NBER WORKING PAPER SERIES

THE RISE, FALL AND STABILIZATION OF U.S. INFLATION: SHIFTING REGIMES AND EVOLVING REPUTATION

Robert G. King
Yang K. Lu

Working Paper 29585
http://www.nber.org/papers/w29585

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2021

We would like to thank the audience at various conferences and seminars for helpful comments and discussions. All errors are ours. We acknowledge financial support from the RGC of HKSAR (GRF HKUST-16504317). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Rise, Fall and Stabilization of U.S. Inflation: Shifting Regimes and Evolving Reputation
Robert G. King and Yang K. Lu
NBER Working Paper No. 29585
December 2021
JEL No. D82,D83,E52

ABSTRACT

The rise, fall, and stabilization of US inflation between 1969 and 2005 is consistent with a model of shifting policy regimes that features a forward-looking New Keynesian Phillips curve, policymakers that can or cannot commit, and private sector learning about policymaker type. Using model-implied inflation forecasting rules to extract state variables from the inflation forecasts in the Survey of Professional Forecasters, we provide evidence that policy regimes without commitment prevailed before 1980 and regimes with commitment prevailed afterward. With theory and quantification, we find that evolution of reputational capital is central to understanding the behavior of inflation.

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

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

"In reality, however, the anchoring of inflation expectations has been a hard-won achievement of monetary policy over the past few decades, and we should not take this stability for granted. [...] a policy of achieving "temporarily" higher inflation over the medium term would run the risk of altering inflation expectations beyond the horizon that is desirable. Were that to happen, the costs of bringing expectations back to their current anchored state might be quite high."

(Donald L. Kohn, "Monetary Policy Research and the Financial Crisis: Strengths and Shortcomings", October 9, 2009)

1 Introduction

The interplay of inflation, policy and expectations is at the heart of modern macroeconomics. Figure 1 plots one-quarter-ahead inflation forecast errors based on the Survey of Professional Forecasters from 1969 to 2005. Around 1980, the average inflation forecast error turns from persistently positive to persistently negative, a feature which is highlighted by an 8 quarter moving average of the errors (the black dashed line).[1] Many authors have used this pattern of runs of forecast errors as evidence for private sector learning with a mis-specified model of the economy.[2]

We instead view the pattern as arising from private sector's gradually learning about the type of policymaker in place.[3] Before 1980, we see a U.S. policymaker that could not commit and thus continually produced positive inflation surprises. After 1980, we see a policymaker committed to carrying out promised inflation plans, but that faced private sector skepticism so that inflation was frequently below forecasts, particularly in the early years.

Building a model to turn these simple ideas into macroeconomic time series is the main research activity reported in this paper. Our model has four

[1]We thank Donghai Zhang for pointing us to this observation.

[2]See G.W. Evans and Honkapohja (2008), Woodford (2013) and Eusepi and Preston (2018) for the surveys of this literature. Relative to our work, this approach captures different aspects of uncertainty regarding economic fundamentals.

[3]Matthes (2015) develops a model in which the private sector learns to differentiate between two competing monetary policymaking approaches, while our setup highlights how purposeful policy depends on private sector learning.

Figure 1: Forecast error of inflation

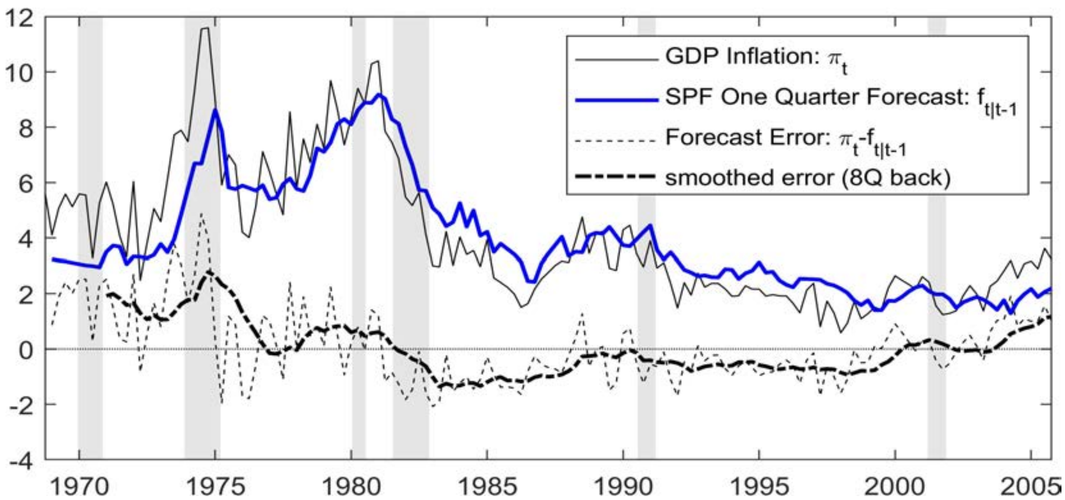

The GDP inflation rate, π_t, and the Survey of Professional Forecasters median inflation forecast rate made one quarter earlier, $f_{t|t-1}$, rise and fall together over 1968Q4 through 2005Q4. Inflation is notoriously difficult to forecast so that the forecasting error, $\pi_t - f_{t|t-1}$, is volatile, although it averages close to zero over the sample period. The errors display serial correlation – lengthy runs of positive and negative forecasting values – that are highlighted by an 8 quarter moving average.

Figure 2: SPF and SPF Spread

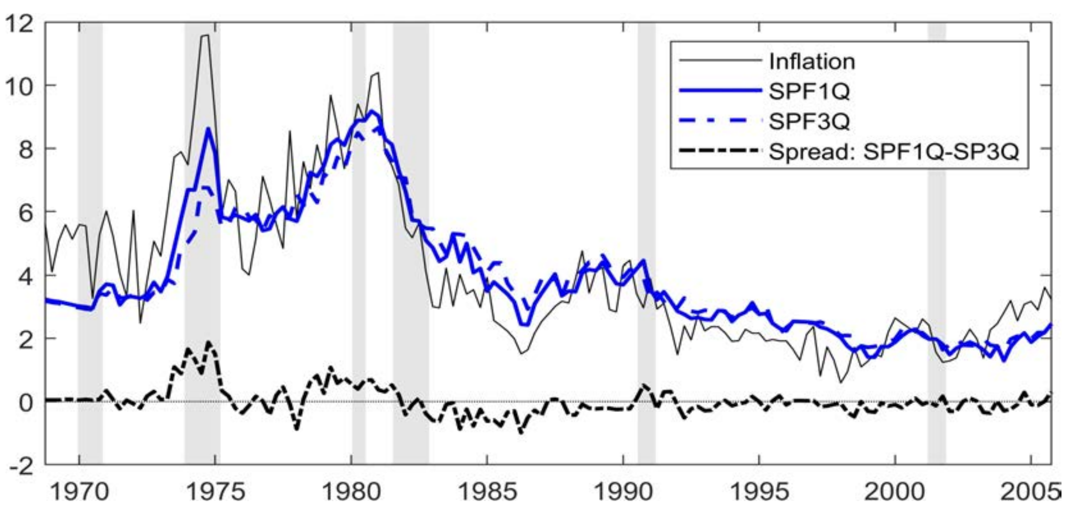

The one quarter ahead forecast $f_{t+1|t}$ and the three quarter ahead forecast $f_{t+3|t}$ rise and fall together. The SPF spread $f_{t+3|t} - f_{t+1|t}$, which will play an important role in the paper, also displays sustained intervals of high or low values. All variables are continuously compounded annualized rates of change. Additional detail on macroeconomic data and the SPF constructions are provided in Appendix C.

main structural elements: a forward-looking New Keynesian Phillips curve, policymakers that can or cannot commit, Bayesian learning by the private sector about policymaker type, and occasional observable changes in regime.[4] Combining our model-implied inflation forecasting rules with the term structure of inflation forecasts in the Survey of Professional Forecasters, we provide evidence that policy regimes without commitment prevailed before 1980 and regimes with commitment prevailed afterward. In our theory and quantification, evolution of reputational capital is central to the behavior of expectations and inflation.

In establishing this understanding of inflation history, we make three contributions. First, we develop a recursive model of optimal intended inflation by a policymaker capable of commitment, facing private sector suspicion that there may be an alternative, optimizing but myopic, policymaker in place.[5] This committed policymaker's reputation - the private sector's likelihood that there is a committed type in place - is a key model state variable that evolves in accordance with optimal Bayesian learning. While we've previously used recursive methods to investigate optimal policy with and without reputation dynamics, the essential new wrinkle is that the alternative policymaker is not a machine, but is purposeful (opportunistic).[6] Formally, this adds an incentive compatibility constraint to our recursive approach: it turns out to be one with substantial content for reputation evolution and inflation.

Second, we show that model-implied inflation forecasts at various horizons are functions of state variables, including the highly persistent reputation state and a more temporary price shock. We then design an empirical strategy to extract these states utilizing SPF forecasts at multiple horizons. Figure 2 highlights the smoother nature of the three-quarter-ahead SPF forecast of inflation relative to the one-quarter-ahead forecast.[7] Taking a cue from literature on the term structure of interest rates, we form an SPF spread, plotted as the black dashed line and defined as the difference between the one and three

[4] Coibion et al. (2018) advocate developing inflation models using survey data like the SPF, focusing on the New Keynesian Phillips curve that confronts our policymaker.

[5] "Intended inflation" captures our policymaker's imperfect control of inflation: see below.

[6] Our label matches Mishkin (2008) and other work on commitment and communication.

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

quarter forecasts.[8] The interest rate term structure analogy suggests that this spread should be positive when there are persistent, but ultimately temporary increases in inflation.[9] Our empirical design extracts unobserved states from the SPF term structure, exploiting the information used by sophisticated forecasters at the time.

Third, as the intended inflation choices of committed and opportunistic policymakers are functions of these state variables, we can trace out the history of these intentions over 1969-2005 without taking a stand on the type of policymaker in place at any date. However, in our theory, observed inflation should differ from intended inflation by a serially uncorrelated shock. Looking at the early history through the end of the 1970s, we find an average deviation of about zero from the opportunistic type's intended inflation but a strongly positive average deviation from the committed type's intended inflation: we thus conclude that the early interval was comprised of regimes without commitment. Looking after 1982, we find the reverse pattern: the deviations average about zero for the committed type and are negative on average for the opportunistic type: we thus conclude that the latter interval was comprised of regimes with commitment.

A roadmap to the paper is as follows. Section 2 locates our work within the literature. Section 3 describes our model economy, while section 4 defines its perfect Bayesian equilibrium. Section 5 highlights the decision problem of the committed type, which is more complicated than in many optimal policy setups due to incentive constraints and reputation dynamics, but shows it fits neatly into a recursive equilibrium. Section 6 begins with model calibration and extraction of latent states. It then shows that our model-implied policy measures capture US inflation's rise, fall, and stabilization between 1970 and 2005.[10] Section 7 develops the interaction of reputation and policy. Section 8 constructs measures of policy credibility and derives a counterfactual inflation history assuming a single committed regime. Section 9 concludes.

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

[9] Conceptually, this description is consistent with a stationary autoregressive component. Empirically, the SPF spread rises in Figure 2 during the 1974-75 inflation surge.

[10] Primiceri (2006) provides an alternative account that focuses on policymaker, rather than private agent, learning.

2 Literatures

Our work has links to history, macroeconomic theory and econometrics.

