Evolving Reputation for Commitment: Understanding Inflation and Inflation Expectations*

Robert G. King† Yang K. Lu‡

May 2025

Abstract

We develop a theory of how inflation expectations respond to monetary policy, emphasizing the role of purposeful policymakers who strategically influence private agents' learning and expectation formation. The central mechanism linking expectations and policy is reputation – private agents' belief in the policymaker's commitment to announced inflation targets. Reputation evolves as agents update their beliefs based on deviations of actual inflation from announced targets. This, in turn, affects their expectations of future inflation and the effectiveness of monetary policy. Optimal policy internalizes this feedback between expectations and policy outcomes. We present a recursive solution and a quantitative implementation of the model calibrated to U.S. inflation history. We also provide empirical evidence supporting the model's prediction of time-varying sensitivity in long-term inflation forecasts to inflation surprises.

Keywords: time inconsistency, reputation game, optimal monetary policy, forward-looking expectations

JEL classifications: E52, D82, D83.

*We thank the participants in various conferences, seminars, and workshops for valuable feedback. Particularly useful advice was provided by Davide Debortoli, Thorsten Drautzburg, Narayana Kocherlakota, Albert Marcet, Elmar Mertens, and Russell Cooper; valuable discussions were made by Donghai Zhang and Shengliang Ou. All errors are ours. We acknowledge the financial support from the RGC of HKSAR (GRF HKUST-16504317).

†Boston University and NBER. Email: ⟨rking@bu.edu⟩

‡Corresponding author, Hong Kong University of Science and Technology. Email: ⟨yanglu@ust.hk⟩

1 Introduction

Managing expectations is central to monetary policy decisions, because inflation expectations are important for economic activities (Coibion et al. (2018b)) and inflation dynamics (Beaudry et al. (2024)), and are responsive to central bank policy communication (Hansen and McMahon (2016), Haldane and McMahon (2018), Blinder et al. (2024)).

This paper develops a theory of how inflation expectations respond to monetary policy, emphasizing the role of purposeful policymakers who strategically influence private-sector learning and expectation formation. We show how to bring the theory to the data and use it to account for the joint dynamics of U.S. inflation and inflation expectations. Empirically, we validate the model's prediction about the evolving sensitivity of long-term inflation forecasts to inflation surprises.

We employ a variant of the textbook New Keynesian (NK) model featuring forward-looking inflation dynamics, purposeful policymakers with a dual mandate to stabilize inflation and output, and stochastic regime changes.[1] The committed policymaker follows an ex-ante optimal, state-contingent inflation plan, while the opportunistic policymaker chooses inflation sequentially to maximize short-term objectives. Private agents do not observe the policymaker's type or chosen inflation directly; instead, they observe noisy realizations of inflation, which they use to update beliefs about the likelihood that the policymaker is committed—this belief constitutes the policymaker's *reputation*—and to form expectations about future inflation.

A key conceptual result is that high reputation narrows the equilibrium policy gap between the two types, as the opportunistic policymaker is less tempted to deviate. In contrast, low reputation widens the gap, as the committed type has stronger incentives to accelerate private-sector learning. We show that this mechanism is essential for quantitatively matching salient features of U.S. inflation history, e.g., the Volcker disinflation, and for explaining the time-varying response of long-term inflation forecast revisions to inflation surprises.

Our modeling approach builds on a mass literature, pioneered by Lucas, Sargent, Kydland and Prescott, that stresses the importance of policymaker commitment capacity in economic policy. Lucas and Sargent (1979) showed that traditional econometric models were inappropriate for analysis of exogenous policy rules when rational expectations is coupled with forward-looking private sector behavior. Kydland and Prescott (1977) took the next step by incorporating purposeful policymakers into theoretical macroeconomic environments,

[1] A regime is a time interval during which outcomes are interpreted as the choices of a single policymaker.

formulated as dynamic games. They stressed the importance of policymaker commitment capacity, showing how its absence could radically change positive and normative outcomes. In the extensive elaboration of these insights over the ensuing decades, there has been growing recognition that private sector learning is important and, indeed, that policymaker commitment capacity is inherently *unobservable*. A substantial body of literature now integrates private sector learning into the theory of economic policy.[2] Yet, an important gap remains as little research features purposeful policymakers who actively seek to steer the learning of private agents.

This paper shows how to close this gap. We use the insights of modern contract theory (mechanism design) to develop a computable recursive equilibrium for a dynamic game with two types of purposeful policymakers, one which can commit and one which cannot, and private agents who learn policymaker type in a Bayesian manner. Forward-looking behavior of private agents, coupled with both types of policymakers being purposeful, necessitates our novel theoretical approach. In recursive equilibrium, reputation – defined as private agents' likelihood that the policymaker can commit – emerges as a key endogenous state variable.

Why new theory is necessary Forward-looking inflation dynamics in New Keynesian (NK) models have largely replaced earlier specifications used by Lucas, Sargent, Kydland, and Prescott, in which private agents form intra-temporal expectations—that is, they expect policy to be chosen in the same period as expectations are formed. In response to supply shocks, forward-looking dynamics increase the gap between optimal inflation policy with and without commitment.[3] Some earlier studies examine the interaction between optimal inflation policy and reputation under intra-temporal expectations, Cukierman and Liviatan (1991); King et al. (2008); Lu (2013); Dovis and Kirpalani (2021), taking advantage of the fact that these expectations allow dynamic games to be solved using backward induction.

When expectations are forward-looking, strategic interactions become intertemporal and the earlier techniques no longer apply. To see why, consider the choice of period-t committed policy: the period-t payoff depends on private agents' expectations, which are affected by future committed policy, future opportunistic policy, and reputation (private agents' likelihood that each policy will take place). But future opportunistic policy cannot be taken as given because it optimally responds to future private agents' expectations that change with how period-t committed and opportunistic policies affect the evolution of reputation.[4]

[2]See for examples: Barro (1986), Backus and Driffill (1985), Phelan (2006), Dovis and Kirpalani (2022).

[3]See, for example, Clarida et al. (1999).

[4]A common way to avoid these strategic interactions is to assume that one type of policymaker being an

Our new mechanism design approach directly tackles these complications. To begin, we recast the equilibrium of the dynamic game as the solution to a dynamic principal-agent problem. The committed policymaker acts as principal to choose state contingent plans for his own policies, the policies of the opportunistic type subject to incentive compatibility constraints, and private agents' expectations subject to rational expectation constraints. We then use the techniques of dynamic contract theory to formulate the principal-agent problem as a recursive optimization with only three state variables including a highly persistent reputation state,[5] a more temporary cost-push shock,[6] and a predetermined pseudo state.[7]

Dynamic theory makes quantitative history feasible Based on the solution to the recursive optimization problem, we construct a calibrated quantitative model that maps structural shocks and latent states to observable macroeconomic data. In particular, we require that the model's inflation expectations match time series from the Survey of Professional Forecasters (SPF), beginning in late 1968. The key identification assumption is that short-term SPF forecasts are more sensitive to temporary factors, such as cost-push shocks, while longer-term forecasts reflect persistent influences, such as the evolution of policymaker reputation. Formally, we exploit the fact that the model's dynamic system implies a nonlinear filtering structure that allows us to jointly identify shocks and states.[8]

Explaining joint dynamics of expected and actual inflation The nonlinear filter produces estimated reputation that exhibits a big swing, declining throughout 1970s to near zero by the end of 1980 and gradually climbing back afterwards. Using the estimated shocks and states, we compute the model-implied inflation values. Remarkably, these values align closely with the observed U.S. inflation, despite the fact that the observed inflation data is not used in estimating shocks and states. To assess the importance of purposeful policy in shaping private-sector learning, we conduct a counterfactual exercise in which a *naive*

automaton (Lu et al. (2016), Amador and Phelan (2021), Morelli and Moretti (2023)), or to assume that the committed policymaker ignores the effect of his policy on private sector learning (Clayton et al. (2025)). However, our analysis below indicates that these assumptions have considerable effects on outcomes, which are – to our minds – undesirable.

[5]Reputation is capital for the committed policymaker but is a martingale in the eyes of private agents.

[6]We use common terminology for this shock, which shifts the policymaker's output-inflation trade-off.

[7]As in other studies of optimal inflation policy, this variable is required to place the committed policy in recursive form, as discussed further below.

[8]The standard Kalman filter is not applicable due to the model's nonlinearity. We instead use a "sigma point" approximation method—the unscented Kalman filter—which has been shown to perform well in nonlinear regime-switching models. See Särkkä and Svensson (2023) for an overview of Gaussian filtering and Binning and Maih (2015), Benigno et al. (2020), and Foerster and Matthes (2022) for recent macroeconomic applications.

committed policymaker optimizes inflation policy but ignores its impact on belief updating. The results show that such a policymaker takes significantly longer to disinflate the economy than what was observed during the Volcker disinflation following 1981.

Forecast revision regressions validate our theory According to our theory, long-term inflation expectations depend on a reputation state that evolves through Bayesian updating of inflation forecast errors. Moreover, the sensitivity of reputation to forecast errors depends on optimal inflation policies and thus on the time-varying reputation state. We test these implications using regressions of SPF long-term forecast revisions on nowcast forecast errors and find that the estimated coefficient on forecast errors indeed is time-varying and in a pattern consistent with theoretical predictions.

Links to the broader literature Our reputational equilibrium analysis adopts one of the two approaches in modern game theory, originated from Milgrom and Roberts (1982) and Kreps and Wilson (1982).[9] Based on Bayesian learning in a noisy environment, our reputational state variable is the likelihood that the current policymaker has commitment capability. Another familiar reputational approach, introduced by Barro and Gordon (1983) 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 (Chari and Kehoe (1990)).[10] However, policymaker reputation does not vary over time in the sustainable plan literature: it is either excellent or nonexistent. Our learning-based framework permits *reputation building* by a policymaker that can commit and *reputation dissipation* by one that can't.

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,[11] 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)

[9]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."

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

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

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 agents' beliefs* and we use the SPF to extract the evolution of such beliefs. In seeking to recover the evolution of private agents' beliefs about the commitment capacity of the Fed, our work is related to Matthes (2015), but policymakers in his study don't purposefully manage private sector learning.[12] 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.

Hazell et al. (2022), Carvalho et al. (2023), and Beaudry et al. (2024) emphasizes the importance of inflation expectation in understanding inflation dynamics. Our work complements theirs by highlighting the role played by monetary policy in determining the dynamics of inflation expectation, and how optimal monetary policy could take into account its role in shaping inflation expectations to improve future output-inflation trade-off.

Use of the SPF also links our research to the large and growing literature on survey measures of inflation (Coibion et al. (2018a)). 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. (2018a)) 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 Section 2 describes 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 perform quantitative analysis of U.S. inflation history. In Section 6, we highlight strategic reputation management as a central feature of our model and explain how it helps to account for the Volcker disinflation and the time-varying sensitivity of long-term inflation forecasts to inflation surprises. Section 7 concludes.

[12] Other papers that investigate U.S. inflation history with private sector learning include Ball (1995), Erceg and Levin (2003), Orphanides and Williams (2005), Goodfriend and King (2005), Cogley et al. (2015), and Melosi (2016).

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 $\varphi\left(\varsigma_{t+1}; \varsigma_t\right)$. 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.[13] There are two types of policymaker. A *committed* type ($\tau = 1$) 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.[14] The committed inflation plan therefore shapes private agents' expected inflation. An *opportunistic* type ($\tau = 0$) makes the same announcements,[15] but chooses intended inflation on a period-by-period basis.

At the start of each period, the policymaker may be replaced through a publicly observed mechanism: the replacement event ($\theta_t = 1$) occurs with probability q. If no replacement

[13] 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 fiscal links to inflation.

[14] We specify intended inflation rather than intended output for analytical convenience, as they are equivalent via the Phillips curve. We also abstract from policy instruments because policy outcome rather than instrument matters in the model, c.f. Faust and Svensson (2001) and Sargent (1999).

[15] The opportunistic type makes the same announcements as the committed type to avoid revelation. In a related fiscal model, Lu (2013) shows that the unique signalling equilibrium involves the 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.

occurs ($\theta_t = 0$), the policymaker type remains unchanged. We discuss the reputation of a new policymaker further below.

Crucially, the private sector does not observe the policymaker's type (τ_t) or his intended inflation, denoted by a_t for the committed type and α_t for the opportunistic type. Yet, it observes an inflation rate π_t which deviates from the policymaker's intention by a random i.i.d. implementation error $v_{\pi,t} \sim g(\cdot)$ with $g(\cdot) = N(0, \sigma^2_{v,\pi})$:[16]

$$\pi_t = \tau_t a_t + (1 - \tau_t)\alpha_t + v_{\pi,t}. \tag{2}$$

The policymaker's momentary objective depends on inflation π and output gap x.

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

There is a long-run inflation target π^* and a strictly positive output target x^*.

The committed type discount factor is β_a; the opportunistic type is myopic.[17]

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, policymaker replacement may or may not occur. If this public event does occur ($\theta_t = 1$), 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] and permits some reputation inheritance. Second, the exogenous cost-push shock ς_t is realized. Third, there is a policy announcement. If the policymaker is new, he announces a new inflation plan that specifies current intended inflation a_t. Otherwise, either type of continuing policymaker reiterates current economic conditions that call for an intended inflation a_t according to the plan announced at the start of the current regime. 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

[16] We interpret random implementation error as a reduced-form representation for all unforeseeable factors that affect inflation beyond 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).

[17] A myopic opportunistic type is the most parsimonious modeling of an optimizing non-committed policymaker. Our framework and recursive method can be extended to a long-lived opportunistic type, but we leave that extension for future research.

Figure 1: Timing of events within a period

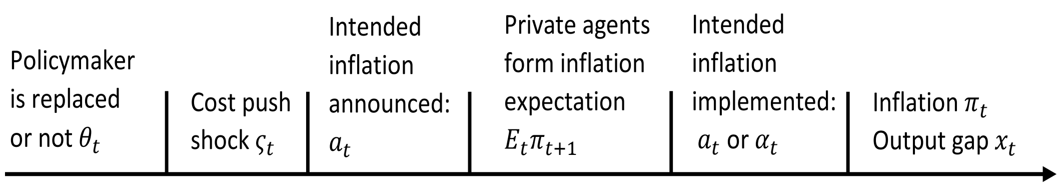

π_t with a density $g(\pi_t|a_t)$ or $g(\pi_t|\alpha_t)$,[18] and an output gap x_t determined by the Phillips curve. New information leads private agents to update their beliefs about policymaker type.

3 Macro Equilibrium in a Dynamic Game

Our economy features random regime switches. In each regime, the policymaker can be one of two types, but their actions do not fully reveal their identity. Private agents form beliefs about the policymaker's type and use these beliefs to forecast future inflation. Following a regime switch, a new inflation plan is initiated and private agents' beliefs about the new policymaker's type are reset randomly to ρ_0. This structure allows each regime to be modeled as a dynamic game with incomplete information. We now describe its equilibrium.

Public Equilibrium Define the public history of the current regime $h_t = \{h_{t-1}, \pi_{t-1}, \varsigma_t\}$ as the collection of all past inflation realizations 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*.[19]

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

3.1 Consistent beliefs: reputation

Consistency of beliefs requires that private agents' assessments of policymaker type are updated according to Bayes' rule (4) which depends on policymakers' equilibrium strategies

[18] With a slight abuse of notation, $g(\pi|\tau a + (1-\tau)\alpha)$ is the density function of $v_\pi = \pi - [\tau a + (1-\tau)\alpha]$.

[19] This restriction is innocuous because: (1) the private sector's strategy is public since its information set is h_t; (2) the committed type's policy is public since it follows the announced policy plan, which needs to be verifiable by the private sector; and (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)).

and observed inflation π_t. Within a regime, private agents' beliefs ρ are updated recursively,

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

If a new regime starts at t, the departing policymaker's reputation $\rho(h_t)$ affects the distribution of the new policymaker's initial reputation: $\rho_0\sim\Xi(\rho_0|\rho(h_t))$.

3.2 Consistent beliefs: inflation expectations

Inflation expectations must be consistent with private agents' beliefs about policymaker type and equilibrium strategies. If a new regime starts at t, the consistent nowcast of inflation is:

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

Within a regime, expectations of future inflation are:

$$e(h_t)=\beta E(\pi_{t+1}|h_t)=\beta\rho(h_t)E(\pi_{t+1}|h_t,\tau_t=1)+\beta(1-\rho(h_t))E(\pi_{t+1}|h_t,\tau_t=0)\tag{6}$$

In both (5) and (6), higher reputation makes expectations more responsive to committed policymaker's future inflation intentions. Turning to the details of (6),

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

$$E(\pi_{t+1}|h_t,\tau_t=0)=\int\sum_{\varsigma_{t+1}}\varphi(\varsigma_{t+1};\varsigma_t)\left[(1-q)\,\alpha(h_{t+1})+qz\left(h_{t+1}\right)\right]g(\pi_t|\alpha(h_t))d\pi_t$$

$E(\pi_{t+1}|h_t,\tau_t=1)$ is conditional on the incumbent policymaker being the committed type. The committed type's intensions lead to stochastic inflation, with density $g(\pi_t|a(h_t))$, contributing to history $h_{t+1}=\{h_t,\pi_t,\varsigma_{t+1}\}$. If the regime continues next period, the expected inflation will be the committed type's intended inflation $a(h_{t+1})$. In the event of a regime change next period, the consistent belief is the history-dependent future nowcast $z(h_{t+1})$. Similarly, $E(\pi_{t+1}|h_t,\tau_t=0)$ is conditional on the incumbent policymaker being the opportunistic type. It 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 $z(h_{t+1})$.

3.3 Sequential rationality of the opportunistic type

Denote $\underline{u}(\alpha, e, \varsigma) \equiv \int u(\pi, x(\pi, e, \varsigma))\, g(\pi|\alpha)\, d\pi$ as the expected momentary objective when the NK Phillips curve (1) is used to replace x with $x(\pi, e, \varsigma) = (\pi - e - \varsigma)/\kappa$. A myopic opportunistic policymaker chooses intended inflation α each period to maximize the expected momentary objective, taking the expected inflation $e_t = e(h_t)$ as given:

$$\alpha(h_t) = \arg\max_{\alpha} \underline{u}(\alpha, e_t, \varsigma_t) \tag{7}$$

The quadratic objective implies a linear best response of α to e and ς.

$$\alpha_t = Ae_t + B(\varsigma_t) \tag{8}$$

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

Since Kydland and Prescott (1977), it has been understood that there is inflation bias when the central bank cannot commit. In our setup with incomplete information, the extent of inflation bias $\alpha_t - \pi^*$ varies with private agents' expected inflation $e_t = \beta(E_t\pi_{t+1})$.

