NBER WORKING PAPER SERIES

EVOLVING REPUTATION FOR COMMITMENT:
THE RISE, FALL AND STABILIZATION OF US INFLATION

Robert G. King
Yang K. Lu

Working Paper 30763
http://www.nber.org/papers/w30763

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2022

The first footnote discloses my co-author's research support from a foundation related to Hong Kong University of Science and Technology. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Evolving Reputation for Commitment: The Rise, Fall and Stabilization of US Inflation
Robert G. King and Yang K. Lu
NBER Working Paper No. 30763
December 2022
JEL No. D82,E52

ABSTRACT

A parsimonious model of shifting policy regimes can simultaneously capture expected and actual US inflation during 1969-2005. Our model features a forward-looking New Keynesian Phillips curve and purposeful policymakers that can or cannot commit. Private sector learning about policymaker type leads to a reputation state variable. We use model inflation forecasting rules to extract state variables from SPF inflation forecasts. US inflation is tracked by optimal policy without commitment before 1981 and by optimal policy with commitment afterward. In theory and quantification, the interaction of private sector learning and optimal policy within regimes is central to expected and actual inflation.

Robert G. King
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215
and NBER
rking@bu.edu

Yang K. Lu
Department of Economics
Hong Kong University of Science and Technology
Clearwater Bay, Kowloon
Boston, MA 02215
United States
yanglu@ust.hk

1 Introduction

Since the 1970s, the distinction between macroeconomic policies with and without commitment has become familiar to macroeconomists and policymakers. In many theoretical models, a policymaker with commitment capability can achieve superior macroeconomic outcomes due to *strategic power* over private agents' expectations, permitting smoothing of real activity and inflation as shocks arise. By contrast, theory suggests that there would be high and volatile inflation if a policymaker can't commit.

How important is this distinction for the behavior of inflation in the United States? Many economists have suggested that inflation policy before 1980 is consistent with lack of commitment and some have suggested that a commitment regime prevailed afterward. But other economists have argued that it is hard to square these polar cases with various aspects of US experience.

We use novel theory and quantification to argue that evolving reputation for commitment is *central* to understanding US inflation history during 1969-2005. Our theory uses a now-standard macroeconomic framework with forward-looking inflation dynamics, purposeful policymaking, and stochastic changes in regime. We depart from that framework by assuming that private agents do not know whether the current policymaker can commit or not, but must learn from observed inflation history. This leads to a reputation measure which is the Bayesian likelihood that the policymaker is of a committed type. In our framework, purposeful policy under commitment manages reputation evolution as well as managing expectations. By contrast, a policymaker without commitment capacity simply responds to private sector expectations.

Even in this simple setting, theory and its quantification is challenging because there is an intertemporal strategic interaction in the dynamic game between two types of policymakers and private agents. We therefore propose a new methodology in which the committed policymaker solves a dynamic principal-agent problem. The resulting recursive equilibrium is quite operational because it has only three state variables including a reputation measure and a cost-push shock.[1] In our setup, reputation is a capital good for the committed policymaker but evolves as a martingale in the eyes of private agents.

[1] The latter shifts the output-inflation trade-off, for which we use common terminology.

Turning to the US historical experience, we focus on inflation starting in late 1968 since this is the start of the Survey of Professional Forecasters (SPF). We reason that short-term SPF forecasts should be more sensitive to temporary factors like cost-push shocks and its longer-term forecasts should better capture persistent factors like reputation. Since our recursive equilibrium spells out how inflation forecasts at various horizons depend on model state variables, we develop a new method of latent state extraction from multi-period SPF forecasts, yielding reputation and cost-push shock time series.

Our model's recursive equilibrium also specifies how optimal intended inflation for each type of policymaker depends on the state variables, so that our extracted states imply time series for intended inflation of each type.[2] Comparing these two policy measures to actual inflation (which is not used in their construction), we find striking support for our quantitative theory: policy without commitment tracks U.S. inflation closely before 1981 and policy with commitment tracks U.S. inflation closely from 1981 to 2005.

Although our model indicates a shift around 1981 from a regime without commitment to a regime with commitment, the optimal policies in each regime reflect private agents uncertainty about policymaker type. During the Great Inflation of the 1970s, though inflation bias arose because of lack of commitment, its extent was initially low and it rose over time as reputation gradually dissipated. The 1980s regime with commitment started with a poor initial reputation, resulting a relatively high optimal committed policy at the start of Volcker's Disinflation. It wasn't until the late-1990s that our committed policymaker built reputation high enough so that his optimal policy came down to a level resembling a standard full commitment solution.

Working under the assumption of a 1981 regime change, we provide our model-based interpretation of U.S. inflation history. We highlight that evolving reputation (i) sheds light on why inflation forecast errors have been serially correlated and (ii) is quantitatively important in helping our model-implied inflation polices to capture key features of U.S. inflation history, including the Great Inflation, the Volcker Disinflation, and the Great Moderation. Finally, we use our quantitative theory to consider recent inflation from the Fed's 2012 inflation target announcement to the present (using the 2022Q3 SPF).

[2] The policymaker sets intended inflation, as he has imperfect control over actual inflation.

Literature Our explanation of the rise, fall, and stabilization of US inflation uses three key ideas from the literature. First, we model inflation as the result of a purposeful policymaker interacting with a private sector with rational expectations. As in Kydland and Prescott (1977) and Barro and Gordon (1983a), inflation bias occurs when the purposeful policymaker cannot commit. However, in our setting where policymaker type is uncertain, inflation bias is shaped by what a committed type would do in similar circumstances and the evolving likelihood that each type is present.

Second, our reputational equilibrium adopts one of the two approaches in modern game theory,[3] but one less standard in macroeconomics: based on Bayesian learning in a noisy environment, our reputational state variable is the likelihood that the current policymaker has commitment capability. Further, the intended inflation of a committed policymaker depends on his reputation and on the behavior of the alternative type in a similar history, since these matter for the inflation expectations that he seeks to manage. The more familiar reputational approach, introduced by Barro and Gordon (1983b) to macroeconomics, demonstrates that reputational forces may substitute for commitment capability, leading a "discretionary" policymaker to behave like a committed one as in the important modern literature on sustainable plans.[4] [5] However, policymaker reputation does not vary over time in most of the substainable plan literature: it is either excellent or nonexistent.[6] Our learning-based framework permits *reputation building* by a policymaker that can commit and

[3] For a general discussion and specific examples see Mailath and Samuelson (2006). These leading theorists advocate for studying reputation as we do, writing "The idea that a player has an incentive to build, maintain, or milk his reputation is captured by the incentive that player has to manipulate the beliefs of other players about his type. The updating of these beliefs establishes links between past behavior and expectations of future behavior. We say 'reputations effects' arise if these links give rise to restrictions on equilibrium payoffs or behavior that do not arise in the underlying game of complete information."

[4] See Chari and Kehoe (1990).

[5] Within the NK framework, optimal policy under commitment involves time-varying inflation when there are Phillips curve shocks: Kurozumi (2008) and Loisel (2008) have shown that a policymaker without commitment capability can be led to follow such a policy so long as he is sufficiently patient and the shocks are not too large.

[6] One exception is Dovis and Kirpalani (2021) which allows for partial reputation – the probability that the central bank is committed. In their analysis, the reputation is either constant if the equilibrium is pooling, or drops to zero if the equilibrium is separating.

reputation dissipation by one that can't, as in Cukierman and Liviatan (1991).[7]

Third, it is well understood that optimal policymaking in a forward-looking New Keynesian setting involves dynamic stabilization absent from the 1980s models.[8] To explore reputation building in the forward-looking NK model, our earlier analysis (Lu et al. (2016)) posited – for tractability – that private agents believed that an alternative policymaker followed a mechanical high inflation rule. But having two optimizing policymakers is crucial for our positive theory of inflation with and without commitment, necessitating the novel theoretical approach that we develop here.

Our paper is related to a large literature studying the rise, fall and stabilization of US inflation, but our approach is quite different. Sargent (1999) stimulated a literature on the role of a purposeful policymaker's beliefs that does not require exogenous regime changes,[9] with Primiceri (2006) extending this approach and quantifying shifts in estimates of the Phillips curve slope and intercept. Bianchi (2013) and Debortoli and Lakdawala (2016) develop and estimate models in which private agents anticipate a possible exogenous policy regime change but do not face a learning problem. Our quantitative theory emphasizes the evolution of *private sector beliefs* and we use the SPF to extract the evolution of such beliefs. Other researchers investigate macroeconomic outcomes with private agent learning, but policymakers in these studies don't purposefully manage private sector learning.[10] Our model features interaction of private sector learning and optimal policies with and without commitment, which we see as essential to matching the pattern of actual inflation and its comovement with the SPF. Carvalho et al. (2022) and Hazell et al. (2022) attribute the Volcker disinflation and the inflation stabilization afterwards to a decline of long-term inflation expectations, highlighting that such expectations are anchored in the 1990s. Our theory rationalizes such long-term inflation expectations behavior as an equilibrium outcome.

[7]These authors used a 1980s-style Phillips curve, rather than an forward-looking specification. We elaborated on this approach in King et al. (2008), highlighting the role of expectations management but maintaining the essentially static output-inflation trade-off.

[8]Various presentations include Clarida et al. (1999), Woodford (2003), and Gali (2015).

[9]See the Riksbank review article by Sargent and Soderstrom (2000) for an introduction.

[10]Examples include Ball (1995), Erceg and Levin (2003), Orphanides and Williams (2005), Goodfriend and King (2005), Cogley et al. (2015), Matthes (2015), and Melosi (2016).

Our use of the SPF also links our research to the large and growing literature on survey measures of inflation (Coibion et al. (2018)). The SPF forecasts systematically underestimated inflation during its rise in the 1970s and then systematically overestimated it during its decline. Our explanation of persistent forecasting errors is consistent with the view that these SPF anomalies arise from agents not knowing the policy regime (Evans and Wachtel (1993), Coibion et al. (2018)) or the model generating the data (Farmer et al. (2021)). Our work differs from the existing literature by having unknown policy optimally evolving over time, rather than being generated by a random process or by exogenous policy rules.

Organization The balance of the paper is as follows. In section 2, we describe the economy. In section 3, we cast the macroeconomic equilibrium in game theoretic terms, defining a Bayesian perfect equilibrium. In section 4, we develop a recursive equilibrium and describe how to solve it. In section 5, we elaborate our new method of latent state extraction from the SPF and use it to construct quantitative measures of policies. Section 6 provides our model-based interpretation of U.S. inflation history and undertakes various exercises to shed light on our model's internal mechanisms. Section 7 concludes.

2 The Economy

A policymaker designs and announces a plan for current and future inflation. A private sector composed of atomistic forward-looking agents is uncertain whether the policymaker can commit or not. Their forward-looking decisions reflect the possibility that an announced policy plan may not be executed.

2.1 Private sector

Private agents' behavior is captured by a standard NK Phillips curve

$$\pi_t = \underbrace{\beta E_t \pi_{t+1}}_{e_t} + \kappa x_t + \varsigma_t, \tag{1}$$

where π_t is inflation, x_t is the output gap, and ς_t is a cost-push shock governed by an exogenous Markov chain with the transition probabilities $\varphi(\varsigma_{t+1}; \varsigma_t)$. Private agents' discount factor is β and $E_t\pi_{t+1}$ is their expectation about the next-period inflation, with e_t shorthand for discounted expected inflation.

2.2 Policymaker

The policymaker is responsible for the inflation rate, π, but cannot control it exactly.[11] There are two types of policymaker. A *committed* type (τ_a) chooses and announces an optimal state-contingent plan for intended inflation at all dates when he first takes office and executes it in all subsequent periods until replaced. [12] [13] The committed inflation plan therefore shapes private sector's expected inflation. An *opportunistic* type (τ_α) makes the same announcements,[14] but may deviate from the announced plan and chooses his own intended inflation on a period-by-period basis.

The private sector does not observe the policymaker's type or his intended inflation, denoted by a_t for the committed type or α_t for the opportunistic type. Yet, it observes an inflation rate π_t that deviates randomly from the policymaker's intention, with a density $g(\pi_t|a_t)$ or $g(\pi_t|\alpha_t)$. We assume that these densities imply zero mean implementation errors which are i.i.d. and independent of the intended inflation:[15]

$$\varepsilon_{at} = \pi_t - a_t \text{ and } \varepsilon_{\alpha t} = \pi_t - \alpha_t. \tag{2}$$

[11] We use "policymaker" rather than "central banker" to recognize that inflation policy may be the result of various actors. For example, DeLong (1996), Levin and Taylor (2013), and Meltzer (2014) stress various political influences on monetary policy outcomes, while other economists see direct connections of fiscal policy to inflation.

[12] We specify intended inflation rather than intended output for analytical convenience. If policy instead controlled intended real aggregate demand $\underline{x}_{\tau t}$ and $x_{\tau t} = \underline{x}_{\tau t} + \sigma_{x\tau}\varepsilon_t$, the Phillips curve $\pi_t = \kappa x_t + e_t + \varsigma_t$ implies that a choice of $\underline{x}_{\tau t} = \frac{1}{\kappa}[a_t - e_t - \varsigma_t]$ leads to identical intended inflation, although certain text expressions – particularly those for inflation expectations – are more cumbersome.

[13] As in some other related studies (see, e.g., Faust and Svensson (2001) and Sargent (1999)), we abstract from policy instruments.

[14] The opportunistic type makes the same announcements as the committed type to avoid revealing his type. This is consistent with a key conclusion made by Lu (2013) in a related fiscal model: the unique signalling equilibrium involves the truth-telling committed type announcing a policy that solves his optimal policy problem and the opportunistic type sending the same message. We therefore abstract from the analysis of signalling equilibria.

[15] We interpret random inflation error as a reduced-form representation for all unforeseeable factors that affect the inflation rate beyond the monetary policy, following Cukierman and Meltzer (1986), Faust and Svensson (2001), Atkeson and Kehoe (2006), etc. There is also ample evidence that realized inflation rates miss the intended inflation target, with examples including Roger and Stone (2005) and Mishkin and Schmidt-Hebbel (2007).

The policymaker has the following momentary objective

$$u(\pi, x, \tau) = -\frac{1}{2}[(\pi - \pi^*)^2 + \vartheta_x (x - x^*)^2] \tag{3}$$

which depends on inflation π and output gap x. There is a long-run inflation target π^* and a strictly positive output target x^*.[16]

The committed type discount factor is β_a; the opportunistic type is myopic.

2.3 Timing of events

Private agents start period t with a probability that the incumbent policymaker is the committed type, which we denote by ρ_t and call *reputation*. The within-period timing is shown in Figure 1. First, with probability q, the current policymaker is replaced via a publicly observed event, in which case the regime clock t is set to zero and the new policymaker's initial reputation ρ_0 is a random draw from the distribution $\Xi(\rho_0|\rho_t)$ with support [0,1].[17] Second, the exogenous cost-push shock ς_t is realized. Third, there is a policy announcement. If there is a new policymaker, he announces a new inflation plan. Otherwise, either type of continuing policymaker simply reiterates that current economic conditions call for an intended inflation a_t. Fourth, private agents form their expectations about the next-period inflation, e_t. Fifth, the policymaker implements intended inflation, a_t or α_t, depending on his type. Sixth, this action leads to a random inflation rate π_t with a density $g(\pi_t|a_t)$ or $g(\pi_t|\alpha_t)$, and an output gap x_t determined by the Phillips curve. New information leads private agents to update their beliefs about policymaker type.

[16] The non-zero inflation target is common in central bank objectives. The output component in the objective can be written as $-\frac{\vartheta_x}{2}[x^2 + (x^*)^2] + (\vartheta_x x^*)x$ highlighting that there is a benefit to an additional unit of output. It is this composite coefficient $(\vartheta_x x^*)$ rather than its components that are important below. Our approach can easily handle publicly observable shocks to the targets π^* and x^*. But since these are not essential to our analysis and have been extensively explored elsewhere, we opt for simplicity in specification.

[17] We allow $\Xi(\rho_0)$ to depend on ρ_t so that the new policymaker can partially inherit his predecessor's reputation.

Figure 1: Timing of events within a period

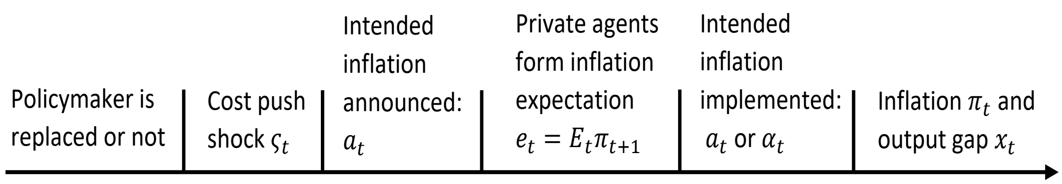

3 Macro Equilibrium in a Dynamic Game

Our economy consists of a private sector and a policymaker that can be one of the two types, but whose actions do not directly reveal his type: a dynamic game with incomplete information. We now describe equilibrium in this game.

3.1 Public Equilibria

Define the public history of the current regime $h_t = \{h_{t-1}, \pi_{t-1}, \varsigma_t\}$ as the collection of all past realizations of inflation rates and exogenous states, with $h_0 = \{\rho_0, \varsigma_0\}$ being the public history of a new regime. We restrict our attention to equilibria in which all strategies depend only on the public history, i.e., "public strategies."[18] We denote the committed and opportunistic policymaker's equilibrium strategies as $\{a(h_t)\}_{t=0}^{\infty}$ and $\{\alpha(h_t)\}_{t=0}^{\infty}$, respectively. Comparably, we can write inflation expectations as $\{e(h_t)\}_{t=0}^{\infty}$.

3.2 Perfect Bayesian Equilibria

We further require the equilibrium of this incomplete information game to be perfect Bayesian. That is, the beliefs of the private sector are consistent and the strategies of the two types of policymakers satisfy sequential rationality.

[18] Such a restriction is innocuous in our equilibrium analysis because: 1) the private sector's strategy has to be public since h_t is its information set; 2) the committed type's policy has to be public since it follows the announced policy plan, which needs to be verifiable by the private sector; 3) given all the other player's strategies are public, it is also optimal for the opportunistic type to choose public strategies (Mailath and Samuelson (2006))

3.2.1 Consistent beliefs: reputation

Consistency of beliefs requires the private sector's assessment of policymaker type is updated according to Bayes' rule (4) which depends on policymakers' equilibrium strategies and the observed inflation π_t. That is, within a regime, the private sector's belief ρ is updated recursively,

$$\text{(4)} \qquad \rho\left(h_{t+1}\right)=\rho\left(h_t, \pi_t\right) \equiv \frac{\rho\left(h_t\right) g(\pi_t|a\left(h_t\right))}{\rho\left(h_t\right) g(\pi_t|a\left(h_t\right))+(1-\rho\left(h_t\right)) g(\pi_t|\alpha\left(h_t\right))}$$

If there is a regime change in period t, the regime clock t is reset to zero and the new policymaker's reputation is $\rho_0 \sim \Xi(\rho_0|\rho(h_t))$, given the inheritance mechanism for reputation discussed above.

3.2.2 Consistent beliefs: inflation expectations

Inflation expectations must be consistent with private sector beliefs about policymaker type and equilibrium strategies. If the policymaker is replaced in period t, the private sector's consistent nowcast of inflation is:

$$\text{(5)} \qquad z(h_t)=\int\left[\rho_0 a(\rho_0, \varsigma_t)+(1-\rho_0) \alpha(\rho_0, \varsigma_t)\right] d\Xi(\rho_0|\rho(h_t)).$$

Expectations of future inflation also reflect unknown policymaker type:

$$\text{(6)} \quad e(h_t)=\beta E(\pi_{t+1}|h_t)=\beta \rho(h_t) E \pi_{t+1}|(h_t, \tau_a)+\beta(1-\rho(h_t)) E \pi_{t+1}|(h_t, \tau_\alpha)$$

$$E \pi_{t+1}|(h_t, \tau_a)=\int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1} ; \varsigma_t)\left[(1-q)\, a(h_{t+1})+q z\left(h_{t+1}\right)\right] g(\pi_t|a(h_t)) d \pi_t$$

$$E \pi_{t+1}|(h_t, \tau_\alpha)=\int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1} ; \varsigma_t)\left[(1-q)\, \alpha(h_{t+1})+q z\left(h_{t+1}\right)\right] g(\pi_t|\alpha(h_t)) d \pi_t$$

Specifically, when private agents form date t inflation expectations, they know that (i) there is a committed type with $\rho_t=\rho(h_t)$,[19] and (ii) the committed type's intentions lead to stochastic inflation, with density $g(\pi_t|a(h_t))$, con-

[19] With a slight abuse of notation, in the start of a new regime, $\rho(h_0)=\rho_0$.

tributing to history $h_{t+1} = \{h_t, \pi_t, \varsigma_{t+1}\}$. Hence, if the regime continues next period, the committed type's intended inflation will be $a(h_{t+1})$. In the event of a regime change next period, the consistent belief is the history-dependent future nowcast $z(h_{t+1})$. Similarly, with probability $1 - \rho_t$, the current policymaker is opportunistic and will generate stochastic inflation π_t with density $g(\pi_t|\alpha(h_t))$ and will implement $\alpha(h_{t+1})$ next period if the regime continues. In the event of a regime change next period, the expected inflation is $z(h_{t+1})$.

3.2.3 Sequential rationality of the committed type

At the beginning of his term, $t = 0$, the committed policymaker selects and announces a state-contingent plan for current and future intended inflation $\{a_t\}_{t=0}^{\infty}$ and then subsequently executes it.

The strategy of the committed type is *sequentially rational* if it maximizes his expected present discounted payoff at the beginning of his term,[20]

$$U_0 = \sum_{t=0}^{\infty} (\beta_a(1-q))^t \sum_{h_t} p(h_t) \underline{u}(a_t, e(h_t), \varsigma_t), \tag{7}$$

where $\underline{u}\,(a, e, \varsigma) \equiv \int u\,(\pi, x(\pi, e, \varsigma))\, g\,(\pi|a)\, d\pi$ is the expected momentary objective when the NK Phillips curve (1) is used to replace x with $x(\pi, e, \varsigma) = (\pi - e - \varsigma)\,/\kappa$. Note that (7) employs the probability of a specific history $h_t = [\varsigma_t, \pi_{t-1}, h_{t-1}]$ when inflation is generated by the committed type, i.e.,[21]

$$p(h_t) = \varphi(\varsigma_t; \varsigma_{t-1}) g(\pi_{t-1}|a(h_{t-1})) p(h_{t-1}) \tag{8}$$

combining the likelihood of the shock ς, the likelihood of inflation π given the committed type's decision, and the probability of the previous history.

In selecting the state-contingent plan at $t = 0$, the committed type takes into account the strategic power of his plan in shaping private sector inflation expectations. We consider this crucial element further below.