2.1 Excessive and rising inflation without commitment

The stagflation of the mid 1970s led to an explosion of new theoretical ideas about macroeconomic policy. Some leading macroeconomic theorists argued that excessive inflation was rooted in the interaction of private agent inflation expectations with a policymaker that couldn't commit to future actions (Kydland and Prescott (1977)). Stagflation itself would arise if inflation bias intensified with a higher natural rate of unemployment (Barro and D.B. Gorden (1983a)). Sargent (1982) argued that regime change was important for stopping sustained inflation, large and small. Yet, while these ideas attained great prominence, many economists have expressed doubt about the historical importance of commitment capacity, reputation evolution, and regime shifts.[11]

Our model provides a positive theory of "The Great Inflation" that was the objective of the early literature and relies on its basic insights, while stressing private sector learning. It does not rely on a time-varying natural rate of unemployment and involves only a small amount of intrinsic inflation bias (a policymaker's choice with well-anchored expected inflation). Its core mechanism for increasing inflation is the positive feedback between expectations and the choices of a policymaker that can't commit: under full information rational expectations, the feedback leads to a bias of 8% or more at the Nash equilibrium in the terminology of Sargent (1999) and others.[12]

But in our model, the rise in inflation can be very gradual as learning slows down the positive feedback mechanism. More specifically, if a policymaker that cannot commit begins with a good reputation, rising actual inflation can have relatively minor effects on expected inflation because the private sector

[11] For example, Blinder (1997) is skeptical about basic inflation bias mechanism. While Parkin (1993) found some support for the Barro-Gordon hypothesis, Levin and Taylor (2013) argued against it by depicting timing misalignment between changes in actual and expected inflation and changes in the natural rate of unemployment. Ireland (2004) found low frequency co-movement between inflation and unemployment consistent with the Barro-Gordon setup, but his econometric analysis didn't support its more detailed time series implications.

[12] That is, if it is evident that a new policymaker is unable to commit, there would be an immediate jump to a high inflation rate until another regime change.

believes that the committed type is most likely in place and expects future anti-inflation policies. Further, faced with inflation expectations that are low and stable, the opportunistic type opts not to increase inflation much because its intrinsic bias is small. In turn, this makes for slow learning and a gradual drift in the level of actual and expected inflation as in Figure 1. Yet, our model also predicts that temporary positive price shocks – such as those at several points in the 1970s – can speed up the erosion of reputation that must ultimately occur when there isn't commitment.

Being myopic, our opportunistic type does not have reputation concerns, but private agents never the less may face difficulty in distinguishing it from a committed policymaker.[13]

2.2 Commitment, disinflation and inflation stabilization

Another strand of literature investigates optimal choices by a policymaker that is endowed with the ability to commit but faces a skeptical private sector. Cukierman and Liviatan (1991) and R.G. King et al. (2005) employ 1980s-type models,[14] and focus on whether a dramatic "cold turkey" or smoother "gradualist" disinflation strategy is desirable for a committed type managing expectations in the face of private sector skepticism and evolving reputation.[15]

Our model features the forward-looking New Keynesian Phillips curve, pri-

[13] Reputational dynamics were also an important element in the 1980s literature, but for a reason very different from our analysis. Barro and D.B. Gorden (1983b), Backus and Driffill (1985b) and Backus and Driffill (1985a), and Barro (1986), showed that reputational forces can substitute for commitment capability, leading a "discretionary" policymaker to behave like a committed one that mechanically adopts a low inflation rule. This early work led to a major literature on "sustainable plans" for patient policymakers in environments without commitment (Chari and Kehoe (1990)). Recently, Dovis and Kirpalani (2021) extends this literature to a situation where there there is uncertainty about whether the policymaker can commit ex post. Insights from these studies will be important to our ongoing work to introduce a long-horizon opportunistic type into our framework.

[14] That is, similar policy objectives and the same Lucas-style Philips curve as used by Kydland and Prescott (1977) and Barro and D.B. Gorden (1983a).

[15] Another dimension of these studies is on signalling equilibria, including the appropriate private sector interpretation of monetary policy announcements when these signals may be sent either by a truth-telling committed type or a dissembling alternative type. The key conclusion – reinforced by the careful work of Lu (2013) on a related fiscal model – was that a signalling equilibrium involves the truth-telling committed type announcing a policy that solves a natural optimal policy problem and the opportunistic type sending the same message. We therefore abstract from the analysis of signalling equilibria.

vate sector learning with imperfect public monitoring, and prospective regime change.[16] We build on our prior work, Lu et al. (2016), that studies optimal reputation building by a committed policymaker with a non-committed policymaker behaving mechanically. In the current setup, both types of policymakers are purposeful. The optimal response of a non-committed policymaker significantly changes the reputation building incentives of the committed policymaker: a new committed policymaker with a good initial reputation will follow an apparently gradualist plan, while one with a poor initial reputation will select more dramatic actions.

2.3 Recursive contracts and optimal policy design

A methodological contribution of this paper is the development of a recursive approach that determines optimal intended inflation of both types of policymakers, when the private sector forms forward-looking rational inflation expectations. We conceive of the committed policymaker as a principal who chooses state-contingent plans for his own actions and those of the two agents – the private sector and the opportunistic policymaker – subject to a rational expectation constraint for the former and an incentive compatibility constraint for the latter. We follow the recursive contracts literature to recast the optimization problem in a recursive form,[17] but modify the standard approach by employing a "change of measure" to tackle a particular challenge in our setup where one agent, the private sector, disagrees with the principal – the committed policymaker – about the probabilities of specific future histories.

2.4 Time series econometrics

There are interesting connections of our work to prominent studies of inflation and inflation forecasting using reduced form and structural models.

Stochastic trends: Many econometric models of inflation contain a stochastic trend.[18] The estimated stochastic trends are sometimes interpreted as time-varying inflation targets, either structurally or informally.

[16] Bianchi (2012), Debortoli and Nunes (2014), Debortoli and Lakdawala (2016) also develop models where agents anticipate a possible policy regime change.

[17] Khan et al. (2003), Golosov et al. (2016) and Marcet and Marimon (2019).

[18] A few examples are Erceg and Levin (2003), Smets and Wouters (2003), Ireland (2004), Stock and Watson (2007), and Cogley et al. (2010).

Our model has only two structural shocks, both stationary, and there is a constant target inflation rate that serves as a long-run objective for the committed type. Yet, within a regime, reputation evolves as a martingale relative to the information set of private agents which includes both actual inflation and price shocks. Reputation is one of our model's state variables – influencing both intended and actual inflation in a nonlinear manner – so that it can potentially impart an apparent stochastic trend to inflation.

Markov switching: A policymaker in our model is one of two types. We assume that this type is not observed by private agents, although the dates of regime switches are public information. Our setup is thus similar to the "Markov switching" models pioneered by Hamilton (1989), but with some important differences. In a standard Markov switching model, agents learn about a hidden state when incoming data is more likely to have come from one state relative to another, so that learning proceeds more rapidly when states are more dispersed. That makes it too easy to learn with the wide variation in observed inflation rates.[19]

Crucially, in our setup, the policy difference between the two regimes varies endogenously with reputation, i.e., the private sector's belief about policymaker type. In particular, the optimal policy difference is small when the private sector believes that the committed type most likely is in place, and is larger when it thinks otherwise. The private sector's belief, in turn, is determined by past policy differences via Bayesian learning. The interplay of optimal policy difference and evolving private sector beliefs allows us to make learning relevant while matching the large inflation swings during the period of the Great Inflation and the Volcker Disinflation.

3 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, and their forward-looking deci-

[19]Combining Markov switching with a stochastic trend, M. Evans and Wachtel (1993) develop a two-regime model to explain U.S. inflation that is immune from this feature, as one regime is a persistent but stationary process while the other is a random walk. They highlight that private agents' learning would lead to runs of forecast errors.

sions reflect the possibility that an announced policy plan may not be executed.

3.1 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 β is their time discount factor, $E_t\pi_{t+1}$ is their expectation about the next-period inflation (with e_t being short-hand for discounted expected inflation), and ς_t is a cost-push shock governed by an exogenous Markov chain with the transition probabilities $\varphi(\varsigma_{t+1}; \varsigma_t)$.

3.2 Policymaker

The policymaker is responsible for the inflation rate, π, but cannot control it exactly.[20] 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.[21] An *opportunistic* type ($\tau = 2$) makes the same announcements, but chooses his own intended inflation on a period-by-period basis.

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

[20] 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.

[21] We specify intended inflation rather than intended output for analytical convenience. If policy instead controlled intended real aggregate demand $\underline{x}_{\tau t}$ and $x_{\tau t} = \underline{x}_{\tau t} + \sigma_{x\tau}\varepsilon_t$, the Phillips curve $\pi_t = \kappa x_t + e_t + \varsigma_t$ implies that a choice of $\underline{x}_{\tau t} = \frac{1}{\kappa}[a_t - e_t - \varsigma_t]$ leads to identical intended inflation, although certain text expressions – particularly those for inflation expectations – are more cumbersome. As in some other related studies (see, e.g., Faust and Svensson (2001) and Sargent (1999)), we abstract from policy instruments. By contrast, Orphanides and Williams (2005), Cogley et al. (2015), and Melosi (2016) study macroeconomomic outcomes with private agent learning under an interest rate rule.

independent of the intended inflation: [22]

$$\varepsilon_{1t} = \pi_t - a_t \text{ and } \varepsilon_{2t} = \pi_t - \alpha_t. \tag{2}$$

The policymaker of type τ has the following momentary objective

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

which depends on inflation π and output gap x. There is a long-run inflation target π^* and a type-specific output target x_τ^*.[23] The committed type has a time discount factor β_1; the opportunistic type is myopic.

3.3 Timing of events

The within-period timing is shown in Figure 3. Private agents start period t with an assessment of the probability that the incumbent policymaker is the committed type, which we denote by ρ_t and call *reputation*. Next, current policymaker is replaced via a publicly observed event that occurs with probability q, in which case the regime clock t is set to zero and the new policymaker partially inherits his predecessor's reputation: $0 < \rho_0 = l(\rho_t) < \rho_t$. Then, the exogenous cost-push shock ς_t is realized. If there is a new policymaker, he announces a new inflation plan. Otherwise, either type of continuing policymaker simply reiterates that current economic conditions call for an intended inflation a_t. Next private agents form their expectations about the next-period inflation, e_t. Then the policymaker implements the intended inflation, a_t or α_t, depending on his type, which leads to a random inflation rate π_t with a density $g(\pi_t|a_t)$ or $g(\pi_t|\alpha_t)$, and an output gap x_t determined by the Phillips

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

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

Figure 3: Timing of events within a period

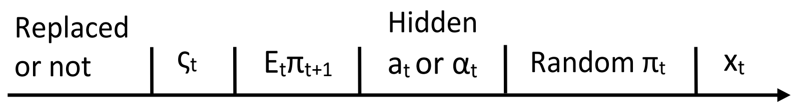

curve. This new information leads private agents to updated their beliefs about policymaker type, as described further below.

4 Macro Equilibrium in a Dynamic Game

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

4.1 Public Equilibria

Define the public history of the current regime $h_t = \{h_{t-1}, \pi_{t-1}, \varsigma_t\}$ as the collection of all past realizations of inflation rates and exogenous states. We restrict our attention to equilibria in which all strategies depend only on the public history, i.e., "public strategies."[24] We denote the committed and opportunistic policymaker's equilibrium strategies as $a(h_t)$ and $\alpha(h_t)$, respectively.