To highlight this unique feature, we rewrite (8) as $\alpha_t - \pi^* = \iota + A\beta(E_t\pi_{t+1} - \pi^*)$ by denoting $\iota \equiv A(\kappa x^* - (1 - \beta)\pi^*)$ and setting $\varsigma_t = 0$,[20] and plots the best response function with the 45 degree line in Figure 2. If private agents expect the inflation to be at target, i.e., $E_t\pi_{t+1} = \pi^*$, the optimal inflation bias is ι; we define this as *intrinsic inflation bias*. If the policy without commitment is fully expected, i.e., $E_t\pi_{t+1} = \alpha_t$, the optimal inflation bias is the intersection of the two lines $\iota/(1 - A\beta)$; it is the well-known *Nash equilibrium (NE) inflation bias*. When $A\beta$ is close to one, NE inflation bias (square marker) can be much larger than intrinsic inflation bias (diamond marker), as it will be in our quantitative model.

Our quantitative model with incomplete information captures the U.S. inflation dynamics of the 1970s as responding to gradually rising expectations as agents learn that the policymaker is opportunistic and follows (8). Foreshadowing this result, Figure 2 shows two one-quarter-ahead inflation forecasts from the Survey of Professional Forecasters (SPF), illustrating the impact of rising expectations on opportunistic policy.

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

Figure 2: Optimal Response of Opportunistic Policy to Inflation Expectations

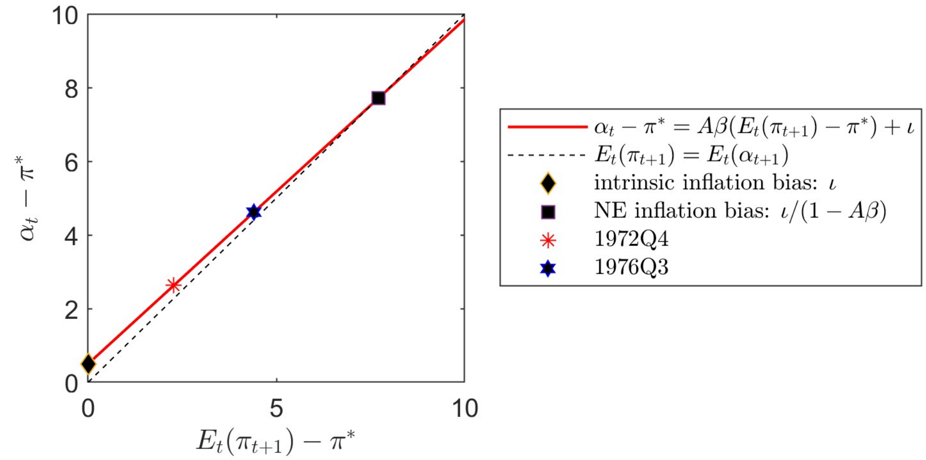

The figure shows how the opportunistic policymaker's optimal inflation bias varies with expected inflation, alongside the 45-degree line. Their intersection (square marker) denotes the *Nash equilibrium (NE) inflation bias*, where policy without commitment is fully anticipated. The diamond marker shows the *intrinsic inflation bias*, assuming agents expect inflation at target. Two SPF one-quarter-ahead forecasts illustrate how rising expectations influence opportunistic policy.

3.4 Sequential rationality of the committed type

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

The strategy of the committed type is *sequentially rational* if he maximizes the expected present discounted payoff at the beginning of his term,[21]

$$U_0 = \sum_{t=0}^{\infty} (\beta_a(1-q))^t \sum_{h_t} p(h_t) \underline{u}(a_t, e(h_t), \varsigma_t), \tag{9}$$

where $\underline{u}(a, e, \varsigma) \equiv \int u(\pi, x(\pi, e, \varsigma))\, g(\pi|a)\, d\pi$ with $x(\pi, e, \varsigma) = (\pi - e - \varsigma)/\kappa$, and

$$p(h_t) = \varphi(\varsigma_t; \varsigma_{t-1}) g(\pi_{t-1}|a(h_{t-1})) p(h_{t-1}) \tag{10}$$

The probability of a specific history $h_t = [\varsigma_t, \pi_{t-1}, h_{t-1}]$ is conditional on inflation being

[21] 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.

generated by the committed type, 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.[22]

In selecting the state-contingent plan $a(h_t)_{t=0}^{\infty}$, the committed policymaker takes into account the *strategic power* of his plan in shaping how private agents' inflation expectations e_t respond to the history h_t. The consistent expectations condition (6) reveals three channels through which this strategic power operates:

i) *Anchoring expectations*: $e(h_t)$ is partially anchored by the next-period committed policy $a(h_{t+1})$. ii) *Managing perceived opportunistic policies*: the opportunistic policymaker chooses $\alpha(h_t)$ as a best response to $e(h_t)$ under sequential rationality. Thus, by influencing expectations, the committed policymaker also indirectly shapes the perceived behavior of the opportunistic type. iii) *Building reputation*: the strength of the anchoring effect depends on $\rho(h_t)$, which evolves based on the history of past committed policies and perceived opportunistic policies, $a(h_{t-1})$ and $\alpha(h_{t-1})$.

3.5 Public Perfect Bayesian Equilibrium

We now define our dynamic game's Public Perfect Bayesian Equilibrium (PBE).

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

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

(ii) given $\alpha(h_t)$, $a(h_t)$, and $z(h_t)$, private agents' beliefs of policymaker type $\rho(h_{t+1})$ are updated according to (4); and their 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 (7);

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 (9), taking into account the strategic power of $\{a(h_t)\}_{t=0}^{\infty}$ on $\{e(h_t)\}_{t=0}^{\infty}$ and $\{\alpha(h_t)\}_{t=0}^{\infty}$.

4 Constructing the Equilibrium

Construction of the Public PBE 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

[22] 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.

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

The strategic power of the committed policy plan $\{a(h_t)\}_{t=0}^{\infty}$ over $\{e(h_t)\}_{t=0}^{\infty}$ and $\{\alpha(h_t)\}_{t=0}^{\infty}$ makes solving the within-regime equilibrium challenging. The committed policymaker's optimal choice depends on the actions of the opportunistic type, as private agents' expectations reflect a weighted average of both types' future policies. At the same time, the committed type cannot treat opportunistic policy as fixed, since the opportunistic type responds to expectations – which are themselves shaped by the committed policy plan.

To address these challenges, we reformulate the within-regime equilibrium as a principal-agent problem. The committed policymaker, acting as the principal, maximizes (9) by choosing state-contingent plans for his own actions and those of two agents – the private sector and the opportunistic policymaker – subject to two incentive compatibility (IC) constraints: (i) consistent beliefs and rational expectations by private agents (4) and (6); and (ii) the opportunistic type's optimal response to expected inflation (8).

This principal-agent problem differs from the standard literature[23] in that one of the agents – the private sector – *disagrees* with the principal in the probability of a specific history. The private sector *thinks* that current inflation could be generated by the opportunistic policymaker, as captured in $E(\pi_{t+1}|h_t, \tau_t = 0)$ in the rational expectations constraint (6). 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 (9).

The disagreement is particularly relevant when the opportunistic policy optimally responds to expected inflation, as different paths of realized inflation lead to different future opportunistic policies. The disagreement in the probability of current inflation thus results in different inflation expectations by the committed policymaker and the private sector.[24]

4.2 Recursive formulation

A key necessary step in recursive formulation is to cast the Lagrangian component associated with the rational expectation constraint (6) into recursive form. Disagreement in inflation

[23] The standard approach to solve Ramsey equilibrium is laid out and advanced by Kydland and Prescott (1980), Chang (1998), Phelan and Stacchetti (2001) and Marcet and Marimon (2019).

[24] The disagreement in the probability of current inflation will be inconsequential if the non-committed type of policymaker follows a policy rule that depends only on exogenous shocks, as in Lu et al. (2016).

expectations between principal and agent poses a challenge in this regard. We overcome it by a "change of measure" that expresses $E(\pi_{t+1}|h_t, \tau_t = 0)$ 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. We then establish:[25]

Proposition 1. Given $z(\varsigma, \rho)$, the within-regime equilibrium is the solution to:

$$W(\varsigma,\rho,\mu) = \min_{\gamma}\max_{a,\alpha,e}\{\underline{u}(a,e,\varsigma) + [\gamma e - \mu\omega(a,\alpha,\rho,\varsigma)] + \beta_a(1-q)\int\sum_{\varsigma'}\varphi(\varsigma';\varsigma)W(\varsigma',\rho',\mu')g(\pi|a)d\pi\}, \tag{11}$$

$$\text{with } \omega(a,\alpha,\rho,\varsigma) \equiv (1-q)a + qz(\varsigma,\rho) + \frac{1-\rho}{\rho}[(1-q)\alpha + qz(\varsigma,\rho)] \tag{12}$$

$$\alpha = Ae + B(\varsigma) \tag{13}$$

$$\mu' = \frac{\beta}{\beta_a(1-q)}\gamma\rho, \text{ given } \mu_0 = 0 \tag{14}$$

$$\rho' = \frac{\rho g(\pi|a)}{\rho g(\pi|a) + (1-\rho)g(\pi|\alpha)} \text{ with prob } g(\pi|a), \text{ given } \rho_0 \tag{15}$$

The component $(\gamma e - \mu\omega)$ arises from the Lagrangian component of the rational expectations constraints (6) expressed in the recursive form. γ is the multiplier to the constraint (6) and is the shadow price of honoring the promised inflation (choice of e).[26] The pseudo state variable μ records past promises made by the committed type (contained in ω).

With incomplete information and stochastic replacement, the composite promise term ω defined in (12) contains more than the committed type's promised a, because the expected inflation by private agents also depends on their perceived inflation α intended by the opportunistic type and their nowcast of inflation z in a new regime.[27] The weight $(1-\rho)/\rho$ attached to $[(1-q)\alpha + qz(\varsigma,\rho)]$ reflects the divergent probability beliefs about inflation π held by the committed policymaker and private agents. Appendix A.9 explains how we eliminate the likelihood ratio λ from the state space.

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

[26] 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.

[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.10 provides a fuller discussion.

4.3 The outer loop fixed point requirement

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

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

$a^*(.)$ and $\alpha^*(.)$ are the within-regime equilibrium strategies obtained in the recursive program (11) given $z^*(\varsigma,\rho)$. The pseudo-state μ_0 in the new regime is zero as a new policymaker is not held accountable for prior commitments made by his predecessor.[28]

4.4 Managing expectations

The recursive formulation simplifies the committed policymaker's management of expectations by reducing the problem to the choice of just two variables: the contemporaneous policy difference $\delta \equiv a - \alpha$, and the future pseudo-state variable μ'.[29]

Lemma 1. Given (ς,ρ) and future equilibrium strategies $a^*(\varsigma',\rho',\mu')$, $\alpha^*(\varsigma',\rho',\mu')$ and $z^*(\varsigma',\rho')$, rationally expected inflation is uniquely determined by δ and μ';

$$\begin{aligned} e = e(\delta,\mu';\varsigma,\rho) &= \beta\rho \textstyle\int M_a(\varsigma, b(v_\pi, v_\pi+\delta,\rho), \mu') g(v_\pi) dv_\pi + \\ &\quad \beta(1-\rho)\textstyle\int M_\alpha(\varsigma, b(v_\pi-\delta, v_\pi,\rho),\mu') g(v_\pi) dv_\pi; \\ \text{where } b(\pi-a,\pi-\alpha,\rho) &\equiv \frac{g(\pi-a)\rho}{g(\pi-a)\rho + g(\pi-\alpha)(1-\rho)}; \\ M_a(\varsigma,\rho',\mu') &\equiv \textstyle\sum_{\varsigma'} \varphi(\varsigma';\varsigma)\left[(1-q)\,a^*(\varsigma',\rho',\mu') + qz^*(\varsigma',\rho')\right]; \\ M_\alpha(\varsigma,\rho',\mu') &\equiv \textstyle\sum_{\varsigma'} \varphi(\varsigma';\varsigma)\left[(1-q)\,\alpha^*(\varsigma',\rho',\mu') + qz^*(\varsigma',\rho')\right]. \end{aligned} \tag{17}$$

The policy difference δ enters the function $b(\cdot)$ – the Bayesian learning rule (4) that describes how ρ' updates in response to the observed inflation outcome π. The distribution of π depends on the policymaker type: $\pi = a + v_\pi$ if the policymaker is committed, and $\pi = \alpha + v_\pi$ if opportunistic. A larger policy difference speeds up private-sector learning about the policymaker's type by increasing the informativeness of inflation outcomes and sharpening the direction of ρ' adjustment, given the actual type.

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

[29] Appendix B.1 contains the proof of the lemma.

The future pseudo-state variable μ' enters the rationally expected future equilibrium policies and thus serves to anchor expectations. The committed policymaker sets μ' through the choice of γ, which represents the shadow value of promising a'. The effect of γ on μ' depends on the current reputation state ρ, which reflects the private sector's skepticism about the likelihood that a' will be delivered.

From expectations to policies: Managing expectations e through the choice of the pair (δ, μ') shapes the inflation policies: $\alpha = Ae + B(\varsigma)$ and $a = \alpha + \delta$. Together, expectations and policies determine the momentary objective function $\underline{u}(\cdot)$ and the composite promise term $\omega(\cdot)$, allowing us to simplify the recursive program (11) from a choice over (γ, a, α, e) to a choice over just (δ, μ'):[30]

Proposition 2. Given $e = e\left(\delta, \mu'; \varsigma, \rho\right)$ and $U^*\left(\varsigma, \rho, \mu\right) = W(\varsigma, \rho, \mu) + \mu\omega^*(\varsigma, \rho, \mu)$,

$$
\begin{aligned}
(18) \qquad W(\varsigma, \rho, \mu) &= \max_{\delta, \mu'} \left[\underline{\underline{u}}(\delta, \mu'; \varsigma, \rho) - \mu \underline{\underline{\omega}}(\delta, \mu'; \varsigma, \rho) + \beta_a \left(1 - q\right) \Omega(\delta, \mu'; \varsigma, \rho) \right] \\
\text{with } \underline{\underline{u}}(\delta, \mu'; \varsigma, \rho) &\equiv \underline{u}(Ae + B(\varsigma) + \delta, e, \varsigma) \\
\underline{\underline{\omega}}(\delta, \mu'; \varsigma, \rho) &\equiv \frac{1}{\rho} \left[\left(1 - q\right) \left(Ae + B(\varsigma)\right) + q z^* \left(\varsigma, \rho\right) \right] + \left(1 - q\right) \delta \\
\Omega(\delta, \mu'; \varsigma, \rho) &\equiv \int \textstyle\sum_{\varsigma'} \varphi(\varsigma'; \varsigma) U^* \left(\varsigma', b(v_\pi, v_\pi + \delta, \rho), \mu'\right) g(v_\pi) dv_\pi
\end{aligned}
$$

Let $s = [\varsigma, \rho, \mu]$ denote the state vector. The optimal choices $\delta^*(s)$ and $\mu'^* = m(s)$ determine the equilibrium objects $\omega^*(s)$, $W(s)$, $a^*(s)$, and $\alpha^*(s)$, which in turn serve as inputs to the rationally expected inflation function $e(\delta, \mu'; \varsigma, \rho)$ in (17) and the value function $U^*(s)$. The program (18) thus defines the inner-loop fixed point problem that solves for the within-regime equilibrium. Combined with the outer-loop fixed point problem (16), the two fixed point problems jointly characterize the Public PBE.

5 Building the Quantitative Model

In this section, we transform the model into a quantitative framework suitable for time series analysis. We begin by converting the within-regime recursive solution into a recursion over calendar time that accommodates random regime switches. After calibrating the model parameters, we implement a state extraction strategy that jointly identifies the latent states and structural shocks by targeting time series data from the Survey of Professional Forecasters

[30] Appendix B6 provides the proof.

(SPF). Finally, using the identified states and shocks, we construct model-implied variables that were not targeted in estimation and compare them to observed data for validation.

5.1 Time series implications of the Public PBE

This juncture marks a shift in how we use t: up to this point, it has served as a regime clock, but from now on it will denote calendar time in the time series analysis. According to Propositions 1 and 2, the state vector $s = [\varsigma, \rho, \mu]$ determines the intended inflation policies, $a(s)$ and $\alpha(s)$, as well as private agents' expected inflation $e(s)$.[31] At the end of each period, inflation is realized as a random variable: $\pi = \tau a(s) + (1-\tau)\alpha(s) + v_\pi$, where $\tau = 1$ indicates a committed regime and $\tau = 0$ an opportunistic regime.

At the start of the next period, a new cost-push shock ς' is drawn from the distribution $\varphi(\varsigma'; \varsigma)$. The evolution of the reputation state ρ' and the pseudo-state μ' depends on two sources of randomness. If no policymaker replacement occurs ($\theta' = 0$), reputation updates via Bayes' rule: $\rho' = b(\pi - a(s), \pi - \alpha(s), \rho)$. Since a and α are equilibrium functions of s, we write $\rho' = b(s, \pi)$ for notational convenience. The pseudo-state evolves deterministically as $\mu' = m(s)$. If a replacement occurs ($\theta' = 1$), the pseudo-state resets to $\mu' = 0$, and reputation is subject to *partial inheritance*. Specifically, inheritance is complete with probability ζ_ρ, and otherwise ρ' is drawn from a Beta distribution with mean $\overline{\rho}$ and standard deviation σ_ρ. Formally, letting $\phi' \sim \text{Bernoulli}(\zeta_\rho)$, reputation evolves according to $\rho' = \phi' b(s, \pi) + (1-\phi')v'_\rho$.

Thus, we can construct an augmented state vector $S = [s, \pi]$ that evolves recursively over time, conditional on the realizations of θ, ϕ, and τ. Private agents observe the outcomes of the replacement events – θ, ϕ, and v_ρ – but do not observe the policymaker type τ.

5.2 Calibration

Table 1 reports the calibrated values of important model parameters. One period is a quarter. 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.[32] The private sector and the committed type share a conventional quarterly discount factor based on a 2% annual real rate.

The slope of the Phillips curve is a central element in any study of optimal 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

[31] We drop the superscript $*$, as we now focus exclusively on equilibrium behavior.

[32] This value matches Shapiro and Wilson (2019) estimate based on FOMC transcripts 2000-2011.

annualized inflation of -1%, a value compatible with diverse empirical evidence.[33]

Table 1: Parameters

π^*	Inflation target	1.5%
β, β_a	Discount factor (private agents, committed type)	0.995
κ	PC output slope	0.08
ϑ_x	Output weight	0.1
x^*	Output target	1.73%
q	Replacement probability	0.03
ζ_ρ	prob of reputation inheritance	0.9
$\overline{\rho}$	mean of reputation draw	0.1
σ_ρ	std of reputation draw	0.05
ξ_ς	Persistence of cost-push shock	0.7
$\sigma_{v,\varsigma}$	Std of cost-push innovation	0.7%
$\sigma_{v,\pi}$	Std of implementation error	1.2%

One period is a quarter. Inflation target π^*, std of cost-push innovation $\sigma_{v,\varsigma}$, and std of implementation error $\sigma_{v,\pi}$ are all annualized rates.

