_t(\mathbf{Y}_t;\theta)|_{\mathbf{P}_t^{e0}} = \int_{\mathbf{Y}_t^0} \nabla_{\mathbf{Y}_t}\mathcal{L}_t(\mathbf{Y}_t^0;\theta)p(\mathbf{Y}_t^0|\mathcal{F}_t)\mathrm{d}\mathbf{Y}_t^0 \ , \qquad \text{(S11)}$$

where $p(\mathbf{Y}_t^0|\mathcal{F}_t)$ is the forecast density under the proposed policy choice $\mathbf{P}_t^{e0}$. Thus, provided the forecast density is available, we can construct the OPP statistic and OPP-based tests as in the main text. The only difference is that there is no closed form expression for the gradient, and numerical or Monte Carlo integration methods will be necessary.

S4 Estimating robust preference parameters

Researchers outside of the policy maker's research staff may not have access to the policy maker's preferences $\mathcal{W}$. For this setting, we outline an approach for conducting *preference robust* OPP inference. The idea is to exploit a sequence of past policy decisions to find the preferences that gives the smallest deviations from optimality on average. This approach can thus be seen as considering a worse-case scenario for rejecting optimality.

Specifically, we write $\omega = \beta \otimes \lambda$, the elements of the preference matrix $\mathcal{W}$, as a function of the $d_\theta \times 1$, parameter vector θ, i.e., $\omega = \omega(\theta)$, with $d_\theta \leq K$, and we estimate θ by numerically solving [5]

$$\hat{\theta} = \arg\min_{\theta \in \Theta} \widehat{S}(\theta) \ , \qquad \widehat{S}(\theta) = \left\| \frac{1}{n} \sum_{s=t_0}^{t} \widehat{\boldsymbol{\delta}}_{a,s}(\theta) \right\|^2 \ , \tag{S12}$$

where $\widehat{\boldsymbol{\delta}}_{a,s}(\theta)$ corresponds to the mean OPP estimate (29) as a function of θ.

With the estimated $\hat{\theta}$ in hand a researcher can base the optimality assessment on the simulated distribution of the OPP given $\hat{\theta}$, which ensures that deviations from optimality are not due to potentially arbitrary choices for the preference parameters.

S5 Additional results for the empirical study

In this section we discuss additional results for our empirical study on testing US monetary policy decisions. These results are complementary to those presented in Section 6 of the main text. In particular, the different subsections discuss the following aspects.

1. Understanding the policy shocks
2. Sensitivity to the preference parameter λ
3. Alternative dynamic causal effect estimates: SVAR inference
4. Testing the stability of the macro environment

[5] To give a specific example, suppose that $M_y = 2$ and $M_p = 2$, then we can take $\theta = (\theta_1, \theta_2)'$ and set $\omega(\theta) = (\theta_1^0, \theta_1^1, \ldots)' \otimes (1, \theta_2)'$, which implies that $\beta = (\theta_1^0, \theta_1^1, \ldots)'$ and $\lambda = (1, \theta_2)'$.

S5.1 Impulse responses to the policy shocks

In our empirical study we use shocks to the policy rate and to the slope of the yield curve. Clearly, this is a subset of all possible policy shocks, and it is of interest to understand the policy experiments that we are considering.

To this extent, Figure S1 plots the full set of impulse responses to (i) innovations to the fed funds rate, and (ii) innovations to the slope of yield curve. Compared to the figures in the main text, Figure S1 also reports the path of the fed funds rate and the path of the slope of the yield curve. We can see that both policy experiments correspond to somewhat persistent changes to the policy instrument, similar to earlier estimates of impulse responses to monetary shocks (e.g., Barnichon and Matthes, 2018).

S5.2 The preference parameter λ

In the main text we used $\lambda = 1$ to compute the OPP-based tests. In this section we compute the OPPs for different choices of λ between $[0.2, 2]$. The results for different choices of λ are shown in Figure S2. We find that the short-rate OPP is not sensitive to the choice for λ. In fact, all of our main findings hold for all choices of λ and the differences are often small. For the slope OPP the findings are a bit different. Here low values of λ, say $\lambda = 0.2$, move the slope OPP towards zero and the findings during the financial crises are no longer significant. The reason is that the causal effects of inflation are estimated with greater uncertainty, such that putting more weight on the inflation mandate (a lower λ) increases the attenuation bias and lessens the power of our optimality test. That said, we stress that the data indicates that such low values for λ are unlikely, as the worst case λ that we computed using (S12) was determined at $\lambda = 0.6$.