[20] We assume the committed policymaker maximizes payoffs within his own term, so his discounting includes both the time discount factor β_a and the replacement probability q.

[21] There is a slight abuse of notation here by using summation Σ over history to capture the joint effects of continuous distribution of π and discrete Markov chain distribution of ς.

3.2.4 Sequential rationality of the opportunistic type

An opportunistic policymaker chooses intended inflation α each period to maximize the expected objective, taking the nature of expected inflation's response to history $\{e(h_t)\}_{t=0}^{\infty}$ as given:

$$\alpha(h_t) = \operatorname*{argmax}_{\alpha} \underline{u}(\alpha, e(h_t), \varsigma_t) \tag{9}$$

where $\underline{u}\,(\alpha, e, \varsigma) \equiv \int u\,(\pi, x(\pi, e, \varsigma))\, g\,(\pi|\alpha)\, d\pi$ with $x(\pi, e, \varsigma) = (\pi - e - \varsigma)/\kappa$. The quadratic objective implies a linear best response of α to e and ς.

$$\alpha_t = Ae_t + B(\varsigma_t) \tag{10}$$

with $A = \frac{\vartheta_x}{\kappa^2+\vartheta_x}$, $B = (1-A)\pi^* + A\kappa x^*$ and $B(\varsigma) = B + A\varsigma_t$.

3.3 Public Perfect Bayesian Equilibrium

We can now define this dynamic game's Public Perfect Bayesian Equilibrium.

Definition 1. A Public Perfect Bayesian Equilibrium is a set of functions $\{z(h_t), e(h_t), \rho(h_t), \alpha(h_t), a(h_t)\}_{t=0}^{\infty}$ such that in each history:

(i) given $\alpha(h_t)$, $a(h_t)$, and $\rho(h_t)$, the private sector's nowcast of inflation $z(h_t)$ conditional on a replacement satisfies (5);

(ii) given $\alpha(h_t)$, $a(h_t)$, and $z(h_t)$, the private sector's belief of policymaker type $\rho(h_{t+1})$ is updated according to (4); and its expected inflation function $e(h_t)$ satisfies (6);

(iii) given the expected inflation function, $e(h_t)$, the action of the opportunistic type policymaker $\alpha(h_t)$ maximizes his expected payoff (9);

and, at the start of a regime (t=0),

(iv) the strategy for the committed type policymaker $\{a(h_t)\}_{t=0}^{\infty}$ maximizes his expected payoff (7), taking into account the strategic power of $\{a(h_t)\}_{t=0}^{\infty}$ on $\{e(h_t)\}_{t=0}^{\infty}$.

By "strategic power" of $\{a(h_t)\}_{t=0}^{\infty}$ on $\{e(h_t)\}_{t=0}^{\infty}$, we refer to the influence that the committed policymaker's state-contingent plan – his *strategy* – has for the response of e_t to the history h_t. In particular, given consistent private sector inflation expectations (6), there are three channels of influence.

First, $e(h_t)$ is partially anchored by future committed policy $a(h_{t+1})$. Second, the extent of this anchoring depends on $\rho(h_t)$ which itself is affected by past committed policy $a(h_{t-1})$. Third, both $e(h_t)$ and $\rho(h_t)$ depend on intended inflation of a possible opportunistic policymaker $\alpha(h_{t+1})$ and $\alpha(h_{t-1})$. Sequential rationality of the opportunistic policymaker makes $\{\alpha(h_t)\}_{t=0}^{\infty}$ a best response to $\{e(h_t)\}_{t=0}^{\infty}$. Therefore, via shaping $\{e(h_t)\}_{t=0}^{\infty}$, the committed state-contingent plan also indirectly determines $\{\alpha(h_t)\}_{t=0}^{\infty}$.[22]

4 Constructing the Equilibrium

Construction of the Public Perfect Bayesian equilibrium is usefully viewed as inner and outer loops of a program. The inner loop builds a within-regime equilibrium $\{e(h_t), \rho(h_t), \alpha(h_t), a(h_t)\}$ taking as given beliefs $z(h_t)$ about the consequences of a regime change. The outer loop adjusts the beliefs $z(h_t)$ to be consistent with future regime outcomes, i.e., to attain a fixed point between $z(h_t)$ and $\{a(h_t), \alpha(h_t), \rho(h_t)\}$.

4.1 Our novel principal-agent approach

Solving the within-regime equilibrium may appear to be a formidable task, due to the strategic power of the committed policy plan $\{a(h_t)\}_{t=0}^{\infty}$ over private sector expectations and opportunistic policies. The optimal choice for a committed policymaker depends on what the opportunistic type would do in the same history since private sector inflation expectations average across both types' future policy choices. However, as the opportunistic type responds to inflation expectations, the committed type's optimization cannot take future opportunistic policy as given since it varies with the committed policy plan.

[22] We assume a myopic opportunistic type because it is the most parsimonious model with an optimizing non-committed policymaker. Our framework and recursive method can be extended to a long-lived opportunistic type whose optimal $\alpha(h_t)$ satisfies the first order condition: $0 = \underline{u}_\alpha(\alpha, e_t, \varsigma_t) + \beta_\alpha(1-q)\int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) V(\varsigma_{t+1}, \pi_t, h_t)\, g_\alpha(\pi_t|\alpha)\, d\pi_t$ and whose value function obeys

$$V(h_t) = \underline{u}(\alpha_t, e_t, \varsigma_t) + \beta_\alpha(1-q)\int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) V(h_{t+1})\, g(\pi_t|\alpha_t)\, d\pi_t$$

when α_t is evaluated at optimal $\alpha(h_t)$. Appendix A.12 further describes the approach.

To tackle these complications, we recast the construction of the within-regime equilibrium as the solution to a principal-agent problem. As principal, the committed policymaker maximizes (7) by choosing state contingent plans for his actions and those of two agents, the private sector and the opportunistic policymaker. Incentive compatibility (IC) constraints of two forms are relevant: (i) private sector consistent beliefs (4) and rational expectations (6); and (ii) opportunistic type optimal response to expected inflation (10).

4.2 Recursive formulation

Our framework is unusual because private agents disagree with the principal – the committed policymaker – in beliefs about the probability of a specific history. The private sector *thinks* that current inflation could be generated by the opportunistic policymaker, as captured in the third line of the expression for expected inflation (6) above. By contrast, the committed policymaker *knows* that current inflation is generated by his policy choices, as reflected in $p(h_t)$ in the intertemporal objective (7). Such disagreement in probability beliefs between principal and agent poses a challenge to casting the Lagrangian component associated with the rational expectation constraint (6) into recursive form.[23] We solve this challenge by a "change of measure". Attaching a multiplier $\gamma(h_t)$ and the committed type's probability of history $p(h_t)$ as weights to the constraint (6), we form the Lagrangian component as:

$$\Psi_0 = \sum_{t=0}^{\infty}(\beta_a(1-q))^t \sum_{h_t} p(h_t)\gamma(h_t)[e_t - e(h_t)], \tag{11}$$

where $e(h_t)$ is given by (6). Then, in (6), we write $E\pi_{t+1}|(h_t, \tau_\alpha)$ in terms of the committed type's probabilities, replacing $g(\pi_t|\alpha(h_t))$ with $\lambda(\pi_t, a_t, \alpha_t)g(\pi_t|a(h_t))$ where $\lambda(\pi_t, a_t, \alpha_t) \equiv g(\pi_t|\alpha_t)/g(\pi_t|a_t)$ is the likelihood ratio. This permits us to express Ψ recursively, so that the dynamic Lagrangian $U_t + \Psi_t$ is also recursive. Defining W_t as the optimized dynamic Lagrangian, we then establish:[24]

[23] Following Kydland and Prescott (1980), Chang (1998), Phelan and Stacchetti (2001) and Marcet and Marimon (2019).

[24] Appendix A provides a detailed derivation of the recursive program.

Proposition 1. The within-regime equilibrium is the solution to a recursive optimization problem, given $z(\varsigma, \rho)$ and the IC constraint $\alpha = Ae + B(\varsigma)$

$$(12) \quad W(\varsigma, \rho, \mu) = \min_{\gamma} \max_{a,\alpha,e} \{\underline{u}(a, e, \varsigma) + (\gamma e + \mu\omega) + \beta_a (1-q) \int \sum_{\varsigma'} \varphi(\varsigma'; \varsigma) W(\varsigma', \rho', \mu') g(\pi|a) d\pi\},$$

with $\omega \equiv -\left\{(1-q) a + qz(\varsigma, \rho) + \frac{1-\rho}{\rho} [(1-q)\alpha + qz(\varsigma, \rho)]\right\}$

$$(13) \quad \mu' = \frac{\beta}{\beta_a (1-q)} \gamma\rho, \text{ with } \mu_0 = 0$$

$$(14) \quad \rho' = \frac{\rho g(\pi|a)}{\rho g(\pi|a) + (1-\rho) g(\pi|\alpha)}, \text{ with prob } g(\pi|a) \text{ given } \rho_0$$

This program enables us to analyze optimal choices of a committed policymaker facing private sector skepticism about his type. The component $(\gamma e + \mu\omega)$ arises from the Lagrangian component of the forward-looking rational expectations constraints (11) expressed in the recursive form.[25] The pseudo state variable μ records past promises (contained in ω) made by the committed type.[26] Next period's pseudo state μ' evolves according (13), keeping track of current promises.

With two possible policymaker types and stochastic replacement, the composite promise term $\omega \equiv -\{(1-q) a + qz(\varsigma, \rho) + \frac{1-\rho}{\rho} [(1-q)\alpha + qz(\varsigma, \rho)]\}$ contains more than the committed type's promised a, because private agents expected inflation also depends on their perceived inflation α intended by the opportunistic type and their nowcast of inflation z in a new regime.[27] The weights attached to a, α, and z reflect the exogenous replacement probability q, the endogenous reputation state ρ, and the divergent probability beliefs

[25] Our rational expectations constraint (6) is equivalent to the Phillips curve. Viewing it as an inequality constraint, with $x_t \leq (\pi_t - \beta E_t \pi_{t+1} - \varsigma_t)/\kappa$, the Phillips curve defines a set of feasible output gaps and inflation rates. Thus, the associated multiplier γ is nonnegative.

[26] The pseudo state variable terminology originates with Kydland and Prescott (1980). A new policymaker isn't held accountable for predecessor promises, so μ is initially zero.

[27] Note $\omega = a$ when $q = 0$, $\beta_a = \beta$, and $\rho = 1$. This is a textbook NK policy problem in recursive form. Appendix A.11 provides a fuller discussion.

about inflation π held by the committed policymaker and the private sector.[28]

It is useful to highlight several implications of Proposition (1). First, it pins down the equilibrium state vector as $s = [\varsigma, \rho, \mu]$, so that optimal decision rules are $a^*(s)$, $\alpha^*(s)$ and $e^*(s)$. Second, taking the first order condition with respect to γ and using an envelope theorem result for W_μ, we recover the rational inflation expectations constraint (6). That is, the optimization imposes the sequence of rational inflation expectations constraints, leading to the following lemma that relates the value function $W(s)$ to the committed policymaker's optimized intertemporal objective $U^*(s)$:[29]

Lemma 1. Let $U^*(s)$ and $\omega^*(s)$ be the intertemporal objective (7) and the composite promise term in (12) evaluated at optimal decision rules, then

$$W(\varsigma, \rho, \mu) = U^*(\varsigma, \rho, \mu) + \mu\omega^*(\varsigma, \rho, \mu) \tag{15}$$

The value function of the committed policymaker is therefore his optimized intertemporal objective net the cost of delivering on his past promises, captured by the term $\mu\omega^*$.

4.3 The PBE fixed point requirement

In a PBE, the nowcast of inflation $z^*(\varsigma, \rho)$ in a new regime must satisfy

$$z^*(\varsigma, \rho) = \int [\rho_0 a^*(\varsigma, \rho_0, 0; z^*(\varsigma, \rho)) + (1-\rho_0)\alpha^*(\varsigma, \rho_0, 0; z^*(\varsigma, \rho))] d\Xi(\rho_0|\rho) \tag{16}$$

with $a^*(.)$ and $\alpha^*(.)$ obtained from the recursive program (12) given $z^*(\varsigma, \rho)$, and $\mu_0 = 0$ as a new policymaker is not held accountable for prior commitments made by his predecessor.[30]

4.4 Key policy trade-offs and computation

The recursive program in Proposition 1 is valuable, as it sheds light on the relevant state variables. But it contains many choice variables, making it

[28] This final feature leads to $(1-\rho)/\rho$ in ω. Appendix A.9 explains how we eliminate the likelihood ratio λ using Bayes' rule.

[29] Appendix B.1 provides the proof.

[30] Schaumburg and Tambalotti (2007) impose a similar fixed point requirement in constructing an equilibrium in which a committed policymaker is randomly replaced.

inefficient for computation and hard to isolate the key trade-offs facing the policymaker. The following Lemma provides both computational and conceptual benefits, by developing implications of the forward-looking rational expectation constraint (6) for an *operational* expectations function.[31]

LEMMA 2. Given the states (ς, ρ) and that future policymakers follow the equilibrium strategies: $a^* (\varsigma', \rho', \mu')$, $\alpha^* (\varsigma', \rho', \mu')$ and $z^* (\varsigma', \rho')$, rationally expected inflation $e(\delta, \mu'; \varsigma, \rho)$ is uniquely determined by the contemporaneous policy difference $\delta = a - \alpha$, and the future pseudo-state variable μ'.

The operational expectations function $e(\delta, \mu'; \varsigma, \rho)$ derives from the committed policymaker's ability to influence expected inflation through two channels. The first is the standard *expectations management* channel, encoded in μ', although we have seen above that this management is complicated by private sector skepticism.[32] The second is a more novel *reputation building* channel: the committed policymaker affects the future reputation variable ρ' by choosing a difference between his intended inflation (a) and the intended inflation of an opportunistic type (α). A larger policy difference $\delta = a - \alpha$ raises the speed of private sector learning about current policymaker type.[33]

Using Lemma 1 and 2, we simplify the recursive program (12), moving from choosing (γ, a, α, e) to merely choosing (δ, μ') to obtain a simplified program.[34]

PROPOSITION 2. Given $z^*(\varsigma, \rho)$ and $U^* (\varsigma, \rho, \mu)$, a simplified program is

$$W(\varsigma, \rho, \mu) = \max_{\delta, \mu'} \left[\underline{\underline{u}}(\delta, \mu'; \varsigma, \rho) + \mu \underline{\underline{\omega}}(\delta, \mu'; \varsigma, \rho) + \beta_a (1 - q) \Omega(\delta, \mu'; \varsigma, \rho) \right]$$

with $\Omega(\delta, \mu'; \varsigma, \rho) = \int \sum_{\varsigma'} \varphi(\varsigma'; \varsigma) U^* (\varsigma', \rho'(\varepsilon_a, \delta, \rho), \mu') \phi_a(\varepsilon_a) d\varepsilon_a$.

[31] For additional details, see Appendix B.2.

[32] The committed policymaker adjusts the future pseudo-state variable μ' by choosing the shadow price of current promises γ. According to (13), a higher γ raises μ' and makes it more expensive for the committed type to have a higher a' next period.

[33] This channel is formalized when we simplify (14) to $\rho' = \rho'(\varepsilon_a, \delta, \rho)$ by replacing $g(\pi|a) = \phi_a(\varepsilon_a)$ and $g(\pi|\alpha) = \phi_\alpha(\pi - a + a - \alpha) = \phi_\alpha(\varepsilon_a + \delta)$, where $\phi_a(\cdot)$ and $\phi_\alpha(\cdot)$ are the densities of ε_a and ε_α respectively.

[34] Specifically, we replace e with $e(\delta, \mu'; \varsigma, \rho)$, α with $Ae + B(\varsigma)$, and a with $\alpha + \delta$ in $\underline{u}(\cdot)$ and $\omega(\cdot)$ of (12) to obtain $\underline{\underline{u}}(\cdot)$ and $\underline{\underline{\omega}}(\cdot)$: Appendix B.2 provides a detailed derivation. Moreover, choosing μ' is a choice of γ in view of (13).

Lemma 2 and Proposition 2 facilitate our computation. With a guessed function $z(\varsigma, \rho)$ specified in the outer loop, we can (i) use $a(\varsigma, \rho, \mu), \alpha(\varsigma, \rho, \mu)$ and $U(\varsigma, \rho, \eta)$ functions to obtain $e(\delta, \mu'; \varsigma, \rho)$ and $\Omega(\delta, \mu'; \varsigma, \rho)$; (ii) optimize over (δ, μ'); (iii) construct new a and α functions from optimal e and δ; and (iv) construct a new U function. Within the inner loop, we iterate until the policy functions converge.[35] We then calculate a new $z(\varsigma, \rho)$ and repeat the process until the outer loop has reached a fixed point in z.

5 Building the quantitative model

We use our framework to construct quantitative measures of intended inflation policy without commitment, α_t, and policy with commitment, a_t, without taking a stand on the type of policymaker that is in place during 1965-2005. These intended inflation measures correspond to the *beliefs* of the private sector about policy during the period and are based on a novel method of latent state extraction from the Survey of Professional Forecasters. We calibrate the various model parameters.

5.1 Inflation bias without commitment

Recall that the opportunistic policymaker chooses α, period by period, as a best response to private sector's expected inflation e. His decision rule is (10): $\alpha(e) = Ae + B(\varsigma)$ where $A = \frac{\vartheta_x}{\kappa^2 + \vartheta_x}$ is between 0 and 1. Denoting $\iota \equiv A(\kappa x^* - (1-\beta)\pi^*)$ and setting $\varsigma = 0$,[36] we can rewrite the best response as $\alpha_t - \pi^* = \iota + A\beta(E_t\pi_{t+1} - \pi^*)$ and plot it in Figure 2.

[35] Bayesian learning makes this not a linear-quadratic problem. In view of Proposition 2, we use direct maximization as part of a projection method to obtain a global solution. Overall, we employ a variant of the "dynamic programming with a rational expectations constraint" as sometimes advocated for calculating optimal policy under commitment.

[36] As is conventional, these inflation bias measures are derived without any shock ς.

Figure 2: Optimal Response of Opportunistic Policy to Inflation Expectations

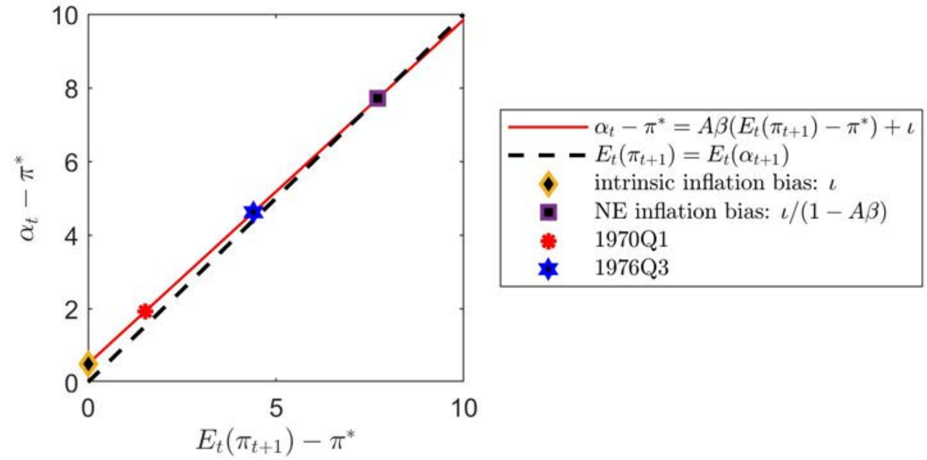

In absence of cost push shocks, the intended inflation α_t of the opportunistic policymaker includes a long-run inflation target π^*, an intrinsic inflation bias ι and a response to the private sector's expected inflation $E_t(\pi_{t+1})$.

The intersection of the best response function with the 45 degree line (the square marker) is the well-known *Nash equilibrium (NE) inflation bias* in which policy without commitment is fully expected (i.e., when $e = \beta\alpha$, $\alpha(e) - \pi^* = \iota/(1 - A\beta)$). But when we introduce uncertainty about policymaker's type, the expected inflation in our model can be lower than $\beta\alpha$. An extreme case is when private sector expects the inflation to be at target, i.e., $e = \beta\pi^*$, the optimal policy without commitment then differs from target only by ι; we define this as *intrinsic inflation bias* and mark it with a diamond in Figure 2.

A key feature of our model is that, with private sector learning, expected inflation varies over time between the two aforementioned extremes. That is, the extent of inflation bias $\alpha - \pi^*$ is time-varying along the red best response line between intrinsic bias ι and NE bias $\iota/(1 - A\beta)$.

Quantitatively, NE bias can be much greater than intrinsic inflation bias when $A\beta$ is close to 1, as in the figure, where $\iota = 0.5\%$ and $\iota/(1-A\beta)$ is around 8%. We also plot two points that correspond to the optimal inflation bias when the expected inflation is set to the SPF one quarter ahead inflation forecast in 1970Q1 and 1976Q3, respectively. Note that as the expected inflation is 3%

higher in 1976Q3 than in 1970Q1, the corresponding optimal policy without commitment is also substantially higher in 1976.

5.2 Additional Parameters

The long-run inflation target π^* is 1.5%, which lies in the 1 to 2 percent range frequently cited by central bankers advocating price stability.[37] Other parameters reported in Table 1 are selected to match some empirical facts and to highlight some model mechanisms.

The private sector and committed type share a conventional quarterly discount factor based on a 2% annual real rate. The replacement probability of $q = .03$ implies an average regime duration of 8 years. The inheritance mechanism for reputation is as follows: the new policymaker's initial reputation ρ_0 is the same as his predecessor's end-of-regime reputation with probability .9, while his initial reputation ρ_0 is otherwise random with a mean of 10%.[38]

The slope of the Phillips curve and the policymaker's concerns about real activity are central elements in any study of inflation policy. In our setup, the PC slope κ relates the output gap x to the quarterly inflation π, holding expected inflation fixed. $\kappa = .08$ implies that an output gap of 3% leads to annualized inflation of -1%, a value compatible with diverse empirical evidence.[39]

Turning to the preference parameters, we set the weight on output ϑ_x to 0.1, which is in the middle of the range used by prominent Fed researchers.[40] This parameter value, together with $\kappa = .08$, implies $A = .94$ according to $A = \vartheta_x/(\vartheta_x + \kappa^2)$. The target output gap x^* is chosen to generate a relatively small intrinsic inflation bias $\iota = .5\%$ so that the range of optimal policy without commitment in our model covers the range of U.S. inflation in the 1970s. Recall that x^* is linked to other parameters via $\iota = A(\kappa x^* - (1-\beta)\pi^*)$. Hence, these

[37] This value matches the estimate of Shapiro and Wilson (2019) in a careful and informative study of FOMC transcripts.