The policymaker's preferences are another key element in the analysis of optimal inflation. We set the weight on output ϑ_x to 0.1, which is in the middle of the range used by prominent Fed researchers.[34] Since $A = \vartheta_x/(\vartheta_x + \kappa^2)$, together with $\kappa = .08$, this implies $A = .94$. The target output gap x^* is chosen to yield a relatively small intrinsic inflation bias $\iota = .5\%$ while yielding a NE bias large enough to capture the magnitude of the Great Inflation: $\iota/(1 - A\beta)$ of around 8%. From above, $\iota = A(\kappa x^* - (1-\beta)\pi^*)$ so that the implied value for $x^* = 1.73\%$.[35]

The replacement probability of $q = .03$ implies an average regime duration of 8 years. We have less empirical guidance about the inheritance mechanism for reputation: we choose parameters so that a new policymaker inherits his predecessor's reputation with probability .9. Otherwise, his initial reputation is random with mean .1 and standard deviation 0.05.[36]

[33] 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.

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

[35] $x^* = 1.73\%$ is equivalent to targeting unemployment about 1% below the natural rate, if we use an Okun's law coefficient of 1.67.

[36] Formally, reputation in th event of replacement next period is governed by $\rho' = \phi' b(s, \pi) + (1 - \phi') v'_\rho$

But our equilibrium policy functions are not sensitive to these choices due to the small replacement probability q.

Beginning in the 1970s, many studies of inflation use an observable "Food and Energy price shock"(FE shock hereafter).[37] We use the FE shock's serial correlation and its standard deviation as the cost-push shock's persistence ξ_ς and innovation volatility $\sigma_{v,\varsigma}$. The transition probability matrix $\varphi(\varsigma';\varsigma)$ is calibrated to approximate $\varsigma' = \xi_\varsigma \varsigma + v_\varsigma$ where $v_\varsigma \sim N(0, \sigma^2_{v,\varsigma})$.[38] To calibrate the standard deviation of implementation errors, we combined the FE shock and the SPF one-quarter-ahead inflation forecast in an initial approximation to opportunistic intended inflation α, estimating the standard deviation of $(\pi - \alpha)$ over 1964Q4-1979Q2.

5.3 State extraction strategy

The model features three structural shocks: to the cost-push process (v_ς), to inflation (v_π), and to reputation in the event of policymaker replacement (v_ρ). It also includes three binary states: policymaker replacement (θ), reputation inheritance (ϕ), and policymaker type (τ). These are supplemented by three continuous state variables in the Public PBE: the highly persistent reputation ρ, the transitory cost-push shock ς, and the predetermined pseudo-state μ. Although unobserved by an outside observer (econometrician), these states and shocks are linked to observable macroeconomic variables through the model's equilibrium functions.

The state extraction strategy uses the term structure of SPF inflation forecasts to identify the latent states. The model implies that longer-term forecasts depend more on the persistent reputation variable ρ_t, while shorter-term forecasts are more responsive to the transitory cost-push shock ς_t. Figure 3 plots the SPF's three-quarter-ahead forecast (SPF3Q) alongside the one-quarter-ahead forecast (SPF1Q), illustrating the smoother behavior of SPF3Q.[39] Drawing from the term structure literature in interest rates, we construct an SPF spread as SPF1Q minus SPF3Q (plotted as a dashed black line). The spread rises notably during the two major oil price shocks of the 1970s (1974–75 and 1978–80), as well as during the recent COVID-19 shock (2021-22).

Given a set of calibrated parameters, the model delivers a function $f(s, j)$ representing

with $\phi' \sim$ Bernoulli(ζ_ρ) and $v'_\rho \sim$ Beta$(\overline{\rho}, \sigma_\rho)$. We set $\zeta_\rho = 0.9$, $\overline{\rho} = 0.1$, and $\sigma_\rho = 0.05$.

[37] 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.

[38] We use the Rouwenhorst (1995) method.

[39] Elmar Mertens pointed us to the use of SPF term structure, as in Mertens and Nason (2020). We do not use SPF4Q due to missing observations, particularly in 1975.

private agents' expectations at each horizon j.[40] Let $f_{t+j|t}$ denote the SPF forecast at horizon j, assumed to differ from the model-implied expectation by a Gaussian observation error: $f_{t+j|t} = f(s_t, j) + \varepsilon_{jt}$.[41]

Following the Markov switching literature (Hamilton (1989), Kim (1994)), the model is cast into a state-space form using the time-series recursion of $S = [s, \pi]$, while the three binary state variables (θ, ϕ, τ) are treated as the outcomes of an unobserved discrete-state Markov process Θ. Six discrete states are defined for Θ: states $\{1, 3, 5\}$ represent a continuing committed policymaker ($\theta = 0, \tau = 1$), a newly appointed committed policymaker with complete reputation inheritance ($\theta = 1, \phi = 1, \tau = 1$), and a new committed policymaker with randomly drawn reputation v_ρ ($\theta = 1, \phi = 0, \tau = 1$); states $\{2, 4, 6\}$ represent the corresponding cases for an opportunistic policymaker. The transition probability matrix for these six states is restricted to reflect the structure of the model – for example, the policymaker type τ can change only when a replacement event occurs. The matrix is presented in Appendix C.3.

5.4 The state space model with Markov-switching

We now detail the evolution of continuous state variables $S_t = [\varsigma_t, \rho_t, \mu_t, \pi_t]$ resulting from $v_t = [v_{\varsigma,t}, v_{\rho,t}, v_{\pi,t}]$, taking as given the discrete state $\Theta_t \in \{1, 2, ..., 6\}$ defined above.

$$S_t = \begin{bmatrix} \xi_\varsigma \varsigma_{t-1} + v_{\varsigma,t} \\ (1 - \theta_t + \theta_t \phi_t) b(s_{t-1}, \pi_{t-1}) + \theta_t (1 - \phi_t) v_{\rho,t} \\ (1 - \theta_t) m(s_{t-1}) \\ \tau_t a(s_t) + (1 - \tau_t) \alpha(s_t) + v_{\pi,t} \end{bmatrix} = F(S_{t-1}, v_t | \Theta_t)$$

The first entry specifies the process for the cost push shock ς. The second entry specifies that ρ_t is determined by the Bayes' rule $b(s_{t-1}, \pi_{t-1})$, if there is no replacement ($\theta = 0$) or if there is reputation inheritance ($\theta = 1$ and $\phi = 1$), while otherwise ρ_t is a random shock $v_{\rho,t}$ with support $[0, 1]$. The third entry indicates that the pseudo state variable evolves according to $\mu_{t+1} = m(s_t)$, except if there is replacement ($\theta = 1$) in which case it is set to zero. The final

[40] Appendix C.2 provides derivations.

[41] If the timing of policymaker replacements were known and model and data expectations were assumed to match exactly, the state variables ς_t and ρ_t could be recovered by inverting the theoretical relationships, since μ_t evolves deterministically. An earlier version of this research, King and Lu (2022), used this approach, which Kollmann (2017) refers to as an "inversion filter" (see also Drautzburg et al. (2022)).

entry captures that inflation π_t depends on the type of policymaker in place.[42]

The one-quarter-ahead and three-quarter-ahead SPF forecasts are taken to be the model inflation forecasts corrupted by Gaussian observation errors ε_1 and ε_3.[43] That is, our observation equations are:

$$Y_t = \begin{bmatrix} f_{t+1|t} \\ f_{t+3|t} \end{bmatrix} = \begin{bmatrix} f(\varsigma_t, \rho_t, \mu_t, 1) \\ f(\varsigma_t, \rho_t, \mu_t, 3) \end{bmatrix} + \begin{bmatrix} \varepsilon_{1t} \\ \varepsilon_{3t} \end{bmatrix} = H(S_t) + \varepsilon_t$$

As our model is not linear, we cannot use the standard Kalman filter. We adopt an efficient unscented Kalman filter – that has been shown to work well in nonlinear regime-switching models. Appendix C.3 provides details on the algorithm, which also employs the collapsing approach of Kim (1994) and Kim and Nelson (2017). For each element of S_t and Θ_t, our approach produces filtered estimates (based on $Y^t = [Y_1, Y_2, ...Y_t]$) and smoothed estimates (based on Y^T).

The smoothed estimates of the cost-push shock $\widehat{\varsigma}_t$ (red) and the reputation state $\widehat{\rho}_t$ (cyan and measured on the right hand axis) are reported in Figure 3.[44] The estimated cost-push shock $\widehat{\varsigma}_t$ covaries positively with the SPF spread (SPF1Q-SPF3Q), consistent with our strategy of exploiting greater sensitivity of near-term forecasts to transitory shocks. The estimated reputation state $\widehat{\rho}_t$ exhibits a big swing, declining from around 0.7 in 1969 to near zero by the end of 1980 and finally climbing back to above 0.8 after 2000.[45]

5.5 Targeted and untargeted time series

We now report the performance of the model-based non-linear Kalman method, in terms of fitting targeted time series, SPF1Q and SPF3Q, and matching untargeted time series.

5.5.1 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 C.6 Figure 10 shows that these additional forecasts lie almost entirely on top of the untargeted SPF2Q and SPF4Q, providing support for our state extraction approach.

[42] The final entry uses $s_t = [\varsigma_t, \rho_t, \mu_t]$ that are shown to be a function of S_{t-1} by the first three entries.

[43] ε_1 and ε_3 are i.i.d. normal random variables with mean zero and standard deviation 0.5% (annualized).

[44] The reported value is the probability-weighted average of smoothed estimates of state variables conditional on being in a discrete state: $\widehat{x}_t = \sum_{i=1}^{6} E(x_t|\Theta_t = i, Y^T) Pr(\Theta_t = i|Y^T)$.

[45] For the estimates of discrete states, please refe to Appendix C.5.

Figure 3: Targeted SPFs, SPF spread, and estimated state variables

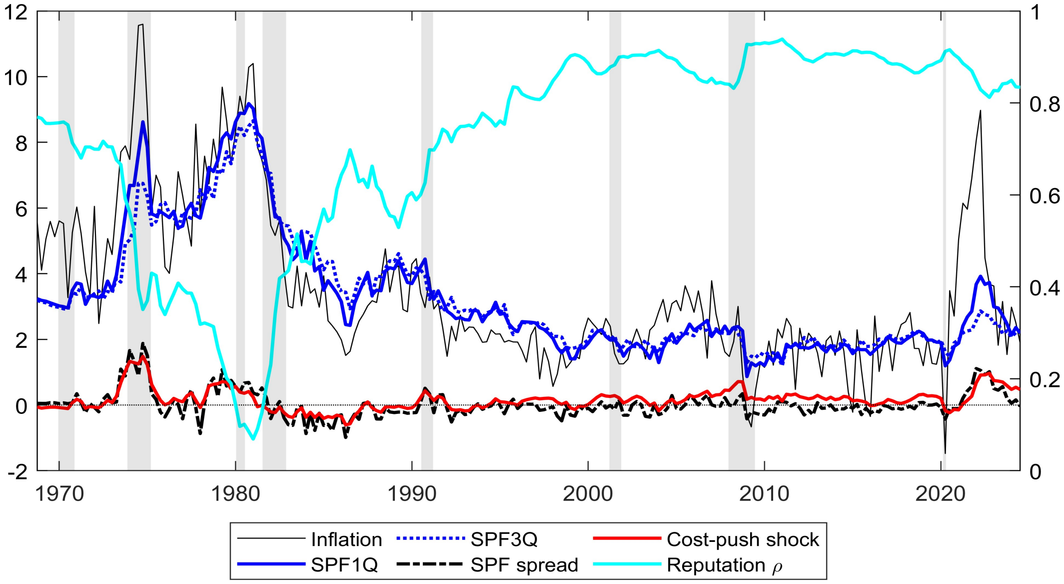

This figure plots selected inputs and outputs of our nonlinear filter. The inputs are the one-quarter-ahead and three-quarter-ahead SPF inflation forecasts from 1968Q4 to 2024Q3. Appendix C provides details on our SPF constructions. The SPF spread (black dashed line) is the difference between SPF1Q and SPF3Q. It rises during the first (1974-75) and the second (1978-80) inflation surges (black solid line). The outputs are smoothed estimates of reputation (cyan and measured on the right axis) and cost-push shocks (red). Consistent with our strategy of exploiting greater sensitivity of near-term forecasts to transitory shocks, the estimated cost-push shock covaries positively with the SPF spread.

5.5.2 Observed and estimated inflation

The state-space model treats inflation as a latent variable and produces both filtered and smoothed estimates of π_t. Since the state extraction relies solely on SPF data, comparing the smoothed estimates $\widehat{\pi}_t$ with observed inflation provides a natural form of model validation. To assess the quality of our estimates, we benchmark them against the SPF's one-quarter-ahead forecast (SPF1Q).[46]

Figure 4 compares observed inflation (blue) with two series: the SPF's one-quarter-ahead forecast made in the current quarter (black) and our model's smoothed estimate (red). The dashed lines plot the deviations of each series from observed inflation, using the

[46] Because our extraction procedure uses SPF1Q and SPF3Q data, one might argue that it is unsurprising the estimates perform well, given the SPF itself tracks observed inflation. Using SPF1Q as the performance benchmark directly addresses this concern.

same respective colors. By two key metrics, our smoothed estimates outperform SPF1Q in tracking actual inflation: (1) they yield a lower mean squared error (MSE) of 1.32 compared to 1.58 for SPF1Q; and (2) their fitting errors exhibit lower persistence (0.52 versus 0.58).

These results suggest that our quantitative model captures U.S. inflation dynamics well, despite inflation not being directly targeted in the estimation.[47]

Figure 4: Inflation history and model-implied inflation

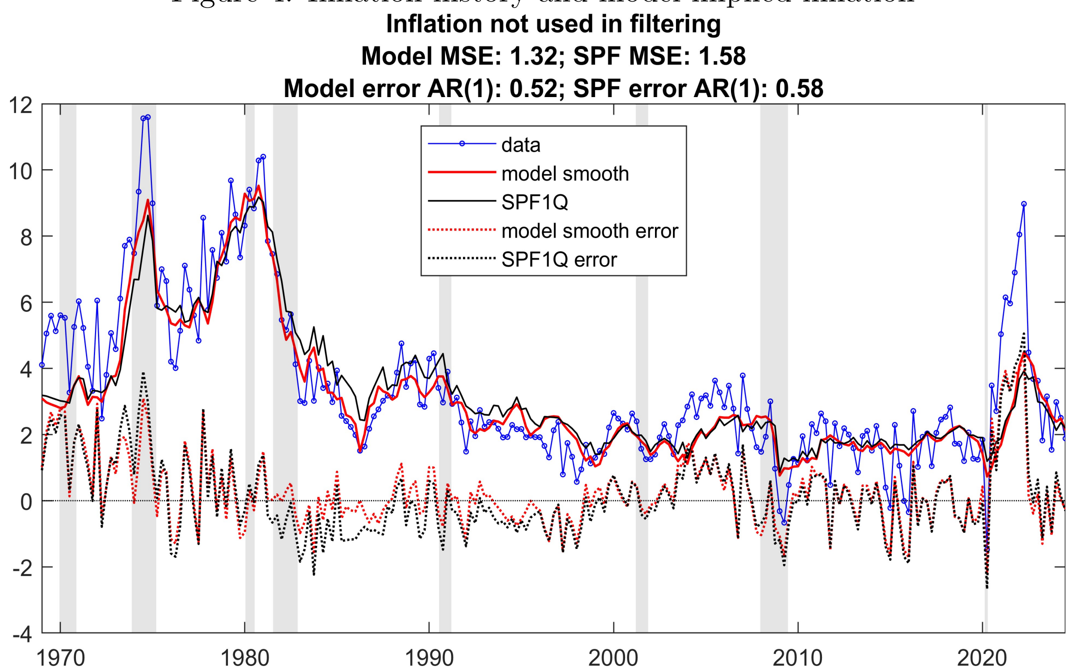

This figure shows that the smoothed inflation estimates (red) produced by our nonlinear filter closely track U.S. inflation data (blue), despite relying solely on SPF data. For comparison, we also plot the one-quarter-ahead SPF forecasts made in the same quarter (SPF1Q, black). The model's fitting error (dashed red) and the SPF1Q error (dashed black) are shown as the respective gaps from observed inflation. The comparison reveals that our smoothed estimates outperform SPF1Q on two fronts: they exhibit a lower mean squared error and less persistent fitting errors.

[47] A skeptical reader may note that our smoothed estimates use the full sample, unlike real-time SPF forecasts. To address this, Appendix C.6, Figure 11, shows one-sided (filtered) estimates. The close match with observed inflation persists even without a full-sample information advantage.

6 Strategic Reputation Management

A central contribution of this paper is to develop a solution method for models in which the committed policymaker takes strategic actions to influence private-sector learning, i.e., his reputation.[48] In this section, we first show how allowing for such strategic behavior makes a difference in the equilibrium policy functions. We then demonstrate why this difference is crucial for understanding the dynamics of the Volcker disinflation episode, and the time-varying sensitivity of long-term inflation forecast revision to inflation surprises.

Figure 5: Equilibrium policy functions

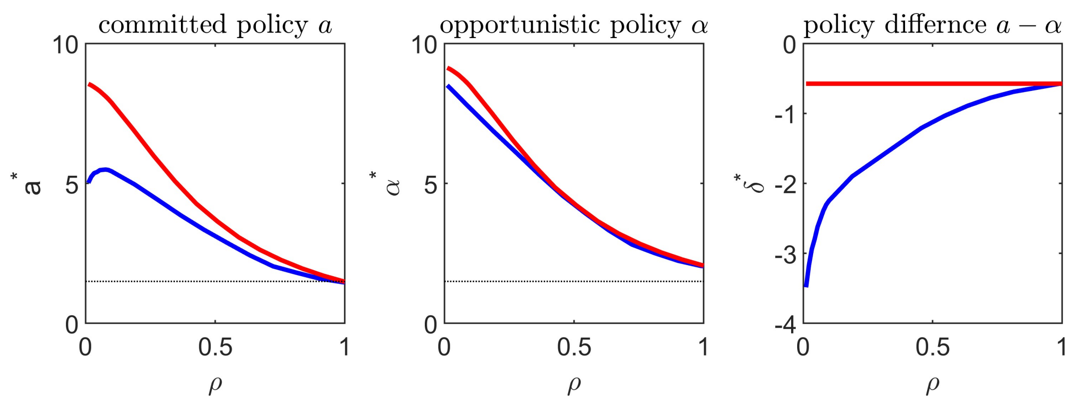

Strategic reputation management makes committed policy less inflationary, particularly when reputation ρ is low, both absolutely and relative to opportunistic policy. Each panel compares policy functions from the benchmark model with strategic reputation management (blue) and the naive model (red).