4.2 Perfect Bayesian Equilibria

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

4.2.1 Consistent beliefs: Reputation Dynamics

Consistency requires that the private sector's belief about policymaker type should be updated according to the equilibrium strategies of intended inflation

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

of both types of policymaker, $a(h_t)$ and $\alpha(h_t)$, the observed inflation π_t, and the Bayes' rule (4). Thus, starting from $\rho(h_0) = \rho_0$, the private sector's belief ρ is updated recursively,

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

The denominator is the private sector's likelihood of a particular inflation rate, so within-regime reputation is a martingale relative to public history,

$$E(\rho(h_{t+1})|h_t) = \int \rho(h_{t+1}) \left[\rho_t g(\pi|a_t) + (1-\rho_t) g(\pi|\alpha_t))\right] d\pi = \rho_t. \tag{5}$$

4.2.2 Consistent beliefs: Inflation Expectations

Consistency further requires the private sector's expectation about the next-period inflation $e_t = \beta E_t(\pi_{t+1})$ to be rational and satisfy:

$$e\left(h_{t}\right)=\beta\left\{\begin{array}{c}\rho\left(h_{t}\right) E\left[(1-q) a\left(h_{t+1}\right)+q z\left(h_{t+1}\right) | h_{t}, \tau_{t}=1\right]+ \\ (1-\rho\left(h_{t}\right)) E\left[(1-q) \alpha\left(h_{t+1}\right)+q z\left(h_{t+1}\right) | h_{t}, \tau_{t}=2\right]\end{array}\right\} \tag{6}$$

This expression is complicated due to the possible future regime change, which occurs with probability q. In the event of a regime change, we use z_{t+1} to denote the private sector's nowcast of inflation in period $t+1$, and z_{t+1} in equilibrium should be

$$z(h_{t+1}) = \rho_0 a(h_0) + (1-\rho_0)\alpha(h_0) \text{ where } \rho_0 = l(\rho_{t+1}) \tag{7}$$

due to the inheritance mechanism for reputation discussed above.

In the event of a regime continuation, the next-period inflation will depend on the type of current policymaker. With probability ρ_t, the current policymaker is committed, who will generate stochastic inflation π_t with density $g(\pi_t|a(h_t))$ and will continue to implement the inflation plan $a(h_{t+1})$ next period. Since $h_{t+1} = \{h_t, \pi_t, \varsigma_{t+1}\}$, the conditional expectation in the first line of (6) is $E\left[\cdot|h_t, \tau_t = 1\right] = \int \sum_{\varsigma_{t+1}} \varphi\left(\varsigma_{t+1}; \varsigma_t\right)(\cdot)\, g\left(\pi_t|a\left(h_t\right)\right) d\pi_t$. Similarly, conditional on the current policymaker being opportunistic, he will generate stochastic inflation π_t with density $g(\pi_t|\alpha(h_t))$ and will implement $\alpha(h_{t+1})$ next period. Hence, $E[\cdot|h_t, \tau_t = 2] = \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t)(\cdot)\, g(\pi_t|\alpha(h_t))\, d\pi_t$.

4.2.3 Sequential rationality of the opportunistic type

An opportunistic policymaker takes private sector expected inflation $e(h_t)$ as given and chooses intended inflation α each period to maximize the expected objective $\int u(\pi, x, \tau = 2)\, g(\pi|\alpha)\, d\pi$, subject to the NK Phillips curve (1). Defining ι as *intrinsic inflation bias* since $\alpha = \pi^* + \iota$ if expected inflation is at target, we write the linear best response function as

$$\alpha(e, \varsigma) = \pi^* + \iota + A[e - \beta\pi^*] + A\varsigma = Ae + B(\varsigma) \tag{8}$$

with $A = \frac{\vartheta_x}{\kappa^2 + \vartheta_x} < 1$, $\iota = A[\kappa x_2^* - (1-\beta)\pi^*]$, and $B(\varsigma) = (1 - A\beta)\pi^* + \iota + A\varsigma$.

The top panel of Figure 4 shows the full information *Nash equilibrium inflation bias*, at the intersection of best response (red solid) and the 45 degree (black dash) lines,

$$\alpha^{NE} - \pi^* = \frac{\iota}{1 - A\beta}, \tag{9}$$

when $\varsigma = 0$: this is greater than the intrinsic bias, particularly when $A\beta$ is close to 1. To draw the figure, we use two conventional parameter values, $\pi^* = 1.5\%$ and $\beta = .995$. Since Blinder (1997) sees little intrinsic bias, we set $\iota = .5\%$. Finally, $A = .94$ leads to a NE bias of 8%.

Next, we construct opportunistic intended inflation α using the same parameters and the SPF one-quarter-ahead forecast as a measure for expected inflation. The bottom panel of Figure 4 shows how this constructed policy comoves with expected inflation, along with actual inflation as a reference point.[25] Our analysis of the rise, fall and stabilization of US inflation incorporates this mechanism.

[25] The gradual increase in e leading to a gradual increase in α is in line with a discussion in Sargent and Soderstrom (2000).

Figure 4: Optimal Response of Opportunistic Policy to Inflation Expectations

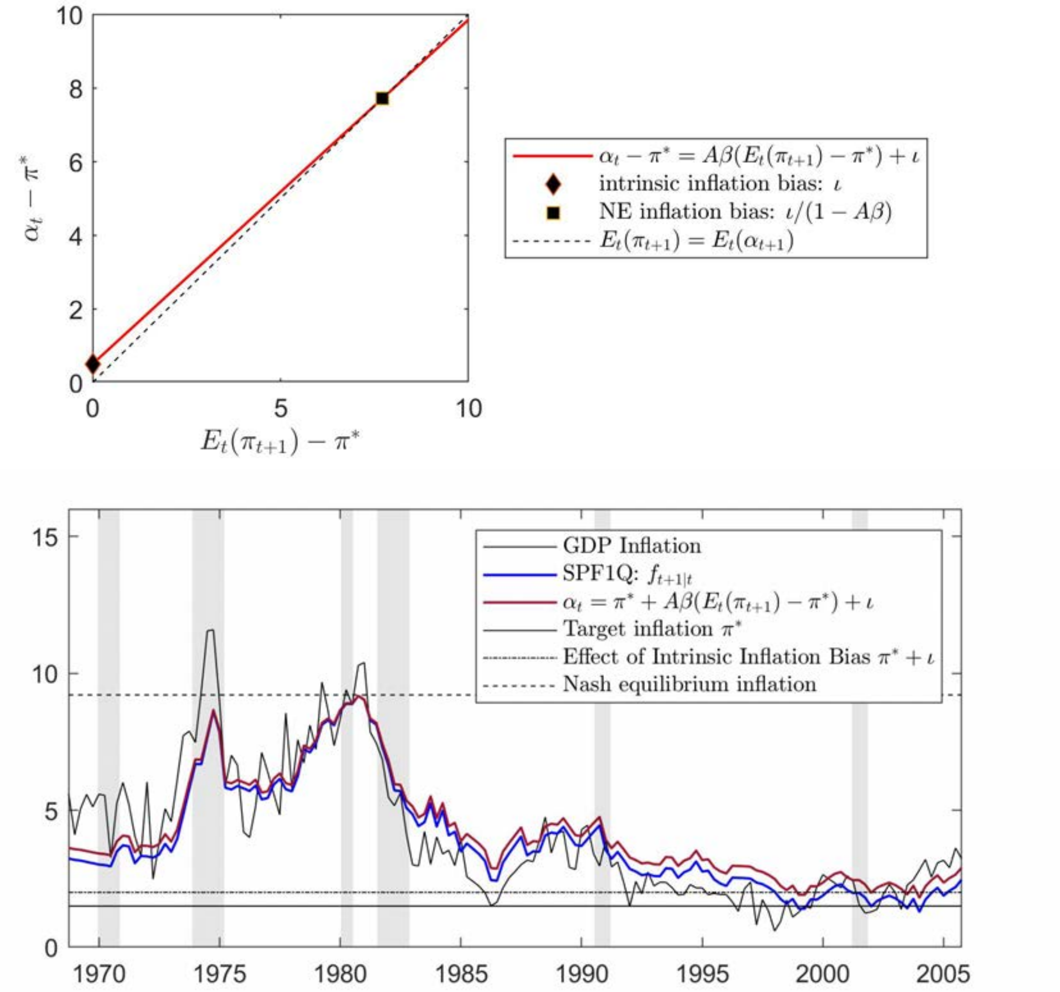

In absence of cost push shocks, the intended inflation α_t of the opportunistic policymaker includes a long-run inflation target π^*, an intrinsic inflation bias ι and a response to the private sector's expected inflation $E_t(\pi_{t+1})$. Top panel: The best response function of α_t to the expected inflation $E_t(\pi_{t+1})$. Bottom panel: The time series of α_t constructed using SPF one quarter forecast as a measure for the private sector's expected inflation $E_t(\pi_{t+1})$.

4.2.4 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 beginning of his term and then subsequently executes it. As the announcement is public information, the committed type has strategic power over the private sector's expectations and thereby over the opportunistic type's intended inflation. In particular, in selecting his state-contingent plan, the committed type takes into account that (i) the private sector's expectation $e\left(h_t\right)$ is based on a *consistent* belief system (6), through which the intended inflation strategies of the committed and opportunistic policymaker $a(h_t)$ and $\alpha(h_t)$ determine how e_t responds to the past history h_t; (ii) the opportunistic intended inflation strategy $\alpha(h_t)$ is sequentially rational, satifying (8), so that it is affected by the expected inflation $e\left(h_t\right)$, and in turn by the committed intended inflation strategy $a(h_t)$.

The strategy of the committed type is *sequentially rational* if it maximizes his expected present discounted payoff at the beginning of his term,[26]

$$U_0 = \sum_{t=0}^{\infty} (\beta_1(1-q))^t \sum_{h_t} p(h_t) \underline{u}(a_t, e(h_t), \varsigma_t, \tau_t = 1), \tag{10}$$

where $\underline{u}\left(a, e, \varsigma, \tau = 1\right) \equiv \int u\left(\pi, x(\pi, e), \varsigma, \tau = 1\right) g\left(\pi | a\right) d\pi$ is the expected momentary objective with x replaced by $x(\pi, e) = \left(\pi - e - \varsigma\right)/\kappa$. Note (10) employs the probability of a specific history $h_t = [\varsigma_t, \pi_{t-1}, h_{t-1}]$ where inflation is generated by the committed type, i.e.,[27]

$$p(h_t) = \varphi(\varsigma_t; \varsigma_{t-1}) g(\pi_{t-1} | a(h_{t-1})) p(h_{t-1}) \tag{11}$$

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.

4.3 Public Perfect Bayesian Equilibrium

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

[26]We assume the committed policymaker maximizes payoffs within his own term, so his discounting includes both the time discount factor β_1 and the replacement probability q.

[27]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.

Definition 1. A Public Perfect Bayesian Equilibrium is a set of functions $e(h_t)$, $\rho(h_t)$, $\alpha(h_t)$, $a(h_t)$ and $z(h_t)$ such that:

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

(ii) given the expected inflation function, $e(h_t)$, the strategy for the opportunistic type policymaker, $\alpha(h_t)$ satisfies (8);

(iii) the strategy for the committed type policymaker, $a(h_t)$, maximizes his expected payoff (10); and

(iv) given $\alpha(h_t)$, $a(h_t)$, and $\rho(h_t)$, the private sector's nowcast of inflation conditional on a replacement, $z(h_t)$, satisfies (7).