[38] We use a beta distribution with parameters 3 and 27.

[39] U.S. data from the 1950s and 1960s suggests that a 1% decrease in unemployment led to about 0.54% - 0.65% increase in inflation. An estimate for Okun's coefficient is about 1.67 using U.S. data prior to 2008, implying a 1% increase in unemployment led to a 1.67% decrease in output. In a structural NKPC, the parameter is also consistent with an adjustment hazard leading to four quarters of stickiness on average and an elasticity of marginal cost with respect to output of unity.

[40] Brayton et al. (2014) after translating time units and using Okun's law.

Table 1: Parameters

β, β_a	Discount factor (private, committed type)	0.995
q	Replacement probability	0.03
κ	PC output slope	0.08
π^*	Inflation target	1.5%
ϑ_x	Output weight	0.1
x^*	Output target	1.73%
ν	Persistence of cost-push shock	0.7
σ_ξ	Std of cost-push innovation	0.7%
σ_ε	Std of implementation error ε_a and ε_α	1.2%

One period is a quarter. Inflation target π^*, std of cost-push innovation σ_ξ, and std of implementation error σ_ε are all annualized rates.

other parameter choices imply $x^* = 1.73\%$.[41]

Beginning in the 1970s, many studies of inflation use an observable "Food and Energy price shock"(FE shock hereafter).[42] We initially used this proxy for ς, but eventually settled on extracting shocks from the SPF because these real time forecasters better capture various events including the 1974 inflation peak.[43] Never the less, we use the FE shock's serial correlation and its standard deviation as the cost-push shock's persistence ν and innovation volatility σ_ξ. We also combined the FE shock and the SPF one-quarter-ahead inflation forecast in an initial approximation to the opportunistic intended inflation α, to obtain the standard deviation of $(\pi - \alpha)$ that prevailed during 1964Q4-1979Q2 and use it as our calibrated standard deviation of implementation errors.

5.3 Extracting reputation and cost-push shocks

We use a novel empirical strategy to extract evolving reputation and cost-push shocks, requiring that our model's expectations match time series from the Survey of Professional Forecasters. More specifically, recall that Proposition

[41]With an Okun's law coefficient of 1.67, $x^* = 1.73\%$ is targeting unemployment about 1% below the natural rate.

[42]See R.J. Gordon (2013) and Watson (2014). It is constructed as the difference between the growth rate of the overall personal consumption deflator and its counterpart excluding food and energy. We display its time series in Appendix C.5.

[43]For additional discussion, see Appendix C.5

1 isolates three state variables $s_t = [\varsigma_t, \rho_t, \mu_t]$, including the highly persistent reputation state ρ, a more temporary cost-push shock ς, and a predetermined pseudo state μ. These states are known to private agents but not to us. At each date, our strategy is to extract the two stochastic unobserved states (reputation and cost-push shock) from the term structure of SPF forecasts.

Figure 3 highlights the smoother nature of the three-quarter-ahead SPF forecast of inflation (SPF3Q) relative to the one-quarter-ahead forecast (SPF1Q).[44] Taking a cue from literature on the term structure of interest rates, we form an SPF spread, plotted as the black dashed line and defined as the difference between the one and three quarter forecasts.[45] The interest rate term structure analogy suggests that this spread should be positive when there are persistent, but ultimately temporary increases in inflation.[46] In other words, longer-term forecasts (SPF3Q) depend more on the persistent reputation variable ρ_t, while shorter-term forecasts are more sensitive to transitory cost-push shocks ς_t.[47]

Our model implies private agents multi-period inflation forecasts are functions of the state variables, which we exploit to extract states by exactly matching the SPF $f_{t+k|t}$ at one quarter and three quarter horizons (k=1,3):[48]

$$f_{t+k|t} = E_t(\pi_{t+k}) = f(\varsigma_t, \rho_t, \mu_t, k), \text{for } k = 1, 3. \tag{17}$$

More specifically, matching SPF1Q and SPF3Q each period allows us to solve for $\widehat{\varsigma}_t$ and $\widehat{\rho}_t$ given the predetermined pseudo state $\widehat{\mu}_t$. With the extracted state $\widehat{\rho}_t$ and the predetermined state $\widehat{\mu}_t$, we determine $\widehat{\mu}_{t+1}$ using the equilibrium decision rule, $\mu'^*(0, \widehat{\rho}_t, \widehat{\mu}_t)$, continuing recursively to calculate a full history of states.[49] The initial value of $\widehat{\mu}$ is set to zero at regime switch dates.

[44] Elmar Mertens guided us to the SPF term structure via Mertens and Nason (2020).

[45] Under the expectations theory, the comparable spread would be the one quarter ahead forward rate less the three quarter ahead forward rate.

[46] Conceptually, this description is consistent with a stationary autoregressive component. Empirically, the SPF spread rises in Figure 3 during the 1974-75 inflation surge.

[47] We do not use SPF4Q due to missing observations, particularly important in 1975.

[48] Appendix C provides recursive forecasting formulae and state extraction details.

[49] We use $\mu'^*(0, \widehat{\rho}_t, \widehat{\mu}_t)$ instead of $\mu'^*(\widehat{\varsigma}_t, \widehat{\rho}_t, \widehat{\mu}_t)$ so that the extracted $\widehat{\varsigma}_t$ will be mean-reverting. Results from using $\mu'^*(\widehat{\varsigma}_t, \widehat{\rho}_t, \widehat{\mu}_t)$ are reported in Appendix C.4.

Figure 3: SPF and extracted states

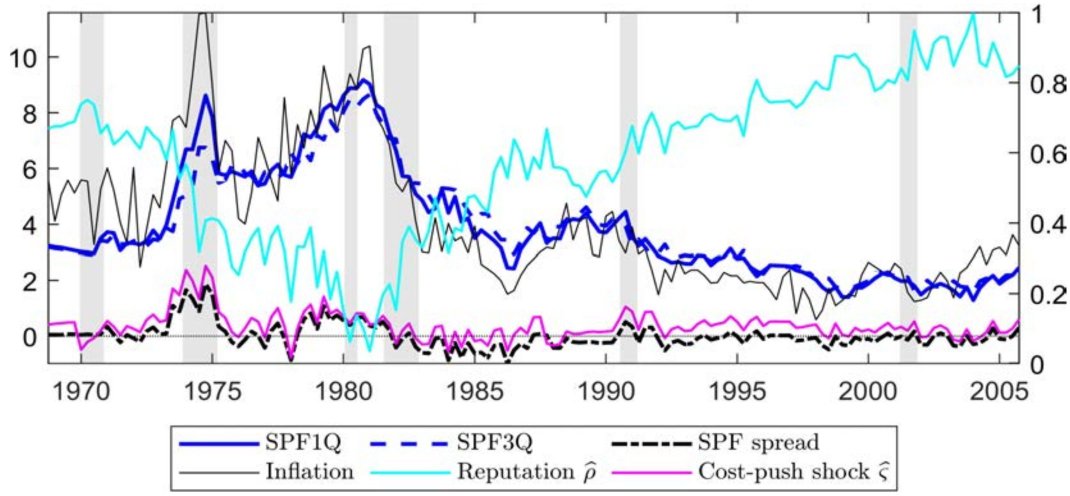

The SPF spread SPF1Q-SPF3Q is the difference between the one and three quarter forecasts. All variables are continuously compounded annualized rates of change. Appendix C provides details on our SPF constructions.

Regimes: Choice of possible regime switch dates is subtle. Many monetary histories highlight the Fed chair's identity and nature, as in Friedman and Schwartz's celebrated Great Contraction chapter. But other histories stress combined efforts of presidential administrations and the central bank (including Meltzer (2014) and Levin and Taylor (2013)). Our benchmark is to specify a new regime with each chairman: 1970Q1 (Burns), 1978Q1 (Miller), 1979Q4 (Volcker's October 1979 announcement of new operating procedures), and 1987Q4 (Greenspan).

Extracted states: Figure 3 also plots the extracted cost-push shock $\widehat{\varsigma}_t$ and the extracted reputation state $\widehat{\rho}_t$ (cyan and measured on the right hand axis). Note first that the extracted cost-push shock $\widehat{\varsigma}$ covaries positively with the SPF spread (SPF1Q-SPF3Q),[50] consistent with state extraction exploiting greater sensitivity of near-term forecasts to transitory shocks. Note also the big swing of the extracted reputation state $\widehat{\rho}$.

[50] Appendix C.5 shows that our extracted cost-push shock covaries well with the "Food and Energy price shock" constructed as the difference between the growth rate of personal consumption deflator and its counterpart excluding food and energy.

Matched and fitted inflation expectations: Our extraction method produces a nearly perfect match for SPF1Q and SPF3Q. Using the extracted states, we can also compute model-implied inflation forecasts at horizons 2 and 4. Appendix Figure 12 shows that these additional forecasts lie almost entirely on top of SPF2Q and SPF4Q, which are not explicitly targeted, providing support for our state extraction approach.

5.4 Model-implied inflation policies and actual inflation

With an extracted state history, we can construct the *intended inflation policy measures* $\widehat{a}_t = a(\widehat{s}_t)$ and $\widehat{\alpha}_t = \alpha(\widehat{s}_t)$. Then, given observed inflation π_t, (2) implies empirical implementation errors, $\widehat{\varepsilon}_{at} = \pi_t - a(\widehat{s}_t)$ and $\widehat{\varepsilon}_{\alpha t} = \pi_t - \alpha(\widehat{s}_t)$. Figure 4 plots these model-implied policies ($\widehat{a}_t$ in green, $\widehat{\alpha}_t$ in red), and their associated implementation errors ($\widehat{\varepsilon}_{at}$ and $\widehat{\varepsilon}_{\alpha t}$ as dash-dotted lines in matching color), with regime switch dates marked by solid vertical lines. Actual inflation π is the black dashed line, but recall that it isn't used to construct $\widehat{a}_t$ or $\widehat{\alpha}_t$.

Seeing the entire history, we have an advantage relative to private agents: they only know events through date t. We use this advantage to break the history at the start of the Reagan administration in 1981Q1 (marked by a dashed vertical line), based on actual inflation and our policy measures. We classify the earlier interval as an opportunistic regime because (i) $\widehat{\varepsilon}_\alpha = \pi - \widehat{\alpha}$ fluctuates around 0, suggesting that actual inflation consistent with opportunistic policy, and (ii) $\widehat{\varepsilon}_a = \pi - \widehat{a}$ is generally positive, suggesting actual inflation inconsistent with committed policy. The situation is clearly reversed in the latter interval, which we classify as a commitment regime. Actual inflation more closely resembles committed policy, with $\widehat{\varepsilon}_a = \pi - \widehat{a}$ fluctuating around zero, whereas actual inflation lies below opportunistic policy, with $\widehat{\varepsilon}_\alpha = \pi - \widehat{\alpha}$ being generally negative.[51]

[51] The mean of $\widehat{\varepsilon}_\alpha = \pi - \widehat{\alpha}$ is 0.16% in the earlier interval, and the mean of $\widehat{\varepsilon}_a = \pi - \widehat{a}$ is only 0.091% in the later interval. By contrast, the mean of $\widehat{\varepsilon}_a$ in the earlier interval and $\widehat{\varepsilon}_\alpha$ in the later interval are 1.6% and -0.97%, respectively. But historians differ in their dating of the regime change on other grounds: see Goodfriend and R.G. King (2005) and Orphanides (2005) for differing views.

Figure 4: Inflation history and model-implied policies

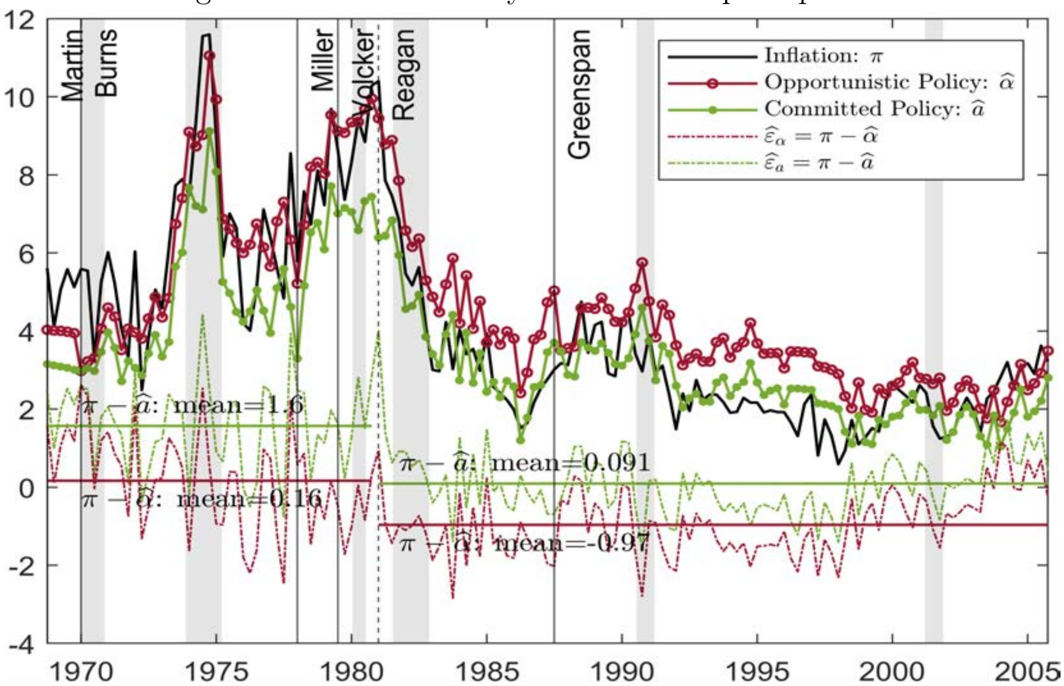

Model-implied policies constructed using calibrated decision rules and extracted states.

6 Evolving reputation: history and prospect

We have seen that our theory provides a potential explanation of the behavior of US inflation over 1968 to 2005, with key ingredients being regime shifts, inability of some policymakers to commit, and private sector learning. We now argue that evolving reputation is essential for this part of US history, highlighting the comovement of inflation and expectations. Drawing on some lessons from history and the features of our theory, we also discuss the current macroeconomic situation and prospects for inflation going forward.

6.1 Interpreting US inflation history 1968-2005

We now examine US inflation assuming an opportunistic policymaker early on and a committed policymaker later. At each date, we display our measures of *actual* policy, which is the current policymaker's intended inflation, and of *shadow* policy, which is private agents' rational belief about an alternative policymaker's behavior if confronted with the same observable history.

Figure 5: Model-based interpretation of US inflation

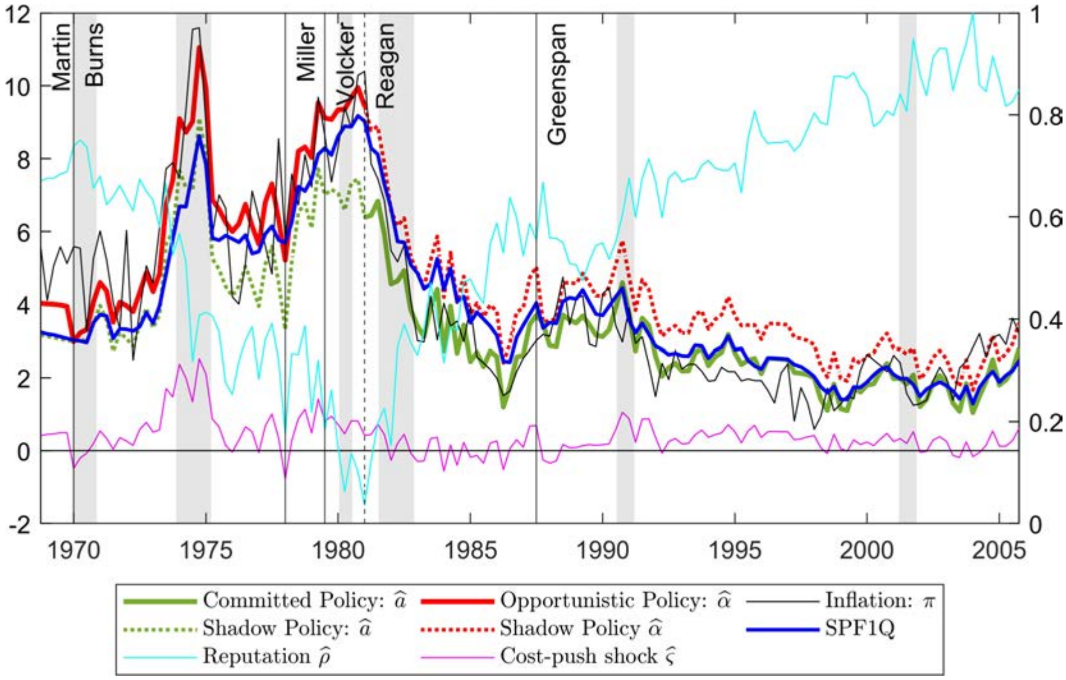

Figure 5 shows our model-based interpretation of US inflation history 1968-2005. Three time series – inflation π (black), model-implied committed policy $\widehat{a}$ (green), and model-implied opportunistic policy $\widehat{\alpha}$ (red) are repeated from Figure 4. But before 1981Q1, an opportunistic policymaker is taken to be generating the observed inflation, so the red line is solid and the green line for shadow policy is dotted. After 1981Q1, the red line is dotted and the green line is solid as a committed policymaker is generating the observed inflation. The blue solid line is the SPF1Q forecast, which is exactly matched by our model expected inflation.

Figure 6: Inflation forecast errors

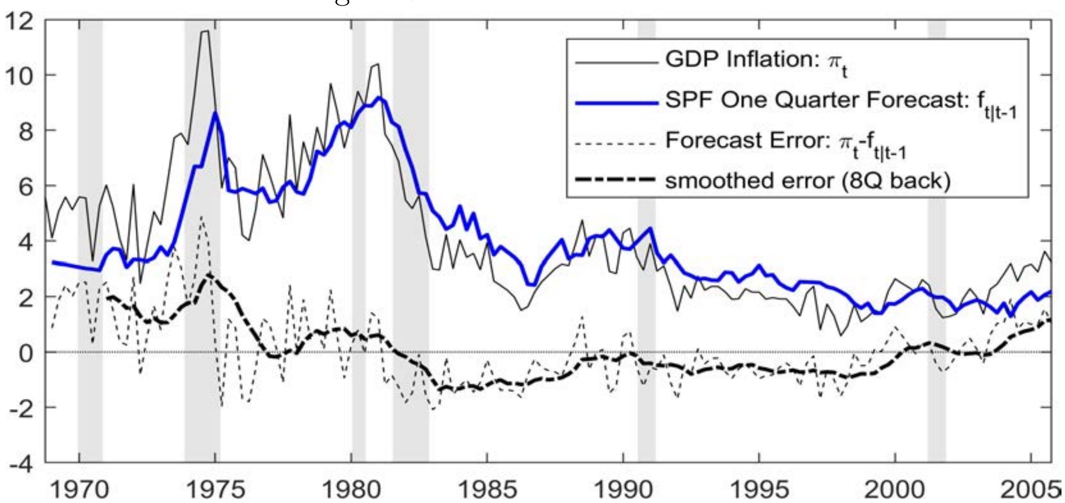

The forecasting error, $\pi_t - f_{t|t-1}$, displays serial correlation – lengthy runs of positive and negative values – that are highlighted by an 8 quarter moving average.

Inflation Forecasting Errors: Our framework thus sheds light on why inflation forecast errors turned from persistently positive to persistently negative around 1980, as highlighted by Figure 6. We saw in Figure 5 that observed inflation was consistent with opportunistic policy intentions before 1981Q1 and committed policy intentions afterward. We also saw that opportunistic intended inflation α is always higher than committed intended inflation a. Finally, expected inflation e is roughly a weighted average of the two. Taking these elements together, before 1981, observed inflation and our opportunistic policy measure $\widehat{\alpha}$ exceed expected inflation (the SPF1Q) so that persistently positive inflation forecast errors arise. After 1981Q1, observed inflation is instead tracked by our committed policy measure $\widehat{a}$, lying below the SPF1Q, yielding persistently negative inflation forecast errors.

Figure 5 includes the extracted cost-push shock $\widehat{\varsigma}$ (magenta) and the extracted reputation state $\widehat{\rho}$ (cyan and measured on the right hand axis). We next utilize these extracted states to facilitate our interpretation of U.S. inflation history.

The Great Inflation: Our model portrays the Great Inflation is a joint product of cost-push shocks and declining reputation. Advocates of the supply shock theory of the Great Inflation, such as Blinder and Rudd (2008), highlight the 1973-1975 surge and decline in inflation. These analysts point out that supply shocks – based on changes in relative prices – necessarily lead to temporary changes in inflation, so that they are well equipped to capture such "hills" as they do in our model.[52] But note that inflation is several percent higher *after* 1973-1975 than it was in the early 1970s. Our framework captures this higher "plateau" of inflation as stemming from a decline in reputation: the observed response to the supply shock makes it more likely that there is a policymaker that can't commit and a policymaker that can't commit responds to higher expected inflation by choosing higher intended inflation. Note that the extracted reputation is about .67 in 1968, falls to around .3 after the 1973-1975 cost-push shocks, and reaches a trough below .1 after another round of major cost-push shocks at the end of 1970s.

Volcker Disinflation: Our model suggests that the "Volcker Disinflation" did not start until 1981. Afterward, both inflation and expected inflation declined, with the former lying below the latter. During the Volcker Disinflation, our extracted reputation rises from a trough level of .1 in early 1981 to about .4 by the end of the recession in November 1982. The process of gaining reputation is incomplete despite a close alignment of actual inflation with the committed policy $\widehat{a}$, because declining inflation expectations also lead to reductions in the opportunistic policy $\widehat{\alpha}$, making it more difficult for private agents to determine the type in place.

The Stabilization of Inflation starts with the end of the major recession in November 1982. From this point on, inflation is well known to be fairly stable and relatively low, particularly during the Greenspan years (1988-2005). Actual inflation roughly tracks the committed policy $\widehat{a}$, even though our state variables were not chosen to produce that result. The shadow opportunistic policy $\widehat{\alpha}$ is also relatively low and stable during this period, but it lies above actual inflation π and the committed policy $\widehat{a}$ in keeping with the inflation

[52] Specifically, there are two relevant cases: (i) a temporary increase in a key price such as energy leads a high inflation period to be followed by a lower inflation period; and (ii) permanent changes in relative prices have at most a temporary effect on inflation.

bias of $\iota = 0.5\%$ per annum that is built into our model. Reputation rises steadily from around .4 in early 1983 to above .9 by the end of 2004.