6.1 Equilibrium policy functions

We start by illustrating how the equilibrium inflation policies vary with reputation – the private agents' belief – in the calibrated benchmark model, where the committed policymaker strategically manages expectations by influencing private-sector learning. We then contrast these policies with those from a model in which the committed policymaker acts *naively* – that is, without attempting to shape reputation, treating it instead as an exogenous factor that determines how much influence policy a has over inflation expectations.[49]

[48] Many prior studies (footnote 12) examine how private agents update their beliefs about the policymaker's type, but typically abstract from the strategic actions a policymaker might take to influence those beliefs.

[49] We thank Davide Debortoli for suggesting the analysis of naive policy. Appendix D details the naive optimization problem, building on Cogley and Sargent (2008) and Kreps (1998).

Figure 5 shows the equilibrium policy functions for the committed policy a (left panel), the opportunistic policy α (middle panel), and their difference $\delta = a - \alpha$ (right panel), comparing results from the benchmark model with strategic reputation management (blue) and the naive model (red).[50]

The difference between the two models is most pronounced in the right panel. Under the benchmark model, the policy gap δ widens as reputation declines: a committed policymaker with low reputation adopts a more aggressive stance to distinguish himself from the opportunistic type and to accelerate reputation building. In contrast, under the naive model, the policy difference remains flat across reputation levels, even though both types optimize.

To understand why this occurs, consider the optimization problem in Proposition 2. First, the optimal δ that maximizes the momentary objective $\underline{\underline{u}}(\cdot)$ is zero, because the opportunistic policy α is already chosen to maximize the same objective as the committed policymaker. Second, when next-period reputation is treated as exogenous to current policies, δ drops out of both the expectation function $e(\cdot)$ and the continuation value term $\Omega(\cdot)$. As a result, the only remaining force generating a nonzero δ is the continuation penalty term $-\mu\underline{\underline{\omega}}(\cdot)$. But since δ enters $\omega(\cdot)$ in a way that is independent of reputation ρ and the cost-push shock ς, the optimal policy difference becomes insensitive to these state variables.

This stark contrast highlights a key implication of strategic reputation management: while belief updating in the standard learning literature typically responds to the policy difference across regimes, here the policy difference itself responds to the private sector's current belief. This feedback loop – where beliefs influence policy, and policy in turn influences beliefs – is essential for capturing the dynamics of episodes like the Volcker disinflation and for explaining the time-varying sensitivity of forecast revisions to inflation surprises.

6.2 Volcker disinflation

We now use the time period between 1979Q3 and 1987Q4 to illustrate how strategic reputation management makes a difference in time series. Our quantitative results show that it is a period when a committed policy regime is more plausible but reputation starts at a low level. We construct the counterfactual dynamics of reputation, inflation, and intended inflation policies under a naive committed policymaker, using the same cost-push shocks,

[50] The policy functions are conditional on the other two state variables, μ and ς, which are held fixed for illustration. Compared to ρ, these states play a less significant role in shaping the policy functions. We set μ at its steady-state level under $\rho = 1$ and $\varsigma = 0$. Appendix Figure 12 plots the equilibrium policy difference, $a - \alpha$, as a function of ρ across various alternative values of the cost-push shock and μ.

implementation errors, and probabilities of states as estimated in our benchmark model.[51] Figure 6 plot them against the benchmark case.

Figure 6: Counterfactual under Volcker

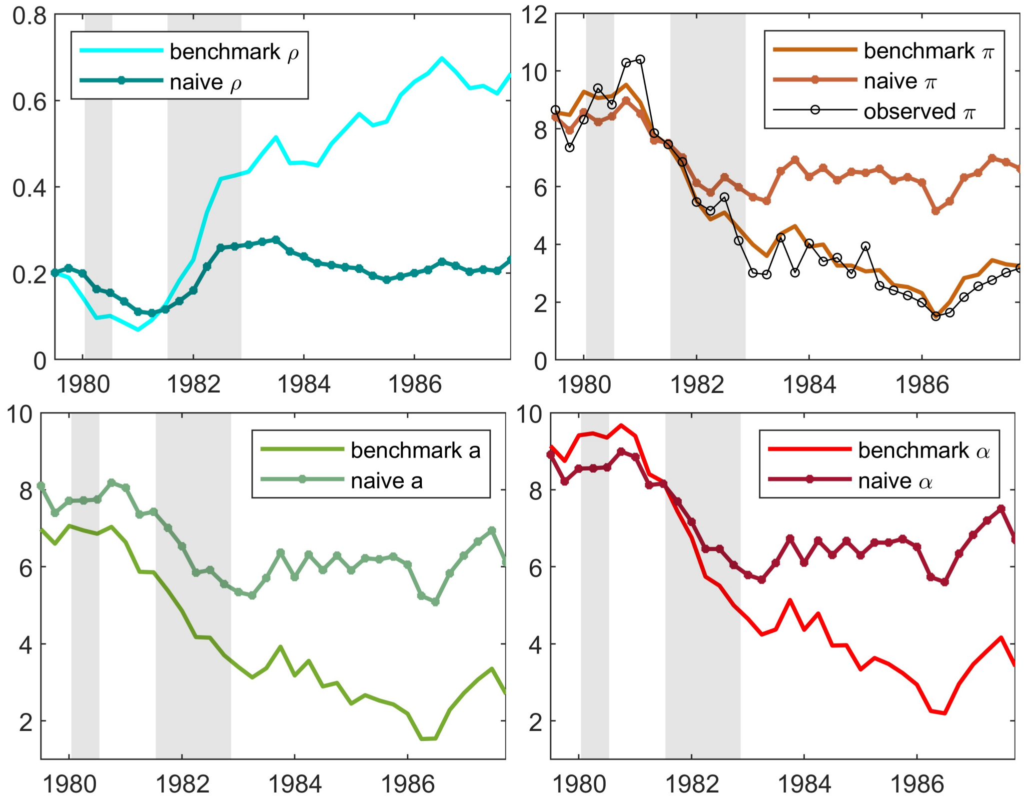

This figure plots the time series of reputation ρ (upper left panel), inflation π (upper right panel), committed policy a (lower left panel), and opportunistic policy α (lower right panel) from 1979Q2 to 1987Q4 when Paul Volcker was the Fed chairman. The solid lines are produced by the benchmark model and the lines with star markers are produced by the naive policy model, using the same cost-push shocks, implementation errors and probabilities of states as estimated in our benchmark model. In the benchmark model, the committed policy after 1981 is much lower than the opportunistic policy, resulting in a rapid increase in reputation. By contrast, in the naive model where the incentives to manage reputation are missing, it takes much longer for the committed policymaker to disinflate the economy after 1981, and reputation remains low for the entire period under Volcker.

In the benchmark model, reputation rises rapidly after 1981—from below 0.1 to over 0.4 by the end of the 1982–83 recession, and above 0.6 by 1987. Model-implied inflation closely tracks the Volcker disinflation, falling from around 10% in 1981 to 4% by the end of the

[51] The initial values for ρ and μ at 1979Q3 are set at the same level as in the benchmark model.

1982–83 recession. This rapid disinflation is driven by aggressive committed policies that remain two to three percentage points below opportunistic policies during 1981–82.

In contrast, under the naive model, reputation gains little during the 1982–83 recession and stays low – around 0.2 – for an extended period. This is because the naive committed policymaker treats reputation as exogenous, leading to a policy path that closely mirrors the opportunistic one. Low reputation sustains high expected inflation, worsening the inflation-output trade-off. As a result, disinflation is much slower, deviating sharply from the post-1981 U.S. inflation experience.

6.3 Long-term inflation forecasts

In this section, we provide empirical evidence for a key implication of strategic reputation management: the sensitivity of long-term inflation forecast revisions to inflation surprises is not constant but varies over time in a manner consistent with our benchmark model. This reflects the feedback loop highlighted earlier – where beliefs influence policy, and policy in turn shapes beliefs. To test this implication, we adopt a reduced-form approach that isolates the model's qualitative predictions without relying on short-term SPF forecasts, which were used in the structural estimation.

6.3.1 Long-term inflation forecast proportional to reputation

Our theory implies that long-term inflation forecasts depend solely on reputation, not on cost-push shocks, because cost-push shocks are stationary and their influence becomes negligible at long horizons. Formally, the period-t long-term inflation forecast is given by:

$$f_{\infty|t} = \rho_t[(1-q)\pi^* + qz(\varsigma = 0, \rho = 1)] + (1-\rho_t)[(1-q)\pi^{NE} + qz(\varsigma = 0, \rho = 0)] \tag{19}$$

which is a weighted average of long-term inflation forecasts conditional on the policymaker's type. With probability ρ_t, the policymaker is committed, and reputation converges to 1 in the long run. In this case, long-term inflation equals the target π^* if the regime persists, or the "startup" inflation associated with a new policymaker following the committed regime. Conversely, with probability $1-\rho_t$, the policymaker is opportunistic, and reputation converges to 0. The corresponding long-term inflation is either the Nash-equilibrium inflation bias π^{NE} if the regime continues, or the startup level associated with a new policymaker following the opportunistic regime. Since inflation is always lower under committed regime, (19) implies that long-term inflation forecasts move inversely with reputation.

6.3.2 Dynamics of reputation

We now examine the dynamics of reputation. According to the Bayes' rule (4), reputation is a function of observed inflation π and evolves as a *martingale* from the perspective of the private sector. This follows from the fact that the likelihood of observing π_t is $\rho_t g(\pi_t|a_t) + (1-\rho_t)g(\pi_t|\alpha_t)$, implying that the expected update in reputation satisfies $E_t\rho_{t+1} = \rho_t$.

Approximating reputation dynamics around the period-t nowcast, $[\rho_t a_t + (1-\rho_t)\alpha_t]$, we obtain the following expression for the change in reputation:[52]

$$\rho_{t+1} - \rho_t \approx k\{\rho_t(1-\rho_t)(a_t - \alpha_t)\}\{\pi_t - [\rho_t a_t + (1-\rho_t)\alpha_t]\} \tag{20}$$

where $k > 0$ depends inversely on the volatility of implementation errors, and the term $\pi_t - [\rho_t a_t + (1-\rho_t)\alpha_t]$ represents the inflation surprise. Under this approximation, the expected change in reputation is zero, and a positive inflation surprise reduces reputation, as the policy difference $a_t - \alpha_t$ is always negative.

Notably, the coefficient on the inflation surprise – the first term in braces – depends on reputation both directly and indirectly through the policy difference. Strategic reputation management implies that lower reputation leads to a wider policy difference. Thus, the overall sensitivity of reputation to inflation surprises increases with $[\rho(1-\rho)](1-\rho)$, where the second $(1-\rho)$ captures the reduced-form effect of reputation on policy difference.

Finally, (19) implies a linear relationship between reputation and long-term inflation forecasts. Consequently, revisions in long-term forecasts inherit the time-varying sensitivity of reputation changes to inflation surprises.

6.3.3 Time variation in sensitivity to inflation news

To study empirically how revisions in long-term inflation forecasts respond to news about inflation, we employ the SPF CPI inflation forecast data that includes both nowcast and 10-year-ahead forecast available starting in 1991Q4. Our sample covers up to 2024Q2.

To construct the "inflation surprise", we utilize the "backcast" estimate of the prior quarter inflation and a "nowcast" estimate of the current quarter inflation:

$$v_t^i = f_{t|t+1}^i - f_{t|t}^i$$

[52] This approximation becomes exact in continuous-time versions of imperfect public monitoring games, such as Eq (2) in Faingold and Sannikov (2011).

The "backcast" $f^i_{t|t+1}$ is the realized CPI inflation in t as perceived by the forecaster i. $f^i_{t|t}$ is the same forecaster's "nowcast" of inflation of the same quarter. For each forecaster, we can also construct 10-year forecast revision as

$$v^i_{40,t} = f^i_{t+41|t+1} - f^i_{t+40|t}.$$

We aggregate individual variables by restricting to forecasters present in both periods t and $t+1$, and then take the mean as the consensus measure: v_t for inflation surprise, $v_{40,t}$ for long-term forecast revision, and $f_{t+40|t}$ for the long-term inflation forecast. We then use $f_{t+40|t}$ to construct a reduced-form measure of reputation, $\widehat{\rho}_t$, based on (19).

We compare three regression models, each corresponding to a class of learning models with different structural features:

$$\text{Model 1: } v_{40,t} = \xi v_t + \epsilon_t$$

$$\text{Model 2: } v_{40,t} = \xi \widehat{\rho}_t (1 - \widehat{\rho}_t) v_t + \epsilon_t$$

$$\text{Model 3: } v_{40,t} = \xi \widehat{\rho}_t (1 - \widehat{\rho}_t)^2 v_t + \epsilon_t$$

Model 1 uses unweighted inflation surprise as the regressor, which would be appropriate if belief updating did not involve learning about the policymaker's type or policy regime. Model 2 weights the surprise by $\widehat{\rho}_t(1 - \widehat{\rho}_t)$, consistent with the idea that forecast revisions are proportional to belief updating – beliefs respond more to shocks when uncertainty about policymaker's type or policy regime is higher. This specification assumes the policy difference across types or regimes is constant. Model 3 modifies the weight to $\widehat{\rho}_t(1-\widehat{\rho}_t)^2$, capturing the interaction between policy and reputation: lower reputation increases the policy difference and accelerates private-sector learning – a mechanism unique to our model with strategic reputation management.

Table 2 reports each model's performance in terms of adjusted R^2 and root mean squared error (RMSE). Compared to the unweighted baseline, the weighted specifications offer substantial gains in explanatory power: adjusted R^2 rises from 0.163 in Model 1 to 0.193 in Model 2, and further to 0.248 in Model 3; RMSE falls from 0.117 to 0.114 and then to 0.110.

Figure 7 visualizes the comparison by plotting the cumulative change in long-term inflation forecasts alongside fitted values from the three models. The black dashed line shows the data; the magenta, blue, and red lines correspond to Models 1, 2, and 3, respectively. Model

3 clearly provides the best fit, especially before 2000, when long-term inflation forecasts were being revised downward—indicating rising reputation.

Table 2: Comparison of three regression models

	ξ	p value	Adjusted R^2	RMSE	N
Model 1	0.038	1.13E-06	0.163	0.117	130
Model 2	0.384	1.04E-07	0.193	0.114	130
Model 3	3.078	1.02E-09	0.248	0.11	130

OLS regressions of long-term CPI inflation forecast revision on inflation surprise. Sample period is from 1991Q4 to 2024Q2. Model 1 uses the unweighted inflation surprise. Models 2 and 3 weight inflation surprise using $\widehat{\rho}_t(1-\widehat{\rho}_t)$ and $\widehat{\rho}_t(1-\widehat{\rho}_t)^2$, respectively, where $\widehat{\rho}_t$ is constructed using the long-term inflation forecast according to (19). RMSE stands for Root Mean Squared Error. All regressions include a constant, but its coefficients are statistically insignificant.

Figure 7: Cumulative change in long-term CPI inflation forecast since 1991Q4

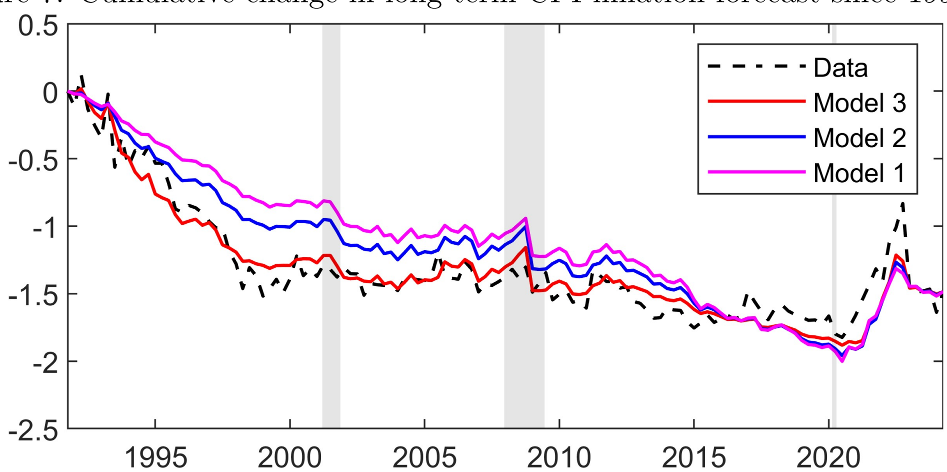

The cumulative change in the long-term CPI inflation forecast from 1991Q4 to 2024Q2 against its fitted counterparts produced by three models. The black dash line is the data, the magenta, blue, and red lines correspond to the fitted line using Models 1, 2, and 3, respectively. Model 1 uses the unweighted inflation surprise. Models 2 and 3 weight inflation surprise using $\widehat{\rho}_t(1-\widehat{\rho}_t)$ and $\widehat{\rho}_t(1-\widehat{\rho}_t)^2$, respectively, where $\widehat{\rho}_t$ is constructed using the long-term inflation forecast according to (19).

7 Summary, Conclusions and Final Remarks

We characterize the recursive equilibrium of a dynamic game that features two types of purposeful policymakers, a committed type which can commit and an opportunistic type which cannot, and private agents who are Bayesian learners about policymaker type and form forward-looking expectations of future policies. In the game, the committed policymaker strategically uses his policy plan to influence private agents' learning and inflation expectations, understanding that (i) private agents inflation expectations include future policy of an opportunistic type; and (ii) an opportunistic type's optimal policy depends on private agents' inflation expectations.

Harnessing the insights of modern contract theory, we develop a computable recursive equilibrium, where the equilibrium policies of both policymaker types and the rational expectations of private agents are shown to be functions of only three state variables, including an important reputation state that captures the evolution of private agents' beliefs about the commitment capacity of current policymaker. Using SPF short-term inflation forecasts data as observables, we extract latent states of the model via a nonlinear filter based on our theoretical model's dynamic system. The model-implied inflation tracks US inflation's rise, fall, and stabilization to a surprising high degree, even though the observed inflation is not used by the nonlinear filter.

Our results reveal that evolving reputation is central to understanding the interplay between inflation expectation and policy. In particular, without the incentives of rebuilding reputation, a switch to committed policy regime in 1981 cannot account for the observed post-1981 Volcker disinflation. Moreover, long-term inflation forecasts depend on reputation that evolves through Bayesian updating of inflation forecast errors. Our model implies a nonlinear relation between reputation and the sensitivity of long-term inflation forecast to forecast errors, which is supported by empirical evidence from regressions of SPF long-term forecast revision on nowcast forecast errors.

Our model is deliberately stark. But it yields results that have surprised us and others. We believe its success in matching U.S. experience of inflation and inflation expectations indicates great promise to further research on leveraging this theory framework to guide how to manage expectations in optimal policy.

References

Amador, Manuel and Christopher Phelan, "Reputation and sovereign default," *Econometrica*, 2021, *89* (4), 1979–2010. Publisher: Wiley Online Library.

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

Backus, David A. and John Driffill, "Inflation and Reputation.," *American Economic Review*, 1985, *75(3)*, 530–538.