S5.3 Alternative dynamic causal effect estimates

In the main text we used LP-IV type estimates for the dynamic causal effects. In this section we estimate the causal effects using the SVAR-IV methodology (e.g. Montiel Olea, Stock and Watson, 2020). The instrumental variables remain identical: high frequency surprises to the fed funds rate and to the slope of the yield curve.

To compute the SVAR-IV dynamic causal effects we use the same specification as for LP-IV and consider a SVAR with unemployment, inflation, the spread between the 10 year and short term interest rates, the short term interest rate and the excess bond yield measure of Gilchrist and Zakrajšek (2012). This implies that we use the same control variables as we used for the LP-IV specification.

The impulse response estimates are shown in Figure S3. We find that the patterns are very close. There are two differences: (i) for the fed funds rate shock the response of inflation is lower for the SVAR-IV method and (ii) for the slope shock the response of unemployment is lower for the SVAR-IV method.

Next, we use the SVAR-IV estimates to compute the OPP statistics. The results are shown in Figure S4. The differences between LP-based and SVAR-based OPPs for the fed funds rate are very small. If anything, the SVAR-based OPP for the fed funds rate is more often significantly different from zero, notably in the early stage of the Great Recession (April 2008) when the OPP is significantly negative even at the 95 percent confidence level. For the SVAR-based OPPs for the slope of the yield curve, the SVAR-based OPP is somewhat muted, still significant at the 68% level in the early stage of the Great Recession but not anymore at the 95% level. This is caused by the fact that SVAR-IV estimates a smaller impulse response of unemployment to a slope shock but a similar response of inflation. As a result, the inflation/unemployment trade-off is less attractive, i.e., there is less of a case for a more aggressive use of slope policy in order to bring down unemployment faster (in exchange for more inflation).

S5.4 Testing the stability of the dynamic causal effects

The retrospective study of past policy decisions is based on the assumption that the dynamic causal effects are stable for some time period (so that we can estimate $\mathcal{R}_a^0$). To verify whether this was the case for our empirical study we consider testing the stability of the dynamic causal effect estimates using the structural change tests for linear models with endogenous variables proposed in Hall, Han and Boldea (2012). We implement the tests using the wild fixed-regressor bootstrap as this allows for heteroskedasticity and an unstable reduced form, see Boldea, Cornea-Madeira and Hall (2019).

The specifications that we consider are given by

$$y_{t+h} = \mathcal{R}_{x,h}^0 x_t + \gamma^y w_t + u_t \qquad y = \pi, u \qquad x = i, s \ , \tag{S13}$$

where s denotes the slope of the yield curve and i the policy rate. The fed funds rate i_t is instrumented by the monetary policy surprises to the fed funds rate measured around the FOMC announcements within a 30 minute window (e.g. Kuttner, 2001). The slope of the yield curve is instrumented by the surprises to the ten-year Treasury yield (orthogonalized with respect to surprises to the current fed funds rate). Note that we conduct the break tests equation-by-equation to study whether there exists breaks at different horizons.

We are interested in testing whether $\mathcal{R}_{x,h}^0$ is stable over the sampling period. Specifically, we test the hypothesis $H_0 : m = 0$ against $H_1 : m = 1$, where m is the number of structural breaks in $\mathcal{R}_{x,h}^0$. The test statistic used is the sup-Wald test, which under homoskedasticity corresponds to the sup-F test of Andrews (1993). We follow Boldea, Cornea-Madeira and Hall (2019) and implement the test using the wild fixed-regressor bootstrap. We refer to their paper for the details.

The resulting bootstrap p-values are reported in Table S1. We find no evidence of parameter instability as we can never reject the null of constant $\mathcal{R}^0_{x,h}$ whether for shocks to the fed funds rate or to the slope of the yield curve. This holds for both inflation and unemployment and for all the horizons considered. We conclude that for both the short-rate OPP and the slope OPP we cannot reject the null hypothesis of parameter stability.

References

Andrews, Donald W. K. 1993. "Tests for Parameter Instability and Structural Change With Unknown Change Point." *Econometrica*, 61(4): 821–856.

Auerbach, Alan J., and Yuriy Gorodnichenko. 2013. "Output Spillovers from Fiscal Policy." *American Economic Review*, 103(3): 141–46.