6.2 Why incomplete information is necessary

It is standard to study inflation policy in two extreme cases, endowing the private sector with full information on whether the policymaker is able to commit or not. In this subsection, we construct two counterfactual inflation polices under the assumption that a single policymaker of a specific type was in place for the entire period 1968Q4 through 2005Q4. We confront a committed and opportunistic policymaker with the same cost push shocks ς as above and, crucially, assume type is known by private agents.[53]

Figure 7 contrasts these counterfactual policies with our model-implied policies to highlight why we depart from the standard assumptions of complete information.

Policy regime without commitment When private agents know there is a policymaker that can't commit, the optimal inflation policy is the dash-dotted red line in Figure 7. It is the quantitative counterpart to the formulae in Clarida et al. (1999): assuming that $e_t = \beta E_t \alpha_{t+1}$, the policy is $\alpha_t = \pi^* + \frac{\iota}{1-\beta A} + \frac{A}{1-\beta\nu A}\varsigma_t$.[54] In the jargon of the literature, this solution includes a Nash equilibrium inflation bias ($\frac{\iota}{1-\beta A}$) and a "stabilization bias." Relative to actual inflation (black line), this counterfactual opportunistic policy is too high and volatile. By contrast, our model-implied opportunistic policy before 1981 tracks the Great Inflation very well due to an important interplay of active and shadow policies in our incomplete information equilibrium. In particular, although the committed policymaker is not active before 1981, his hypothetical behavior a (green dotted line) holds down inflation expectations e when private sector does not know the policymaker's type. This in turn holds down the optimal policy without commitment (red solid line).

[53] We have also incorporated ε shocks in this exercise, but suppress these in this discussion because (i) they have no bearing on optimal policy under complete information; and (ii) the figure is less cluttered when they are not present.

[54] Given the cost push shock ς follows a Markov process with persistence ν, $E_t\varsigma_{t+1} = \nu\varsigma_t$.

Figure 7: Counterfactual inflation with and without commitment

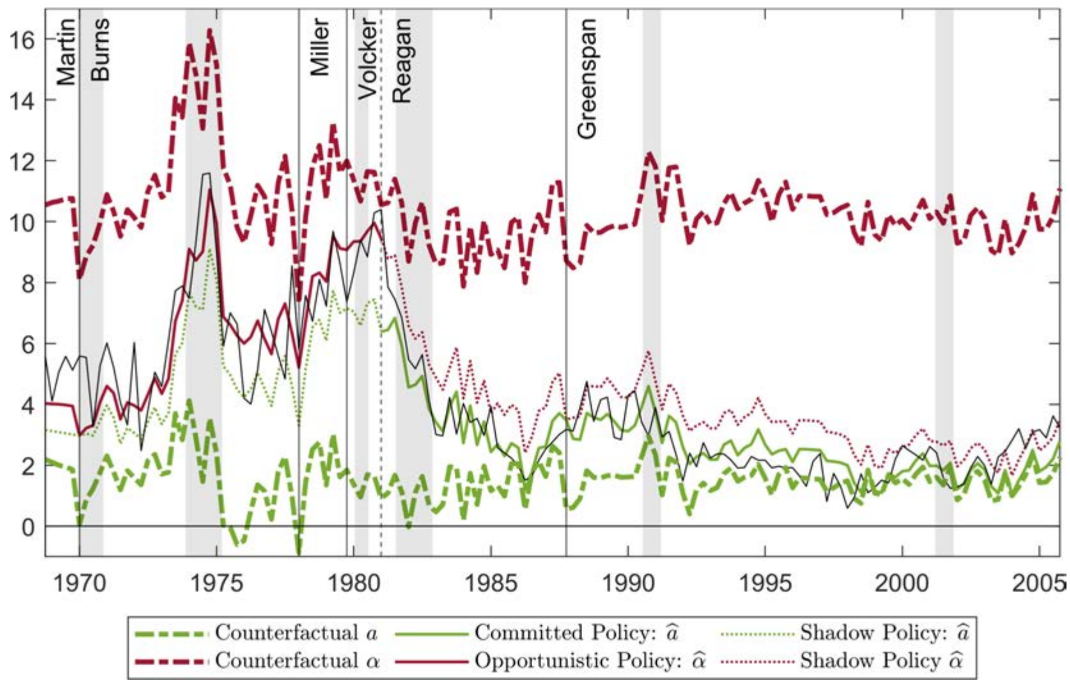

Counterfactual rates are computed assuming a policymaker of known type in place for the entire period, faced with the same extracted ς shocks. Inflation is plotted in the background using the black solid line.

Policy regime with commitment When private agents know that they are facing a committed policymaker, the optimal committed inflation policy is the dash-dotted green line in Figure 7.[55] In line with the "flexible inflation targeting" concept of Svensson (1999) and the discussion of Clarida et al. (1999), the dramatic shocks of the early 1970s lead the counterfactual committed policy a to rise well above the long-run target π^* of 1.5%. But this increased counterfactual a is soon followed by an interval of policy below the long-run target, so that the deviations from target average out.

For most of 1968-2005, the counterfactual committed policy is significantly

[55] This is $a^*(\widehat{\varsigma}, \rho = 1, \mu)$ with the same extracted cost push shocks and an appropriately calculated pseudo state μ, according to its equilibrium evolution equation: $\mu' = \mu'^*(\widehat{\varsigma}, 1, \mu)$. The initial value of μ is set to zero at 1968Q4.

lower than our model-implied committed policy when the committed policymaker is facing imperfect reputation (the dashed green line before 1981 and the solid green line after 1981). Confronted with less-than-perfect reputation, a committed policymaker recognizes that expectations management is more difficult, so that a greater accommodation of shocks and inflation expectations is desirable. As reputation evolves to a high level after the mid 1990s, the difference between the two commitment policies becomes smaller.

6.3 Why evolving reputation is essential

Our framework specifies that $\widehat{a}$ and $\widehat{\alpha}$ depend on the extracted states $(\widehat{\varsigma}, \widehat{\rho}, \widehat{\mu})$. To isolate the quantitative importance of time-varying reputation state ρ for optimal policies, we construct new policy measures by evaluating our optimal policy rules holding reputation at a constant value $\underline{\rho} = 0.6$,[56] while keeping cost-push shock $\widehat{\varsigma}_t$ and pseudo-state variable $\widehat{\mu}_t$ at the extracted levels.[57] The gap between constant-ρ and original policies then measures the effect of time-varying reputation. Figure 8 contrasts these constant-ρ policies (dash-dotted line) constructed using $(\widehat{\varsigma}_t, \underline{\rho}, \widehat{\mu}_t)$, with original policies (solid line) constructed using $(\widehat{\varsigma}_t, \widehat{\rho}_t, \widehat{\mu}_t)$. Observed inflation (black line) facilitates assessment of how much time variation in ρ helps our model match the US inflation experience: the bottom line is that it is essential.

The top panel of Figure 8 reveals that even if equipped with the same price shocks and regime changes, the constant-ρ policy without commitment misses two key features of the Great Inflation. First, it is almost 2% short of the inflation around the 1974Q3 peak. Second, it is unable to capture the post-1975 higher plateau of inflation. The bottom panel reveals that extracted reputation $\widehat{\rho}$ has fallen below its historical average since 1974Q1. In our model, lower reputation leads to higher expected inflation and a higher e, which in turn leads to a higher α^* given the best response function (10). In other words, it is the decline of reputation that enables our model-implied opportunistic policy to capture the Great Inflation.

[56] $\underline{\rho} = 0.6$ is the sample average of extract reputation state $\widehat{\rho}_t$ during 1968Q4-2005Q4.

[57] An alternative decomposition would based on holding ρ fixed and recomputing the decision rules. See Appendix A.10 for how to compute the decision rules with constant ρ.

Figure 8: Historical decomposition: effect of ρ constant at historical average

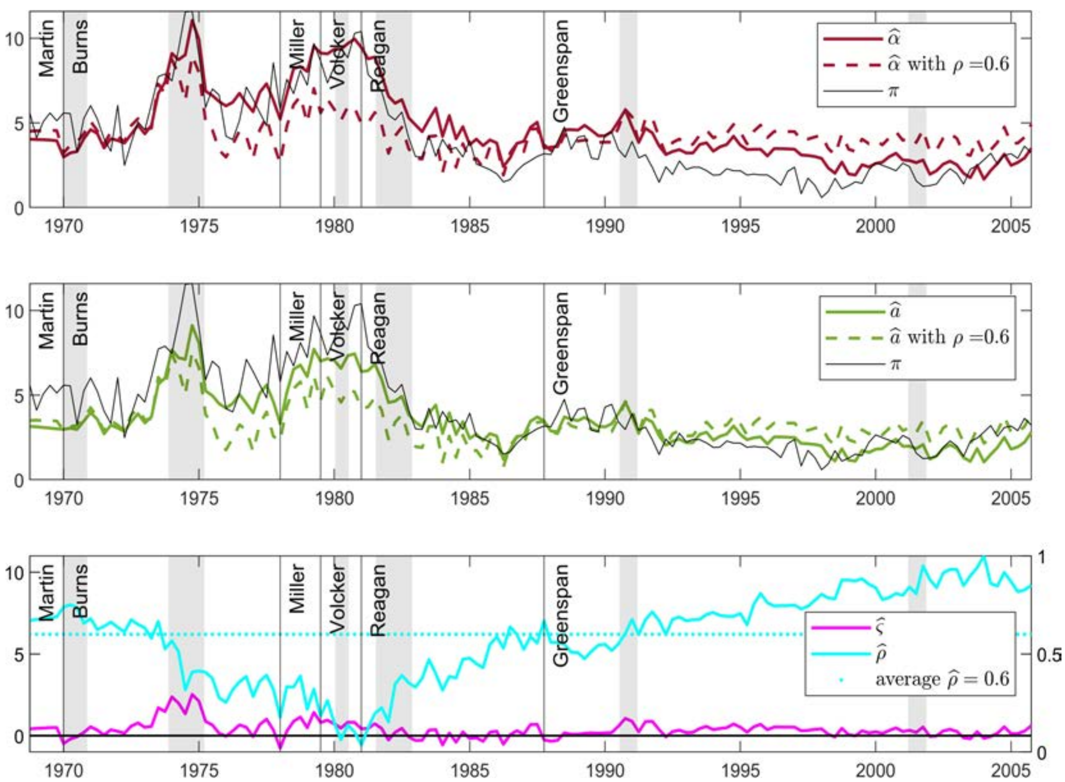

The middle panel contrasts the constant-ρ committed policy with our model-implied committed policy that depends on time-varying reputation. Looking before 1981, we can see that the committed policy exceeds π^*. From Section 6.2, we know this is partly due to price shocks, but it mainly occurs because expectations management is less effective when reputation is low. In particular, the committed authority sees expectations as less directly responsive to a and as varying in response to private sector beliefs about what an opportunistic policymaker would do.[58] After 1981, the constant-ρ committed policy follows a disinflation path, which is a gradual one due to imperfect expectations management. But it is too moderate relative to actual inflation and our time-varying-ρ committed policy. We think that the historical record

[58]Recall that expectations depend on future a, future α and ρ in Equation 6.

shows that Volcker placed substantial weight on expectations management and reputation. The bottom panel shows that the extracted time-varying reputation gradually rises from its trough after 1981Q1. According to Proposition 2, optimal policy under commitment is based in part on reputation building. Thus, again, it is time-varying reputation that enables our model-implied committed policy to quantitatively match the Volcker Disinflation.

Finally, when inflation is relatively stable after 1990, the constant-ρ committed policy becomes too high relative to our model-based policies with time-varying reputation. That is, our model would miss the Great Moderation too if it omitted time-varying reputation.[59]

6.4 Contrasting Reputation and Credibility

Macroeconomists frequently discuss policymaker reputation, as we have above, and the credibility of a specific announcement or program,[60] as we have not. In a prelude to a survey of macroeconomists and central bankers about credibility, Blinder (2000) remarks that his "own favorite definition involves matching deeds to words: a central bank is credible if people believe it will do what it says." While we also like this definition,[61] it is incomplete because it does not allow for imperfect credibility. We agree with M. King (2005) that "credibility is not an all-or-nothing matter. Policy is neither credible nor incredible. It is, as we say in economics, a continuous variable."

Measuring credibility One natural measure of partial credibility of a committed policymaker's policy plan $a(s)$ is the *credibility gap*, defined as $a(s) - z(s)$, where $z(s)$ is the private sector nowcast of inflation.[62] This measure in our setup is just $(1-\rho)\delta(s)$. Another is the *degree of credibility*, defined as the private sector's probability that inflation will fall in a band around target

[59]This is because since 1990, the time-varying reputation has risen above its historical average. That makes our model-implied committed policy less accommodating, which is closer to actual inflation, than the constant-ρ committed policy. Appendix D explains in detail the effects of reputation on equilibrium decision rules.

[60]For example, some economists point to a country's long-term interest rate as a measure of credibility for low inflation, presuming it dominated by long-term inflation expectations, e.g., M. King (2005) and Goodfriend (1993).

[61]See the opening discussion of "Managing Expectations" (R.G. King et al. (2008))

[62]Cukierman and Meltzer (1986) define credibility as: "the absolute distance between the policymaker's plans and agents beliefs about those plans."

(e.g., $a - \theta \leq \pi \leq a + \theta$) relative to the committed type's probability. In our setup, this probability reflects implementation errors with standard deviation σ and the private sector's lack of knowledge about policymaker type.[63]

Figure 9 displays these credibility measures along with extracted reputation $\widehat{\rho}$.[64] A strand of empirical central banking literature seeks to measure credibility rather than reputation: the correlation between credibility and reputation measures shown in Figure 9 suggests that this work may shed light on our theoretical reputation mechanism.[65]

Figure 9: Credibility and Reputation

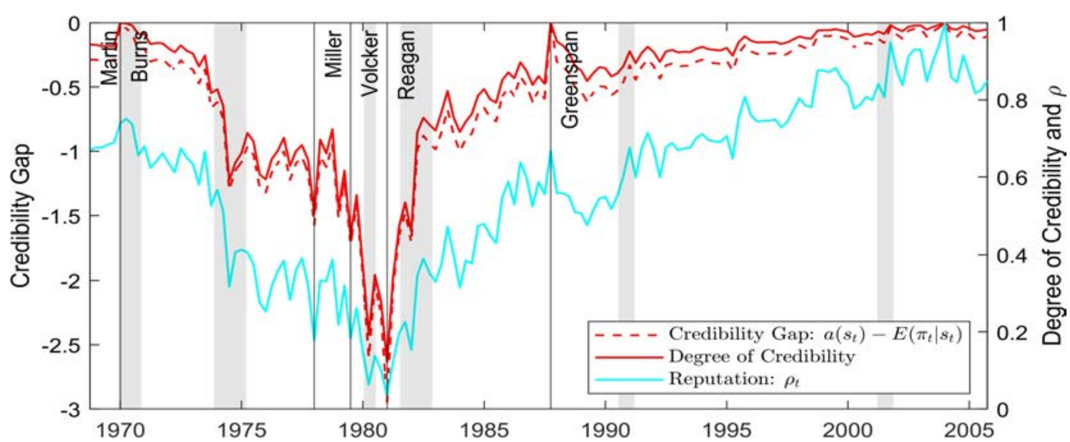

The inflation credibility gap is defined as $a(s) - E(\pi|s) = (1-\rho)[a(s) - \alpha(s))] = (1-\rho)\delta$. The degree of credibility is defined as the ratio of the private sector's probability that $a - \theta \leq \pi \leq a + \theta$ relative to the committed type's probability.

Long-term credibility and reputation. Both of these evolving partial credibility measures depend on near-term intended inflation. Under commitment, though, a long-lasting regime will attain $\rho = 1$ and intended inflation

[63] Additional details on these measures are provided in Appendix E.

[64] The degree of credibility is based on $\theta = \sigma$.

[65] The strong association of reputation and credibility combined with better expectations management in our model when reputation is high echos many ideas regularly discussed by practical macroeconomists and central bankers, such as "greater credibility improves the short-run inflation-unemployment trade-off," "greater credibility brings down the cost of reducing inflation" and "once low inflation has been achieved, a more credible central bank is better able to maintain low inflation." These quotes are from Blinder (2000), p. 145.

will have a stationary distribution with $E(a) = \pi^*$. Hence, ρ_t is a measure of longer-term credibility and, in particular, of the date t likelihood that the current regime will achieve "price stability." In this sense, our model captures the views of some Blinder (2000) survey participants: "a central bank can raise the public's subjective probability that it is 'tough' by keeping inflation low. This probability is, in turn, taken as a measure of the bank's credibility." It is also consistent with his summary "that many central bankers take the degree of dedication to price stability as synonymous with credibility."

6.5 Looking Forward

Our quantitative analysis has so far focused on 1968Q4 through 2005Q4. We made this choice for several reasons. First, the sample matches that of leading studies of U.S. inflation's rise, fall and stabilization that were mainly undertaken prior to the Global Financial Crisis.[66] Second, our analysis abstracted from monetary policy instruments and fiscal actions, even as these varied through US history, because we viewed these as subordinated to intended inflation. Yet, the now-standard theory of policy with short-term interest rates at zero requires an aggregate demand specification and imposes additional constraints on our policy problem, so we avoided these complications.[67]

However, our refraining from interpreting the longer U.S. inflation history with our model does not mean it performs poorly beyond 2005. In fact, Appendix F shows that our model performs well in matching the SPF term structure and observed inflation through 2020Q1.

Since 2020Q4, the U.S. inflation has transitioned from ranging around 1.5 to 2 percent to a sustained increase that takes it to over 8 percent in 2022Q4. This recent development has prompted many to make comparisons with the 1973-4 upswing in inflation.[68]

We thus use a segment of our more lengthy model-based interpretation of

[66]Examples include Sargent (1999) and Primiceri (2006) but there are many others.

[67]See for example Eggertsson and Woodford (2003).

[68]A Fed staffer from that period recalls Arthur Burn's fixation on special factors in inflation that were assumed to be transitory. He observes that "The US Federal Reserve is insisting that recent increases in (prices of specific goods) reflect transitory factors that will quickly fade with post-pandemic normalization. But what if they are a harbinger, not a "noisy" deviation?" Roach 2021.

Figure 10: Model-based interpretation of recent US inflation

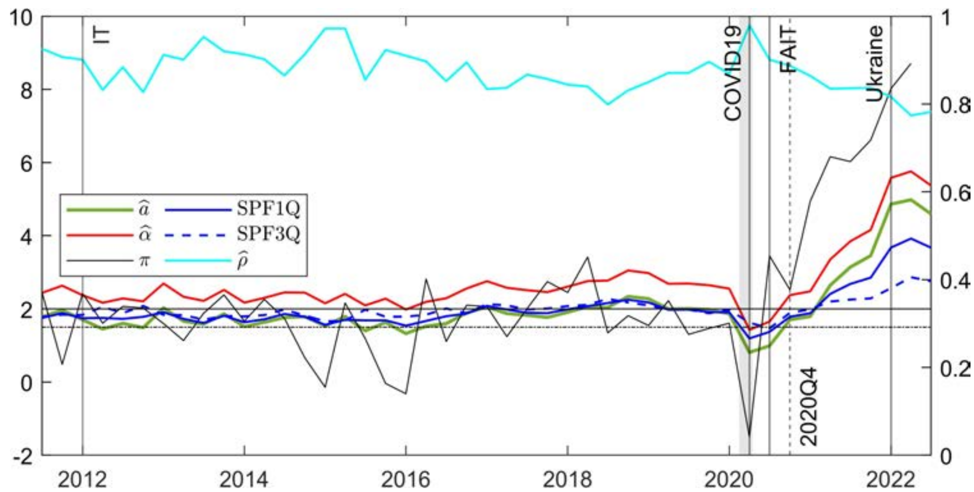

US inflation history (Figure 10) to consider our model's relevance for today's policy and the years ahead. We start in 2011Q4, shortly before the Fed announced an inflation target of 2 percent in January 2012. We display the SPF forecasts at the 1 and 3 quarter horizons which, for much of 2011Q4 to 2022Q3, were generally close to each other and mainly between the announced target of 2 percent and our π^* of 1.5 percent.

Figure 10 also contains our SPF-based reputation measure $\widehat{\rho}$ (cyan and measured on the right hand axis) and the intended inflation measures $\widehat{a}$ (in green) and $\widehat{\alpha}$ (in red) produced by our model. Through 2019, extracted reputation remains high and actual inflation π fluctuates around the relatively constant commitment policy (and the SPF expectations) as it would under optimal policy with high or perfect reputation.

During 2020 through 2022, though, three events took place that are not present in our model: the winter 2020 onset of the COVID-19 pandemic, the summer 2020 Fed announcement of a shift to a flexible average inflation targeting (FAIT) approach with a short-run inflation target above 2 percent, and the early 2022 onset of the Ukraine war. In 2020Q4, extracted reputation $\widehat{\rho}$ is roughly unchanged from its the 2011 level, the SPF forecasts are close to each other and to a, midway between 1.5 and 2 percent. However, after 2020Q4, inflation moved up to 2.5 percent and then rose to 8 percent.

In the 6 quarters after 2020Q4, the SPF1Q increased by more than SPF3Q: our extraction approach interprets as a transitory positive cost-push shock. But since both expectations measures have risen, we extract declining reputation. Accordingly, both of our policy measures rise but α by more than a, as a combination of inflation and stabilization bias.

To this writing, though, the movements in expected inflation and model-based policy responses are still small relative to the 1970s. Given inflation is at 8 percent in 2022Q2 while committed policy is slightly below 5 percent and opportunistic policy is around 5.5 percent, the optimistic vision is that the temporary rise of inflation is principally just due to a series of positive, transitory shocks ϵ and ς, with these shocks so far not having triggered the sort of dramatic decline in reputation experienced in 1974-1975. But, if inflation continues to rise, our model predicts the reputation will further decline, resulting in a high plateau of high inflation and expected inflation even after these transitory shocks pass, just as was the case after 1973-1975. At that point, disinflation will be costly as it was in the 1980s.

7 Summary, Conclusions and Final Remarks

We present a parsimonious model that can simultaneously capture the expected and actual inflation in the U.S. Our setup features a standard forward-looking New Keynesian Phillips curve, policy regime changes with policymakers differing in commitment capacity, and Bayesian learning by the private sector about policymaker type.

Both types of policymakers, committed and opportunistic, behave purposefully in our model. The committed policymaker strategically uses its policy plan to influence private sector's learning and inflation expectations, understanding that (i) private sector inflation expectations include future policy of an opportunistic type; and (ii) an opportunistic type's optimal policy depends on private sector inflation expectations. We adopt a mechanism design approach to formulate the problem of the committed type in a compact recursive form with only three state variables. This permits calculation of decision rules of both types of policymakers and the rational expectations of the private sector in the Bayesian perfect equilibrium.