5 Constructing the Equilibrium

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

5.1 The principal-agent approach

Solving the within-regime equilibrium may appear a formidable task, due to two dynamic game elements. First, the policymaker and the public are connected intertemporally: (i) forward-looking expectations make future policies enter the policymaker's current payoffs, thus affecting his current policy choice and (ii) the policymaker's current choices enter the private sector's belief updating, thus affecting its future expectations and in turn the policymaker's policy choices. Second, interactions between the two policymaker types arise via private sector expectations: even though one policymaker type is in charge in each period, an optimal choice depends on what the other type would do since private expectations average across both types' policy choices.

To tackle these complications, we recast the construction of the within-regime equilibrium as the solution to a principal-agent problem. As principal, the committed policymaker maximizes (10) by choosing state contingent plans for his current and future actions and those of two agents, the private sector and the opportunistic policymaker. Two forms of incentive compatibility (IC) constraints are relevant: (i) private sector consistent beliefs (4) and rational expectations (6); and (ii) opportunistic type optimal response to expected inflation (8). We then develop a recursive form for the principal's problem.

Relative to a standard dynamic principal-agent problem, we encounter an unusual challenge to constructing a recursive optimization problem for the principal, in that one agent – the private sector – disagrees with the principal – the committed policymaker – in its belief about the probability of a specific history. The private sector *thinks* that current inflation could be generated by the opportunistic policymaker, as reflected by the second line in expected inflation (6), whereas the committed policymaker *knows* that current inflation is generated by his policy choices, as reflected in $p(h_t)$ in the intertemporal objective (10). Such disagreement in probability beliefs between the principal and the agent creates difficulty in putting a Lagrangian component associated with the rational expectation constraint (6) in a recursive form, following the approach laid out by Marcet and Marimon (2019) and others.[28]

We solve this challenge by a "change of measure". Attaching a multiplier $\gamma(h_t)$ and the committed type's probability of history $p(h_t)$ as weights to the constraint (6), we form the Lagrangian component as:

$$\Psi_0 = \sum_{t=0}^{\infty} (\beta_1(1-q))^t \sum_{h_t} p(h_t)\gamma(h_t)[e_t - e(h_t)], \tag{12}$$

Then, we rewrite $E\left[\cdot|h_t, \tau_t = 2\right]$ in (6) as $\int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1};\varsigma_t)[\cdot]\lambda(\pi_t, a_t, \alpha_t) g(\pi_t|a(h_t))d\pi_t$ where $\lambda(\pi_t, a_t, \alpha_t) \equiv g(\pi_t|\alpha_t)/g(\pi_t|a_t)$ is the likelihood ratio. This allows us to express (12) recursively, which leads to the following proposition.[29]

[28]See also Kydland and Prescott (1980), Chang (1998) and Phelan and Stacchetti (2001).

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

PROPOSITION 1. Given $z(\varsigma, \rho)$, the within-regime equilibrium is the solution to the following recursive optimization problem, subject to the IC constraint $\alpha = Ae + B(\varsigma)$

$$(13) \quad W(\varsigma, \rho, \mu) = \min_{\gamma} \max_{a,\alpha,e} \{\underline{u}(a, e, \varsigma, \tau = 1) + (\gamma e + \mu\omega) + \beta_1 (1-q) \int \sum_{\varsigma'} \varphi(\varsigma'; \varsigma) W(\varsigma', \rho', \mu') g(\pi|a) d\pi\},$$

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

$$(14) \quad \mu' = \frac{\beta}{\beta_1 (1-q)} \gamma\rho, \text{ with } \mu_0 = 0$$

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

In this program, some components are familiar from prior research using the recursive approach,[30] but some are not due to the unique features of our model. The component $(\gamma e + \mu\omega)$ arises from the forward-looking rational expectations constraint (6). The pseudo state variable μ records past promises made by the committed type in his management of expectations. Next period's pseudo state μ' evolves according (14), keeping track of the shadow price of current promise γ and the effect of a on past expected inflation measured by ρ, as well as adjusting for differing discount factors. A new policymaker is not held accountable for his predecessor's promises, so the initial value of μ is zero.

With two policymaker types and stochastic replacement, the term ω con-

[30]In fact, if $q = 0$, $\beta_1 = \beta$, and $\rho = 1$ always, this is 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 along Marcet and Marimon (2019) lines 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 $\mu_0 = 0$. The presence of both min and max stems from the fact that the optimum is a saddlepoint as in the Kuhn-Tucker theorem.

tains more than the promised a because private sector's expected inflation also depends on the opportunistic policymaker's intended inflation α and the inflation z in a new regime. The weights attached to a, α, and z reflect the exogenous replacement probability q, the endogenous reputation state ρ, and the divergent probability beliefs about inflation π held by the committed policymaker and the private sector. This final feature leads to $(1-\rho)/\rho$ in ω.[31]

5.2 The PBE fixed point requirement

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

$$z^*(\varsigma, \rho) = \rho_0 a^*\left(\varsigma, \rho_0, 0; z^*(\varsigma, \rho)\right) + (1-\rho_0)\, \alpha^*\left(\varsigma, \rho_0, 0; z^*(\varsigma, \rho)\right) \tag{16}$$

with $a^*(.)$ and $\alpha^*(.)$ obtained from the recursive program (13) given $z^*(\varsigma, \rho)$, $\rho_0 = l(\rho)$ being partially inherited initial reputation, and $\mu_0 = 0$ as prior commitments are no longer binding in a new regime.[32]

5.3 Focusing on policy trade-offs and computation

The recursive program in Proposition 1 is valuable, as it sheds light on the relevant state variables. But it is inefficient for computation because there are many choice variables. Further, it can be hard to isolate the key trade-offs facing the policymaker. The following Lemma provides both computational and conceptual benefits, by developing implications of the forward-looking rational expectation constraint (6).[33]

Lemma 1. Given (ς, ρ) and that future policymakers follow the equilibrium strategies: $a^*\left(\varsigma', \rho', \mu'\right)$, $\alpha^*\left(\varsigma', \rho', \mu'\right)$ and $z^*\left(\varsigma', \rho'\right)$, rationally expected inflation $e(\delta, \mu'; \varsigma, \rho)$ is uniquely determined by the contemporaneous policy difference $\delta = a - \alpha$, and the future pseudo-state variable μ'.

This lemma stems from the fact that the committed policymaker can influence expected inflation through two channels. Via the *learning channel*, the

[31] Appendix A.9 eliminates the likelihood ratio λ using Bayes' rule.

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

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

committed policymaker affects the future reputation variable ρ', as a larger policy difference δ raises the speed of private sector learning about current policymaker type. This channel is formalized when we simplify (15) to $\rho' = b(\varepsilon_1, \delta, \rho)$ by replacing $g(\pi|a) = \phi_1(\varepsilon_1)$ and $g(\pi|\alpha) = \phi_2(\pi - a + a - \alpha) = \phi_2(\varepsilon_1 + \delta)$, where $\phi_1(\cdot)$ and $\phi_2(\cdot)$ are the densities of ε_1 and ε_2 respectively. Via the *expectation anchoring channel*, the committed policymaker adjusts the future pseudo-state variable, with a higher μ' lowering next-period committed intended inflation essentially by making it more expensive.

Using Lemma 1, we simplify the recursive program (13), moving from choosing (γ, a, α, e) to merely choosing (δ, μ'). Specifically, we replace e with $e(\delta, \mu'; \varsigma, \rho)$, α with $Ae + B(\varsigma)$, and a with $\alpha + \delta$ in $\underline{u}(\cdot)$ and $\omega(\cdot)$ of (13) to obtain $\underline{\underline{u}}(\cdot)$ and $\underline{\underline{\omega}}(\cdot)$ in a simplified program:[34]

Proposition 2. Given $z^*(\varsigma, \rho)$ and $U^*(\varsigma, \rho, \mu)$, the recursive optimization (13) reduces to

$$W(\varsigma, \rho, \mu) = \max_{\delta, \mu'} \left[\underline{\underline{u}}(\delta, \mu') + \mu \underline{\underline{\omega}}(\delta, \mu') + \beta_1 (1 - q) \Omega(\delta, \mu'); \varsigma, \rho \right] \tag{17}$$

with $\Omega(\delta, \mu') = \int \sum_{\varsigma'} \varphi(\varsigma'; \varsigma) U^*(\varsigma', b(\varepsilon_1, \delta, \rho), \mu') \phi_1(\varepsilon_1) d\varepsilon_1$.

The equilibrium $U^*(\varsigma, \rho, \mu)$ satisfies the following functional fixed point

$$U^*(\varsigma, \rho, \mu) = W(\varsigma, \rho, \mu) - \mu \underline{\underline{\omega}}(\delta^*, \mu'^*) \tag{18}$$

where $W(\varsigma, \rho, \mu)$, δ^* and μ'^* are the solution to the simplified recursive program (17) conditional on $U^*(\varsigma, \rho, \mu)$.

Lemma 1 and Proposition 2 facilitate our computation. With a guessed function $z(\varsigma, \rho)$ specified in the outer loop, we (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 new U function. Within the inner loop, we iterate until the policy functions converge.[35] We then calculate a new $z(\varsigma, \rho)$ and repeat the process

[34] Appendix B provides detailed derivation of this simplified recursive program.

[35] The problem is not linear quadratic due to Bayesian learning. We therefore use a projection method to obtain a global solution.

until the outer loop has reached a fixed point in z.

6 Inflation Regimes and Reputation

We now move to exploring the positive implications of our theory.

6.1 Linking the theory to the data

Quantification requires regime change dates, state variables and parameters.

Regimes: Choice of possible regime switch dates is subtle. Many monetary histories highlight the Fed chair's identity and nature, as in Friedman and Schwartz's celebrated Great Contraction chapter. But other histories stress combined efforts of presidential administrations and the central bank (including Meltzer (2014), Levin and Taylor (2013), and Binder and Spindel (2017)). Our benchmark is to specify a new regime with each chairman: 1970Q1 (Burns), 1978Q1 (Miller), 1979Q4 (Volcker's October 1979 announcement of new operating procedures), and 1987Q4 (Greenspan).

Reputation and cost-push shocks: Proposition 1 highlights three state variables $s_t = [\varsigma_t, \rho_t, \mu_t]$, known to private agents but not to us. We exploit our model's implication that private agents multi-period inflation forecasts are functions of $E_t(\pi_{t+k}) = f(\varsigma_t, \rho_t, \mu_t, k)$.[36] We choose states to exactly match the SPF data $f_{t+k|t}$ at one quarter and three quarter horizons (k=1,3) [37]:

$$f_{t+k|t} = E_t(\pi_{t+k}) = f(\varsigma_t, \rho_t, \mu_t, k), \text{for } k = 1, 3. \tag{19}$$

With the extracted state $\widehat{\rho}_t$ and the predetermined state $\widehat{\mu}_t$, we determine $\widehat{\mu}_{t+1}$ using the equilibrium decision rule, $\mu'^*(0, \widehat{\rho}_t, \widehat{\mu}_t)$, continuing recursively to calculate a full history of states.[38] At regime switch dates $\widehat{\mu}$ is set to zero.

Following the term structure intuition discussed earlier, longer-term forecasts (SPF3Q) depend more on the persistent reputation variable ρ_t, while shorter-term forecasts are more sensitive to transitory price shocks ς_t, as illus-

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

[37] This allows us to solve for $\widehat{\varsigma}_t$ and $\widehat{\rho}_t$ given the predetermined pseudo state $\widehat{\mu}_t$.