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

Barro, Robert J., "Reputation in a model of monetary policy with incomplete information," *Journal of Monetary Economics*, January 1986, *17* (1), 3–20.

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

Beaudry, Paul, Chenyu Hou, and Franck Portier, "The Dominant Role of Expectations and Broad-Based Supply Shocks in Driving Inflation," in "NBER Macroeconomics Annual 2024, volume 39," University of Chicago Press, May 2024.

Benigno, Gianluca, Andrew Foerster, Christopher Otrok, and Alessandro Rebucci, "Estimating Macroeconomic Models of Financial Crises: An Endogenous Regime-Switching Approach," April 2020.

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.

Binning, Andrew and Junior Maih, "Sigma Point Filters for Dynamic Nonlinear Regime Switching Models," *SSRN Electronic Journal*, 2015.

Blinder, Alan S., Michael Ehrmann, Jakob de Haan, and David-Jan Jansen, "Central Bank Communication with the General Public: Promise or False Hope?," *Journal of Economic Literature*, June 2024, *62* (2), 425–457.

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," 2023, *15* (1), 1–47.

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.

Clayton, Christopher, Amanda Dos Santos, Matteo Maggiori, and Jesse Schreger, "Internationalizing Like China," *American Economic Review*, March 2025, *115* (3), 864–902.

Cogley, Timothy and Thomas J. Sargent, "Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making," *International Economic Review*, 2008, *49* (1), 185–221.

— , 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.

— , — , and Saten Kumar, "How Do Firms Form Their Expectations? New Survey Evidence," *American Economic Review*, September 2018, *108* (9), 2671–2713.

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.

— and — , "Reputation, Bailouts, and Interest Rate Spread Dynamics," *American Economic Journal: Macroeconomics*, July 2022, *14* (3), 411–449.

Drautzburg, Thorsten, Jesus Fernandez-Villaverde, Pablo Guerron-Quintana, and Dick Oosthuizen, "Filtering with limited information," 2022.

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.

Faingold, Eduardo and Yuliy Sannikov, "Reputation in Continuous-Time Games," *Econometrica*, May 2011, *79* (3), 773–876.

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.

Foerster, Andrew and Christian Matthes, "Learning About Regime Change," *International Economic Review*, 2022, *63* (4), 1829–1859.

Goodfriend, Marvin 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.

Haldane, Andrew and Michael McMahon, "Central Bank Communications and the General Public," *AEA Papers and Proceedings*, May 2018, *108*, 578–583.

Hamilton, James D., "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 1989, *57* (2), 357–384.

Hansen, Stephen and Michael McMahon, "First Impressions Matter: Signalling as a Source of Policy Dynamics," *The Review of Economic Studies*, February 2016, *83* (4), 1645–1672.

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.

Kim, Chang-Jin, "Dynamic linear models with Markov-switching," *Journal of Econometrics*, January 1994, *60* (1-2), 1–22.

_ and Charles R. Nelson, *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*, The MIT Press, November 2017.

King, Robert G. and Yang K. Lu, "Evolving Reputation for Commitment: The Rise, Fall and Stabilization of US Inflation," December 2022.

— , — , **and Ernesto S. Pastén**, "Managing Expectations," *Journal of Money, Credit and Banking*, December 2008, *40* (8), 1625–1666.

Kollmann, Robert, "Tractable likelihood-based estimation of non-linear DSGE models," *Economics Letters*, December 2017, *161*, 90–92.

Kreps, David M., "Anticipated Utility and Dynamic Choice (1997)," in Donald P. Jacobs, Ehud Kalai, Morton I. Kamien, and Nancy L. Schwartz, eds., *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983–1997*, Econometric Society Monographs, Cambridge: Cambridge University Press, 1998, pp. 242–274.

Kreps, David M and Robert Wilson, "Reputation and imperfect information," *Journal of Economic Theory*, August 1982, *27* (2), 253–279.

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.

Lucas, Robert E. and Thomas J. Sargent, "After Keynesian macroeconomics," *Quarterly Review, Federal Reserve Bank of Minneapolis*, 1979, *3* (Spr). Publisher: Federal Reserve Bank of Minneapolis.

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.

Milgrom, Paul and John Roberts, "Limit pricing and entry under incomplete information: An equilibrium analysis," *Econometrica*, 1982, *50* (2), 443–460.

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

Morelli, Juan M. and Matias Moretti, "Information Frictions, Reputation, and Sovereign Spreads," *Journal of Political Economy*, March 2023, pp. 000–000. Publisher: The University of Chicago Press.

Orphanides, Athanasios 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, "Public trust and government betrayal," *Journal of Economic Theory*, September 2006, *130* (1), 27–43.

— **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.

Rouwenhorst, K. Geert, "Asset Pricing Implications of Equilibrium Business Cycle Models," in Thomas F. Cooley, ed., *Frontiers of Business Cycle Research*, Princeton University Press, 1995, pp. 294–330.

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.

Särkkä, Simo and Lennart Svensson, *Bayesian Filtering and Smoothing*, 2023 ed., Cambridge University Press, 2023.

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

Appendices

A Recursive optimal policy design

The optimal policy problem for the committed type at the start of his 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 derives the recursive program stated in Proposition 1. Key elements from the main text are repeated to ensure the appendix is self-contained. The derivation proceeds step by step, accommodating readers with varying familiarity with recursive optimal policy design. A distinctive feature of this application is the "change of measure" in the expectations constraint faced by the committed policymaker, which arises because private agents understand that inflation may result from an optimizing opportunistic type.

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 his replacement. 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 impose this requirement in Section 4.3 of the main text.

A.1 Intended and actual inflation

At each date, the policymaker chooses intended inflation, denoted as a for the committed type ($\tau = 1$) and α for the opportunistic type ($\tau = 0$). Intended inflation is not observed by the private sector. Actual inflation is randomly distributed around this intention:

$$\pi_t = \tau_t a_t + (1 - \tau_t)\alpha_t + v_{\pi,t}. \tag{A21}$$

where $v_{\pi,t}$ is an i.i.d. implementation error and $v_{\pi,t} \sim g(\cdot)$ with $g(\cdot) = N(0, \sigma^2_{v,\pi})$. With a slight abuse of notation, we use $g(\pi|a)$ and $g(\pi|\alpha)$ to denote the density of inflation conditional on the intended inflation.

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) and the expectations term e in the output-inflation trade-off, $\pi = e + \kappa x + \varsigma$, are both conditioned on this information. 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. 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 opportunistic type which will choose α_t. Accordingly, their rational likelihood of the outcome π_t is

$$g(\pi_t|a_t)\rho_t + g(\pi_t|\alpha_t)(1 - \rho_t) \quad \text{(A22)}$$

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

(A23) $$\rho(h_t^+) = \frac{g(\pi_t|a(h_t))\rho(h_t)}{g(\pi_t|a(h_t))\rho(h_t) + g(\pi_t|\alpha(h_t))(1-\rho(h_t))}$$

As no information about type is revealed by ς_{t+1}, $\rho(h_{t+1}) = \rho(h_t^+)$. This updating may be written as

(A24) $$\rho(h_{t+1}) \quad = \quad \frac{\rho(h_t)}{\rho(h_t) + \lambda(\pi_t, h_t)(1-\rho(h_t))}$$

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 private agents' expected inflation, $E_t\pi_{t+1}$, 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 ($\theta_{t+1} = 1$) with probability q and won't ($\theta_{t+1} = 0$) with probability $1 - q$.

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

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

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

$$E(\pi_{t+1}|h_{t+1}, \tau_{t+1} = 0) = \alpha([\varsigma_{t+1}, h_t^+])$$

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

(A25) $$E(\pi_{t+1}|h_{t+1}, \theta_{t+1} = 0) = \rho(h_{t+1})a(h_{t+1}) + (1-\rho(h_{t+1}))\alpha(h_{t+1})$$

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

(A26) $$E(\pi_{t+1}|h_{t+1}, \theta_{t+1} = 1) = z(h_{t+1})$$

as private agents' expectation of inflation conditional on a replacement.

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

$$\text{(A27)}\quad E(\pi_{t+1}|h_t) = \rho(h_t)\int\sum_{\varsigma_{t+1}}\varphi(\varsigma_{t+1};\varsigma_t)\left[(1-q)\,a(h_{t+1}) + qz\,(h_{t+1})\right]g(\pi_t|a(h_t))d\pi_t$$
$$+\,(1-\rho\,(h_t))\int\sum_{\varsigma_{t+1}}\varphi(\varsigma_{t+1};\varsigma_t)\left[(1-q)\,\alpha(h_{t+1}) + qz\,(h_{t+1})\right]g(\pi_t|\alpha(h_t))d\pi_t$$

There may appear to be a conflict between this expression and (A25) that contains reputation at t+1. But there is not. Weighting (A25) and (A26) by $(1-q)$ and q and then integrating over private agents' belief about inflation (A22) leads directly to it. The simplicity arises because (A22) also occurs in the denominator of the Bayesian updating expression (A23).

A.6 Intertemporal objective

We assume that the policymaker's intertemporal objective involves discounting at $\beta_a(1-q)$, where β_a is his 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

$$\text{(A28)}\qquad U_t = \sum_{j=0}^{\infty}(\beta_a(1-q))^j\sum_{h_{t+j}}\frac{p(h_{t+j})}{p(h_t)}\underline{u}(a(h_{t+j}), e(h_{t+j}), \varsigma(h_{t+j}))$$

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:

$$\text{(A29)}\qquad p(h_{t+j}) = \varphi(\varsigma_{t+j};\varsigma_{t+j-1}) \times g(\pi_{t+j-1}|a(h_{t+j-1})) \times p(h_{t+j-1})$$

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.[1]

[1] 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.

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

$$\text{(A30)} \qquad \Psi_t = \sum_{j=0}^{\infty}(\beta_a(1-q))^j \sum_{h_{t+j}} \frac{p(h_{t+j})}{p(h_t)}\gamma(h_{t+j})[e(h_{t+j}) - \beta E(\pi_{t+j+1}|h_{t+j})]$$

and then express it recursively. We detailed $E(\pi_{t+1}|h_t)$ in (A27), but the expression involved the probability of inflation under the opportunistic type. So, we undertake a "change of measure" and rewrite it as

$$\text{(A31)} \qquad \rho(h_t)\int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1};\varsigma_t)[\beta(1-q)a(h_{t+1}) + \beta q z(h_{t+1})]g(\pi|a(h_t))d\pi$$

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

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

$$\text{(A32)} \qquad \frac{g(\pi_t|\alpha(h_t))}{g(\pi_t|a(h_t))} = \lambda(h_t^+) = \lambda(h_{t+1})$$

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 (A30) and replace $E(\pi_{t+1}|h_t)$ with the expression in (A31). 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 discount factors, we can write (A30) as

$$\text{(A33)} \qquad \Psi_t = E_t^c[\sum_{j=0}^{\infty}(\beta_a(1-q))^j \psi_{t+j}]$$

with

$$\text{(A34)} \qquad \psi_t = \gamma_t e_t - \frac{\beta}{\beta_a(1-q)}\gamma_{t-1}\{\rho_{t-1}[(1-q)a_t + qz_t] + (1-\rho_{t-1})\lambda_t[(1-q)\alpha_t + qz_t]\}$$

This latter expression captures past commitments about current state-contingent decisions as these were relevant to past expectations of inflation.[2] 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

$$
\begin{aligned}
\text{(A35)}\quad W(\varsigma_t, \eta_t, \rho_{t-1}, \lambda_t) = \min_{\gamma} \max_{a,\alpha,e} \{ & \underline{u}(a_t, e_t, \varsigma_t) + \gamma_t e_t \\
& - \eta_t[\rho_{t-1}((1-q)a_t + qz_t) + (1-\rho_{t-1})\lambda_t((1-q)\alpha_t + qz_t)] \\
& + \beta_a(1-q) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) W(\varsigma_{t+1}, \eta_{t+1}, \rho_t, \lambda_{t+1}) g(\pi_t|a_t) d\pi_t \}
\end{aligned}
$$

subject to the IC constraint

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

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

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

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 (A34) as

$$\text{(A36)}\qquad \psi_t = \gamma_t e_t - \frac{\beta}{\beta_a(1-q)}\gamma_{t-1}\rho_{t-1}\{[(1-q)a_t + qz_t] + \frac{(1-\rho_{t-1})\lambda_t}{\rho_{t-1}}[(1-q)\alpha_t + qz_t]\}$$

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

$$\text{(A37)}\qquad \psi_t = \gamma_t e_t - \frac{\beta}{\beta_a(1-q)}\gamma_{t-1}\rho_{t-1}\{[(1-q)a_t + qz_t] + \boxed{\frac{(1-\rho_t)}{\rho_t}}[(1-q)\alpha_t + qz_t]\}$$

[2] Our short hand notation replaces $\lambda(h_t)$ with λ_t. Given (A32), the likelihood ratio λ_t is predetermined in period t by actions and inflation outcome in period $t-1$.

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:

$$\text{(A38)} \qquad \mu_{t+1} = \frac{\beta}{\beta_a(1-q)}\gamma_t\rho_t \text{ with } \mu_0 = 0$$

$$\text{(A39)} \qquad \rho_{t+1} = \frac{\rho_t g(\pi_t|a_t)}{\rho_t g(\pi_t|a_t) + (1-\rho_t)g(\pi_t|\alpha_t)} \text{ with probability } g(\pi_t|a_t)$$

The recursive optimization (A35) can now be written with only three state variables $(\varsigma_t, \rho_t, \mu_t)$ as stated in Proposition 1.

Proposition 1. Given $z(\varsigma, \rho)$, the within-regime equilibrium is the solution to:

$$\text{(A40)} \qquad W(\varsigma, \rho, \mu) = \min_{\gamma}\max_{a,\alpha,e}\{\underline{u}(a, e, \varsigma) + [\gamma e - \mu\omega(a, \alpha, \rho, \varsigma)] + \beta_a(1-q)\int\sum_{\varsigma'}\varphi(\varsigma';\varsigma)\,W(\varsigma', \rho', \mu')\,g(\pi|a)\,d\pi\},$$

$$\text{(A41)} \qquad \text{with } \omega(a, \alpha, \rho, \varsigma) \equiv (1-q)\,a + qz(\varsigma, \rho) + \frac{1-\rho}{\rho}\left[(1-q)\,\alpha + qz(\varsigma, \rho)\right]$$

$$\text{(A42)} \qquad \alpha = Ae + B(\varsigma)$$

$$\text{(A43)} \qquad \mu' = \frac{\beta}{\beta_a(1-q)}\gamma\rho, \text{ given } \mu_0 = 0$$

$$\text{(A44)} \qquad \rho' = \frac{\rho g(\pi|a)}{\rho g(\pi|a) + (1-\rho)\,g(\pi|\alpha)} \text{ with prob } g(\pi|a), \text{ given } \rho_0$$

A.10 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 $\mu_0 = 0$.

B Consolidation

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 inefficient for computation. This appendix explains how we consolidate choice variables by exploring the implications of private agents' rational expectation constraint.

B.1 The rational expectation function

We now show that imposing the rational expectation constraint (A31) on the choice of e_t implies an operational expectation function:

Lemma 1. Given (ς, ρ) and future equilibrium strategies $a^*(\varsigma', \rho', \mu')$, $\alpha^*(\varsigma', \rho', \mu')$ and $z^*(\varsigma', \rho')$, rationally expected inflation is uniquely determined by δ and μ';

(B1)
$$\begin{aligned}
e = e(\delta, \mu'; \varsigma, \rho) &= \beta\rho \int M_a(\varsigma, b(v_\pi, v_\pi + \delta, \rho), \mu') g(v_\pi) dv_\pi + \\
&\quad \beta(1-\rho) \int M_\alpha(\varsigma, b(v_\pi - \delta, v_\pi, \rho), \mu') g(v_\pi) dv_\pi; \\
\text{where } b(\pi - a, \pi - \alpha, \rho) &\equiv \frac{g(\pi - a)\rho}{g(\pi - a)\rho + g(\pi - \alpha)(1-\rho)}; \\
M_a(\varsigma, \rho', \mu') &\equiv \textstyle\sum_{\varsigma'} \varphi(\varsigma'; \varsigma) \left[(1-q) a^*(\varsigma', \rho', \mu') + q z^*(\varsigma', \rho')\right]; \\
M_\alpha(\varsigma, \rho', \mu') &\equiv \textstyle\sum_{\varsigma'} \varphi(\varsigma'; \varsigma) \left[(1-q) \alpha^*(\varsigma', \rho', \mu') + q z^*(\varsigma', \rho')\right].
\end{aligned}$$

Proof. Before taking a "change of measure", the rational expectation constraint on e is:

(B2)
$$e = \beta\rho \int \sum_{\varsigma'} \varphi(\varsigma'; \varsigma) \left[(1-q) a' + q z'\right] g(\pi | a) d\pi$$
$$+ \beta(1-\rho) \int \sum_{\varsigma'} \varphi(\varsigma'; \varsigma) \left[(1-q) \alpha' + q z'\right] g(\pi | \alpha) d\pi$$

with a', α', and z' determined by the three states $(\varsigma', \rho', \mu')$ through the equilibrium strategies: $a^*(\cdot)$, $\alpha^*(\cdot)$, and $z^*(\cdot)$.

According to the Bayes' rule (A23), ρ' is the function of π, where $\pi = a + v_\pi$ under the committed type and $\pi = \alpha + v_\pi$ under the opportunistic type, with v_π being zero mean random variables. We can therefore rewrite ρ' as

(B3)
$$\rho' = b(\pi - a, \pi - \alpha, \rho) = \begin{cases} b(v_\pi, v_\pi + \delta, \rho) & \text{if } \tau = 1 \\ b(v_\pi - \delta, v_\pi, \rho) & \text{if } \tau = 0 \end{cases}$$

Replacing $g(\pi|a)$ and $g(\pi|\alpha)$ in (B2) with $g(v_\pi)$, ρ' with (B3), and realizing choosing γ is equivalent to choosing μ' due to $\mu' = \frac{\beta}{\beta_a(1-q)}\gamma\rho$, we obtain the rational expectation function in (B1). ∎

B.2 Relation between W and U

We now show that the committed policymaker's value function equals his optimized intertemporal objective minus the cost of honoring past promises, captured by the term $\mu\omega^*$.