Putting our theory to work, we show how to extract these state variables from just the SPF inflation data. We use these extracted states to construct time series of optimal committed and opportunistic policy without using actual inflation. Yet, when we assume regimes with opportunistic policymakers before 1981 and with committed policymakers afterward, the corresponding optimal policy tracks US inflation's rise, fall, and stabilization between 1970 and 2005.

Our quantitative exercise reveals that evolving reputation is very important in accounting for actual inflation. In particular, endogenous policy differences help explain why private sector learning is slow in early 1970s; why cost-push shocks in the mid 1970s sped up learning, intensifying and prolonging the Great Inflation; and why the Volcker disinflation may be understood as a committed policymaker rebuilding reputation lost during the Great Inflation.

These lessons from the 1970s and 1980s appear particularly relevant for the ongoing fight against inflation in the U.S. Small but persistent deviations of inflation from targets can eventually lead to run-away inflation expectations, even though such expectations may appear very sticky early on. Explicitly committing to inflation targets – including flexible inflation targeting – helps the central bank to acquire and maintain credibility for attaining its monetary policy objectives. Our theory highlights that reputation for commitment, a measure of long-term credibility, can be gained or lost. We think reputations evolution is important for US macroeconomic history, past and future.

Our model is deliberately stark. But it yields results that have surprised us and others. We believe its success in matching U.S. time series indicates great promise to further research on models that feature agents learning about the commitment capacity of purposeful policymakers.

References

Atkeson, Andrew and Patrick J. Kehoe, "The advantage of transparency in monetary policy instruments," *Federal Reserve Bank of Minneapolis Staff Report 297*, 2006.

Ball, Laurence, "Disinflation with Imperfect Credibility," *Journal of Monetary Economics*, 1995, *35* (1), 5–23.

Barro, Robert J. and David B. Gordon, "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Political Economy*, 1983, *91* (4), 589–610.

_ **and** _ , "Rules, discretion and reputation in a model of monetary policy," *Journal of Monetary Economics*, 1983, *12* (1), 101–121.

Bianchi, Francesco, "Regime Switches, Agents' Beliefs, and Post-World War II U.S. Macroeconomic Dynamics," *The Review of Economic Studies*, oct 2013, *80* (2), 463–490.

Blinder, Alan and Jeremy Rudd, "The Supply-Shock Explanation of the Great Stagflation Revisited," Technical Report w14563, National Bureau of Economic Research, Cambridge, MA December 2008.

Blinder, Alan S., "Central-Bank Credibility: Why Do We Care? How Do We Build It?," *The American Economic Review*, 2000, *90* (5), 1421–1431.

Brayton, Flint, Thomas Laubach, and David Reifschneider, "Optimal-Control Monetary Policy in the FRB/US Model," *FEDS Notes*, November 2014, *2014* (0035).

Carvalho, Carlos, Stefano Eusepi, Emanuel Moench, and Bruce Preston, "Anchored Inflation Expectations," *American Economic Journal: Macroeconomics*, 2022.

Chang, Roberto, "Credible Monetary Policy in an Infinite Horizon Model: Recursive Approaches," *Journal of Economic Theory*, 1998, *81*, 431 – 461.

Chari, V. V. and Patrick J. Kehoe, "Sustainable Plans," *Journal of Political Economy*, 1990, *98* (4), 783–802.

Clarida, Richard, Jordi Gali, and Mark Gertler, "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, 1999, *37* (4), 1661–1707.

Cogley, Timothy, Christian Matthes, and Argia M. Sbordone, "Optimized Taylor rules for disinflation when agents are learning," *Journal of Monetary Economics*, may 2015, *72*, 131–147.

Coibion, Olivier, Yuriy Gorodnichenko, and Rupal Kamdar, "The Formation of Expectations, Inflation, and the Phillips Curve," *Journal of Economic Literature*, December 2018, *56* (4), 1447–1491.

Cukierman, Alex and Allan H. Meltzer, "A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information," *Econometrica*, September 1986, *54* (5), 1099.

_ **and Nissan Liviatan**, "Optimal accommodation by strong policymakers under incomplete information," *Journal of Monetary Economics*, February 1991, *27* (1), 99–127.

Debortoli, Davide and Aeimit Lakdawala, "How Credible Is the Federal Reserve? A Structural Estimation of Policy Re-Optimizations," *American Economic Journal: Macroeconomics*, jul 2016, *8* (3), 42–76.

DeLong, J. Bradford, "America's Only Peacetime Inflation: The 1970s," Technical Report h0084, National Bureau of Economic Research, Cambridge, MA May 1996.

Dovis, Alessandro and Rishabh Kirpalani, "Rules without Commitment: Reputation and Incentives," *The Review of Economic Studies*, feb 2021.

Eggertsson, Gauti B. and Michael Woodford, "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 2003, *2003* (1), 139–211.

Erceg, Christopher J. and Andrew T. Levin, "Imperfect credibility and inflation persistence," *Journal of Monetary Economics*, May 2003, *50* (4), 915–944.

Evans, Martin and Paul Wachtel, "Inflation Regimes and the Sources of Inflation Uncertainty," *Journal of Money, Credit and Banking*, 1993, *25* (3), 475–511.

Farmer, Leland, Emi Nakamura, and Jón Steinsson, "Learning About the Long Run," Working Paper 29495, National Bureau of Economic Research November 2021.

Faust, Jon and Lars E. O. Svensson, "Transparency and Credibility: Monetary Policy with Unobservable Goals," *International Economic Review*, 2001, *42* (2), 369–397.

Gali, Jordi, *Monetary Policy, Inflation, and the Business Cycle*, Princeton Univers. Press, June 2015.

Goodfriend, Marvin, "Interest Rate Policy and the Inflation Scare Problem: 1979-1992," *Federal Reserve Bank of Richmond Economomic Quarterly*, 1993.

_ **and Robert G. King**, "The incredible Volcker disinflation," *Journal of Monetary Economics*, July 2005, *52* (5), 981–1015.

Gordon, Robert, "The Phillips Curve is Alive and Well: Inflation and the NAIRU During the Slow Recovery," Technical Report aug 2013.

Hazell, Jonathon, Juan Herreño, Emi Nakamura, and Jón Steinsson, "The Slope of the Phillips Curve: Evidence from U.S. States*," *The Quarterly Journal of Economics*, 02 2022, *137* (3), 1299–1344.

King, Mervyn, "Credibility and Monetary Policy: Theory and Evidence," *Bank of England Quarterly Bulletin*, March 2005.

King, Robert G., Yang K. Lu, and Ernesto S. Pastén, "Managing Expectations," *Journal of Money, Credit and Banking*, December 2008, *40* (8), 1625–1666.

Kurozumi, Takushi, "Optimal sustainable monetary policy," *Journal of Monetary Economics*, 2008, *55* (7), 1277–1289.

Kydland, Finn E. and Edward C. Prescott, "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, 1977, *85* (3), 473–491.

_ **and** _ , "Dynamic optimal taxation, rational expectations and optimal control," *Journal of Economic Dynamics and Control*, January 1980, *2*, 79–91.

Levin, Andrew and John B. Taylor, "Falling Behind the Curve: A Positive Analysis of Stop-Start Monetary Policies and the Great Inflation," in "The Great Inflation: The Rebirth of Modern Central Banking," University of Chicago Press, June 2013, pp. 217–244.

Loisel, Olivier, "Central bank reputation in a forward-looking model," *Journal of Economic Dynamics and Control*, 2008, *32* (11), 3718–3742.

Lu, Yang K., "Optimal policy with credibility concerns," *Journal of Economic Theory*, September 2013, *148* (5), 2007–2032.

_ **, Robert G. King, and Ernesto Pasten**, "Optimal reputation building in the New Keynesian model," *Journal of Monetary Economics*, December 2016, *84*, 233–249.

Mailath, George J. and Larry Samuelson, *Repeated Games and Reputations: Long-run Relationships*, Oxford University Press, July 2006.

Marcet, Albert and Ramon Marimon, "Recursive Contracts," *Econometrica*, 2019, *87* (5), 1589–1631.

Matthes, Christian, "Figuring Out the Fed-Beliefs about Policymakers and Gains from Transparency," *Journal of Money, Credit and Banking*, January 2015, *47* (1), 1–29.

Melosi, Leonardo, "Signalling Effects of Monetary Policy," *The Review of Economic Studies*, sep 2016, p. rdw050.

Meltzer, Allan H., *A history of the Federal Reserve. vol. 2, book 2: 1970 - 1986*, paperback ed ed., Chicago: Univ. of Chicago Press, 2014.

Mertens, Elmar and James M. Nason, "Inflation and professional forecast dynamics: An evaluation of stickiness, persistence, and volatility," *Quantitative Economics*, November 2020, *11* (4), 1485–1520.

Mishkin, Frederic S. and Klaus Schmidt-Hebbel, "Does Inflation Targeting Make a Difference?," NBER Working Papers 12876 January 2007.

Orphanides, Athanasios, "Comment on: "The incredible Volcker disinflation"," *Journal of Monetary Economics*, July 2005, *52* (5), 1017–1023.

_ **and John C. Williams**, "The decline of activist stabilization policy: Natural rate misperceptions, learning, and expectations," *Journal of Economic Dynamics and Control*, 2005, *29* (11), 1927–1950.

Phelan, Christopher and Ennio Stacchetti, "Sequential Equilibria in a Ramsey Tax Model," *Econometrica*, 2001, *69* (6), 1491–1518.

Primiceri, Giorgio E, "Why Inflation Rose and Fell: Policy-Makers's Beliefs and U. S. Postwar Stabilization Policy," *Quarterly Journal of Economics*, aug 2006, *121* (3), 867–901.

Roger, Scott and Mark R. Stone, "On Target? the International Experience with Achieving Inflation Targets," *IMF Working Paper No. 05/163*, 2005.

Sargent, Thomas J., *The conquest of American inflation*, 1st print ed., Princeton, NJ: Princeton University Press, 1999.

_ **and Ulf Soderstrom**, "The conquest of American inflation: A summary," *Sveriges Riksbank Economic Review*, 2000, *3*, 12–45.

Schaumburg, Ernst and Andrea Tambalotti, "An investigation of the gains from commitment in monetary policy," *Journal of Monetary Economics*, 2007, *54* (2), 302–324.

Shapiro, Adam and Daniel J. Wilson, "The Evolution of the FOMC's Explicit Inflation Target," *FRBSF Economic Letter*, April 2019.

Svensson, Lars E. O., "Inflation Targeting: Some Extensions," *The Scandinavian Journal of Economics*, 1999, *101* (3), 337–361.

Watson, Mark W., "Inflation Persistence, the NAIRU, and the Great Recession," *American Economic Review*, May 2014, *104* (5), 31–36.

Woodford, Michael, *Interest and prices*, Princeton, NJ [u.a.]: Princeton Univ. Press, 2003.

Appendices

A Recursive optimal policy design

The optimal policy problem for the committed type at the start of its tenure involves forward-looking constraints, which must be transformed to yield a recursive specification. Conceptually, this involves casting Lagrangian components in recursive form, relying on (i) application of the law of iterated expectation and (ii) appropriate rearrangement of expected discounted sums. In the current model, the transformation to recursive form must also take into account that the committed policymaker and the private sector have different discount factors and probability beliefs, so that the law of iterated expectation must be applied carefully.

This appendix's derivation of the recursive program in Proposition 1 incorporates three structural features described in section 2 of the text: (1) informational subperiods; (2) different information sets for the committed policymaker and the private sector; and (3) private sector learning. It also generalizes the section 2 framework so that (a) it can be used with constant reputation or a mechanical alternative type; (b) it can be used when the opportunistic type of policymaker is forward-looking with time discount factor β_α. Various elements from the main text are repeated, so that the appendix may be read separately.

The detailed derivation of the recursive form is a slow-moving proof, designed for readers with various degrees of prior exposure to recursive optimal policy design. A key new feature relative to other macro applications is a "change of measure" in the expectations constraint on the committed policymaker, which arises because private agents understand that inflation may come from the decisions of an optimizing alternative type.[1]

As we develop the optimal policy for the committed type, we assume that the committed type takes as given a function governing private agents' expected inflation in the event of its replacement, which may depend on events during its tenure and, in particular, on its terminal reputation. But in the background, there is an equilibrium requirement that private agents form rational beliefs about inflation in the event of a replacement next period. We discuss imposing this requirement at the end of this appendix.

A.1 Intended and actual inflation

At each date, the policymaker chooses intended inflation, denoted as a for the committed type (τ_a) and α for the alternative type (τ_α). Intended inflation is not observed by the

[1] This feature will play an even more important role in future research that makes the alternative type care more about the future than in the current case of a myopic alternative.

private sector. Actual inflation is randomly distributed around this intention, with density $g(\pi|a)$ if there is a committed type and $g(\pi|\alpha)$ if there is an alternative type. We assume

$$\begin{aligned} a &= \int \pi g(\pi|a) d\pi \\ \alpha &= \int \pi g(\pi|\alpha) d\pi \end{aligned}$$

Implementation errors are $\varepsilon_a = \pi - a$ and $\varepsilon_\alpha = \pi - \alpha$ for the two types. While we allow for different continuous distributions on the same range of inflation outcomes, we do not separately include type τ as an argument to avoid notation clutter in the balance of this appendix (i.e., we write $g(\pi|a)$ and $g(\pi|\alpha)$).

A.2 Measures of history

We use period t as the time index within a regime, so period 0 is the date of last regime change. The committed type begins with a reputation, ρ_0, known to private agents.

Private agents at the end of period t know the entire history of inflation (π), output (x), and inflation shocks (ς) since period 0 (the last regime change date). After the next period starts, the ς shock is realized. The policymaker's intended inflation (a or α) is conditioned on this information, as is the expectations shifter in the output-inflation trade-off, e. We write the information history as

$$h_t = [\varsigma_t, \{\varsigma_{t-s}\}_{s=1}^{t}, \{\pi_{t-s}\}_{s=1}^{t}]$$

After the policymaker chooses his intended inflation, actual inflation and output are realized. Other variables, notably private agents' updated belief about policymaker type, are conditioned on this extended information,

$$h_t^+ = [\pi_t, h_t].$$

Note that

$$h_{t+1} = [\varsigma_{t+1}, h_t^+] = [\varsigma_{t+1}, \pi_t, h_t]$$

A word on notation: In the Public Perfect Bayesian Equilibrium of our dynamic game, variables depend just on the relevant history (e.g., $a(h_t)$) and not separately on the date (e.g., $a_t(h_t)$). To further streamline some formulas, we will sometimes condense variables even further, writing $a(h_t)$ as a_t.

A.3 Beliefs about current inflation

Although private agents do not know the type of policymaker that is in place, at the start of period t, they have a prior belief ρ_t that there is a committed type which will choose a_t and a complementary prior belief $1-\rho_t$ that there is an alternative type which will choose α_t. Accordingly, their rational likelihood of the outcome π_t is

$$g(\pi_t|a_t)\rho_t + g(\pi_t|\alpha_t)(1-\rho_t) \tag{A18}$$

A.4 Beliefs about policymaker type

On observing inflation within a regime, private agents use Bayes' law to update their conditional probability that the current policymaker is the committed type

$$\begin{aligned} \rho(h_t^+) &= \frac{g(\pi_t|a(h_t))\rho(h_t)}{g(\pi_t|a(h_t))\rho(h_t) + g(\pi_t|\alpha(h_t))(1-\rho(h_t))} \\ &\equiv b(\pi_t, a(h_t), \alpha(h_t), \rho(h_t)) \end{aligned} \tag{A19}$$

where the b function is a convenient short-hand and $h_t^+ = [\pi_t, h_t]$. As there is no information about type revealed by ς_{t+1}, $\rho(h_{t+1}) = \rho(h_t^+)$. This updating may be written

$$\rho(h_t^+) = \frac{\rho(h_t)}{\rho(h_t) + \lambda(\pi_t, h_t)(1-\rho(h_t))} \tag{A20}$$

using the likelihood ratio $\lambda(\pi_t, h_t) \equiv \frac{g(\pi_t|\alpha(h_t))}{g(\pi_t|a(h_t))}$.

A.5 Constructing expected inflation

We now construct the private sector's expected inflation, $E\pi_{t+1}$, working backwards from the start of next period to the start of this period. We take into account that there will be a regime change ($n_{t+1}=1$) with probability q and won't ($n_{t+1}=0$) with probability $1-q$.

If the committed type is known to be in place, with decision rule $a([\varsigma_{t+1}, h_t^+])$, then

$$E(\pi_{t+1}|h_{t+1}, \tau_a) = a([\varsigma_{t+1}, h_t^+])$$

since intended inflation is the mean of realized inflation. Similarly,

$$E(\pi_{t+1}|h_{t+1}, \tau_\alpha) = \alpha([\varsigma_{t+1}, h_t^+])$$

Since the private sector will not know the type of policymaker in place at the start of next period, expected inflation will be

$$E(\pi_{t+1}|h_{t+1}, n_{t+1} = 0) = \rho(h_{t+1})a(h_{t+1}) + (1 - \rho(h_{t+1}))\alpha(h_{t+1}) \tag{A21}$$

if there isn't a regime change. Without taking a stand on the details of reputation inheritance, we simply define

$$E(\pi_{t+1}|h_{t+1}, n_{t+1} = 1) = z(h_{t+1}) \tag{A22}$$

as the private sector's expectation of inflation conditional on a replacement.

Stepping back now to period t, expected inflation conditional on h_t is

$$\begin{aligned} E(\pi_{t+1}|h_t) = \rho(h_t) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) \left[(1-q)\, a(h_{t+1}) + qz\,(h_{t+1})\right] g(\pi_t|a(h_t)) d\pi_t \\ + (1 - \rho\,(h_t)) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) \left[(1-q)\, \alpha(h_{t+1}) + qz\,(h_{t+1})\right] g(\pi_t|\alpha(h_t)) d\pi_t \end{aligned} \tag{A23}$$

There may appear to be a conflict between this expression and (A21) that contains reputation at t+1. But there is not. Weighting (A21) and (A22) by $(1 - q)$ and q and then integrating over the private sector's belief about inflation (A18) leads directly to it. The simplicity arises because (A18) also occurs in the denominator of the Bayesian updating expression (A19).

A.6 Intertemporal objective

We assume that the policymaker's intertemporal objective involves discounting at $\beta_a(1-q)$, where β_a is its structural discount factor and $(1-q)$ reflects discounting due to replacement.

$$U_t = \underline{u}(a_t, e_t, \varsigma_t) + (\beta_a(1 - q))E_t^c U_{t+1}$$

where $\underline{u}\,(a, e, \varsigma) \equiv \int u\,(\pi, x(\pi, e), \varsigma)\, g\,(\pi|a)\, d\pi$ is the expected momentary objective with x replaced by $x(\pi, e) = (\pi - e - \varsigma)\,/\kappa$, and the conditional expectation operator $E_t^c(\cdot)$ is using the committed type's probability $p(h_{t+j})$ of a specific history h_{t+j} when his actions generate inflation.

More specifically, at any date t given the history h_t, the intertemporal objective is

$$U_t = \sum_{j=0}^{\infty} (\beta_a(1 - q))^j \sum_{h_{t+j}} \frac{p(h_{t+j})}{p(h_t)} \underline{u}(a(h_{t+j}), e(h_{t+j}), \varsigma(h_{t+j})) \tag{A24}$$

Given $h_{t+j} = [\varsigma_{t+j}, \pi_{t+j-1}, h_{t+j-1}]$, the committed type's probability of a specific history is:

$$p(h_{t+j}) = \varphi(\varsigma_{t+j}; \varsigma_{t+j-1}) \times g(\pi_{t+j-1}|a(h_{t+j-1})) \times p(h_{t+j-1}) \tag{A25}$$

That is, it combines the likelihood of inflation π given the committed type's decision, the likelihood of the shock ς and the probability of the previous history.[2]

A.7 Rational expectations constraint

To develop the desired recursive form, we construct the Lagrangian component using the committed type's probabilities as weights on the multipliers

$$\Psi_t = \sum_{j=0}^{\infty} (\beta_a(1-q))^j \sum_{h_{t+j}} \frac{p(h_{t+j})}{p(h_t)} \gamma(h_{t+j})[e(h_{t+j}) - \beta E(\pi_{t+j+1}|h_{t+j})] \tag{A26}$$

and then express it recursively. We detailed $E(\pi_{t+1}|h_t)$ in (A23), but the expression involved the probability of inflation under the alternative type. So, we undertake a "change of measure" and rewrite it as

$$\begin{aligned} &\rho(h_t) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t)[\beta(1-q)a(h_{t+1}) + \beta q z(h_{t+1})] g(\pi|a(h_t)) d\pi \\ &+(1-\rho(h_t)) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t)[\beta(1-q)\alpha(h_{t+1}) + \beta q z(h_{t+1})] \boldsymbol{\lambda}(\mathbf{h_{t+1}}) g(\pi|a(h_t)) d\pi \end{aligned} \tag{A27}$$

where $\lambda(h_{t+1})$ is the likelihood ratio discussed above in the context of Bayesian updating.

$$\frac{g(\pi_t|\alpha(h_t))}{g(\pi_t|a(h_t))} = \lambda(h_t^+) = \lambda(h_{t+1}) \tag{A28}$$

As the notations emphasize, this is a random variable from the standpoint of h_t but it is known as of $h_t^+ = [\pi_t, h_t]$ and $h_{t+1} = [\varsigma_{t+1}, h_t^+]$.

We now return to (A26) and replace $E(\pi_{t+1}|h_t)$ with the expression in (A27). Note that $a(h_{t+1}), \alpha(h_{t+1})\lambda(h_{t+1})$, and $z(h_{t+1})$ are multiplied by $\varphi(\varsigma_{t+1}; \varsigma_t) g(\pi|a(h_t)) p(h_t)$ and by $\gamma(h_t)$, which is $p(h_{t+1})\gamma(h_t)$. So, just as in simpler models, it is possible to eliminate expectations at future dates, essentially by applying the law of iterated expectation. Adjusting for different

[2] We ask for the reader's patience in using a sum over histories to capture the joint effects of the possibly continuous distribution of π and the discrete Markov chain distribution for ς.

discount factors, we can write (A26) as

$$\Psi_t = E_t^c[\sum_{j=0}^{\infty}(\beta_a(1-q))^j \psi_{t+j}] \tag{A29}$$

with

$$\psi_t = \gamma_t e_t - \frac{\beta}{\beta_a(1-q)}\gamma_{t-1}\{\rho_{t-1}[(1-q)a_t + qz_t] + (1-\rho_{t-1})\lambda_t[(1-q)\alpha_t + qz_t]\} \tag{A30}$$

This latter expression captures past commitments about current state-contingent decisions as these were relevant to past expectations of inflation.[3] Note that at the start of the regime, when $t = 0$, $\gamma_{t-1} = 0$ by assumption. The initial condition on reputation specifies ρ_0.