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

trated by the spread between SPF1Q and SPF3Q in Figure 2.[39]

Parameters: Table 1 parameters are selected to match some data and to highlight some model mechanisms. The private sector and committed type share a conventional quarterly discount factor based on a 2% annual real rate. The replacement probability of $q = .03$ implies an average regime duration of 8 years. A new policymaker inherits $\rho_0 = .01 + 0.9\rho_{-1}$, implying initial reputation ranging between 1% and 91%.[40] The 1.5% long-run inflation target lies in the 1 to 2 percent range sometimes cited by central bankers advocating price stability.[41] As in Section 4.2.3, we posit a small intrinsic inflation bias, $\iota = .5\%$ annual rate and set the reduced form parameter $A = .94$ to produce a Nash Equilibrium inflation bias of 8% annual rate at the peak of the Great Inflation. The PC slope κ relates the output gap x to the quarterly inflation π, holding expected inflation fixed, so that $\kappa = .08$ means that an output gap of 3% leads to annualized inflation of -1%, a value compatible with diverse empirical evidence.[42] Given $A = \vartheta_x/(\vartheta_x + \kappa^2)$, matching $A = .94$ requires $\vartheta_x = 0.1$, within the range used by prominent Fed researchers.[43] Finally, all these parameter choices imply that the opportunistic type's target output gap is $x_2^* = 1.75\%$.[44]

Beginning in the 1970s, many studies of inflation use an observable "Food and Energy price shock"(FE shock hereafter).[45] We initially used this proxy

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

[40] This inheritance mechanism would capture private agents rational expectation if a new policymaker's type is the same as his predecessor with probability .9 and is otherwise a random draw with the chance of a committed type equal to 10%.

[41] See Shapiro and Wilson (2019).

[42] 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.

[43] Brayton et al. (2014) and Orphanides and Williams (2013) after translating time units and using Okun's law.

[44] Section 4.2.3 links x_2^* to other parameters via $\iota = A(\kappa x_2^* - (1-\beta)\pi^*)$. With an Okun's law coefficient of 1.67, $x_2^* = 1.75\%$ is targeting unemployment about 1% below the natural rate. A x_1^* is set slightly below x_2^* for computational reasons.

[45] See R.J. Gorden (2013) and Watson (2014). It is constructed as the difference between the growth rate of the overall personal consumption deflator and its counterpart excluding

Table 1: Parameters

β, β_1	Discount factor (private, committed type)	0.995
q	Replacement probability	0.03
ρ_0	Initial reputation after replacement	$1\% + 0.9\rho_{-1}$
κ	PC output slope	0.08
π^*	Inflation target	1.5%
ϑ_x	Output weight	0.1
x_1^*	Committed type's output target	1.7%
x_2^*	Opportunistic type's output target	1.75%
ν	Persistence of cost-push shock (not δ)	0.7
σ_ξ	Std of cost-push innovation	0.7%
σ_ε	Std of implementation error ε_1 and ε_2	1.2%

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

for ς, but eventually settled on extracting shocks from the SPF because these real time forecasters appear to better capture various events including the 1974 inflation peak.[46] The FE shock's serial correlation and its standard deviation are never the less used to determine ν and σ_ξ. We also combined the FE shock and the SPF1Q in an initial approximation to the opportunistic intended inflation α, generalizing the approach behind Figure 4, to obtain the standard deviation of $(\pi - \alpha)$ that prevailed during 1964Q4-1979Q2 and use it as our calibrated standard deviation of implementation errors.

6.2 Inflation history and model-based inflation policies

With an extracted state history, we can construct the *intended inflation policy measures* $\widehat{a}_t = a(\widehat{s}_t)$ and $\widehat{\alpha}_t = \alpha(\widehat{s}_t)$. Then, given observed inflation π_t, (2) implies empirical implementation errors, $\widehat{\varepsilon}_{1t} = \pi_t - a(\widehat{s}_t)$ and $\widehat{\varepsilon}_{2t} = \pi_t - \alpha(\widehat{s}_t)$.

Figure 5 plots these model-based policies ($\widehat{a}_t$ in green, $\widehat{\alpha}_t$ in red), and their associated implementation errors ($\widehat{\varepsilon}_{1t}$ and $\widehat{\varepsilon}_{2t}$ with dash-dotted lines in matching color), with regime switch dates marked by solid vertical lines. Within our model, $\widehat{a}_t$ and $\widehat{\alpha}_t$ are private agent beliefs about the intended inflation of each

food and energy. We display its time series in Appendix C.5.

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

policymaker type and are constructed using states extracted from the SPF forecasts conditional on the set of regime switch dates. Actual inflation π is the black dashed line, but recall that it doesn't enter construction of $\widehat{a}_t$ or $\widehat{\alpha}_t$.

Seeing the entire history of these series, we have an advantage relative to private agents: they only know events through date t. Using this advantage, we identify the Burns-Miller interval 1971Q1-1979Q2 as an opportunistic regime because (i) $\widehat{\varepsilon}_2 = \pi - \widehat{\alpha}$ fluctuates around 0, suggesting that actual inflation consistent with opportunistic policy, and (ii) $\widehat{\varepsilon}_1 = \pi - \widehat{a}$ is generally positive, suggesting actual inflation inconsistent with committed policy.

By the 1982Q4 recession, the situation is clearly reversed. Actual inflation more closely resembles committed policy, with $\widehat{\varepsilon}_1 = \pi - \widehat{a}$ fluctuating around zero, whereas actual inflation lies below opportunistic policy, with $\widehat{\varepsilon}_2 = \pi - \widehat{\alpha}$ being generally negative. We thus identify the bulk of the Volcker regime and the full Greenspan regime as involving commitment policy. Perhaps more controversially,[47] Figure 5 breaks the history at the start of the Reagan administration in 1981Q1 (marked by a dashed vertical line): the average mean of $\widehat{\varepsilon}_2 = \pi - \widehat{\alpha}$ is 0.13% in the earlier interval, and the average mean of $\widehat{\varepsilon}_1 = \pi - \widehat{a}$ is only 0.019% in the later interval. By contrast, the average means of $\widehat{\varepsilon}_1$ in the earlier interval and $\widehat{\varepsilon}_2$ in the later interval are 1.5% and -1%, respectively.

6.3 Interpreting US inflation history 1968-2005

We now examine US inflation history assuming there is an opportunistic policymaker early on and a committed policymaker later. At any point in time, one policy is the current policymaker's intended inflation and the other policy is private agents' rational belief about an alternative policymaker's behavior if confronted with the same observable history.

Figure 6 shows our model-based interpretation of US inflation history 1968-2005. Three time series – inflation π (black), model-implied committed policy $\widehat{a}$ (green), and model-implied opportunistic policy $\widehat{\alpha}$ (red) are repeated from Figure 5. But before 1981Q1, the red line is solid and the green line is dotted because an opportunistic policymaker is taken to be generating the observed inflation. After 1981Q1, the red line is dotted and the green line is solid as a

[47] See Goodfriend and R.G. King (2005) and Orphanides (2005)

Figure 5: Inflation history and model-implied policies

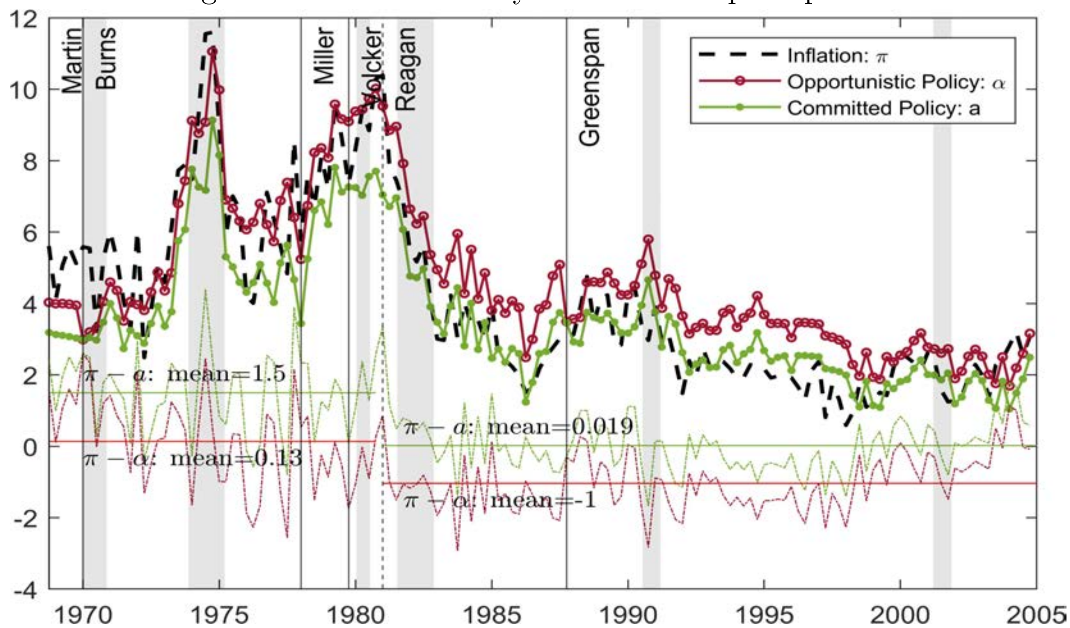

committed policymaker is generating the observed inflation. Figure 6 also plots the SPF1Q forecast (blue), exactly matched by our model expected inflation.

Our framework thus sheds light on why inflation forecast errors turned from persistently positive to persistently negative around 1980, as highlighted by Figure 1. Opportunistic intended inflation α is always higher than committed intended inflation a in our model and expected inflation e is roughly a weighted average of the two. Observed inflation before 1981Q1 is tracked by our opportunistic policy measure $\widehat{\alpha}$ so it exceeds expected inflation – the SPF1Q, hence persistently positive inflation forecast errors arise. After 1981Q1, observed inflation is instead tracked by our committed policy measure $\widehat{a}$, lying below the SPF1Q, yielding persistently negative inflation forecast errors.

Figure 6 also plots the extracted cost-push shock $\widehat{\varsigma}$ (magenta) and the extracted reputation state $\widehat{\rho}$ (cyan and measured on the right hand axis). Note first that the extracted cost-push shock $\widehat{\varsigma}$ covaries strongly with the SPF spread (SPF1Q-SPF3Q plotted in black dotted line), consistent with state

Figure 6: Model-based interpretation of US inflation history

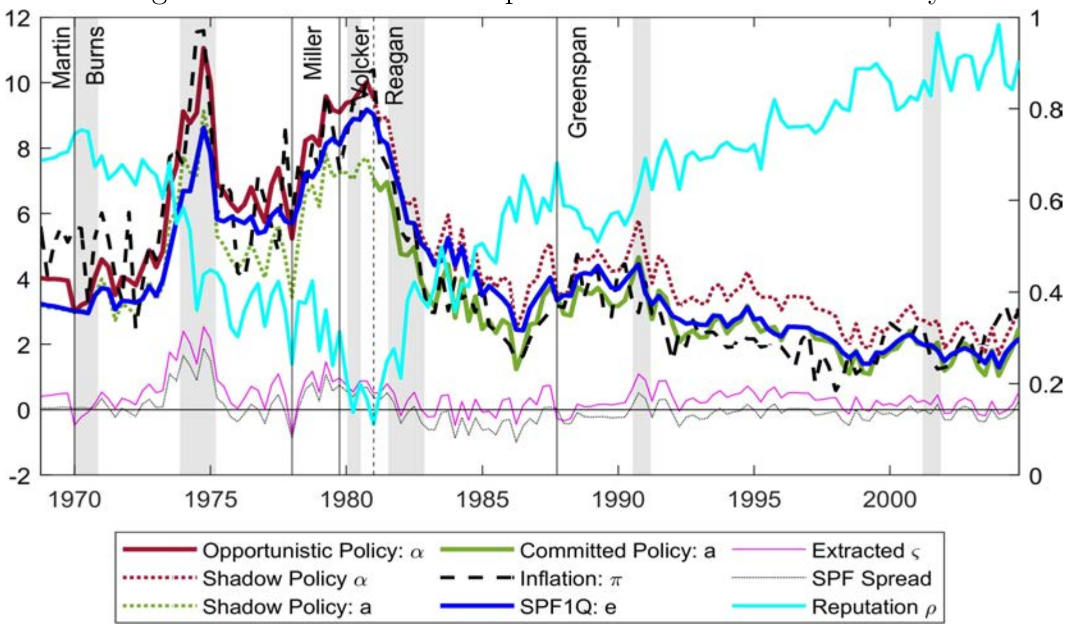

extraction exploiting greater sensitivity of near-term forecasts to transitory shocks. Note next the extracted reputation's big swing: $\widehat{\rho}$ starts from .7 in 1968, decreases through the 1970s to a 1981Q1 trough at .1, and increases afterwards to above .9 in 2005. These reputation dynamics are quantitatively important for our model-implied policy measures, as we will show next.