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

Bernanke, Ben. 2007. "Inflation expectations and inflation forecasting." Board of Governors of the Federal Reserve System (US).

Boldea, Otilia, Adriana Cornea-Madeira, and Alastair R. Hall. 2019. "Bootstrapping structural change tests." *Journal of Econometrics*, 213(2): 359–397.

Dennis, John E., and Robert B. Schnabel. 1996. *Numerical Methods for Unconstrained Optimization and Nonlinear Equations.* Siam.

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

Gilchrist, Simon, and Egon Zakrajšek. 2012. "Credit Spreads and Business Cycle Fluctuations." *American Economic Review*, 102(4): 1692–1720.

Hall, Alastair R., Sanggohn Han, and Otilia Boldea. 2012. "Inference regarding multiple structural changes in linear models with endogenous regressors." *Journal of Econometrics*, 170(2): 281–302.

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.

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

Montiel Olea, José L., James H. Stock, and Mark W. Watson. 2020. "Inference in Structural Vector Autoregressions identified with an external instrument." *Journal of Econometrics.* forthcoming.

Figure S1: IMPULSE RESPONSES TO POLICY INNOVATIONS

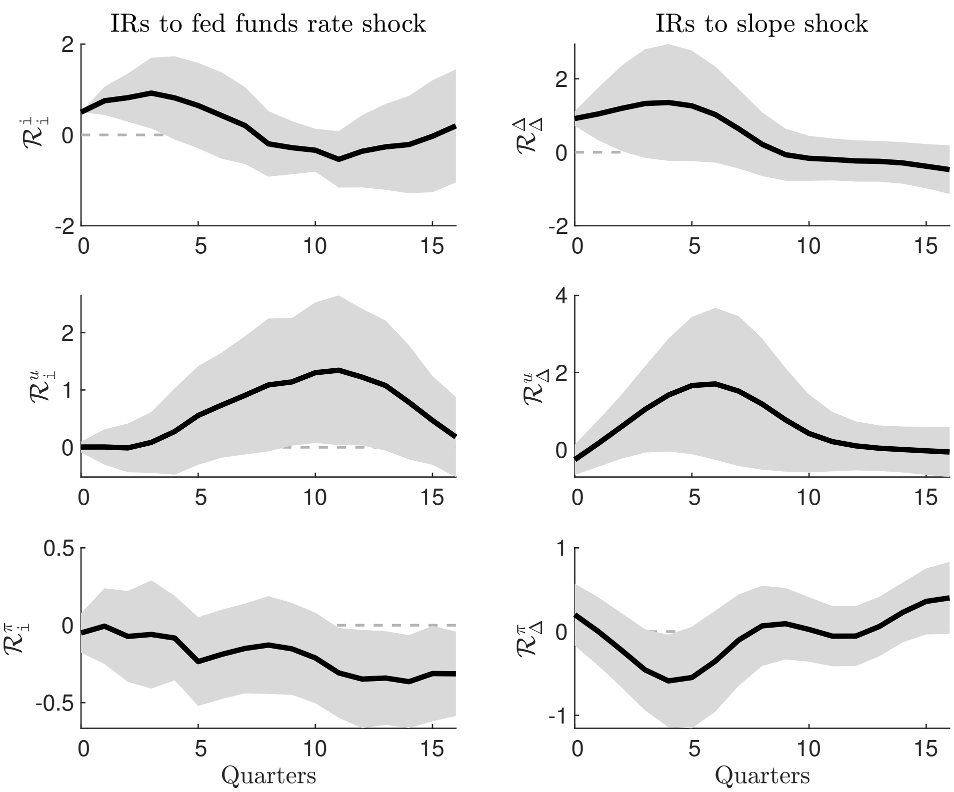

Notes: Left panel: impulse responses (IRs) of the fed funds rate, inflation and unemployment gaps to a fed funds rate shock. Right panel: impulse responses (IRs) of the slope of the yield curve, inflation and unemployment gaps to a slope policy shock. Shaded bands denote the 95 percent confidence intervals.