A.8 The basic recursive specification

The preceding derivations suggest a recursive version of $U_t + \Psi_t$ with states $(\varsigma_t, \gamma_{t-1}, \rho_{t-1}, \lambda_t)$. For algebraic convenience, we define $\eta_t = \frac{\beta}{\beta_a(1-q)}\gamma_{t-1}$. Then, the recursive form as in Marcet and Marimon (2019) is

$$\begin{aligned} W(\varsigma_t, \eta_t, \rho_{t-1}, \lambda_t) = \min_{\gamma} \max_{a,\alpha,e} \{ & \underline{u}(a_t, e_t, \varsigma_t) + \gamma_t e_t \\ & - \eta_t[\rho_{t-1}((1-q)a_t + qz_t) + (1-\rho_{t-1})\lambda_t((1-q)\alpha_t + qz_t)] \\ & + \beta_a(1-q)\int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1};\varsigma_t) W(\varsigma_{t+1}, \eta_{t+1}, \rho_t, \lambda_{t+1}) g(\pi_t|a_t) d\pi_t \} \end{aligned} \tag{A31}$$

subject to the IC constraint

$$\alpha_t = Ae_t + B(\varsigma_t)$$

with state dynamics (from the perspective of the committed type)

$$\begin{aligned} \eta_{t+1} &= \frac{\beta}{\beta_a(1-q)}\gamma_t \text{ with } \gamma_{-1} = 0 \\ \rho_t &= \frac{\rho_{t-1}}{\rho_{t-1} + (1-\rho_{t-1})\lambda_t} \text{ given } \rho_0 \\ \lambda_{t+1} &= \lambda(\pi_t, a_t, \alpha_t) \text{ with probability } g(\pi_t|a_t) \end{aligned}$$

Defining $S_t = [\varsigma_t, \eta_t, \rho_{t-1}, \lambda_t]$, this program delivers optimal choices $a^*(S)$, $\alpha^*(S)$, $e^*(S)$, $\gamma^*(S)$ with optimal state evolution induced by these decision rules. As is standard in recursive

[3] Our short hand notation replaces $\lambda(h_t)$ with λ_t. Given (A28), the likelihood ratio λ_t is predetermined in period t by actions and inflation outcome in period $t-1$.

systems, these rules also imply a value of the objective, $U^*(S)$.

A.9 State space reduction

For computational and analytical benefits, it is desirable to reduce the state space. We now show how to eliminate the likelihood ratio (λ) from the state vector so that we only need three state variables instead of four. Start by rewriting (A30) as

$$\text{(A32)} \qquad \psi_t = \gamma_t e_t - \frac{\beta}{\beta_a(1-q)}\gamma_{t-1}\rho_{t-1}\{[(1-q)a_t + qz_t] + \frac{(1-\rho_{t-1})\lambda_t}{\rho_{t-1}}[(1-q)\alpha_t + qz_t]\}$$

Then, note that $\rho_t = \frac{\rho_{t-1}}{\rho_{t-1}+(1-\rho_{t-1})\lambda_t}$ implies that $\frac{(1-\rho_{t-1})\lambda_t}{\rho_{t-1}} = \frac{1-\rho_t}{\rho_t}$ so that Bayes' rule can be used to eliminate λ_t. Substitution of this expression into that above yields

$$\text{(A33)} \qquad \psi_t = \gamma_t e_t - \frac{\beta}{\beta_a(1-q)}\gamma_{t-1}\rho_{t-1}\{[(1-q)a_t + qz_t] + \boxed{\frac{(1-\rho_t)}{\rho_t}}[(1-q)\alpha_t + qz_t]\}$$

which indicates that the states $(\varsigma_t, \eta_t, \rho_{t-1}, \lambda_t)$ can be reduced to ς_t, $\mu_t = \frac{\beta}{\beta_1(1-q)}\gamma_{t-1}\rho_{t-1}$ and ρ_t with the following transition rules for the endogenous states given ρ_0:

$$\text{(A34)} \qquad \mu_{t+1} = \frac{\beta}{\beta_a(1-q)}\gamma_t\rho_t \text{ with } \mu_0 = 0$$

$$\text{(A35)} \qquad \rho_{t+1} = b(\pi_t, a_t, \alpha_t, \rho_t) \text{ with probability } g(\pi_t|a_t)$$

A.10 Extended recursive program

The recursive optimization (A31) can now be written with only three state variables. While doing so, we extend the program to make it easy to shut down each of the two key mechanisms: endogenous reputation and optimizing behavior by the alternative type.

$$\text{(A36)} \qquad \begin{aligned} W(\varsigma_t, \rho_t^s, \mu_t) &= \min_{\gamma}\max_{a,\alpha,e}\{\underline{u}(a_t, e_t, \varsigma_t) + \gamma_t e_t + \mu_t\omega_t \\ &\quad + \beta_a(1-q)\int\sum_{\varsigma_{t+1}}\varphi(\varsigma_{t+1};\varsigma_t)W(\varsigma_{t+1}, \rho_{t+1}^s, \mu_{t+1})g(\pi;a_t)d\pi\} \end{aligned}$$

where

$$\omega_t = -\{[(1-q)a_t + qz_t] + \boxed{\frac{(1-\rho_t^s)}{\rho_t^s}}[(1-q)\alpha_t + qz_t]\}$$

subject to the IC constraint

$$\alpha_t = \begin{cases} Ae_t + B(\varsigma_t) & \text{if optimizing alternative type} \\ \underline{\alpha}(\varsigma_t) & \text{if mechanical alternative type} \end{cases}$$

with state dynamics allowing exogenous reputation ($\mathbf{1}_{endo} = 0$) or endogenous reputation ($\mathbf{1}_{endo} = 1$)

$$
\begin{aligned}
\mu_{t+1} &= \frac{\beta}{\beta_a(1-q)}\gamma_t\rho_t \\
\rho_{t+1} &= \mathbf{1}_{endo}\rho^s_{t+1} + (1-\mathbf{1}_{endo})\,\rho \\
\rho^s_{t+1} &= b(\pi_t, a_t, \alpha_t, \rho_t)
\end{aligned}
$$

The recursive program here is written in a general form that allows (i) optimizing or mechanical alternative type and (ii) endogenous or exogenous reputation. The program in Proposition 1 of the main text is a special form of (A36) where there is an optimizing alternative type and endogenous reputation. Hence, in that setting, there is no need to distinguish ρ^s from ρ.

A.11 A special case

If $q = 0$, $\beta_a = \beta$, and $\rho = 1$ always, our recursive program collapses to a textbook NK policy problem in recursive form. For example, in Clarida et al. (1999), the policymaker maximizes $E_0\sum_{t=0}^{\infty}\beta^t u(\pi_t, x_t)$ subject to $\pi_t = \kappa x_t + \beta E_t\pi_{t+1} + \varsigma_t$.

To create a dynamic Lagrangian one attaches $E_0\sum_{t=0}^{\infty}\beta^t\gamma_t[\pi_t - \kappa x_t - \beta E_t\pi_{t+1} - \varsigma_t]$ to the objective. The law of iterated expectation and rearrangement of terms allow this expression to be written as $E_0\sum_{t=0}^{\infty}\beta^t\{(\gamma_t - \gamma_{t-1})\pi_t - \gamma_t\kappa x_t - \gamma_t\varsigma_t\}$ with $\gamma_{-1} = 0$. Defining the pseudo state variable $\mu_t = \gamma_{t-1}$, the recursive optimization problem is

$$
W(\varsigma_t, \mu_t) = \min_{\gamma_t}\max_{\pi_t, x_t}\{u(\pi_t, x_t) + \gamma_t(\pi_t - \kappa x_t - \varsigma_t) - \mu_t\pi_t + \beta E_t W(\varsigma_{t+1}, \mu_{t+1})\}
$$

with $\mu_{t+1} = \gamma_t$ and $\mu_0 = 0$.

A.12 Forward-looking opportunistic type

Suppose that the opportunistic type of policymaker has a time discount factor β_α, this section derives an extended version of Proposition 1 for the committed policymaker's recursive optimization.

Given $h_{t+1} = [\varsigma_{t+1}, h_t^+] = [\varsigma_{t+1}, \pi_t, h_t]$, the optimization problem of the opportunistic type can be written as

$$
\alpha(h_t) = \arg\max_{\alpha_t} \underline{u}(\alpha_t, e_t, \varsigma_t) + \beta_\alpha(1-q)\int\sum_{\varsigma_{t+1}}\varphi(\varsigma_{t+1};\varsigma_t)V(\varsigma_{t+1}, \pi_t, h_t)\, g(\pi_t|\alpha_t)\, d\pi_t
$$

where $\underline{u}(\alpha_t, e_t, \varsigma_t) = \int u(\pi_t, e_t, \varsigma_t) g\left(\pi_t | \alpha_t\right) d\pi_t$. The FOC of α is

$$0 = \underline{u}_\alpha(\alpha_t, e_t, \varsigma_t) + \beta_\alpha(1-q) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) V\left(\varsigma_{t+1}, \pi_t, h_t\right) g_\alpha\left(\pi_t | \alpha_t\right) d\pi_t \tag{A37}$$

when α_t is evaluated at optimal $\alpha\left(h_t\right)$.

The implied value function is

$$V\left(h_t\right) = \underline{u}(\alpha_t, e_t, \varsigma_t) + \beta_\alpha(1-q) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) V\left(h_{t+1}\right) g\left(\pi_t | \alpha_t\right) d\pi_t \tag{A38}$$

when α_t is evaluated at optimal $\alpha\left(h_t\right)$.

A.12.1 Recursive formulation of the sequential rationality constraints

The optimal path of (α_t, V_t) chosen by the committed policymaker needs to satisfy the two sequential rationality constraints (A37) and (A38). Similar to the rational expectation constraint, we construct the Lagrangian components of the sequential rationality constraints using the committed type's probabilities as weights on the multipliers, where $p(h_{t+j+1}) = p\left(h_{t+j}\right) \varphi(\varsigma_{t+j+1}; \varsigma_{t+j}) g(\pi_{t+j} | a_{t+j})$.

The Lagrangian component for the FOC of α is

$$\begin{aligned}
& \Lambda_t \\
= \; & \sum_{j=0}^{\infty} (\beta_a(1-q))^j \sum_{h_{t+j}} \frac{p(h_{t+j})}{p(h_t)} \phi(h_{t+j}) [\underline{u}_\alpha(\alpha_t, e_t, \varsigma_t)] \\
& + \sum_{j=0}^{\infty} (\beta_a(1-q))^j \sum_{h_{t+j}} \frac{p(h_{t+j})}{p(h_t)} \phi(h_{t+j}) [\beta_\alpha(1-q) \int \sum_{\varsigma_{t+j+1}} \varphi(\varsigma_{t+j+1}; \varsigma_{t+j}) V\left(h_{t+j+1}\right) g_\alpha\left(\pi_{t+j} | \alpha_{t+j}\right) d\pi_{t+j}] \\
= \; & \sum_{j=0}^{\infty} (\beta_a(1-q))^j \sum_{h_{t+j}} \frac{p(h_{t+j})}{p(h_t)} \phi(h_{t+j}) [\underline{u}_\alpha(\alpha_t, e_t, \varsigma_t)] \\
& + \sum_{j=0}^{\infty} (\beta_a(1-q))^{j+1} \sum_{h_{t+j+1}} \frac{p(h_{t+j+1})}{p(h_t)} \left[\frac{\beta_\alpha}{\beta_a} \phi(h_{t+j}) \frac{g_\alpha\left(\pi_{t+j} | \alpha_{t+j}\right)}{g\left(\pi_{t+j} | a_{t+j}\right)} \right] V\left(h_{t+j+1}\right) \\
= \; & \phi(h_t) \; \underline{u}_\alpha(\alpha_t, e_t, \varsigma_t) + \frac{\beta_\alpha}{\beta_a} \phi(h_{t-1}) \frac{b_\alpha\left(\pi_{t-1} | \alpha_{t-1}\right)}{g\left(\pi_{t-1} | a_{t-1}\right)} V\left(h_t\right) + \beta_1(1-q) E_t^c \Lambda_{t+1}
\end{aligned}$$

with the understanding that $\phi_{-1} = 0$.

Note that there are two conceptual elements of the term $\frac{g_\alpha(\pi_{t-1}|\alpha_{t-1})}{g(\pi_{t-1}|a_{t-1})}$, which we can separate by writing this term as $\frac{g_\alpha(\pi_{t-1}|\alpha_{t-1})}{g(\pi_{t-1}|\alpha_{t-1})} \frac{g(\pi_{t-1}|\alpha_{t-1})}{g(\pi_{t-1}|a_{t-1})} = \kappa_t \lambda_t$. One is the change of measure necessary to transform the opportunistic type's probabilities into those of the committed type, i.e., λ_t using the notation established above. The other arises from the implication that a marginally higher α has on the likelihood of a specific inflation history, captured by $\kappa_t = \frac{g_\alpha(\pi_{t-1}|\alpha_{t-1})}{g(\pi_{t-1}|\alpha_{t-1})}$. Accordingly, we have a convenient recursive expression, which includes a lagged multiplier, for this Lagrangian component

$$\Lambda_t = \phi_t \; \underline{u}_\alpha(\alpha_t, e_t, \varsigma_t) + \frac{\beta_\alpha}{\beta_1} \phi_{t-1}\kappa_t\lambda_t V_t + \beta_a(1-q)E_t^c\Lambda_{t+1}. \tag{A39}$$

The understanding is that the initial condition is $\phi_{-1} = 0$.

The Lagrangian component for the value function of the opportunistic type is

$$\begin{aligned}
&\Theta_t \\
= &\sum_{j=0}^{\infty}(\beta_a(1-q))^j \sum_{h_{t+j}} \frac{p(h_{t+j})}{p(h_t)}\chi(h_{t+j})\left[\underline{u}(\alpha_t, e_t, \varsigma_t) - V\left(h_{t+j}\right)\right] \\
&+\sum_{j=0}^{\infty}(\beta_a(1-q))^j \sum_{h_{t+j}} \frac{p(h_{t+j})}{p(h_t)}\chi(h_{t+j})[\beta_\alpha(1-q)\int \sum_{\varsigma_{t+j+1}} \varphi(\varsigma_{t+j+1};\varsigma_{t+j})V\left(h_{t+j+1}\right) g\left(\pi_{t+j}|\alpha_{t+j}\right) d\pi_{t+j}] \\
= &\sum_{j=0}^{\infty}(\beta_a(1-q))^j \sum_{h_{t+j}} \frac{p(h_{t+j})}{p(h_t)}\chi(h_{t+j})\left[\underline{u}(\alpha_t, e_t, \varsigma_t) - V\left(h_{t+j}\right)\right] \\
&+\sum_{j=0}^{\infty}(\beta_a(1-q))^{j+1} \sum_{h_{t+j}} \frac{p(h_{t+j+1})}{p(h_t)}\left[\frac{\beta_\alpha}{\beta_a}\chi(h_{t+j})\frac{g\left(\pi_{t+j}|\alpha_{t+j}\right)}{g\left(\pi_{t+j}|a_{t+j}\right)}\right] V\left(h_{t+j+1}\right) \\
= &\;\chi(h_t)\; \underline{u}(\alpha_t, e_t, \varsigma_t) - \chi(h_t)V\left(h_t\right) + \frac{\beta_\alpha}{\beta_1}\chi(h_{t-1})\frac{g\left(\pi_{t-1}|\alpha_{t-1}\right)}{g\left(\pi_{t-1}|a_{t-1}\right)}V\left(h_t\right) + \beta_a(1-q)E_t^c\Theta_{t+1}
\end{aligned}$$

Note that the change of measure applies to this value recursion, so that we can replace $\frac{g(\pi_{t-1}|\alpha_{t-1})}{g(\pi_{t-1}|a_{t-1})}$ with λ_t using our notation from above. We have a convenient recursive expression for this Lagrangian component

$$\Theta_t = \chi_t \; \underline{u}(\alpha_t, e_t, \varsigma_t) - \chi_t V_t + \frac{\beta_\alpha}{\beta_a}\chi_{t-1}\lambda_t V_t + \beta_a(1-q)E_t^c\Theta_{t+1} \tag{A40}$$

which also contains a lagged multiplier. The understanding is that the initial condition is $\chi_{-1} = 0$.

A.12.2 The basic recursive specification

The preceding derivations suggest a recursive version of $U_t + \Psi_t + \Lambda_t + \Theta_t$ with states $(\varsigma_t, \gamma_{t-1}, \rho_{t-1}, \phi_{t-1}, \chi_{t-1}, \lambda_t, \kappa_t)$. This is

$$
\begin{aligned}
& W(\varsigma_t, \gamma_{t-1}, \rho_{t-1}, \phi_{t-1}, \chi_{t-1}, \lambda_t, \kappa_t) \\
= & \min_{\gamma,\phi,\chi} \max_{a,\alpha,e,V} \{\underline{u}(a_t, e_t, \varsigma_t) \\
& +\gamma_t e_t - \frac{\beta}{\beta_a(1-q)}\gamma_{t-1}[\rho_{t-1}((1-q)a_t + qz_t) + (1-\rho_{t-1})\lambda_t((1-q)\alpha_t + qz_t)] \\
& +\phi_t \ \underline{u}_\alpha(\alpha_t, e_t, \varsigma_t) + \chi_t \underline{u}(\alpha_t, e_t, \varsigma_t) - \chi_t V_t \\
& +\frac{\beta_\alpha}{\beta_a}\left[\phi_{t-1}\kappa_t\lambda_t + \chi_{t-1}\lambda_t\right] V_t \\
& +\beta_a(1-q)\int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1};\varsigma_t) W(\varsigma_{t+1}, \gamma_t, \rho_t, \phi_t, \chi_t, \lambda_{t+1}, \kappa_{t+1}) g(\pi_t|a_t) d\pi_t\}
\end{aligned}
$$

with state dynamics (from the perspective of the committed type)

$$
\begin{aligned}
\rho_t &= \frac{\rho_{t-1}}{\rho_{t-1} + (1-\rho_{t-1})\lambda_t} \text{ given } \rho_0 \\
\lambda_{t+1} &= \lambda(\pi_t, a_t, \alpha_t) \text{ with probability } g(\pi_t|a_t) \\
\kappa_{t+1} &= \kappa(\pi_t, a_t, \alpha_t) \text{ with probability } g(\pi_t|a_t)
\end{aligned}
$$

allowing for direct adjustment for the lagged multiplier psuedo states γ, ϕ, χ from initial conditions $\gamma_{-1} = 0, \phi_{-1} = 0$ and $\chi_{-1} = 0$.

A.12.3 State reduction

Using the same technique in Section A.9, $(\eta_t, \rho_{t-1}, \lambda_t)$ can be reduced to $\mu_t = \frac{\beta}{\beta_a(1-q)}\gamma_{t-1}\rho_{t-1}$ and ρ_t. Further note that ϕ_{t-1}, χ_{t-1} and κ_t enter only in the bracketed expression on the fourth line of the basic recursive specification, so we can define $y_t = \frac{\beta_\alpha}{\beta_1}\left[\phi_{t-1}\kappa_t\lambda_t + \chi_{t-1}\lambda_t\right]$ to consolidate states.

These observations establish the following extended version of our Proposition 1.

Proposition 1. Given $z(\varsigma_t, \rho_t)$, the within-regime equilibrium is the solution to the following recursive optimization problem

$$
\begin{aligned}
W(\varsigma_t, \mu_t, \rho_t, y_t) = & \min_{\gamma,\phi,\chi} \max_{a,\alpha,e,V} \{\underline{u}(a_t, e_t, \varsigma_t) \\
& +\gamma_t e_t \mu_t[(1-q)a_t + qz_t + \frac{(1-\rho_t)}{\rho_t}[(1-q)\alpha_t + qz_t] \\
& +\phi_t \ \underline{u}_\alpha(\alpha_t, e_t, \varsigma_t) + \chi_t \underline{u}(\alpha_t, e_t, \varsigma_t) + (y_t - \chi_t)V_t \\
& +\beta_a(1-q)\int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1};\varsigma_t) W(\varsigma_{t+1}, \mu_{t+1}, \rho_{t+1}, y_{t+1}) g(\pi_t|a_t) d\pi_t\}
\end{aligned}
$$

having state dynamics

$$\begin{aligned}
\mu_{t+1} &= \frac{\beta}{\beta_a(1-q)}\gamma_t\rho_t \text{ with } \mu_0 = 0 \\
\rho_{t+1} &= b(\pi_t, a_t, \alpha_t, \rho_t) \text{ with probability } g(\pi_t|a_t) \\
y_{t+1} &= \frac{\beta_\alpha}{\beta_a}\left[\phi_t \frac{g_\alpha(\pi_t|\alpha_t)}{g(\pi_t|a_t)} + \chi_t \frac{g(\pi_t|\alpha_t)}{g(\pi_t|a_t)}\right] \text{ with } y_0 = 0 \text{ and probability } g(\pi_t|a_t)
\end{aligned}$$

B Consolidation

This appendix explains how to simplify the recursive program in Proposition 1 to the one in Proposition 2, via the implications of private sector's rational expectation constraint.

B.1 Proof of Lemma 1: Relationship between U and W

If $W(\cdot)$ in (A36) is differentiable, there are two notable implications of this structure.

The **envelope theorem implication** for μ is

$$W_\mu(\varsigma_t, \rho_t^s, \mu_t) = -\{[(1-q)a_t + qz_t] + \boxed{\frac{(1-\rho_t^s)}{\rho_t^s}}[(1-q)\alpha_t + qz_t]\} = \omega_t$$

The **first order necessary condition for** γ_t is

$$\begin{aligned}
0 &= e_t + \beta_a(1-q)\sum_{\varsigma_{t+1}}\varphi(\varsigma_{t+1};\varsigma_t)\int W_\mu(\varsigma_{t+1}, \rho_{t+1}, \mu_{t+1})\frac{\partial\mu_{t+1}}{\partial\gamma_t}g(\pi_t|a_t)d\pi_t \\
&= e_t + \beta\sum_{\varsigma_{t+1}}\varphi(\varsigma_{t+1};\varsigma_t)\int W_\mu(\varsigma_{t+1}, \rho_{t+1}, \mu_{t+1})\rho_t g(\pi_t|a_t)d\pi_t
\end{aligned}$$

where the state evolution equation (A34) implies $\partial\mu_{t+1}/\partial\gamma_t = \rho_t\beta/(\beta_a(1-q))$.