Lemma 2. Let $U^*(s)$ and $\omega^*(s)$ be the intertemporal objective (A28) and the composite promise term in (A41) evaluated at optimal decision rules, then

$$W\left(\varsigma,\rho,\mu\right) = U^*\left(\varsigma,\rho,\mu\right) - \mu\omega^*\left(\varsigma,\rho,\mu\right) \tag{B4}$$

Proof. In the recursive optimization problem (A40), the envelope theorem implies:

$$W_\mu(\varsigma,\rho,\mu) = -\{[(1-q)a + qz] + \boxed{\tfrac{(1-\rho)}{\rho}}[(1-q)\alpha + qz]\} = -\omega$$

The first order necessary condition for γ is

$$\begin{aligned} 0 &= e + \beta_a(1-q)\sum_{\varsigma'}\varphi(\varsigma';\varsigma)\int W_\mu(\varsigma',\rho',\mu')\frac{\partial\mu'}{\partial\gamma}g(\pi|a)d\pi \\ &= e + \beta\sum_{\varsigma'}\varphi(\varsigma';\varsigma)\int W_\mu(\varsigma',\rho',\mu')\rho g(\pi|a)d\pi \end{aligned}$$

where the state evolution equation (A38) implies $\partial\mu'/\partial\gamma = \rho\beta/(\beta_a(1-q))$.

When combined with an updated version of the envelope theorem implication, this FOC recovers private agents' rational expectation constraint as in (A31):

$$e = \beta\int\sum_{\varsigma'}\varphi(\varsigma';\varsigma)\left[[(1-q)a' + qz'] + \boxed{\tfrac{(1-\rho')}{\rho'}}[(1-q)\alpha' + qz']\right]\rho g(\pi|a)d\pi$$

where

$$\frac{1-\rho'}{\rho'} = \frac{(1-\rho)\,\lambda'}{\rho}.$$

Hence, in equilibrium where the rational expectation constraint must hold, we obtain

$$e^* = \beta\int\sum_{\varsigma'}\varphi(\varsigma';\varsigma)\omega^{*\prime}\rho g(\pi|a^*)d\pi$$

Utilizing this equilibrium condition, we now show by "guess and verify" that in equilibrium: $W(\varsigma,\rho,\mu) = U^*(\varsigma,\rho,\mu) - \mu\omega^*$. The following recursion must hold:

$$\begin{aligned} W(\varsigma,\rho,\mu) + \mu\omega^* &= \underline{u}(a^*,e^*,\varsigma) + \gamma e^* \\ &\quad + \beta_a(1-q)\int\sum_{\varsigma'}\varphi(\varsigma';\varsigma)W(\varsigma',\rho',\mu')g(\pi|a^*)d\pi \end{aligned} \tag{B5}$$

Suppose $W(\varsigma', \rho', \mu') = -\mu'\omega^{*\prime} + U^*(\varsigma', \rho', \mu')$, the right hand side can be written as

$$
\begin{aligned}
&\underline{u}(a^*, e^*, \varsigma) + \gamma e^* - \beta_a(1-q)\int \sum_{\varsigma'} \varphi(\varsigma';\varsigma)[\frac{\beta}{\beta_a(1-q)}\gamma\rho\omega^{*\prime}]g(\pi|a^*)d\pi \\
&+ \beta_a(1-q)\int \sum_{\varsigma'} \varphi(\varsigma';\varsigma)U^*(\varsigma', \rho', \mu')g(\pi|a^*)d\pi \\
=&\underline{u}(a^*, e^*, \varsigma) + \gamma[e^* - \beta\int \sum_{\varsigma'} \varphi(\varsigma';\varsigma)\omega^{*\prime}\rho g(\pi|a^*)d\pi] \\
&+ \beta_a(1-q)\int \sum_{\varsigma'} \varphi(\varsigma';\varsigma)U^*(\varsigma', \rho', \mu')g(\pi|a^*)d\pi \\
=&\underline{u}(a^*, e^*, \varsigma) + \beta_a(1-q)\int \sum_{\varsigma'} \varphi(\varsigma';\varsigma)U^*(\varsigma', \rho', \mu')g(\pi|a^*)d\pi \\
=&U^*(\varsigma, \rho, \mu)
\end{aligned}
$$

which implies $W(\varsigma, \rho, \mu) = U^*(\varsigma, \rho, \mu) - \mu\omega^*$. ∎

B.3 Simplified recursive program

Using Lemma 1 and 2, we simplify the recursive program (A40), moving from choosing (γ, a, α, e) to choosing (δ, μ'):

Proposition 2. Given $e = e(\delta, \mu'; \varsigma, \rho)$ and $U^*(\varsigma, \rho, \mu) = W(\varsigma, \rho, \mu) + \mu\omega^*(\varsigma, \rho, \mu)$,

(B6) $$
\begin{aligned}
W(\varsigma, \rho, \mu) = \max_{\delta,\mu'} &\left[\underline{\underline{u}}(\delta, \mu'; \varsigma, \rho) - \mu\underline{\underline{\omega}}(\delta, \mu'; \varsigma, \rho) + \beta_a(1-q)\,\Omega(\delta, \mu'; \varsigma, \rho)\right] \\
\text{with } \underline{\underline{u}}(\delta, \mu'; \varsigma, \rho) &\equiv \underline{u}(Ae + B(\varsigma) + \delta, e, \varsigma) \\
\underline{\underline{\omega}}(\delta, \mu'; \varsigma, \rho) &\equiv \frac{1}{\rho}\left[(1-q)\left(Ae + B(\varsigma)\right) + qz^*(\varsigma, \rho)\right] + (1-q)\,\delta \\
\Omega(\delta, \mu'; \varsigma, \rho) &\equiv \int \textstyle\sum_{\varsigma'} \varphi(\varsigma';\varsigma)U^*(\varsigma', b(v_\pi, v_\pi + \delta, \rho), \mu')\, g(v_\pi)dv_\pi
\end{aligned}
$$

Proof. Lemma 2 implies that the objective of the recursive optimization (A40) can be reduced to

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

Lemma 1 implies that (δ, μ') determines $e = e(\delta, \mu'; \varsigma, \rho)$, $\alpha = Ae + B(\varsigma)$, and $a = \alpha + \delta$. Because $\underline{u}(\cdot)$ and $\omega(\cdot)$ are functions of (e, α, a), they can be written as functions of (δ, μ'):

(B7) $$\underline{\underline{u}}(\delta, \mu') \equiv \underline{u}(Ae + B(\varsigma) + \delta, e, \varsigma)$$

(B8) $$\underline{\underline{\omega}}(\delta, \mu') \equiv \frac{1}{\rho}\left[(1-q)\left(Ae + B(\varsigma)\right) + qz^*(\varsigma, \rho)\right] + (1-q)\,\delta$$

The optimization is from the perspective of the committed policymaker so that $\pi = a + v_\pi$. Therefore, $\rho' = b(v_\pi, v_\pi + \delta, \rho)$ as defined in (B3) and $g(\pi|a) = g(v_\pi)$. We then obtain the simplified program. ∎

Computation: Lemma 1 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.[3] We then calculate a new $z(\varsigma, \rho)$ and repeat the process until the outer loop has reached a fixed point in z.

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

[3] 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.

compounded inflation rate is less prone than is the FRBP inflation rate.

Figure 8 contrasts the two measures.

Figure 8: Contrasting median inflation and change in median price

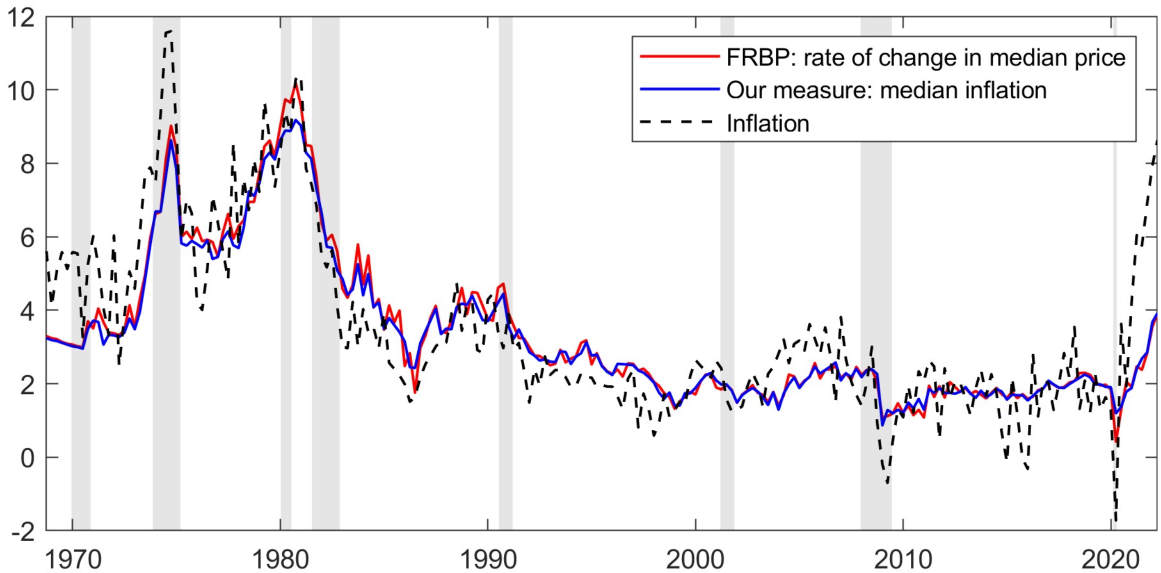

This figure compares two measures of the median SPF inflation forecast for the GDP deflator. The red line, from the summary data file ("annualized percent change of median response") constructed by the FRBP, is obtained by first finding the median price forecast at each horizon and then computing the annualized inflation rate from these median price forecasts. The blue line, based on the individual responses file, calculates annualized inflation for each individual forecast and then takes the median across individuals. We adopt the latter measure in the paper as it is less prone to transitory outliers.

C.2 Recursive forecasting in our theory

This appendix describes the calculation of private agents' expectations of inflation at each horizon j: $E(\pi_{t+j}|h_t)$.

The information set is assumed to be the start of period information of the private sector, $s_t = (\varsigma_t, \rho_t, \mu_t)$. We denote the forecast function using $f(s_t, j) = E(\pi_{t+j}|s_t)$.

Given s_t, private agents know the intended inflation policies of the committed and the opportunistic policymakers: $a(\varsigma_t, \rho_t, \mu_t)$ and $\alpha(\varsigma_t, \rho_t, \mu_t)$. Because implementation errors have mean zero, the private agents' "nowcast" of inflation is

$$f(\varsigma_t, \rho_t, \mu_t, 0) = \rho_t a(\varsigma_t, \rho_t, \mu_t) + (1 - \rho_t)\alpha(\varsigma_t, \rho_t, \mu_t)$$

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:

$$E(\pi_{t+j}|s_t) = f(\varsigma_t, \rho_t, \mu_t, j) = E[E(\pi_{t+j}|s_{t+1})|s_t] = E[f(\varsigma_{t+1}, \rho_{t+1}, \mu_{t+1}, j-1)|s_t] \tag{C1}$$

The pseudo state variable μ_{t+1} evolves according to:

$$\mu_{t+1} = \begin{cases} \mu'^{*}(\varsigma_t, \rho_t, \mu_t) & \text{with prob } 1-q \\ 0 & \text{with prob } q \end{cases}$$

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

$$\rho_{t+1} = \begin{cases} b(v_\pi, v_\pi + \delta, \rho_t) & \text{with prob } (1-q)\rho_t \\ b(v_\pi - \delta, v_\pi, \rho_t) & \text{with prob } (1-q)(1-\rho_t) \\ \phi_{t+1} b(v_\pi, v_\pi + \delta, \rho_t) + (1-\phi_{t+1}) v_{\rho,t+1} & \text{with prob } q\rho_t \\ \phi_{t+1} b(v_\pi - \delta, v_\pi, \rho_t) + (1-\phi_{t+1}) v_{\rho,t+1} & \text{with prob } q(1-\rho_t) \end{cases}$$

where $\phi_{t+1} \sim \text{Bernoulli}(\zeta_\rho)$ and $v_{\rho,t+1} \sim \text{Beta}(\overline{\rho}, \sigma_\rho)$. Therefore:

$$\begin{aligned} & f(\varsigma_t, \rho_t, \mu_t, j) = \sum \varphi(\varsigma_{t+1}; \varsigma_t) \Big\{ q(1-\zeta_\rho) \int f(\varsigma_{t+1}, v_\rho, 0, j-1) d\text{Beta}(v_\rho | \overline{\rho}, \sigma_\rho) \\ & (1-q)\rho_t \int f\big(\varsigma_{t+1}, b(v_\pi, v_\pi + \delta, \rho_t), \mu'^{*}(\varsigma_t, \rho_t, \mu_t), j-1\big) g(v_\pi) dv_\pi \\ & +(1-q)(1-\rho_t) \int f\big(\varsigma_{t+1}, b(v_\pi - \delta, v_\pi, \rho_t), \mu'^{*}(\varsigma_t, \rho_t, \mu_t), j-1\big) g(v_\pi) dv_\pi \\ & +q\rho_t\zeta_\rho \int f\big(\varsigma_{t+1}, b(v_\pi, v_\pi + \delta, \rho_t), 0, j-1\big) g(v_\pi) dv_\pi \\ & +q(1-\rho_t)\zeta_\rho \int f\big(\varsigma_{t+1}, b(v_\pi - \delta, v_\pi, \rho_t), 0, j-1\big) g(v_\pi) dv_\pi \Big\} \end{aligned}$$

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 the dates of policymaker replacement and the policymaker type within each regime: we have just three random inputs – cost-push shocks ς_t, implementation errors $v_{\pi,t}$, and reputation shocks $v_{\rho,t}$ – that drive many observable macro time series, including the policies a_t and α_t, inflation π_t, and, as we just discussed, expectations at various horizons $E_t(\pi_{t+j})$.

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.[4] 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.

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 private agents' inflation expectations at various horizons, the latter of which are measured by the SPF.

The state-space representation of our model can be written as follows

$$\text{(C2)} \qquad S_t = [\varsigma_t, \rho_t, \mu_t, \pi_t]' = F(S_{t-1}, v_t | \theta_t, \phi_t, \tau_t)$$

$$= \begin{bmatrix} \xi_\varsigma \varsigma_{t-1} + v_{\varsigma,t} \\ (1-\theta_t+\theta_t\phi_t) b(\varsigma_{t-1,}\rho_{t-1}, \mu_{t-1}, \pi_{t-1}) + \theta_t(1-\phi_t) v_{\rho,t} \\ (1-\theta_t) m(\varsigma_{t-1,}\rho_{t-1}, \mu_{t-1}) \\ \tau_t a(\varsigma_t, \rho_t, \mu_t) + (1-\tau_t)\alpha(\varsigma_t, \rho_t, \mu_t) + v_{\pi,t} \end{bmatrix}$$

$$\text{(C3)} \qquad Y_t = \begin{bmatrix} f_{t+1|t} \\ f_{t+3|t} \end{bmatrix} = \begin{bmatrix} f(\varsigma_t, \rho_t, \mu_t, 1) \\ f(\varsigma_t, \rho_t, \mu_t, 3) \end{bmatrix} + \begin{bmatrix} \varepsilon_{1t} \\ \varepsilon_{3t} \end{bmatrix} = H(S_t) + \varepsilon_t$$

The state vector collects the three state variables $(\varsigma_t, \rho_t, \mu_t)$ identified in Proposition 1 and inflation π_t. The state evolution equations are the stochastic processes of shocks and the equilibrium policy functions, conditional on $(\theta_t, \phi_t, \tau_t)$, representing the event of policymaker replacement ($\theta_t = 1$), continuing type in a new regime ($\phi_t = 1$), and committed type in place ($\tau_t = 1$).

The observable vector consists of the SPF at one quarter and three quarter horizons ($f_{t+j|t}$, j=1,3). The measurement equations are model-implied one-period and three-period ahead inflation forecasts by private agents. $\varepsilon_{j,t}$ is the normal measurement error with mean zero and standard deviation 0.5% at annualized rate.

We model $(\theta_t, \phi_t, \tau_t)$ as the outcome of an unobserved discrete-state Markov process

[4] 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.

Θ_t, with six discrete states:[5] $\{(\theta_t = 0, \tau_t = 1), (\theta_t = 0, \tau_t = 0), (\theta_t = 1, \phi_t = 1, \tau_t = 1), (\theta_t = 1, \phi_t = 1, \tau_t = 0), (\theta_t = 1, \phi_t = 0, \tau_t = 1), (\theta_t = 1, \phi_t = 0, \tau_t = 0)\}$. The transitional probability matrix $T_{i,j} = Pr(\Theta_t = j|\Theta_{t-1} = i)$ is determined by the structure of our model: 1) when $\theta_t = 0$, i.e., no replacement of policymaker, the policymaker type remains the same in period $t-1$ and t; 2) when $\theta_t = 1$ and $\phi_t = 1$, i.e., there is a new policymaker whose type is the same as his predecessor, the probability that a committed type will be in place in period t is the private agents' posterior belief at the end of period $t-1$, $b^i_{t-1} \equiv b(s_{t-1}, \pi_{t-1}|\Theta_{t-1} = i)$; 3) when $\theta_t = 1$ and $\phi_t = 0$, i.e., there is a new policymaker whose type is a random draw, the probability that a committed type will be in place in period t is the unconditional mean $\overline{\rho}$ of the reputation shock v_ρ.

(C4)
$$T = \begin{bmatrix} 1-q & 0 & \zeta_\rho b^{i=1}_{t-1} q & \zeta_\rho(1-b^{i=1}_{t-1})q & (1-\zeta_\rho)\overline{\rho}q & (1-\zeta_\rho)(1-\overline{\rho})q \\ 0 & (1-q) & \zeta_\rho b^{i=2}_{t-1} q & \zeta_\rho(1-b^{i=2}_{t-1})q & (1-\zeta_\rho)\overline{\rho}q & (1-\zeta_\rho)(1-\overline{\rho})q \\ 1-q & 0 & \zeta_\rho b^{i=3}_{t-1} q & \zeta_\rho(1-b^{i=3}_{t-1})q & (1-\zeta_\rho)\overline{\rho}q & (1-\zeta_\rho)(1-\overline{\rho})q \\ 0 & (1-q) & \zeta_\rho b^{i=4}_{t-1} q & \zeta_\rho(1-b^{i=4}_{t-1})q & (1-\zeta_\rho)\overline{\rho}q & (1-\zeta_\rho)(1-\overline{\rho})q \\ 1-q & 0 & \zeta_\rho b^{i=5}_{t-1} q & \zeta_\rho(1-b^{i=5}_{t-1})q & (1-\zeta_\rho)\overline{\rho}q & (1-\zeta_\rho)(1-\overline{\rho})q \\ 0 & (1-q) & \zeta_\rho b^{i=6}_{t-1} q & \zeta_\rho(1-b^{i=6}_{t-1})q & (1-\zeta_\rho)\overline{\rho}q & (1-\zeta_\rho)(1-\overline{\rho})q \end{bmatrix}$$

C.4 Unscented Kalman filter with Markov-switching

This subsection describes the detailed algorithm we employ to obtain filtered and smoothed estimates of latent states in the state space model (C2) and (C3). Relative to a standard nonlinear system with additive Gaussian errors, our model has three complications.