Figure S2: OPP for different λ, 1990-2019

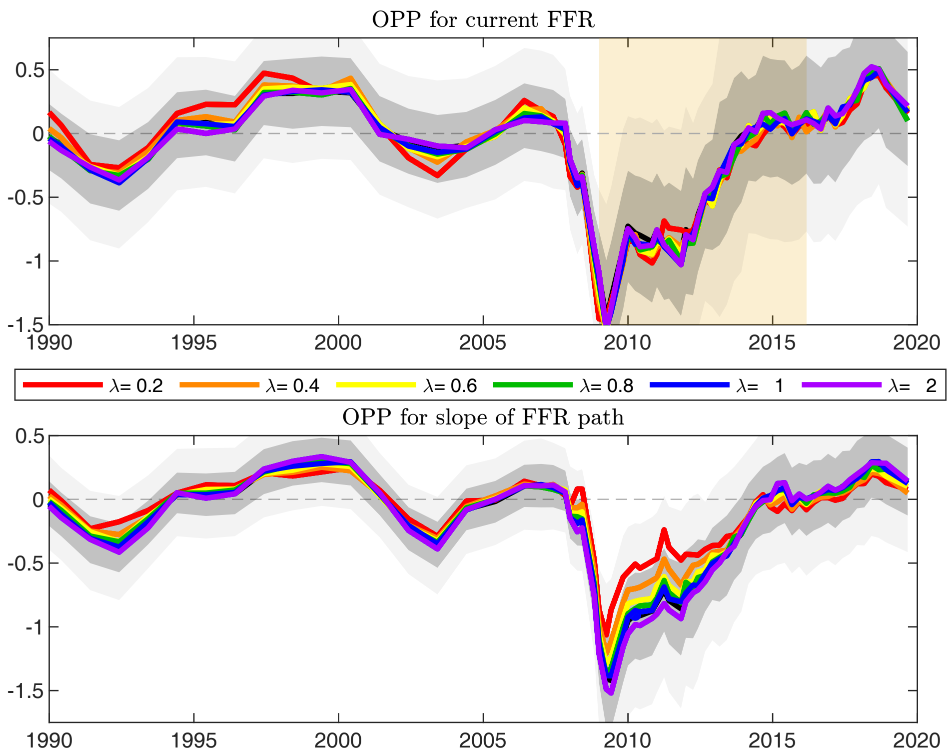

Notes: Top panel: short-rate OPP. The colored lines correspond to the OPP with λ's between 0.2 and 2. Bottom panel: slope OPP. The colored lines correspond to the OPP with λ's between 0.2 and 2. The grey areas capture impulse response and model uncertainty at respectively 68% (darker shade) and 95% (lighter shade) confidence when using $\lambda = 1$ as in the main text.

Figure S3: LP-IV vs SVAR-IV

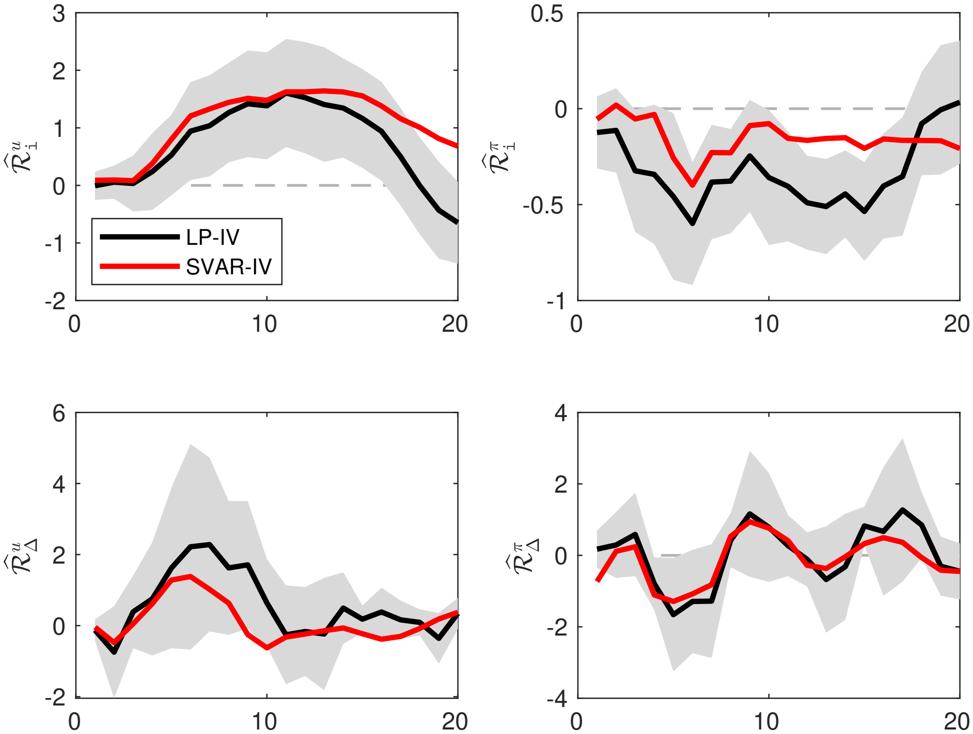

Notes: The panels compare the LP-IV and SVAR-IV estimates of the impulse responses of inflation (left column) or unemployment (right column) to an innovation to the fed funds rate (top row) or to the slope of the yield curve (bottom row).