When combined with an updated version of the envelope theorem implication, this FOC recovers the private sector's rational expectation constraint as in (A27):

$$e_t = \beta\int\sum_{\varsigma_{t+1}}\varphi(\varsigma_{t+1};\varsigma_t)\left[[(1-q)a_{t+1} + qz_{t+1}] + \boxed{\frac{(1-\rho_{t+1}^s)}{\rho_{t+1}^s}}[(1-q)\alpha_{t+1} + qz_{t+1}]\right]\rho_t g(\pi_t|a_t)d\pi_t$$

where

$$\frac{1-\rho_{t+1}^s}{\rho_{t+1}^s} = \frac{(1-\rho_t)\lambda_{t+1}}{\rho_t}.$$

Hence, in equilibrium where the rational expectation constraint must hold, we obtain the following relationship between the value function $W(\cdot)$ and the optimized objective $U^*(\cdot)$: $W(\varsigma_t, \rho_t^s, \mu_t) - \mu_t\omega_t^* = U^*(\varsigma_t, \rho_t^s, \mu_t)$, where $\omega_t^* = -\{[(1-q)a_t^* + qz_t^*] + \boxed{\frac{(1-\rho_t^s)}{\rho_t^s}}[(1-q)\alpha_t^* + qz_t^*]\}$.

To see why it holds, notice in equilibrium

$$W(\varsigma_t, \rho^s_t, \mu_t) - \mu_t \omega^*_t = \underline{u}(a^*_t, e^*_t, \varsigma_t) + \gamma_t e^*_t + \beta_a(1-q) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) W(\varsigma_{t+1}, \rho^s_{t+1}, \mu_{t+1}) g(\pi_t | a^*_t) d\pi_t \tag{B1}$$

Suppose $W(\varsigma_{t+1}, \rho^s_{t+1}, \mu_{t+1}) = \mu_{t+1}\omega^*_{t+1} + U^*(\varsigma_{t+1}, \rho^s_{t+1}, \mu_{t+1})$, the right hand side can be written as

$$\begin{aligned}
&\underline{u}(a^*_t, e^*_t, \varsigma_t) + \gamma_t e^*_t + \beta_a(1-q) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) [\frac{\beta}{\beta_a(1-q)} \gamma_t \rho_t \omega^*_{t+1}] g(\pi_t | a^*_t) d\pi_t \\
&+ \beta_a(1-q) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) U^*(\varsigma_{t+1}, \rho^s_{t+1}, \mu_{t+1}) g(\pi_t | a^*_t) d\pi_t \\
=&\underline{u}(a^*_t, e^*_t, \varsigma_t) + \gamma_t [e^*_t + \beta \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) \omega^*_{t+1} \rho_t g(\pi_t | a^*_t) d\pi_t] \\
&+ \beta_a(1-q) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) U^*(\varsigma_{t+1}, \rho^s_{t+1}, \mu_{t+1}) g(\pi_t | a^*_t) d\pi_t \\
=&\underline{u}(a^*_t, e^*_t, \varsigma_t) + \beta_a(1-q) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) U^*(\varsigma_{t+1}, \rho^s_{t+1}, \mu_{t+1}) g(\pi_t | a^*_t) d\pi_t \\
=&U^*(\varsigma_t, \rho^s_t, \mu_t)
\end{aligned}$$

which implies $W(\varsigma_t, \rho^s_t, \mu_t) - \mu_t \omega^*_t = U^*(\varsigma_t, \rho^s_t, \mu_t)$.

B.2 Proof of Lemma 2: Further consolidation

We now show that imposing the rational expectation constraint (A27) on the choice of e_t implies Lemma 2, which allows us to further reduce the recursive program in Proposition 1 to the one in Proposition 2. The key idea is that only the policy difference $\delta = a - \alpha$ matters rather than the levels of a and α.

Recall that (A27) comes from (A23) before undertaking a "change of measure". So the original form of the rational expectation constraint on e_t is:

$$\begin{aligned}
e_t = \beta \rho_t \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) \left[(1-q)\, a_{t+1} + q z_{t+1}\right] g(\pi_t | a_t) d\pi_t \\
+ \beta\,(1-\rho_t) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) \left[(1-q)\, \alpha_{t+1} + q z_{t+1}\right] g(\pi_t | \alpha_t) d\pi_t
\end{aligned} \tag{B2}$$

with a_{t+1}, α_{t+1}, and z_{t+1} determined by the three states $(\varsigma_{t+1}, \rho_{t+1}, \mu_{t+1})$ through the equilibrium strategies: $a^*(\cdot)$, $\alpha^*(\cdot)$, and $z^*(\cdot)$.

Recall $\rho_{t+1} = b(\pi_t, a_t, \alpha_t, \rho_t)$ from (A35) and $b(\cdot)$ is the Bayes' learning rule specified in (A19). The inflation distribution is $\pi = a + \varepsilon_1$ under the committed type and $\pi = \alpha + \varepsilon_2$ under the opportunistic type, with ε_1 and ε_2 being zero mean random variables. We can therefore rewrite the Bayes' learning rule (A19) as

$$\text{(B3)} \qquad \begin{aligned} \rho_{t+1} &= \frac{\phi_1(\pi_t - a_t)\rho_t}{\phi_1(\pi_t - a_t)\rho_t + \phi_2(\pi_t - \alpha_t)(1-\rho_t)} \\ &\equiv b(\pi_t - a_t, \pi_t - \alpha_t, \rho_t) \end{aligned}$$

where $g(\pi|a) = \phi_1(\pi - a)$ and $g(\pi|\alpha) = \phi_2(\pi - \alpha)$, and the b function is a version of our general convenient short-hand which is identified by its three argument nature.

Then, in terms of the policy difference $\delta = a - \alpha$, future reputation is

$$\text{(B4)} \qquad \rho' = b(\varepsilon_1, \varepsilon_1 + \delta, \rho) \text{ conditional on } \tau_a$$

$$\text{(B5)} \qquad \rho' = b(\varepsilon_2 - \delta, \varepsilon_2, \rho) \text{ conditional on } \tau_\alpha$$

Replacing $g(\pi|a)$ and $g(\pi|\alpha)$ in (B2) with $\phi_1(\pi - a)$ and $\phi_2(\pi - \alpha)$, ρ_{t+1} with (B4) and (B5), and realizing choosing γ_t is equivalent to choosing μ_{t+1} due to $\mu_{t+1} = \frac{\beta}{\beta_a(1-q)}\gamma_t\rho_t$, we obtain Lemma 2 with the added details as follows:

Lemma 2. Given (ς, ρ), and that future policymakers follow the equilibrium strategies $a^*(\varsigma', \rho', \mu')$, $\alpha^*(\varsigma', \rho', \mu')$ and $z^*(\varsigma', \rho')$, rationally expected inflation is uniquely determined by the contemporaneous policy difference $\delta = a - \alpha$, and the future pseudo-state variable μ'.

$$\begin{aligned} e &= e(\delta, \mu'; \varsigma, \rho) = \beta\rho \int \widehat{M}_1(\varsigma, b(\varepsilon_1, \varepsilon_1 + \delta, \rho), \mu')\phi_1(\varepsilon_1)\, d\varepsilon_1 + \\ &\quad \beta(1-\rho)\int \widehat{M}_2(\varsigma, b(\varepsilon_2 - \delta, \varepsilon_2, \rho), \mu')\phi_2(\varepsilon_2)\, d\varepsilon_2; \end{aligned}$$

where $\phi_1(\cdot)$ and $\phi_2(\cdot)$ denote the density functions of ε_1 and ε_2;

$$\begin{aligned} \widehat{M}_1(\varsigma, \rho', \mu') &:= \sum_{\varsigma'} \varphi(\varsigma'; \varsigma)\left[(1-q)\, a^*(\varsigma', \rho', \mu') + q z^*(\varsigma', \rho')\right]; \\ \widehat{M}_2(\varsigma, \rho', \mu') &:= \sum_{\varsigma'} \varphi(\varsigma'; \varsigma)\left[(1-q)\, \alpha^*(\varsigma', \rho', \mu') + q z^*(\varsigma', \rho')\right]; \end{aligned}$$

Lemma 1 and 2 enable us to simplify the recursive program in Proposition 1, moving from choosing (γ, a, α, e) to merely choosing (δ, μ'). More specifically, once $e(\delta, \mu'; \varsigma, \rho)$ is chosen via the choices of (δ, μ'), we can obtain α from $Ae + B(\varsigma)$, and a from $\alpha + \delta$.

Furthermore, the relationship between U and W specified in (B1) implies that the objective of the recursive optimization can be reduced to

$$\underline{u}(a, e, \varsigma) + \mu\omega(a, \alpha) + \beta_a(1-q)\int \sum_{\varsigma'} \varphi(\varsigma';\varsigma)U^*(\varsigma', \rho', \mu')g(\pi|a)d\pi$$

where $U^*(\varsigma, \rho, \mu) = W(\varsigma, \rho, \mu) - \mu\omega(a^*, \alpha^*)$.

Replacing (e, α, a) in $\underline{u}(\cdot)$ and $\omega(\cdot)$ with $e(\delta, \mu'; \varsigma, \rho)$, $Ae + B(\varsigma)$, and $\alpha + \delta$ makes u and ω only depend on (δ, μ'):

$$\text{(B6)} \qquad \underline{\underline{u}}(\delta, \mu') := \underline{u}(Ae + B(\varsigma), e, \varsigma)$$

$$\text{(B7)} \qquad \underline{\underline{\omega}}(\delta, \mu') := -\frac{1}{\rho}\left[(1-q)\left(Ae + B(\varsigma)\right) + qz^*\left(\varsigma, \rho\right)\right] - (1-q)\,\delta$$

where $e = e(\delta, \mu'; \varsigma, \rho)$. Replacing ρ' in $U^*(\cdot)$ with (B4) and $g(\pi|a)$ with $\phi_1(\varepsilon_1)$, we then arrive at the recursive program in Proposition 2.

C Forecasting Functions and Matching the SPF

C.1 SPF Data

We construct the SPF inflation data from "individual responses" file for the *level* of the GDP deflator available at https://www.philadelphiafed.org/surveys-and-data/pgdp. The sample starts from the fourth quarter of 1968.

In the middle of each quarter, each survey participant submits a forecast for the price level in that quarter and the next four. We first calculate inflation forecasts for each individual forecaster j, using the continuously compounded growth rate: $400 \times \ln(P^j_{t+k|t}/P^j_{t+k-1|t})$. We then take the median of these inflation forecasts.

Alternatively, one can use the summary data files constructed by the Federal Reserve Bank of Philadelphia, particularly the "annualized percent change of median responses" file from https://www.philadelphiafed.org/surveys-and-data/pgdp, as a measure for the SPF inflation data. This file includes an inflation "nowcast" and forecasts at the 1,2,3, and 4 quarter horizons. The nature of these inflation series is explained by Stark (2010). The FRBP first constructs a median price level for each horizon from "individual responses", say $P_{t+k|t}$ for k=0,1,...4. It then constructs an annualized percentage growth rate using the formula $100 \times ([P_{t+k|t}/P_{t+k-1|t}]^4 - 1)$.

Our procedure yields time series that are less prone to transitory outliers than the standard FRBP constructions. Each difference matters, i.e., (i) the median of the inflation rates is less prone than is the change in the median price level; and (ii) the continuously compounded inflation rate is less prone than is the FRBP inflation rate.

Figure 11 contrasts the two measures.

Figure 11: Contrasting median inflation and change in median price

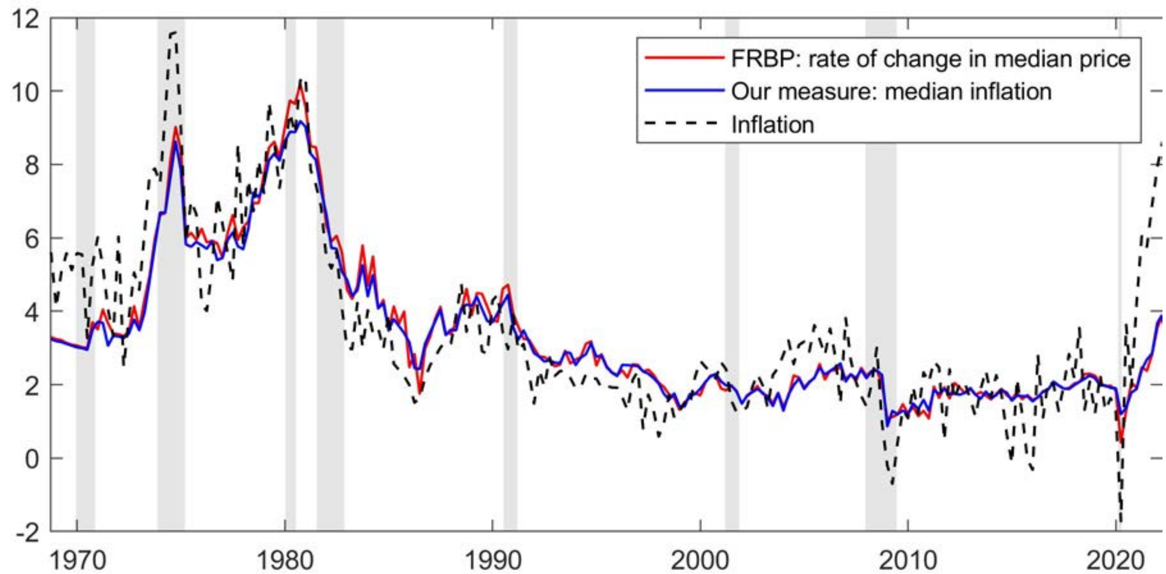

C.2 Recursive forecasting in our theory

The SPF contains multiperiod forecasts of inflation. Real and nominal interest rates contain multiperiod forecasts of output and inflation. This appendix describes the calculation of such forecasts. We specialize the inflation distributions to

$$\pi_t = a_t + \sigma\varepsilon_t \quad \text{and} \quad \pi_t = \alpha_t + \sigma\varepsilon_t \tag{C1}$$

with a density $\phi(\varepsilon)$ compatible with a zero mean and a unit standard deviation such as the standard normal.[4]

The information set is assumed to be the start of period information of the private sector, $(\varsigma_t, \rho_t, \mu_t)$. Generally, our approach is applicable to forecasting any variable v_{t+k} which has

[4]The recipe allows for type dependent parameters σ_1 and σ_2 but we use the common σ assumption for simplicity in this discussion.

a functional solution

$$v(\varsigma_t, \rho_t, \mu_t)$$

that is known to private agents and our specific applications are to inflation and output.

C.2.1 Forecasting inflation

Let us start with forecasting inflation k steps ahead, which we denote $f_{t+k|t}$.[5] Private agents know the intended inflation functions of the two policymakers:

$$a(\varsigma_t, \rho_t, \mu_t)$$
$$\alpha(\varsigma_t, \rho_t, \mu_t)$$

Accordingly, given that implementation errors have mean zero, the private sector "nowcast" of inflation is

$$f_{t|t} = f(\varsigma_t, \rho_t, \mu_t, 0) = \rho_t a(\varsigma_t, \rho_t, \mu_t) + (1 - \rho_t)\alpha(\varsigma_t, \rho_t, \mu_t)$$

Utilizing the law of iterated expectation, today's forecast of π_{t+j} is today's forecast of tomorrow's forecast of π_{t+j}. We can compute multistep forecasts of inflation recursively building up $f_{t+j|t}$ from $f_{t+j|t+1}$:

$$f_{t+j|t} = f(\varsigma_t, \rho_t, \mu_t, j) = E_t(f_{t+j|t+1}) = E_t[f(\varsigma_{t+1}, \rho_{t+1}, \mu_{t+1}, j-1)] \tag{C2}$$

The state variables ρ_{t+1} and μ_{t+1} evolve as follows.

With probability $1 - q$ there is no regime change. The pseudo-state variable is evolves according to:

$$\mu_{t+1} = \mu'^{*}(\varsigma_t, \rho_t, \mu_t).$$

The reputation state variable ρ_{t+1} evolves according to:

$$\rho_{t+1} = b(a_t + \sigma\varepsilon_t, a_t, \alpha_t, \rho_t) \qquad \text{with prob } \rho_t$$
$$\rho_{t+1} = b(\alpha_t + \sigma\varepsilon_t, a_t, \alpha_t, \rho_t) \quad \text{with prob } 1 - \rho_t$$

With probability q, there is a regime change in which case $\mu_0 = 0$ and the new policymaker's initial reputation ρ_0 is a random draw from the distribution $\Xi(\rho_0|\rho_{t+1})$, where ρ_{t+1} is

[5] The model solution already contains a one-step ahead forecast for inflation as a function of the state, i.e, $f_{t+1|t} = f(\varsigma_t, \rho_t, \mu_t, 1) = e^{*}(\varsigma_t, \rho_t, \mu_t)/\beta$. Our concern here is longer-term inflation.

what the reputation would have been if there was no replacement. Then, we can determine

$$\text{(C3)} \qquad \begin{aligned} & f_{t+j|t} = f(\varsigma_t, \rho_t, \mu_t, j) = \sum \varphi(\varsigma_{t+1}; \varsigma_t)\{ \\ & (1-q)\rho_t \int f[\varsigma_{t+1}, b(a_t + \sigma\varepsilon, a_t, \alpha_t, \rho_t), \mu_{t+1}, j-1]\phi(\varepsilon)d\varepsilon \\ & +(1-q)(1-\rho_t) \int f[\varsigma_{t+1}, b(\alpha_t + \sigma\varepsilon, a_t, \alpha_t, \rho_t), \mu_{t+1}, j-1]\phi(\varepsilon)d\varepsilon \\ & +q\rho_t \int \{ \int f[\varsigma_{t+1}, \rho_0, 0, j-1] d\Xi(\rho_0 | b(a_t + \sigma\varepsilon, a_t, \alpha_t, \rho_t))\}\phi(\varepsilon)d\varepsilon \\ & +q(1-\rho_t) \int \{ \int f[\varsigma_{t+1}, \rho_0, 0, j-1] d\Xi(\rho_0 | b(\alpha_t + \sigma\varepsilon, a_t, \alpha_t, \rho_t))\}\phi(\varepsilon)d\varepsilon \} \end{aligned}$$

C.2.2 Forecasting output

We now turn to forecasting output, determined by

$$x_t = \frac{1}{\kappa}[\pi_t - \beta f(\varsigma_t, \rho_t, \mu_t, 1) - \varsigma_t]$$

so that a "nowcast" of output is

$$\widehat{x}_0(\varsigma_t, \rho_t, \mu_t) = \frac{1}{\kappa}[f(\varsigma_t, \rho_t, \mu_t, 0) - \beta f(\varsigma_t, \rho_t, \mu_t, 1) - \varsigma_t)]$$

Hence, we can use the same recipe for multistep forecasts:

$$\widehat{x}_{j+1}(\varsigma_t, \rho_t, \mu_t) = E_t[\widehat{x}_j(\varsigma_{t+1}, \rho_{t+1}, \mu_{t+1})]$$

recursively building up $\widehat{x}_{j+1}$ from $\widehat{x}_j$.

C.3 Matching the SPF: motivation and mechanics

From the standpoint of modern econometrics, our theory is a very simple one that is easily rejected: conditional on regime change dates and the identification of policymaker type within each regime: we have just two random inputs – price shocks ς_t and implementation errors ε_t – that drive many observable macro time series. To review, there are three state variables $s_t = [\varsigma_t, \rho_t, \mu_t]$, governed by a Markov process with a special form

$$\begin{aligned} \varsigma_t &= \nu\varsigma_{t-1} + \xi_t \\ \rho_{t+1} &= b(\pi_t, a^*(s_t), \alpha^*(s_t), \rho_t) \\ \mu_{t+1} &= \mu'^*(\varsigma_t, \rho_t, \mu_t) \end{aligned}$$

Many variables depend just on these states, including the policies a_t and α_t and, as we just discussed, expectations at various horizons $f_{t+k|t}$. Others, including inflation π_t and real activity x_t, also depend on ε_t.

Our work in this paper is quantitative theory and, following early RBC analyses, we fix model parameters and use a transparent strategy for extracting the unobserved states. Then, with the states in hand, we calculate the historical behavior of observables.[6] But the literature has stressed that one of the difficulties with this RBC strategy is that the technology state is measured by the Solow residual, which is based on observable variables (output, capital, and labor) whose behavior is ultimately to be explored.

C.3.1 The strategy for extracting states

We therefore develop a strategy for extracting state information that does not use the behavior of the GDP deflator. It relies on the fact that our model provides a mapping between states and inflation expectations at various horizons:

$$f_{t+k|t} = f(\varsigma_t, \rho_t, \mu_t, k).$$

Since the pseudo state μ_t is predetermined, we can solve for $\widehat{\varsigma}_t$ and $\widehat{\rho}_t$ from two elements of the SPF term structure, if we identify model expectations at horizon k with the k-quarter-ahead SPF inflation forecast.

With the date t extracted states $\widehat{\varsigma}_t$ and $\widehat{\rho}_t$ and the predetermined state $\widehat{\mu}_t$ in hand, we can create $\widehat{\mu}_{t+1}$ using the third transition rule, $\mu'^*(\cdot)$, continuing recursively to calculate a full history of states. It is true that we need an initial condition on μ, but that is supplied by specifying a set of regime switch dates at which μ is set zero.

In the main text, we use $\mu'^*(0, \widehat{\rho}_t, \widehat{\mu}_t)$ to determine $\widehat{\mu}_{t+1}$ recursively, instead of $\mu'^*(\widehat{\varsigma}_t, \widehat{\rho}_t, \widehat{\mu}_t)$. We do so because it is a natural way to preserve the mean-reverting property of ς shock in the extract $\widehat{\varsigma}_t$ series. We nonetheless redo our quantitative fitting exercise with a version of state extraction using $\widehat{\mu}_{t+1} = \mu'^*(\widehat{\varsigma}_t, \widehat{\rho}_t, \widehat{\mu}_t)$. The model's fitting to the U.S. inflation is similar to that reported in the main text. The results are reported in Section C.4.

[6] Prescott (1986) constructs Solow residuals as productivity indicators and then calculates moment implications for many variables of a model with calibrated parameters. Our work is closer to Plosser (1989), who uses the Solow residual time series and a basic calibrated model to construct time series of many variables, including consumption, investment and so on.