First, the shocks v_ς and v_ρ enter the evolution equation of π nonlinearly because the policy function $a(\cdot)$ and $\alpha(\cdot)$ are nonlinear functions of ς and ρ. Moreover, the shock v_ρ follows a Beta distribution instead of a Gaussian one. Following Särkkä and Svensson (2023), this complication can be dealt with by: 1) approximating the Beta random variable v_ρ using a nonlinear transformation of a Gaussian random variable $\tilde{v}_\rho$:

$$v_\rho = R(\tilde{v}_\rho) = \frac{exp(\tilde{v}_\rho)}{1 + exp(\tilde{v}_\rho)}$$

2) forming sigma points for the state vector augmented by v_ς and $\tilde{v}_\rho$.

Second, the reputation state ρ is bounded between 0 and 1. To enforce the boundary

[5] In general, there will be eight discrete states constructed from combinations of three binary variables. In this case, the state ϕ_t is only relevant in a new regime, i.e., $\theta_t = 1$.

condition, we use "constrained unscented Kalman filter" (Kandepu et al. (2008), Rouhani and Abur (2018)) that projects the sigma points outside the feasible region to the nearest points within the region.

Third, the state evolution equations depend on the outcome of an unobserved discrete-state Markov process Θ_t. We follow Kim (1994) and Kim and Nelson (2017) to obtain the conditional probability of Θ_t being in each discrete state and to collapse state estimate and covariance.

To ease the notation, we rewrite the state space model (C2) and (C3) as follows:

$$
\begin{aligned}
S_t &= F_{\Theta_t}\big(S_{t-1}, [v_{\varsigma,t}, \tilde{v}_{\rho,t}]\big) + [0, 0, 0, v_{\pi,t}]' \\
Y_t &= H(S_t) + \varepsilon_t
\end{aligned}
$$

where $\Theta_t \in \{1, ..., 6\}$ corresponding to $\{(\theta_t = 0, \tau_t = 1), (\theta_t = 0, \tau_t = 0), (\theta_t = 1, \phi_t = 1, \tau_t = 1), (\theta_t = 1, \phi_t = 1, \tau_t = 0), (\theta_t = 1, \phi_t = 0, \tau_t = 1), (\theta_t = 1, \phi_t = 0, \tau_t = 0)\}$ with transitional probability matrix $T_{i,j} = Pr(\Theta_t = j|\Theta_{t-1} = i)$ defined in (C4).

Notation:

- Covariance of $[0, 0, 0, v_{\pi,t}]'$: Q
- Covariance of measurement noise ε: R
- Mean of the shock vector $[v_{\varsigma,t}, \tilde{v}_{\rho,t}]'$: $\hat{v} = [0, \tilde{\rho}]'$
- Covariance of the shock vector $[v_{\varsigma,t}, \tilde{v}_{\rho,t}]'$: $V = diag(\sigma_\varsigma^2, \tilde{\sigma}_\rho^2)$
- Initial state estimate: $\hat{s}_0^j$, $j = 1, 2, ..., 6$
- Initial state covariance: P_0^j, $j = 1, 2, ..., 6$

Parameters related to sigma points

- Number parameter: L
- Scaling parameters: α, β, κ
- Weight parameter: $\lambda = \alpha^2(L + \kappa) - L$
- $w_{m,0} = \frac{\lambda}{L+\lambda}$, $w_{m(n)} = \frac{1}{2(L+\lambda)}$, $n = 1, ..., 2L$
- $w_{c,0} = \frac{\lambda}{L+\lambda} + (1 - \alpha^2 + \beta)$, $w_{c(n)} = \frac{1}{2(L+\lambda)}$, $n = 1, ..., 2L$

Prediction Step: $L = 6$, conditional on $\Theta_{t-1} = i$, $\Theta_t = j$:

- Augment the state vector:

$$\hat{x}^i_{t-1} = \begin{bmatrix} \hat{s}^i_{t-1} \\ \hat{v} \end{bmatrix}; \quad \tilde{P}^i_{t-1} = \begin{bmatrix} P^i_{t-1} & 0 \\ 0 & V \end{bmatrix}$$

- Generate $2L + 1$ sigma points:
  - $X^i_{t-1,(0)} = \hat{x}^i_{t-1}$
  - $X^i_{t-1,(n)} = \hat{x}^i_{t-1} + \sqrt{(L+\lambda)}\big[\sqrt{\tilde{P}^i_{t-1}}\big]_n$
  - $X^i_{t-1,(n+L)} = \hat{x}^i_{t-1} - \sqrt{(L+\lambda)}\big[\sqrt{\tilde{P}^i_{t-1}}\big]_n, \quad n = 1, ..., L$
- Propagate sigma points through the state transition function:
  - $S'^{(i,j)}_{t(n)} = F_j(X^i_{t-1,(n)}), \quad n = 0, ..., 2L$
- Compute the predicted state estimate:
  - $\hat{s}^{-(i,j)}_t = \sum_{n=0}^{2L} w_{m(n)} S'^{(i,j)}_{t(n)}$
- Compute the predicted state covariance:
  - $P^{-(i,j)}_t = \sum_{n=0}^{2L} w_{c(n)} (S'^{(i,j)}_{t(n)} - \hat{s}^{-(i,j)}_t)(S'^{(i,j)}_{t(n)} - \hat{s}^{-(i,j)}_t)^\top + Q$

Update Step: $L = 4$, conditional on $\Theta_{t-1} = i$, $\Theta_t = j$:

- Generate sigma points:
  - $S^{-(i,j)}_{t,(0)} = \hat{s}^{-(i,j)}_t$
  - $S^{-(i,j)}_{t,(n)} = \hat{s}^{-(i,j)}_t + \sqrt{(L+\lambda)}\big[\sqrt{P^{-(i,j)}_t}\big]_n$
  - $S^{-(i,j)}_{t,(n+L)} = \hat{s}^{-(i,j)}_t - \sqrt{(L+\lambda)}\big[\sqrt{P^{-(i,j)}_t}\big]_n, \quad n = 1, ..., L$
- Propagate sigma points through the measurement function:
  - $Y^{-(i,j)}_{t(n)} = H(S^{-(i,j)}_{t(n)}), \quad n = 0, ..., 2L$
- Compute the predicted measurement mean and covariance:
  - $\hat{y}^{-(i,j)}_t = \sum_{n=0}^{2L} w_{m(n)} Y^{-(i,j)}_{t(n)}$

- $P_{yy,t}^{-(i,j)} = \sum_{n=0}^{2L} w_{c(n)}(Y_{t(n)}^{-(i,j)} - \hat{y}_t^{-(i,j)})(Y_{t(n)}^{-(i,j)} - \hat{y}_t^{-(i,j)})^\top + R$

- Compute the cross-covariance between state and measurement:
  - $P_{sy,t}^{-(i,j)} = \sum_{n=0}^{2L} w_{c(n)}(S_{t(n)}^{-(i,j)} - \hat{s}_t^{-(i,j)})(Y_{t(n)}^{(i,j)} - \hat{y}_t^{-(i,j)})^\top$
- Compute the Kalman gain:
  - $K_t^{(i,j)} = P_{sy,t}^{-(i,j)}(P_{yy,t}^{-(i,j)})^{-1}$
- Update the state estimate:
  - $\hat{s}_t^{(i,j)} = \hat{s}_t^{-(i,j)} + K_t^{(i,j)}(Y_t - \hat{y}_t^{-(i,j)})$
- Update the state covariance:
  - $P_t^{(i,j)} = P_t^{-(i,j)} - K_t^{(i,j)} P_{yy,t}^{-(i,j)} (K_t^{(i,j)})^\top$

Conditional Probability Step:

- Start from $Pr(\Theta_{t-1} = i | Y^{t-1})$
  - $Pr(\Theta_{t-1} = i, \Theta_t = j | Y^{t-1}) = Pr(\Theta_t = j | \Theta_{t-1} = i) Pr(\Theta_{t-1} = i | Y^{t-1})$
- Update using Bayes' rule

$$Pr(\Theta_{t-1} = i, \Theta_t = j | Y^t) = \frac{f(Y_t | \Theta_{t-1} = i, \Theta_t = j, Y^{t-1}) Pr(\Theta_{t-1} = i, \Theta_t = j | Y^{t-1})}{\sum_{j=1}^{6} \sum_{i=1}^{6} f(Y_t, \Theta_{t-1} = i, \Theta_t = j | Y^{t-1})}$$

  where $f(Y_t | \Theta_{t-1} = i, \Theta_t = j, Y^{t-1}) \sim N(\hat{y}_t^{-(i,j)}, P_{yy,t}^{-(i,j)})$

- Collapse $Pr(\Theta_t = j | Y^t) = \sum_{i=1}^{6} Pr(\Theta_{t-1} = i, \Theta_t = j | Y^t)$

Collapse Step:

$$\hat{s}_t^j = \frac{\sum_{i=1}^{6} Pr(\Theta_{t-1} = i, \Theta_t = j | Y^t) \hat{s}_t^{(i,j)}}{Pr(\Theta_t = j | Y^t)}$$

$$P_t^j = \frac{\sum_{i=1}^{6} Pr(\Theta_{t-1} = i, \Theta_t = j | Y^t) \{P_t^{(i,j)} + (\hat{s}_t^j - \hat{s}_t^{(i,j)})(\hat{s}_t^j - \hat{s}_t^{(i,j)})^\top\}}{Pr(\Theta_t = j | Y^t)}$$

Smooth Step:

- Initialize the smoothed state estimate and covariance at the last time step:
  - $\hat{s}_T^{s,j} = \hat{s}_T^j$
  - $P_T^{s,j} = P_T^j$
  - $Pr(\Theta_T = j|Y^T)$
- Smooth probability for $\Theta_t = j$ and $\Theta_{t+1} = k$ from $t = T-1, ..., 1$:

$$\begin{aligned}
&Pr(\Theta_t = j, \Theta_{t+1} = k|Y^T) \\
&\quad = Pr(\Theta_{t+1} = k|Y^T) Pr(\Theta_t = j|\Theta_{t+1} = k, Y^T) \\
&\quad \approx Pr(\Theta_{t+1} = k|Y^T) Pr(\Theta_t = j|\Theta_{t+1} = k, Y^t) \\
&\quad = \frac{Pr(\Theta_{t+1} = k|Y^T) Pr(\Theta_t = j, \Theta_{t+1} = k|Y^t)}{Pr(\Theta_{t+1} = k|Y^t)} \\
&\quad = Pr(\Theta_{t+1} = k|Y^T) \frac{Pr(\Theta_t = j|Y^t) Pr(\Theta_{t+1} = k|\Theta_t = j)}{\sum_{j=1}^{6} Pr(\Theta_t = j|Y^t) Pr(\Theta_{t+1} = k|\Theta_t = j)}
\end{aligned}$$

- Smooth probability for $\Theta_t = j$ for $t = T-1, ..., 1$:

$$Pr(\Theta_t = j|Y^T) = \sum_{k=1}^{6} Pr(\Theta_t = j, \Theta_{t+1} = k|Y^T)$$

- Perform the smoothing recursion from $t = T-1, ..., 1$, conditional on $\Theta_t = j, \Theta_{t+1} = k$:
  - Augment the state vector:

$$\hat{x}_t^j = \begin{bmatrix} \hat{s}_t^j \\ \hat{v} \end{bmatrix}; \quad \tilde{P}_t^j = \begin{bmatrix} P_t^j & 0 \\ 0 & V \end{bmatrix}$$

  - Generate $2L+1$ sigma points given $L = 6$:
    * $X_{t,(0)}^j = \hat{x}_t^j$
    * $X_{t,(n)}^j = \hat{x}_t^j + \sqrt{(L+\lambda)}\big[\sqrt{\tilde{P}_t^j}\big]_n$
    * $X_{t,(n+L)}^j = \hat{x}_t^j - \sqrt{(L+\lambda)}\big[\sqrt{\tilde{P}_t^j}\big]_n, \quad n = 1, ..., L$
  - Propagate sigma points through the state transition function:
    * $S_{t+1,(n)}^{\prime(j,k)} = F_k(X_{t,(n)}^j)$

– Compute the predicted state mean and covariance:

* $\hat{s}_{t+1}^{-(j,k)} = \sum_{n=0}^{2L} w_{m(n)} S'^{(j,k)}_{t+1,(n)}$
* $P_{t+1}^{-(j,k)} = \sum_{n=0}^{2L} w_{c(n)} (S'^{(j,k)}_{t+1,(n)} - \hat{s}_{t+1}^{-(j,k)})(S'^{(j,k)}_{t+1,(n)} - \hat{s}_{t+1}^{-(j,k)})^\top + Q$

– Compute the cross-covariance:

* $D_{t+1}^{-(j,k)} = \sum_{n=0}^{2L} w_{c(n)} (X_{t,(n)}^{j,S} - \hat{s}_t^j)(S'^{(j,k)}_{t+1,(n)} - \hat{s}_{t+1}^{-(j,k)})^\top$
* where $X_{t,(n)}^{j,S}$ denotes the part of sigma point n which corresponds to S_t

– Compute the smoothed state gain:

* $K_t^{s,(j,k)} = D_{t+1}^{-(j,k)} (P_{t+1}^{-(j,k)})^{-1}$

– Compute the smoothed state estimate:

* $\hat{s}_t^{s,(j,k)} = \hat{x}_t^j + K_t^{s,(j,k)} (\hat{s}_{t+1}^{s,k} - \hat{s}_{t+1}^{-(j,k)})$

– Compute the smoothed state covariance:

* $P_t^{s,(j,k)} = P_t^j + K_t^{s,(j,k)} (P_{t+1}^{s,k} - P_{t+1}^{-(j,k)})(K_t^{s,(j,k)})^\top$

– Collapse the smoothed state estimate and covariance

$$\hat{s}_t^{s,j} = \frac{\sum_{k=1}^{6} Pr(\Theta_t = j, \Theta_{t+1} = k|Y^T)\hat{s}_t^{s,(j,k)}}{Pr(\Theta_t = j|Y^T)}$$

$$P_t^{s,j} = \frac{\sum_{k=1}^{6} Pr(\Theta_t = j, \Theta_{t+1} = k|Y^T)\{P_t^{s,(j,k)} + (\hat{s}_t^{s,j} - \hat{s}_t^{s,(j,k)})(\hat{s}_t^{s,j} - \hat{s}_t^{s,(j,k)})^\top\}}{Pr(\Theta_t = j|Y^T)}$$

Filtered and smoothed estimates of states and observables:

- Filtered estimates for states conditional on $\Theta_{t-1} = i, \Theta_t = j$

$$\hat{s}_t = \sum_{j=1}^{6} Pr(\Theta_t = j|Y^t)\hat{s}_t^j$$

$$P_t = \sum_{j=1}^{6} Pr(\Theta_t = j|Y^t)\{P_t^j + (\hat{s}_t - \hat{s}_t^j)(\hat{s}_t - \hat{s}_t^j)^\top\}$$

- Filtered estimates for observables

  – Generate sigma points, $L = 4$:

  * $S_{t,(0)}^{(i,j)} = \hat{s}_t^{(i,j)}$
  * $S_{t,(n)}^{(i,j)} = \hat{s}_t^{(i,j)} + \sqrt{(L+\lambda)}\big[\sqrt{P_t^{(i,j)}}\big]_n$

    * $S^{(i,j)}_{t,(n+L)} = \hat{s}^{(i,j)}_t - \sqrt{(L+\lambda)}\big[\sqrt{P^{(i,j)}_t}\big]_n, \quad n = 1,...,L$

- Propagate sigma points through the measurement function:
    * $Y^{(i,j)}_{t(n)} = H(S^{(i,j)}_{t(n)}), \quad n = 0,...,2L$
- Compute the predicted measurement mean and covariance:
    * $\hat{y}^{(i,j)}_t = \sum_{n=0}^{2L} w_{m(n)} Y^{(i,j)}_{t(n)}$
    * $P^{(i,j)}_{yy,t} = \sum_{n=0}^{2L} w_{c(n)} (Y^{(i,j)}_{t(n)} - \hat{y}^{(i,j)}_t)(Y^{(i,j)}_{t(n)} - \hat{y}^{(i,j)}_t)^\top + R$
- Collapse

$$\hat{y}_t = \sum_{i=1}^{6}\sum_{j=1}^{6} Pr(\Theta_{t-1} = i, \Theta_t = j|Y^t)\hat{y}^{(i,j)}_t$$

$$P_{yy,t} = \sum_{i=1}^{6}\sum_{j=1}^{6} Pr(\Theta_{t-1} = i, \Theta_t = j|Y^t)\{P^{(i,j)}_{yy,t} + (\hat{y}_t - \hat{y}^{(i,j)}_t)(\hat{y}_t - \hat{y}^{(i,j)}_t)^\top\}$$

- Smoothed estimates for states

$$\hat{s}^s_t = \sum_{j=1}^{6} Pr(\Theta_t = j|Y^T)\hat{s}^{s,j}_t$$

$$P^s_t = \sum_{j=1}^{6} Pr(\Theta_t = j|Y^T)\{P^{s,j}_t + (\hat{s}^s_t - \hat{s}^{s,j}_t)(\hat{s}^s_t - \hat{s}^{s,j}_t)^\top\}$$

- Smoothed estimates for observables conditional on $\Theta_t = j, \Theta_{t+1} = k$
    - Generate sigma points, $L = 4$:
        * $S^{s,(j,k)}_{t,(0)} = \hat{s}^{s,(j,k)}_t$
        * $S^{s,(j,k)}_{t,(n)} = \hat{s}^{s,(j,k)}_t + \sqrt{(L+\lambda)}\big[\sqrt{P^{s,(j,k)}_t}\big]_n$
        * $S^{s,(j,k)}_{t,(n+L)} = \hat{s}^{s,(j,k)}_t - \sqrt{(L+\lambda)}\big[\sqrt{P^{s,(j,k)}_t}\big]_n, \quad n = 1,...,L$
    - Propagate sigma points through the measurement function:
        * $Y^{(i,j)}_{t(n)} = H(S^{s,(j,k)}_{t(n)}), \quad n = 0,...,2L$
    - Compute the predicted measurement mean and covariance:
        * $\hat{y}^{s,(j,k)}_t = \sum_{n=0}^{2L} w_{m(n)} Y^{s,(j,k)}_{t(n)}$
        * $P^{s,(j,k)}_{yy,t} = \sum_{n=0}^{2L} w_{c(n)} (Y^{s,(j,k)}_{t(n)} - \hat{y}^{s,(j,k)}_t)(Y^{s,(j,k)}_{t(n)} - \hat{y}^{s,(j,k)}_t)^\top + R$

– Collapse

$$\hat{y}_t^s = \sum_{j=1}^{6}\sum_{k=1}^{6} Pr(\Theta_t = j, \Theta_{t+1} = k|Y^T)\hat{y}_t^{s,(j,k)}$$

$$P_{yy,t}^s = \sum_{j=1}^{6}\sum_{k=1}^{6} Pr(\Theta_t = j, \Theta_{t+1} = k|Y^T)\{P_{yy,t}^{s,(j,k)} + (\hat{y}_t^s - \hat{y}_t^{s,(j,k)})(\hat{y}_t^s - \hat{y}_t^{s,(j,k)})^\top\}$$

C.5 Estimates of discrete states

As an example of estimated conditional probabilities, Figure 9 plots the smoothed probabilities of a committed policy regime (blue) and of a policymaker replacement (red) in each period. The probability of a committed policy regime echos the dynamics of the estimated reputation state $\widehat{\rho}_t$ in Figure 3.[6] The probability is close to zero after 1975 and sharply increases to close to one in 1981-1982, suggesting that the most likely discrete state consistent with the observed SPF data switches from $\tau = 0$ (an opportunistic policy regime) to $\tau = 1$ (a committed policy regime). According to the model, policymaker's type can only switch in the event of a policymaker replacement. Our estimated probability of a replacement event peaks in the first quarter of 1982.