Figure S4: OPPs for fed monetary policy, Using SVAR-IV to estimate $\mathcal{R}_a^0$

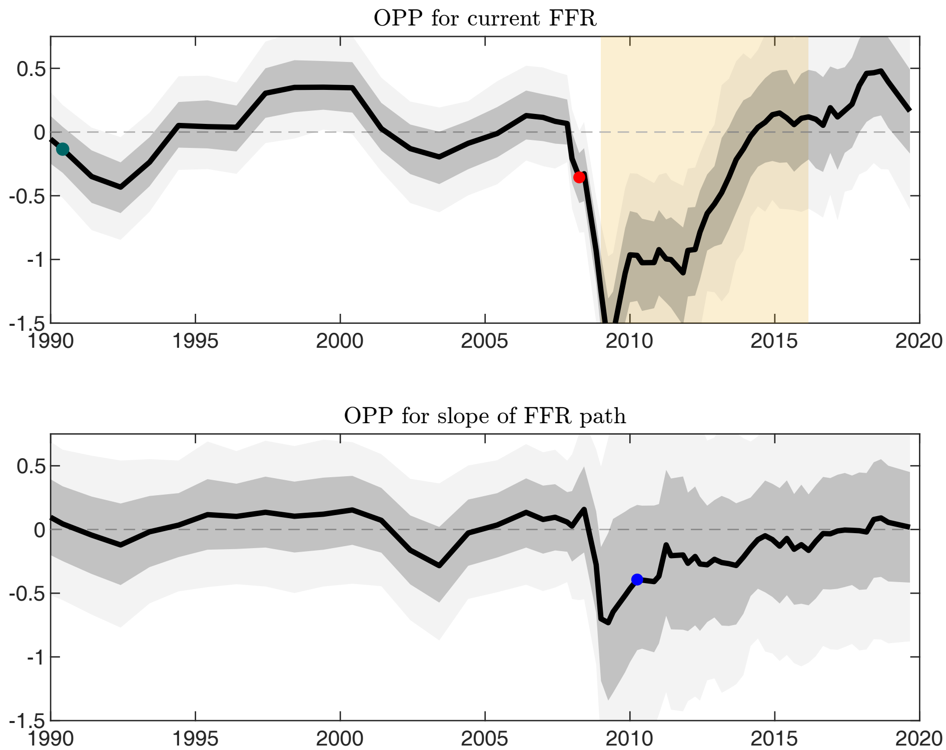

Notes: OPP sequences over 1990-2019. Top panel: fed funds rate OPP estimated using SVAR-IV estimates for $\mathcal{R}_a^0$. Bottom panel: slope OPP estimated using SVAR-IV estimates for $\mathcal{R}_a^0$.

Table S1: Structural Break Tests for $\mathcal{R}^0_{x,h}$

FFR: 1990-2007

h	π	u
0	0.49	0.57
5	0.43	0.10
10	0.58	0.27
15	0.44	0.29
20	0.60	0.62

Slope: 2008-2018

h	π	u
0	0.71	0.62
5	0.72	0.36
10	0.96	0.47
15	0.81	0.34
20	0.50	0.79

Notes: We report the fixed-regressor bootstrap p-values for testing the hypothesis $H_0 : m = 0$ vs $H_1 : m = 1$, where m is the number of breaks in the causal effect $\mathcal{R}^0_{x,h}$ of an innovation to x (the fed funds rate or the slope of the yield curve) on inflation (π) or unemployment (u) after $h = 0, 5, 10, 15, 20$ quarters. In the top panel, the fed funds rate is instrumented with the high frequency monetary policy surprises of Kuttner (2001). In the bottom panel the slope of the yield curve is instrumented with surprises to the ten-year Treasury yield (orthogonalized with respect to surprises to the current fed funds rate). The bootstrap was implemented following Boldea, Cornea-Madeira and Hall (2019).