Figure 12: Model-implied and SPF forecasts of inflation

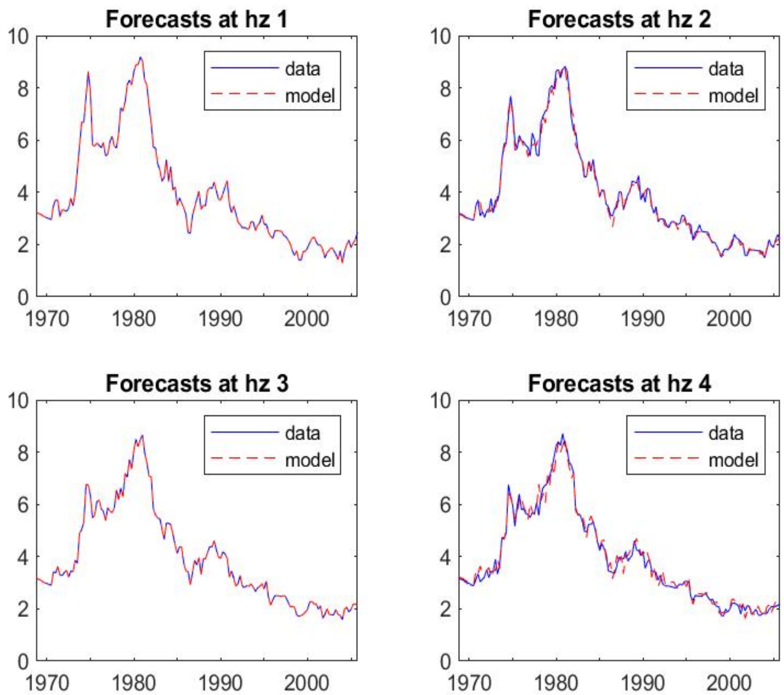

C.3.2 Application and fitting performance

As discussed in section 5.3 of the main text, we extract latent states by matching model-implied inflation forecasts at horizons 1 and 3 with SPF one-quarter-ahead and three-quarter-ahead forecasts. The left panels in Figure 12 shows our match is nearly perfect given that we choose two state variables ς and ρ each period to match two data points SPF1Q and SPF3Q. Using the extracted states, we can also compute model-implied inflation forecasts at horizons 2 and 4, and compare them with SPF two-quarter-ahead and four-quarter-ahead forecasts. The comparison is shown in the right panels of Figure 12. It is notable that our model-implied forecasts lie almost entirely on top of the SPF data for both forecasting horizons, which are not explicitly targeted. We view this figure as evidence in support of our state extraction approach.

C.4 Results without imposing mean-reverting on extracted $\widehat{\varsigma}$

The results in the main text are based on using a decision rule $\widehat{\mu}' = \mu'^{*}(0, \widehat{\rho}, \widehat{\mu})$ rather than $\widehat{\mu}' = \mu'^{*}(\widehat{\varsigma}, \widehat{\rho}, \widehat{\mu})$. That is, we do not allow the extracted shock to influence the dynamics of the pseudo state variable. Figure 13 displays the results when we alternatively allow this influence. The main messages from the text are maintained.

Figure 13: Inflation history and model-implied policies

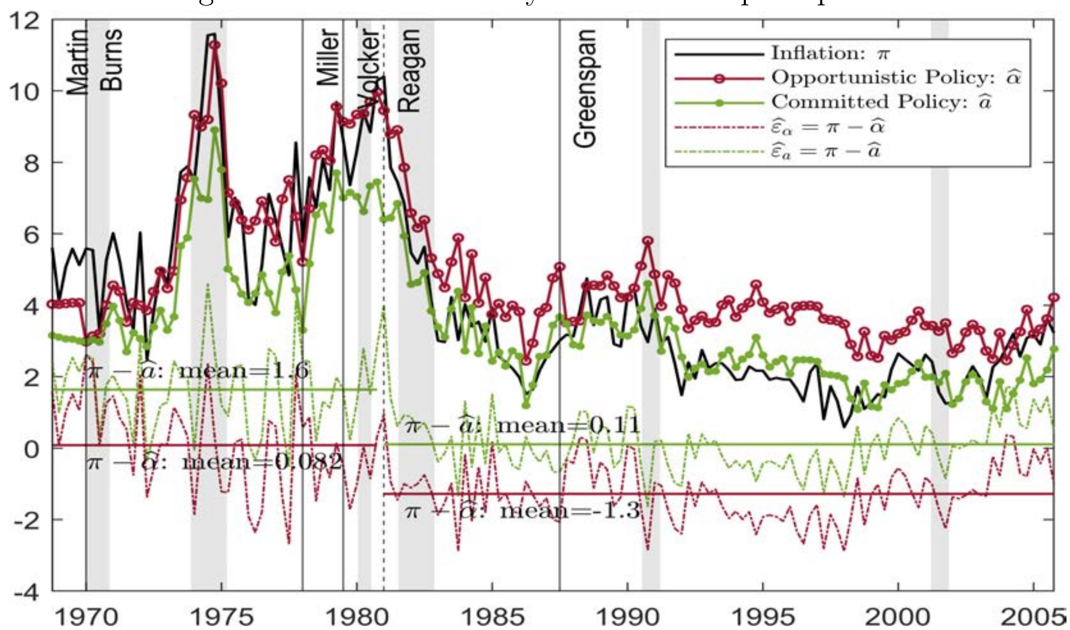

Note: Model-implied policies are based on extracted states produced using $\widehat{\mu}' = \mu'^*(\widehat{\varsigma}, \widehat{\rho}, \widehat{\mu})$ recursively.

C.5 Shock Comparisons

We have explored four indicators of the cost-push shock. First, there is a food and and energy shock constructed along the lines of Watson (2014). Second, there is the SPF spread. Third, there is the extracted shock series from the main text. Fourth, there is the extracted shock using the procedure that we just discussed. Figure 14 displays these alternative series. Note first that all measures rise dramatically during the famous "oil price shock" of late 1973 and early 1974 and also during the late 1970s interval that preceded Volcker's appointment. Note next that the extracted shocks and the SPF spread are more persistent during the earlier episode. Contemporary sources, such as the January 1975 Economic Report of the President prepared by Alan Greenspan and his CEA colleagues, point to other price shocks in addition to oil during the preceding year. Econometric studies such as those of R.J. Gordon (2013) and Watson (2014) estimate price shocks, including those from price decontrols in the 1970s, of more lasting form. So, on this basis, we are led to prefer extracted shocks as a parsimonious approach. Note further that the extracted shocks depart from each other toward the end of the period, which is the motivation for us to adopt the extraction strategy employed in the main text rather than that discussed in the prior section. Our extraction procedure is a straightforward and transparent way to induce the extracted price shocks to be mean-reverting, but does not explicitly impose the requirement that extracted price shocks are stationary. More sophisticated methods, applicable to hidden Markov models such as ours, would impose that requirement.

Figure 14: Various indicators for cost-push shocks

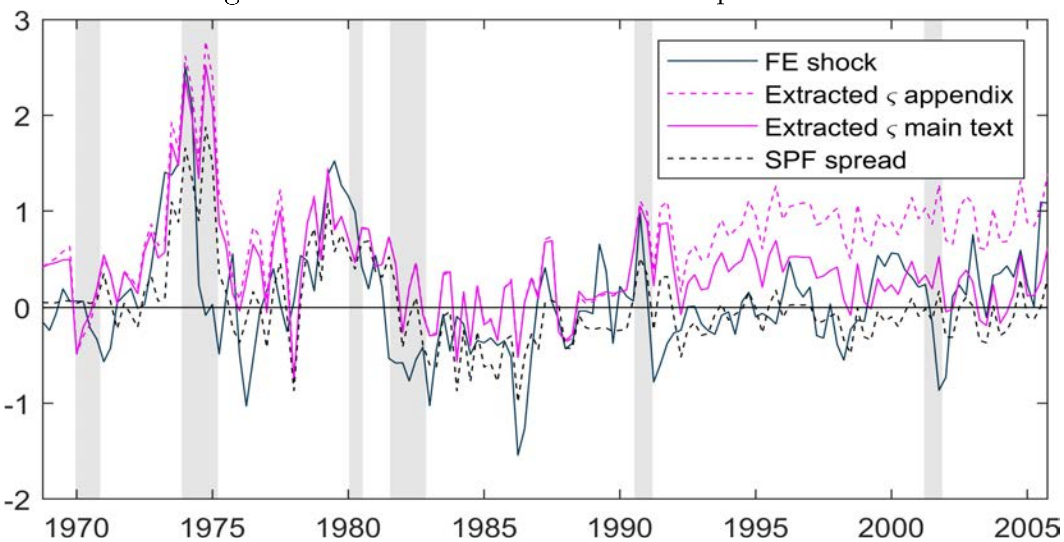

Note: Comparing $\widehat{\varsigma}$ extracted using $\widehat{\mu}' = \mu'^*(\widehat{\varsigma}, \widehat{\rho}, \widehat{\mu})$ to its counterpart in main text using $\widehat{\mu}' = \mu'^*(0, \widehat{\rho}, \widehat{\mu})$, the two series behave similarly before 1990, but the former is higher than the latter after 1990. The FE shock is the "Food and Energy price shock," constructed as the difference between the growth rate of the overall personal consumption deflator and its counterpart excluding food and energy. SPF spread is SPF1Q-SPF3Q.

D Effects of reputation on equilibrium decision rules

This section reports equilibrium decision rules to help understand the Figure 8 historical decomposition, focusing on the positive correlation between $\widehat{\delta}$ and $\widehat{\rho}$, and the effects of evolving reputation on equilibrium policies. In the process, we highlight the crucial role played by a purposeful non-committed policymaker.

Recall from Section 4.4 that the committed type's choice problem can be simplified to choosing (δ, μ'), where the policy difference δ determines reputation evolution and the future pseudo-state μ' controls the level at which inflation expectations are anchored.

Adopting this perspective and using equilibrium (δ^*, μ'^*), inflation expectations e^* are given by the operational expectation function $e(\delta^*, \mu'^*; \varsigma, \rho)$. Inflation expectations then result in an equilibrium opportunistic policy via the best response function $\alpha^* = Ae^* + B(\varsigma)$. Finally, equilibrium committed policy a^* is the sum of opportunistic policy α^* and the policy difference δ^*.

These equilibrium decisions $\{\delta^*, \alpha^*, a^*\}$ depend on reputation ρ, along with μ and ς. The blue lines in Figure 15 reveal how the reputation state ρ affects each of these equilibrium decisions, conditional on $\varsigma = 0$. The columns contrast two levels of the pseudo state μ.[7] In the middle and bottom panels, we also plot the inflation target $\pi^* = 1.5\%$ (black dotted line).

Equilibrium policy difference is displayed in the top panels. Notice first that $\delta^* \leq 0$: equilibrium committed policy is always lower than equilibrium opportunistic policy. This is because the committed policy has strategic power on anchoring the inflation expectations. The strategic power is stronger with higher reputation ρ and the incentive of anchoring expectations is stronger with higher μ. Notice next that δ^* increases with ρ for $\rho > 0$, indicating shrinking policy difference as reputation improves. A larger policy difference helps the committed type to distinguish itself from the opportunistic type but also costs more output to implement. Better reputation reduces the incentive of the committed policymaker to invest further in reputation. Third, at high reputation ρ, δ^* is either zero or close to zero,[8] which makes it difficult for the private sector to distinguish policymaker type.

Equilibrium opportunistic policy shown in the middle panels highlight that α^* decreases with reputation ρ. It is intuitive given that higher reputation yields lower expected inflation and a lower e leads to a lower α^* according to the best response function (10). With very low ρ, α approaches the Nash equilibrium inflation bias of around 8% above π^* in both panels. With ρ close to 1, opportunistic policy is higher in the left panel than the right, because the opportunistic policymaker responds positively to the expected "startup inflation" generated by the committed policymaker in the absence of his incentive of anchoring expectations ($\mu = 0$).

Equilibrium committed policy is displayed in the bottom row. We start by noticing that a^* in the right panel is π^* when $\rho = 1$, highlighting that $\mu = 0.025$ corresponds to a

[7]The two values are zero (the initial value at the regime switch date) and the certainty steady state value when $\rho = 1$. We choose these levels because most equilibrium values of μ lie between them in absence of the cost-push shock.

[8]That is, small relative to 1.2% standard deviation of ε_1 and ε_2.

Figure 15: Effects of ρ on equilibrium policies

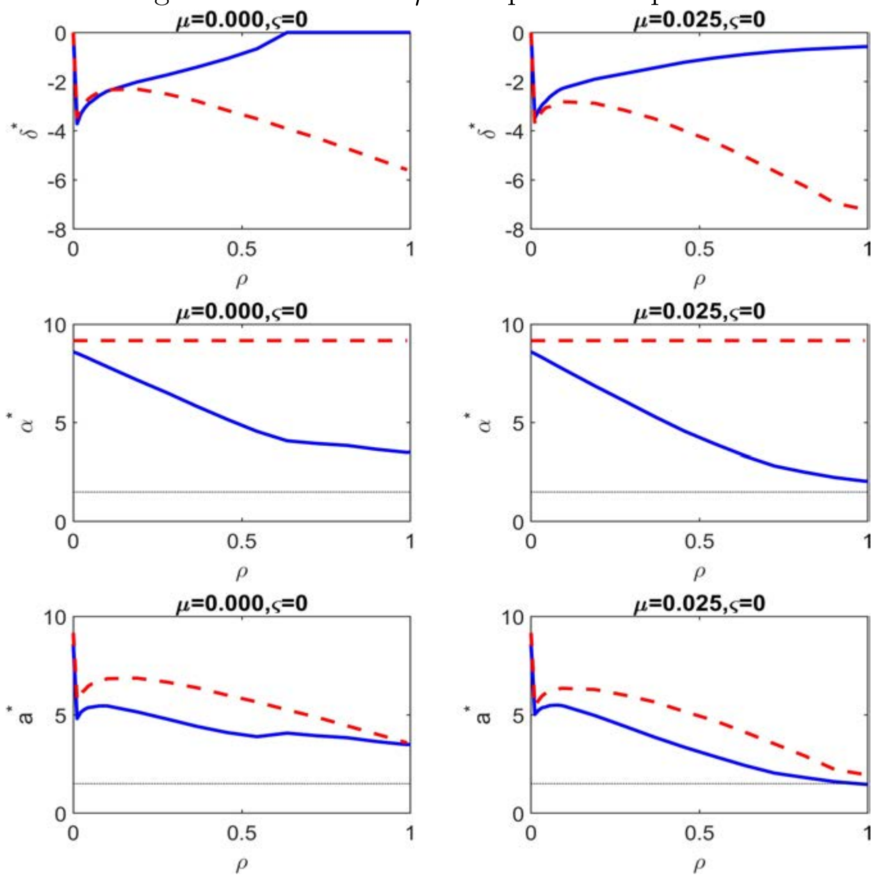

Equilibrium decision rules: top panels: policy difference $\delta^* = a^* - \alpha^*$; middle panels: intended inflation of non-committed policymaker α^*; bottom panels: intended inflation of committed policymaker a^*. Blue solid lines are decision rules in our model where the non-committed policymaker optimally responds to the expected inflation. Red dashed lines are decision rules in a model where the non-committed policymaker mechanically adopts a policy rule that would be optimal if $\rho = 0$.

full commitment steady state. With $\rho = 1$ but $\mu = 0$, a^* in left panel shows that perfect reputation policy involves "start up inflation" higher than π^* as discussed above. When $\rho = 0$, there is no difference between the committed policy and the opportunistic policy. Hence, a^* in both panels is the same as α^*.

With $0 < \rho < 1$, the committed policymaker has two additional considerations. First, leverage over inflation expectations via future intended inflation a^*_{t+1} is reduced, making it more desirable to accommodate shifts in expectations. Second, reputation building is important and this requires optimal a to be different from optimal α. Since a^* is the sum of α^* and δ^*, the net effect of these considerations can be understood as the relative strength of reputation effects on α^* versus on δ^*. In our calibration, the Nash Equilibrium inflation bias (α^* at $\rho = 0$) is much higher than the intrinsic inflation bias (α^* at $\rho = 1$), resulting in a dominant effect of ρ on α^*. In turn, a^* is generally decreasing in ρ, with a flatter slope than α^*.

Purposeful versus mechanical non-committed type An important new element, relative to our prior work (Lu et al. (2016)), is a purposeful, if myopic, policymaker rather than a mechanical alternative type. If we instead assume that the non-committed policymaker mechanically adopts a policy rule that would be optimal if $\rho = 0$ – incorporating the Nash Equilibrium inflation bias – then matters are very different: the results are the red dashed lines in Figure 15. The most salient implication is for the policy difference δ^*. Comparing the red dashed lines with the blue lines in the bottom panels, we find that at majority values of ρ, the policy difference is much larger than when the non-committed policymaker is purposeful. With such a mechanical alternative policymaker, the large δ^* means that private agents learn about policymaker type so fast that we lose the time-varying reputation shown above to crucial for capturing many elements of the US inflation experience.[9]

E Details about credibility construction

Credibility gap in inflation units One intuitive measure is the distance between the $a(s)$ and the private sector's nowcast of inflation $E(\pi|s)$, i.e.,

$$a(s) - E(\pi|s) = (1 - \rho)[a(s) - \alpha(s)] = (1 - \rho)\delta. \tag{C1}$$

so that it depends only on reputation and the policy difference δ.

Degree of credibility In an inflation targeting context, credibility is sometimes related to the private sector's probability that inflation will fall in a band around the target, e.g., $a-\theta \leq \pi \leq a+\theta$. In our setup, this probability reflects implementation errors and the private sector's lack of knowledge about policymaker type. We now assume normal implementation errors and let $N(\cdot, \bar{\pi}, \sigma)$ be the normal cdf with mean $\bar{\pi}$ and standard deviation σ. Normal

[9] Recall the standard deviation of implementation error in our calibration is 1.2%. When the equilibrium policy difference δ^* is as large as three or four times 1.2%, as the red line indicates at majority values of ρ, the policymaker's type will be revealed immediately.

errors imply the private sector's probability of $a-\theta \leq \pi \leq a+\theta$ is

$$\begin{aligned}
& \int_{a-\theta}^{a+\theta} [\rho n(\pi, a, \sigma) + (1-\rho) n(\pi, \alpha, \sigma)] d\pi \\
= \; & \rho[N(a+\theta, a, \sigma) - N(a-\theta, a, \sigma)] + (1-\rho)[N(a+\theta, \alpha, \sigma) - N(a-\theta, \alpha, \sigma)].
\end{aligned}$$

Expressing this as a ratio to $[N(a+\theta, a, \sigma) - N(a-\theta, a, \sigma)]$ leads to our second credibility measure:[10]

$$\text{(C2)} \qquad \psi(a, \alpha, \theta, \rho, \sigma) = \rho + (1-\rho) \frac{[N(a+\theta, \alpha, \sigma) - N(a-\theta, \alpha, \sigma)]}{[N(a+\theta, a, \sigma) - N(a-\theta, a, \sigma)]}$$

which is the ratio of the private sector's probability that inflation falls within the band relative to the committed policymaker's probability. Note that the denominator expression is constant across $a(s)$, while the numerator may be written to stress the policy difference, $N(\delta+\theta, 0, \sigma) - N(\delta-\theta, 0, \sigma)$. That is, our second credibility measure also depends on reputation ρ and the policy difference δ.

[10]This measure is readily generalized to an asymmetric band and type-specific implementation error volatility.

F Model performance with a longer sample

Figure 16: Model-implied and SPF forecasts of inflation

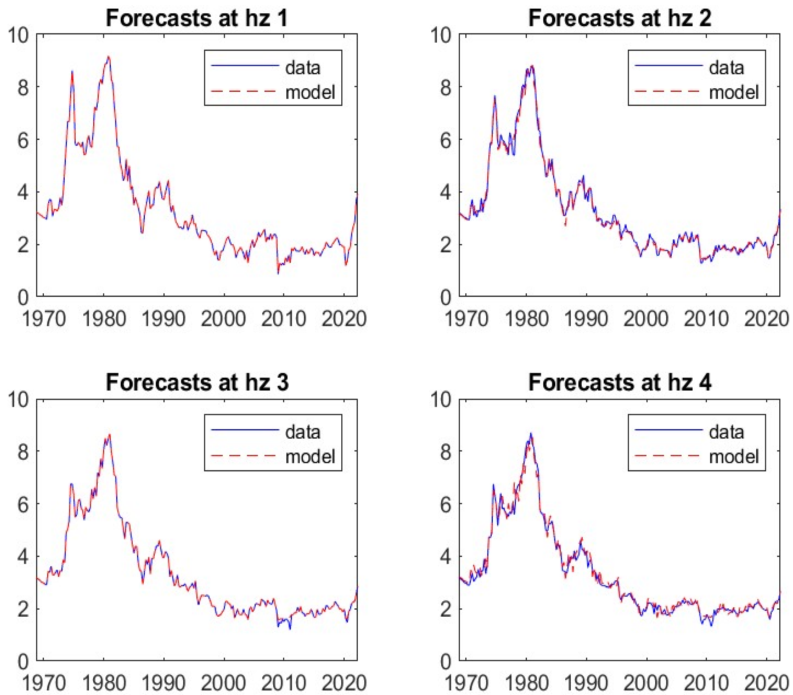

Figure 17: Inflation history and model-implied policies

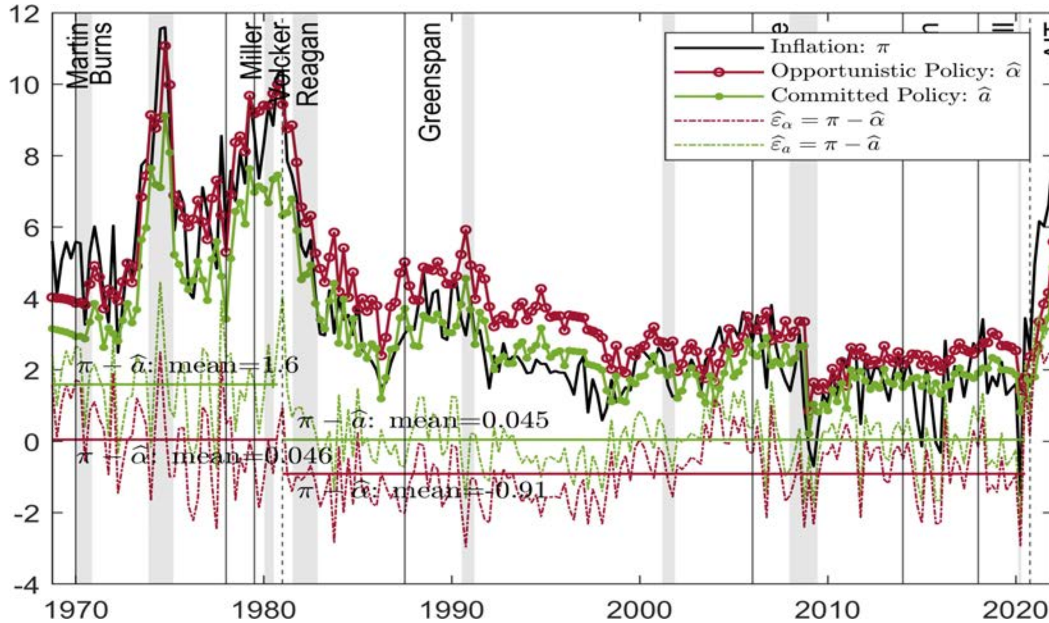

This is computed with only one regime change at 1981Q1. The mean error during the latter regime does not include data from and after 2020Q3, when AIT is announced.

Figure 18: Model-based interpretation of US inflation history