7 Reputation and Policy

We start with a historical decomposition illustrating the quantitative importance of time-varying reputation in aligning our model-implied policies with historical inflation. We then use equilibrium decision rules to explain how reputation affects optimal policies, and in the process, highlight the crucial role of having a purposeful non-committed policymaker to time-varying reputation.

7.1 Historical Decomposition

Our framework permits a *historical decomposition* of $\widehat{a}$ and $\widehat{\alpha}$ into parts attributable to each state (ς, ρ, μ). To focus on the importance of the reputation state ρ for optimal policies, we construct new policy measures with reputation at a reference value $\underline{\rho}$, but maintaining the cost-push shock $\widehat{\varsigma}_t$ and the pseudo-state variable $\widehat{\mu}_t$ at extracted levels. The gap between constant-ρ and original policies measures the effect of time-varying reputation.[48] Figure 7 contrasts these constant-ρ policies (dash-dotted line) constructed using $(\widehat{\varsigma}_t, \underline{\rho}, \widehat{\mu}_t)$, with original policies (solid line) constructed using $(\widehat{\varsigma}_t, \widehat{\rho}_t, \widehat{\mu}_t)$. Observed inflation (black line) facilitates assessments on how much time variation in ρ helps our model match the US inflation experience.

Figure 7: Historical decomposition: effect of ρ constant at historical average.

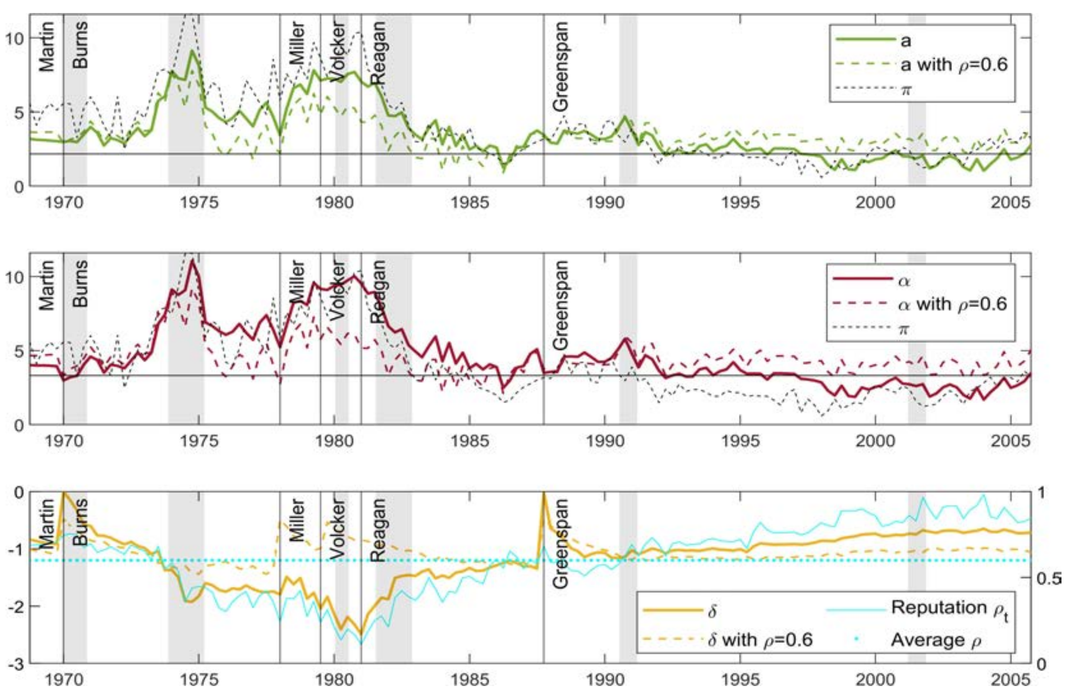

The top two panels in Figure 7 reveal just how important time-varying ρ is for the intended inflation measures $\widehat{a}$ and $\widehat{\alpha}$. Between 1974Q1 and 1985Q4, the constant-ρ policies lie below the original policies, with particularly large gaps during the Great Inflation and the Volcker Disinflation. Our model would

[48] We hold the reputation state constant at $\underline{\rho} = 0.6$ – the sample average of extract reputation state $\widehat{\rho}_t$ during 1968Q4-2005Q4.

badly miss these two important episodes of inflation history without time-varying reputation, even if equipped with the same price shocks and regime changes. When inflation is relatively stable after 1990, our model-based policies are quite close to observed inflation, but the constant-ρ policies are uniformly greater. That is, our model would miss the Great Moderation too if it omitted time-varying reputation.

Section 5.3 showed how to formulate the committed type's choice problem in terms of the policy difference $\delta = a - \alpha$, which is key to Bayesian learning and expected inflation. The model-implied policy difference $\widehat{\delta} = \widehat{a} - \widehat{\alpha}$ (solid line) and its constant-ρ counterpart (dash-dotted line),[49] are shown in the bottom panel of Figure 7 along with the reputation state $\widehat{\rho}$ (cyan solid line) and its historical average (cyan dotted line) measured on the right hand axis. Three notable features shed light on the large gaps between constant-ρ policies and original policies (a, α) in the top two panels. First, the policy difference $\widehat{\delta}$ moves closely with the extracted reputation state $\widehat{\rho}$. Second, when $\widehat{\rho}$ is lower than its historical average, e.g., between 1974Q1 and 1985Q4, both $\widehat{a}$ and $\widehat{\alpha}$ rise above their constant-ρ counterparts, with the policy difference $\widehat{\delta}$ larger than its constant-ρ counterpart. Third, when $\widehat{\rho}$ is higher than its historical average, e.g., after 1990, both $\widehat{a}$ and $\widehat{\alpha}$ fall relative to their constant-ρ counterparts, with the policy difference $\widehat{\delta}$ smaller than its constant-ρ counterpart.

7.2 Effects of reputation on equilibrium decision rules

Equilibrium decisions $\{a^*, \alpha^*\}$ and their difference $\delta^* = a^* - \alpha^*$ depend on reputation ρ, along with the other two more standard states. The three panels of Figure 8 display how each choice (blue line) depends on ρ, conditional on the cost-push shock ς being zero and two levels of the pseudo state μ.[50] In the top and middle panels, we also plot the inflation target $\pi^* = 1.5\%$ (black dotted line): note that a^* in the upper right panel is at this level when $\rho = 1$.

The middle panels highlight that the equilibrium opportunistic policy α^* decreases with reputation ρ, which is intuitive given that higher reputation

[49] $\delta(\widehat{\varsigma}_t, \underline{\rho}, \widehat{\mu}_t) = a(\widehat{\varsigma}_t, \underline{\rho}, \widehat{\mu}_t) - \alpha(\widehat{\varsigma}_t, \underline{\rho}, \widehat{\mu}_t)$ is the constant-ρ policy difference.

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

Figure 8: Effects of ρ on equilibrium policies

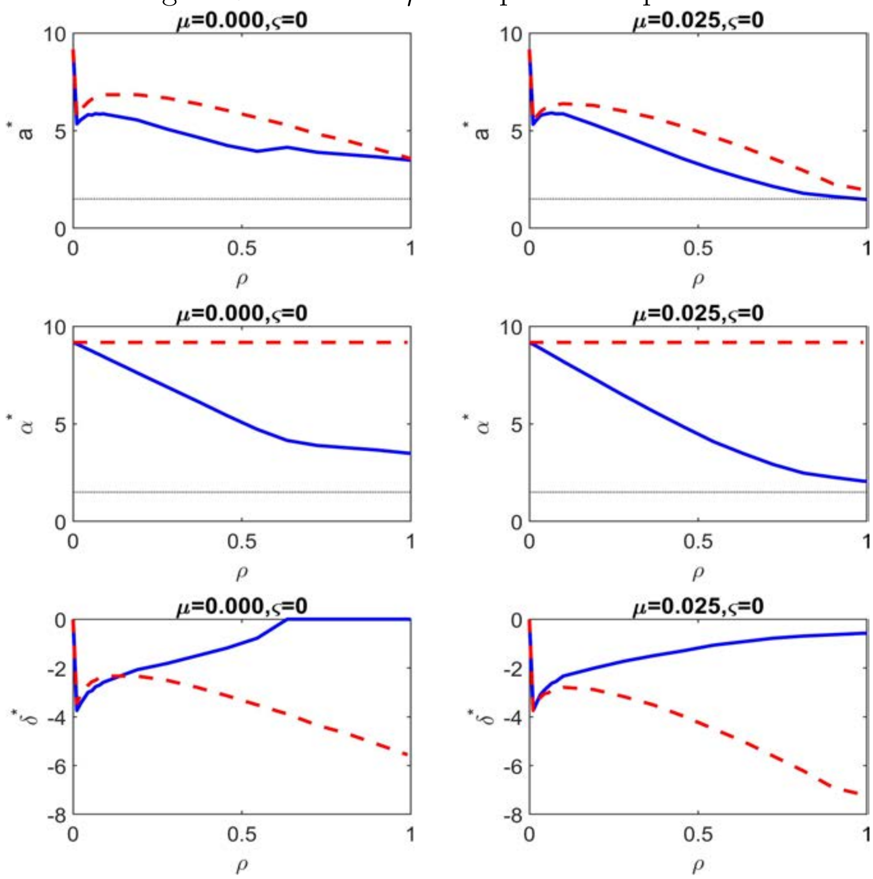

Equilibrium decision rules: top panels: intended inflation of committed policymaker a^*; middle panels: intended inflation of non-committed policymaker α^*; bottom panels: policy difference $\delta^* = a^* - \alpha^*$. Blue solid lines are decision rules in our model where the non-committed policymaker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

yields lower expected inflation and the Section 4.2.3 link between α^* and e.

The consequences of reputation for the policy difference, δ^*, are shown in the bottom panels. Notice first that $\delta^* \leq 0$: equilibrium committed policy is always lower than equilibrium opportunistic policy. Intuitively, the committed policymaker invests in reputation when $\rho < 1$. Notice next that δ^* increases with ρ for $\rho > 0$, indicating diminishing returns to investing in reputation.[51] Third, at high reputation, δ^* is either zero or close to zero.[52] Consequently, the learning speed of private agents is zero or very slow when the policymaker is likely to be the committed type, because the observed inflation is only informative about type when the two types behave differently.