Figure 9: Smoothed probabilities

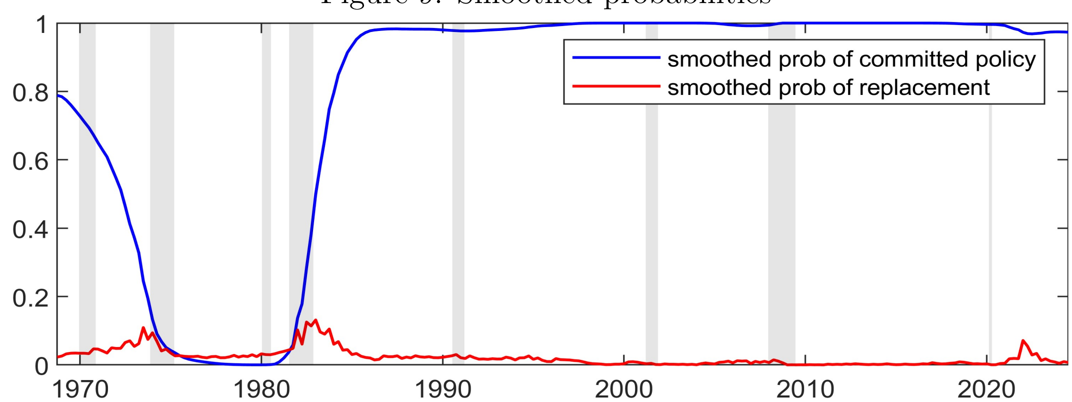

The figure plots the smoothed probabilities of discrete Markov states estimated by the nonlinear filter. The smoothed probability of committed policy $Pr(\tau_t = 1|Y^T)$ is the sum of three smoothed probabilities $Pr(\theta_t = 0, \tau_t = 1|Y^T)$ and $Pr(\theta_t = 1, \phi_t = 0 \text{ or } 1, \tau_t = 1|Y^T)$. The smoothed probability of replacement $Pr(\theta_t = 1|Y^T)$ is the sum of four smoothed probabilities $Pr(\theta_t = 1, \phi_t = 0 \text{ or } 1, \tau_t = 0 \text{ or } 1|Y^T)$.

[6] The smoothed estimate of ρ_t is different from the smoothed probability of a committed policy regime. Our filter calculates the optimal estimates of $s_t = (\varsigma_t, \rho_t, \mu_t)$ for fitting the observed SPF data, given the assumption of being in a specific policy regime. Subsequently, it applies these regime-specific estimates to obtain the probability of that particular policy regime, taking into account the structure of shocks. The smoothed estimate of ρ_t is derived as a probability-weighted average of these regime-specific estimates.

C.6 Fitting performance

Inflation expectations: As discussed in Section 5.4 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 10 shows our match is nearly perfect for 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 10. 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.

Figure 10: Model-implied and SPF forecasts of inflation

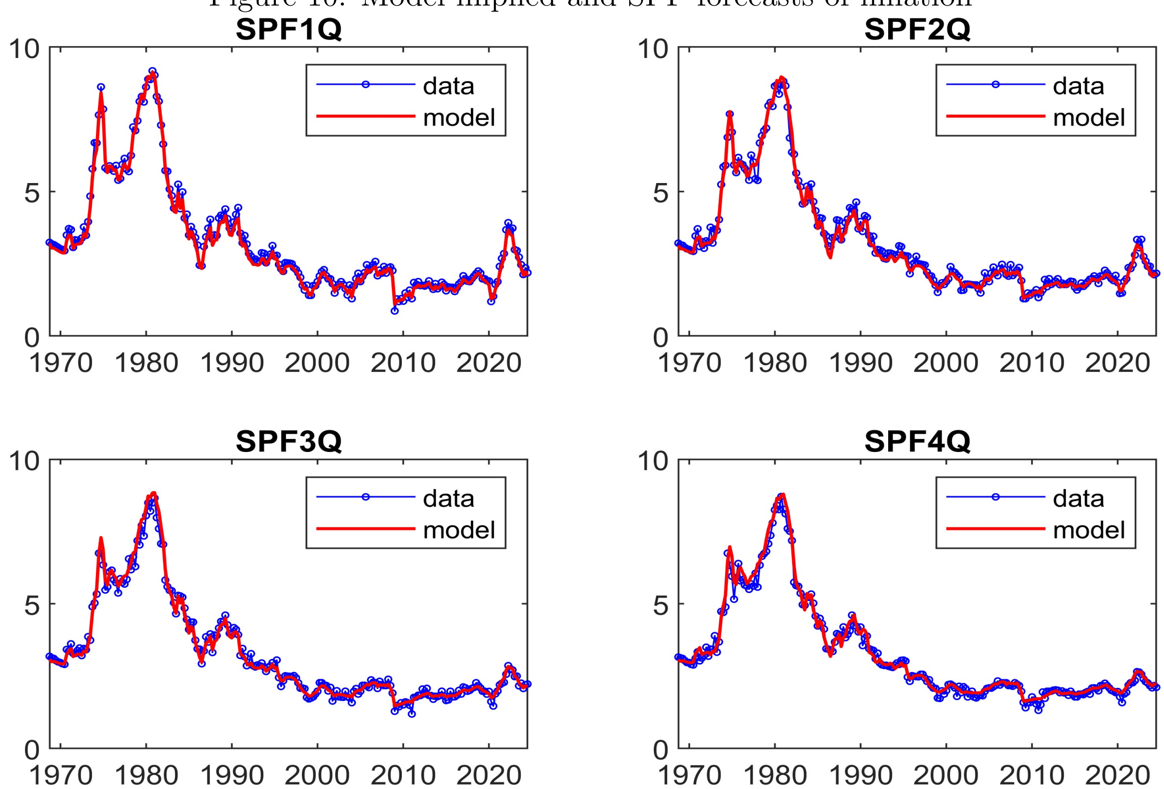

This figure compares the smoothed estimates of model inflation forecast produced by our nonlinear filter with the SPF inflation forecast of the same forecasting horizon. SPF1Q and SPF3Q are targeted by our nonlinear filter; SPF2Q and SPF4Q are not targeted.

Filtered inflation Section 5.5.2 demonstrates that the smoothed estimates of inflation by our state-space model fit the observed U.S. inflation well without explicitly targeting it. The benchmark we use to measure the fitting performance is to compare the smoothed estimates with the SPF1Q, as shown in Figure 4. A skeptical reader may concern that our smoothed measure performs better simply because it is based on the full sample of SPF, while the SPF1Q is prepared with information up to the period t. We therefore provide a filtered version Figure 11, where no information after the period t is used to obtain the period-t filtered measure. Our filtered estimates for inflation continue to outperform SPF1Q in both measures of fit: lower persistence of fitting error and lower mean-squared error.

Figure 11: Untargeted inflation: filtered result

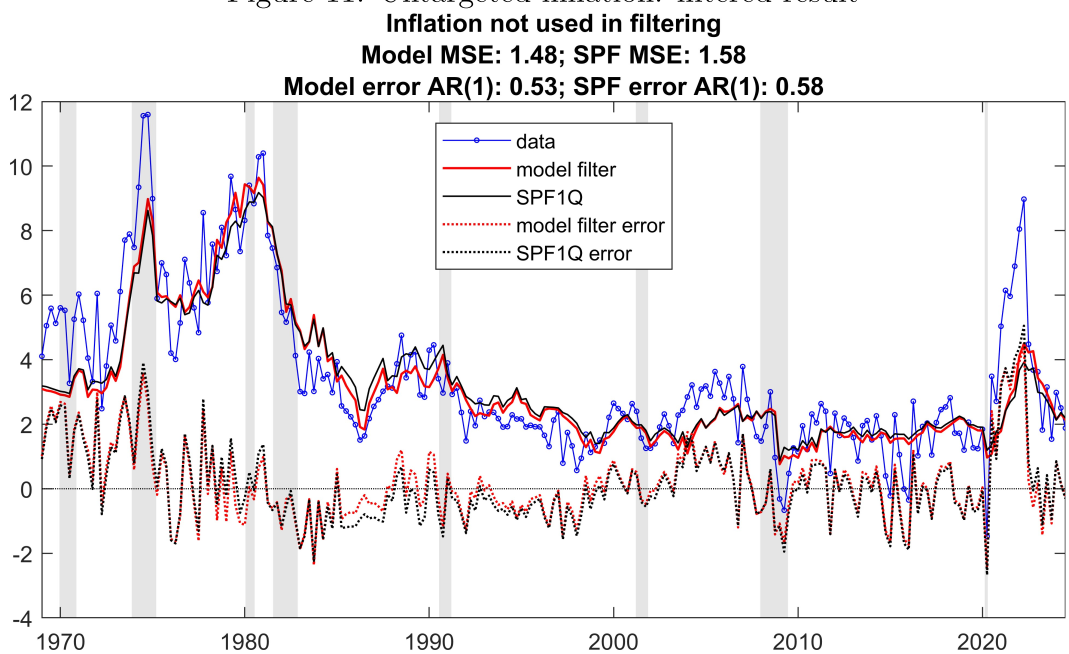

This figure is the counterpart of Figure 4 except that we replace the smoothed estimates for inflation with the filtered estimates that only use information up to period t. Our filtered estimates for inflation continue to outperform SPF1Q with lower persistence of fitting error and lower mean-squared error.

D Counterfactual with Naive Committed Policy

D.1 Optimization of a naive committed policymaker

The key departure from the benchmark model is that the committed type optimizes as if the reputation is a given parameter ρ. When the reputation is no longer a function of the inflation shock π (at least in the committed type's optimization), there is no channel for the current π_t to affect future state variables.[7]

This observation helps us to reduce the forwarding expectation constraint to:

$$e_t = \beta E_t \pi_{t+1} = \beta (1-q) \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) [\rho a(h_{t+1}) + (1-\rho) \alpha(h_{t+1})] + \beta q \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) z(h_{t+1})$$

because a_{t+1}, α_{t+1}, and z_{t+1} are independent of π_t. As a result, we avoid carrying the likelihood ratio $\lambda(h_{t+1}) \equiv \frac{g(\pi_t|\alpha_t)}{g(\pi_t|a_t)}$ as a state variable.

The recursive form of the naive optimization of the committed policymaker is

$$\begin{aligned} W(\varsigma_t, \eta_t; \rho) &= \min_{\gamma} \max_{a,e} \underline{u}(a_t, e_t, \varsigma_t) + \gamma_t e_t - (1-q) \eta_t [\rho a_t + (1-\rho) \alpha_t] - q \eta_t z(\varsigma_t, \rho) \\ &\quad + \beta_a (1-q) \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t) W(\varsigma_{t+1}, \eta_{t+1}; \rho) \end{aligned}$$

subject to

$$\alpha_t = A e_t + B(\varsigma_t)$$

with

$$\eta_{t+1} = \frac{\beta}{\beta_a (1-q)} \gamma_t \text{ with } \gamma_{-1} = 0.$$

Given $z(\varsigma_t, \rho)$, the optimization yields the following policy rules: $a(\varsigma_t, \eta_t; \rho)$, $e(\varsigma_t, \eta_t; \rho)$, and $\gamma(\varsigma_t, \eta_t; \rho)$. The fixed point requires

$$z(\varsigma_t, \rho) = \rho a(\varsigma_t, 0; \rho) + (1-\rho) [A e(\varsigma_t, 0; \rho) + B(\varsigma_t)]$$

The policy function under the setup of naive committed policymaker are denoted by

$$\begin{aligned} & a^N(\varsigma, \rho, \mu) \\ & \alpha^N(\varsigma, \rho, \mu) \\ & \mu'^N(\varsigma, \rho, \mu) \end{aligned}$$

D.2 Constructing counterfactual time series

Initialization step for $t = 1$: $\rho_1^{N,j} = \hat{\rho}_1^j$ and $\mu_1^{N,j} = \hat{\mu}_1^j$ for $\Theta_1 = j$. $\{\hat{\varsigma}_t^j\}_{t=1}^T$, $\{\hat{v}_{\pi,t}^j\}_{t=1}^T$, and $\{Pr(\Theta_t = j|Y^T)\}_{t=1}^T$ are smoothed estimates of cost-push shocks, implementation errors, and smoothed probabilities of $\Theta_t = j$ obtained from the benchmark model.

[7] Recall that the lagrangian multiplier γ_t is chosen before the realization of π_t and it will determine the next-period pseudo state variable.

Conditional on $\Theta_t = j$ and $\Theta_{t+1} = k$, we obtain

$$a_t^{N,j} = a^N(\hat{\varsigma}_t^j, \rho_t^{N,j}, \mu_t^{N,j})$$

$$\alpha_t^{N,j} = \alpha^N(\hat{\varsigma}_t^j, \rho_t^{N,j}, \mu_t^{N,j}))$$

$$\rho_{t+1}^{N,(j,k)} = \begin{cases} b(\pi_t^j; a_t^{N,j}, \alpha_t^{N,j}, \rho_t^{N,j}) & \text{if } k = 1,2,3,4 \\ \hat{\rho}_{t+1}^k & \text{if } k = 5,6 \end{cases}$$

$$\mu_{t+1}^{N,(j,k)} = \begin{cases} \mu'^N(\hat{\varsigma}_t^j, \rho_t^{N,j}, \mu_t^{N,j}) & \text{if } k = 1,2 \\ 0 & \text{if } k = 3,4,5,6 \end{cases}$$

where $\pi_t^j = a_t^{N,j} + \hat{v}_{\pi,t}^j$ if $j = 1,3,5$ and $\pi_t^j = \alpha_t^{N,j} + \hat{v}_{\pi,t}^j$ if $j = 2,4,6$.

We then perform the collapsing step:

$$\rho_{t+1}^{N,k} = \sum_{j=1}^{6} \rho_{t+1}^{N,(j,k)} Pr(\Theta_t = j | \Theta_{t+1} = k, Y^T)$$

$$\mu_{t+1}^{N,k} = \sum_{j=1}^{6} \mu_{t+1}^{N,(j,k)} Pr(\Theta_t = j | \Theta_{t+1} = k, Y^T)$$

where

$$Pr(\Theta_t = j | \Theta_{t+1} = k, Y^T) = \frac{Pr(\Theta_t = j, \Theta_{t+1} = k | Y^T)}{\sum_{j=1}^{6} Pr(\Theta_t = j, \Theta_{t+1} = k | Y^T)}$$

$$= \frac{Pr(\Theta_{t+1} = k | \Theta_t = j) Pr(\Theta_t = j | Y^T)}{\sum_{j=1}^{6} Pr(\Theta_{t+1} = k | \Theta_t = j) Pr(\Theta_t = j | Y^T)}$$

The transitional probability $Pr(\Theta_{t+1} = k | \Theta_t = j)$ are the same as the one in the benchmark model (C4) except that b_{t-1}^i is replaced with the naive-policy version $b(\pi_t^j; a_t^{N,j}, \alpha_t^{N,j}, \rho_t^{N,j})$.

The reported counterfactual time series t=1,...T are constructed as follows:

$$\rho_t^N = \sum_{j=1}^{6} \rho_t^{N,j} Pr(\Theta_t = j | Y^T)$$

$$a_t^N = \sum_{j=1}^{6} a_t^{N,j} Pr(\Theta_t = j | Y^T)$$

$$\alpha_t^N = \sum_{j=1}^{6} \alpha_t^{N,j} Pr(\Theta_t = j | Y^T)$$

Figure 12: Policy Difference

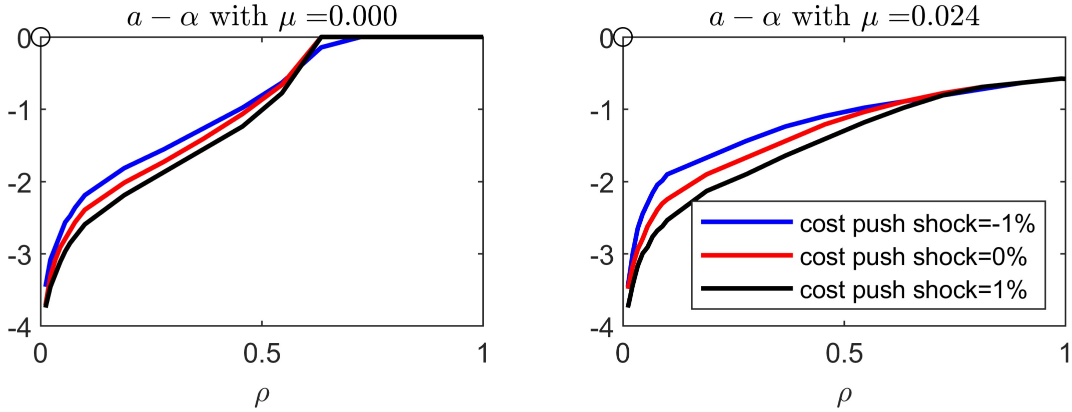

Difference between equilibrium committed policy and opportunistic policy is larger at lower levels of reputation. This property holds for various levels of cost push shock and pseudo state μ. The two values of μ are representative as they correspond to the steady state μ when $\rho = 0$ and $\rho = 1$, respectively.

Appendix References

Kandepu, Rambabu, Lars Imsland, and Bjarne A. Foss, "Constrained state estimation using the Unscented Kalman Filter," in "2008 16th Mediterranean Conference on Control and Automation" June 2008, pp. 1453–1458.

Kim, Chang-Jin, "Dynamic linear models with Markov-switching," *Journal of Econometrics*, January 1994, *60* (1-2), 1–22.

_ **and Charles R. Nelson**, *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*, The MIT Press, November 2017.

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

Rouhani, Alireza and Ali Abur, "Constrained Iterated Unscented Kalman Filter for Dynamic State and Parameter Estimation," *IEEE Transactions on Power Systems*, May 2018, *33* (3), 2404–2414.

Stark, Tom, "Realistic Evaluation of Real-Time Forecasts in the Survey of Professional Forecasters," Special Report, Federal Reserve Bank of Philadelphia 2010.

Särkkä, Simo and Lennart Svensson, *Bayesian Filtering and Smoothing*, 2023 ed., Cambridge University Press, 2023.