The equilibrium committed policy a^*, shown in the top panels, can be understood as the sum of α^* and δ^*. Therefore, the effect of reputation ρ on a^* depends on the relative strength of effects on α^* versus on δ^*. In our calibration, the Nash Equilibrium inflation bias (α^* at $\rho = 0$) is much higher than the intrinsic inflation bias (α^* at $\rho = 1$), resulting in a dominant effect of ρ on α^*. In turn, a^* is generally decreasing in ρ, with a flatter slope than α^*.

These decision rules help us understand the historical decomposition in Figure 7: since a^* and α^* are both decreasing in ρ, our model-implied $\widehat{a}$ and $\widehat{\alpha}$ are higher than their constant-ρ counterparts when extracted $\widehat{\rho}$ is below its historical average, and vice versa.

An important new element, relative to our prior work (Lu et al. (2016)), is a purposeful, if myopic, policymaker rather than a mechanical alternative type. If we instead assume that the non-committed policymaker mechanically adopts a policy rule that would be optimal if $\rho = 0$ – incorporating the Nash Equilibrium inflation bias – then matters are very different: the results are the red dashed lines in Figure 8. The most salient implication is for the policy difference δ^*. Comparing the red dashed lines with the blue lines in the bottom panels, we find that at majority values of ρ, the policy difference is much larger than when the non-committed policymaker is purposeful. With such a mechanical alternative policymaker, the large δ^* means that private agents learn about policymaker type so fast that we lose the time-varying reputation shown above to crucial for capturing many elements of the US inflation experience.[53]

[51] Moreover, it becomes harder to distinguish between the two policy regimes when the private sector attaches a higher likelihood that it is the committed policymaker in place.

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

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

8 Exploring Credibility and Counterfactuals

Our framework sheds light on the much-discussed idea of credibility and permits us to undertake an important counterfactual.

8.1 Credibility and Reputation

Macroeconomists frequently discuss policymaker reputation, as we have above, and the credibility of a specific announcement or program,[54] as we have not. As a prelude to a survey of macroeconomists and central bankers about credibility, Blinder (2000) remarks that his "own favorite definition involves matching deeds to words: a central bank is credible if people believe it will do what it says." While we also like this definition,[55] it is incomplete because it does not allow for partial, but not perfect, credibility.[56] Practical macroeconomists and central bankers regularly discuss ideas such as "greater credibility improves the short-run inflation-unemployment trade-off," "greater credibility brings down the cost of reducing inflation" and "once low inflation has been achieved, a more credible central bank is better able to maintain low inflation."[57] We now describe two measures of partial credibility of a committed policymaker's policy plan $a(s)$ and display these along with reputation in Figure 9.

Credibility gap in inflation units One intuitive measure is the distance between the $a(s)$ and the private sector's nowcast of inflation $E(\pi|s)$,[58] i.e.,

$$a(s) - E(\pi|s) = (1-\rho)[a(s) - \alpha(s)] = (1-\rho)\delta. \tag{20}$$

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

[54]For example, some point to a country's long-term interest rate, presuming it dominated by inflation expectations, as a a measure of credibility for low inflation. M. King (2005) interprets international cross-section of nominal rates in this way. He highlights shifts in nominal and real yields during notable U.K. events, while Goodfriend (1993) links U.S. long-term nominal interest rate to inflation scares, and evolving credibility.

[55]See the opening discussion of "Managing Expectations" (R.G. King et al. (2005))

[56]As M. King (2005) puts it "credibility is not an all-or-nothing matter. Policy is neither credible nor incredible. It is, as we say in economics, a continuous variable."

[57]These quotes are from Blinder (2000), p. 145.

[58]Cukierman and Meltzer (1986) define credibility as: "the absolute distance between the policymaker's plans and agents beliefs about those plans." If a=0 and $\alpha > 0$ is constant, $a(s) - E(\pi|s)$ depends only on ρ and varies inversely with it.

Degree of credibility In an inflation targeting context, credibility is sometimes related to the private sector's probability that inflation will fall in a band around the target, e.g., $a - \theta \leq \pi \leq a + \theta$. In our setup, this probability reflects implementation errors and the private sector's lack of knowledge about policymaker type.[59] We now assume normal implementation errors and let $N(\cdot, \bar{\pi}, \sigma)$ be the normal cdf with mean $\bar{\pi}$ and standard deviation σ. Our second credibility measure is

$$\psi(a, \alpha, \theta, \rho, \sigma) = \rho + (1 - \rho) \frac{[N(a + \theta, \alpha, \sigma) - N(a - \theta, \alpha, \sigma)]}{[N(a + \theta, a, \sigma) - N(a - \theta, a, \sigma)]} \tag{21}$$

which is the ratio of the private sector's probability that inflation falls within the band relative to the committed policymaker's probability. Note that the denominator expression is constant across $a(s)$, while the numerator may be written to stress the policy difference, $N(\delta + \theta, 0, \sigma) - N(\delta - \theta, 0, \sigma)$. That is, our second credibility measure also depends on reputation ρ and the policy difference δ. Figure 9 displays this credibility measure for $\theta = \sigma$, along with the credibility gap and reputation. There is a strikingly high correlation between the credibility gap in inflation units and the degree of credibility.

Credibility and reputation Both of these evolving, partial credibility measures depend on near-term inflation. Under commitment, though, a long-lasting regime will attain $\rho = 1$ and intended inflation will have a stationary distribution with $E(a) = \pi^*$. Hence, ρ_t is a measure of longer-term credibility and, in particular, of the date t likelihood that the current regime will achieve "price stability." In this sense, our model captures the views of some academicians in the Blinder (2000) survey: "a central bank can raise the public's subjective probability that it is 'tough' by keeping inflation low. This probability is, in turn, taken as a measure of the bank's credibility." It is also

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

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

Expressing this as a ratio to $[N(a + \theta, a, \sigma) - N(a - \theta, a, \sigma)]$ leads to (21). This measure is readily generalized to an asymmetric band and type-specific implementation error volatility.

Figure 9: Credibility and Reputation

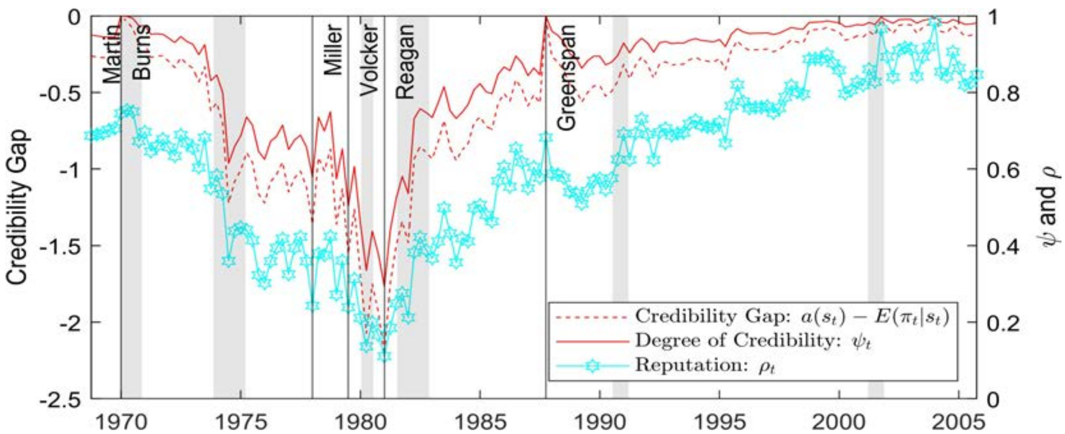

Two measures of the short-term policy credibility are closely associated: the inflation credibility gap defined as $a(s) - E(\pi|s) = (1-\rho)[a(s) - \alpha(s))] = (1-\rho)\delta$ and the degree of credibility defined as the ratio of the private sector's probability that $a - \theta \leq \pi \leq a + \theta$ relative to the committed type's probability. Further, these two measures also rise and fall over time with reputation, ρ, which can be viewed as the likelihood that inflation will be at π^* if the current regime continuous for a long time.

consistent with his summary "that many central bankers take the degree of dedication to price stability as synonymous with credibility."

8.2 US inflation under a commitment regime

We have seen that our theory provides a potential explanation of the behavior of US inflation over 1968 to 2005, with key ingredients being regime shifts, inability of some policymakers to commit, and private sector learning. We now consider how inflation would have evolved if there had been a single committed policymaker in place for the entire period, faced with the same series of ς and ε shocks, and with his type known by private agents.

Figure 10 displays the answer which is striking. The green line labelled "commitment policy" is $\widehat{a}^c = a^*(\widehat{\varsigma}, \rho = 1, \mu)$ so that it evolves with extracted price shocks, while the black solid line marked "commitment inflation" is $\widehat{\pi}^c = \widehat{a}^c + \widehat{\varepsilon}$ so that it contains the same extracted implementation errors.[60] These

[60] $\widehat{\varepsilon}_t = \widehat{\varepsilon}_{2,t}$ before 1981Q1 and $\widehat{\varepsilon}_t = \widehat{\varepsilon}_{1,t}$ afterwards.

Figure 10: Counterfactual Inflation Under Commitment

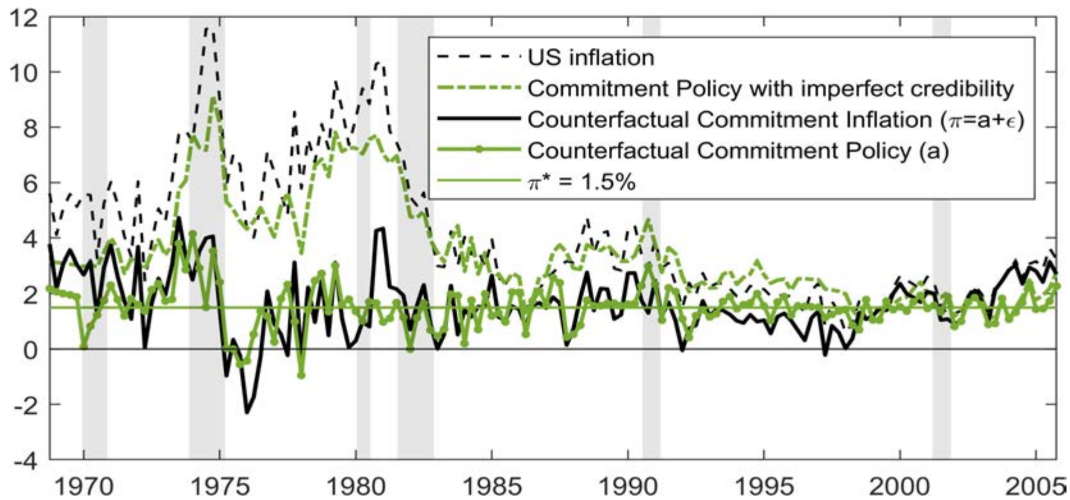

Counterfactual rates are computed assuming a single committed policymaker in place for the entire period, faced with the same shocks and with his type known by private agents. The counterfactual US inflation is dramatically different from observed inflation, even with large temporary "price shocks" in the 1970s. While the policymaker permits inflation to rise temporarily in response to such shocks, there is a subsequent reversal reflecting the desirability of price level targeting in this New Keynesian model.

two series are governed by general NK policy principles. First, our model's timing assumptions imply that policy under commitment will not respond to ε, because these one-time disturbances do not affect intertemporal trade-offs. Second, in line with the "flexible inflation targeting" analytics of Clarida et al. (1999), the dramatic price shocks of the 1970s lead to an increase in actual and intended inflation, rising to about 4% relative to a long-run target π^* of 1.5%. But this above-average inflation is soon followed by an interval of inflation below the long-run target and the large price shocks lead to deflation in 1975-1976.

For most of 1968-2005, intended inflation is very different in this full commitment counterfactual than it would have been with a series of committed policymakers facing evolving reputation (the dashed green line) as in our earlier historical decomposition. However, after the mid 1990s, there is a relatively small difference as reputation is at a high level.

9 Conclusions and Final Remarks

We show that a monetary regime shift model can capture the main features of U.S. inflation between the late 1960s and the mid 2000s. Our setup features a standard forward-looking New Keynesian Phillips curve, policymakers of differing commitment capacity, Bayesian learning by the private sector policymaker type, and occasional observable changes in regime.

Both types of policymakers, committed and opportunistic, behave purposefully in line with the earlier 1980s literature on monetary policy and inflation bias. To this end, we construct a Bayesian perfect equilibrium, in which policymakers and private agents rationally anticipate future regime change. Within a regime, a committed policymaker solves a recursive optimization problem that generalizes a now-standard approach in two necessary and important ways. First, the committed policymaker takes into account the effect of policy actions – intended inflation – on reputation, defined as the private sector's rational belief that a committed type is in place. Second, the committed policymaker understands that (i) private sector inflation expectations include future behavior of an opportunistic type; and (ii) an opportunistic type's intended inflation depends on private sector inflation expectations. A compact representation of this optimization problem permits calculation of decision rules and construction of time series within the Bayesian perfect equilibrium.

Our framework has state variables, observed by the policymaker and private agents, but not by us. We use the inflation forecasts from the Survey of Professional Forecasters to extract these state variables (reputation and a cost-push shock). These state variables allow us to construct time series of intended inflation for committed and opportunistic policymakers from the SPF data without using actual inflation. Yet, when we assume regimes with opportunistic policymakers before 1981 and regimes with committed policymakers afterward, the corresponding intended inflation tracks US inflation's rise, fall, and stabilization between 1970 and 2005.

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

References

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.

_ **and** _ , "Rational Expectations and Policy Credibility Following a Change in Regime," *The Review of Economic Studies*, April 1985, *52* (2), 211–221.

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**, "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Political Economy*, 1983, *91* (4), 589–610.

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

Bianchi, Francesco, "Regime Switches, Agents' Beliefs, and Post-World War II U.S. Macroeconomic Dynamics," *The Review of Economic Studies*, oct 2012, *80* (2), 463–490.

Binder, Sarah A. and Mark Spindel, *The myth of independence: how Congress governs the Federal Reserve*, Princeton ; Oxford: Princeton University Press, 2017.

Blinder, Alan S, "Distinguished Lecture on Economics in Government: What Central Bankers Could Learn from Academics—and Vice Versa," *Journal of Economic Perspectives*, May 1997, *11* (2), 3–19.

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

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

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

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

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

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

_ , Giorgio E. Primiceri, and Thomas J. Sargent, "Inflation-Gap Persistence in the US," *American Economic Journal: Macroeconomics*, 2010, *2* (1), 43–69.

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

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

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

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

_ and Rivcardo Nunes, "Monetary Regime Switches and Central Bank Preferences," *Journal of Money, Credit and Banking*, 2014, *46* (8), 1591–1625.

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.

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

Eusepi, Stefano and Bruce Preston, "The Science of Monetary Policy: An Imperfect Knowledge Perspective," *Journal of Economic Literature*, March 2018, *56* (1), 3–59.

Evans, George W. and Seppo Honkapohja, "Expectations, Learning, And Monetary Policy: An Overview Of Recent Research," Working Papers Central Bank of Chile 501, Central Bank of Chile October 2008.

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

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

Golosov, M., A. Tsyvinski, and N. Werquin, "Recursive Contracts and Endogenously Incomplete Markets," in "Handbook of Macroeconomics," Elsevier, 2016, pp. 725–841.

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

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

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

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

Ireland, Peter N., "Technology Shocks in the New Keynesian Model," *Review of Economics and Statistics*, November 2004, *86* (4), 923–936.

Khan, Aubhik, Robert G. King, and Alexander L. Wolman, "Optimal Monetary Policy," *Review of Economic Studies*, October 2003, *70* (4), 825–860.

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

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

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.

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

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

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

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

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

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

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

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

Mishkin, Frederic S., “Central Bank Commitment and Communication,” April 2008.

— and Klaus Schmidt-Hebbel, “Does Inflation Targeting Make a Difference?,” NBER Working Papers 12876 January 2007.

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

— and John C. Williams, “Imperfect Knowledge, Inflation Expectations, and Monetary Policy,” in “The Inflation-Targeting Debate,” University of Chicago Press, 2005, pp. 201–245.

_ **and** _ , "Monetary Policy Mistakes and the Evolution of Inflation Expectations," in "The Great Inflation: The Rebirth of Modern Central Banking," University of Chicago Press, June 2013.

Parkin, Michael, "Inflation in North America," in "Price Stabilization in the 1990s," Palgrave Macmillan UK, 1993, pp. 47–93.

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

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

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

Sargent, Thomas J., "The Ends of Four Big Inflations," in "Inflation: Causes and Consequences," Univ. of Chicago Press, 1982, pp. 41–98.

_ , *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.

Smets, Frank and Raf Wouters, "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area," *Journal of the European Economic Association*, September 2003, *1* (5), 1123–1175.

Stock, James H. and Mark W. Watson, "Why Has U.S. Inflation Become Harder to Forecast?," *Journal of Money, Credit and Banking*, February 2007, *39* (s1), 3–33.

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

Woodford, Michael, "Macroeconomic Analysis Without the Rational Expectations Hypothesis," *Annual Review of Economics*, August 2013, *5* (1), 303–346.

For Online Publication

Appendices

A Recursive optimal policy design

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

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

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

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

A.1 Intended and actual inflation

At each date, the policymaker chooses intended inflation, denoted as a for the committed type ($\tau = 1$) and α for the alternative type ($\tau = 2$). Intended inflation is not observed by the

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

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

$$\begin{aligned} a &= \int \pi g(\pi|a, \tau = 1) d\pi \\ \alpha &= \int \pi g(\pi|\alpha, \tau = 2) d\pi \end{aligned}$$

Implementation errors are $\varepsilon_1 = \pi - a$ and $\varepsilon_2 = \pi - \alpha$ for the two types. While we allow for different continuous distributions on the same range of inflation outcomes, we do not separately include type τ as an argument to avoid notation clutter in the balance of this appendix (i.e., we write $g(\pi|a)$ and $g(\pi|\alpha)$).

A.2 Measures of history

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

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

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

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

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

Note that

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

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

A.3 Beliefs about current inflation

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

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

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

$$\begin{aligned} \rho(h_t^+) &= \frac{g(\pi_t|a(h_t))\rho(h_t)}{g(\pi_t|a(h_t))\rho(h_t) + g(\pi_t|\alpha(h_t))(1-\rho(h_t))} \\ &\equiv b(\pi_t, a(h_t), \alpha(h_t), \rho(h_t)) \end{aligned} \tag{A23}$$

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

$$\rho(h_t^+) = \frac{\rho(h_t)}{\rho(h_t) + \lambda(\pi_t, h_t)(1-\rho(h_t))} \tag{A24}$$

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

A.5 Constructing expected inflation

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

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

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

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

$$E(\pi_{t+1}|h_{t+1}, \tau_{t+1} = 2) = \alpha([\varsigma_{t+1}, h_t^+])$$

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

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

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

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

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

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

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

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 the private sector's 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 $\beta_1(1-q)$, where β_1 is its structural discount factor and $(1-q)$ reflects discounting due to replacement.

$$U_t = \underline{u}(a_t, e_t, \varsigma_t, \tau = 1) + (\beta_1(1-q))E_t^c U_{t+1}$$

where $\underline{u}\,(a, e, \varsigma, \tau = 1) \equiv \int u\,(\pi, x(\pi, e), \varsigma, \tau = 1)\, g\,(\pi|a)\, d\pi$ is the expected momentary objective with x replaced by $x(\pi, e) = (\pi - e - \varsigma)\,/\kappa$, and the conditional expectation operator $E_t^c(\cdot)$ is using the committed type's probability $p(h_{t+j})$ of a specific history h_{t+j} when his actions generate inflation.

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

$$U_t = \sum_{j=0}^{\infty} (\beta_1(1-q))^j \sum_{h_{t+j}} \frac{p(h_{t+j})}{p(h_t)} \underline{u}(a(h_{t+j}), e(h_{t+j}), \varsigma(h_{t+j}), \tau = 1) \tag{A28}$$

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

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

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

A.7 Rational expectations constraint

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

$$\Psi_t = \sum_{j=0}^{\infty} (\beta_1(1-q))^j \sum_{h_{t+j}} \frac{p(h_{t+j})}{p(h_t)} \gamma(h_{t+j})[e(h_{t+j}) - \beta E(\pi_{t+j+1}|h_{t+j})] \tag{A30}$$

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 alternative type. So, we undertake a "change of measure" and rewrite it as

$$\begin{aligned} \text{(A31)} \quad & \rho(h_t) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t)[\beta(1-q)a(h_{t+1}) + \beta q z(h_{t+1})] g(\pi|a(h_t)) d\pi \\ & + (1-\rho(h_t)) \int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1}; \varsigma_t)[\beta(1-q)\alpha(h_{t+1}) + \beta q z(h_{t+1})] \boldsymbol{\lambda}(\mathbf{h_{t+1}}) g(\pi|a(h_t)) d\pi \end{aligned}$$

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

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

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

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

discount factors, we can write (A30) as

$$\Psi_t = E_t^c[\sum_{j=0}^{\infty}(\beta_1(1-q))^j \psi_{t+j}] \tag{A33}$$

with

$$\psi_t = \gamma_t e_t - \frac{\beta}{\beta_1(1-q)}\gamma_{t-1}\{\rho_{t-1}[(1-q)a_t + qz_t] + (1-\rho_{t-1})\lambda_t[(1-q)\alpha_t + qz_t]\} \tag{A34}$$

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

A.8 The basic recursive specification

The preceding derivations suggest a recursive version of $U_t + \Psi_t$ with states $(\varsigma_t, \gamma_{t-1}, \rho_{t-1},\lambda_t)$. For algebraic convenience, we define $\eta_t = \frac{\beta}{\beta_1(1-q)}\gamma_{t-1}$. Then, the recursive form as in Marcet and Marimon (2019) is

$$\begin{aligned} W(\varsigma_t, \eta_t, \rho_{t-1}, \lambda_t) = \min_{\gamma}\max_{a,\alpha,e}\{&\underline{u}(a_t, e_t, \varsigma_t, \tau = 1) + \gamma_t e_t \\ &- \eta_t[\rho_{t-1}((1-q)a_t + qz_t) + (1-\rho_{t-1})\lambda_t((1-q)\alpha_t + qz_t)] \\ &+ \beta_1(1-q)\int \sum_{\varsigma_{t+1}} \varphi(\varsigma_{t+1};\varsigma_t)W(\varsigma_{t+1}, \eta_{t+1}, \rho_t, \lambda_{t+1})g(\pi_t|a_t)d\pi_t\} \end{aligned} \tag{A35}$$

subject to the IC constraint

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

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

$$\begin{aligned} \eta_{t+1} &= \frac{\beta}{\beta_1(1-q)}\gamma_t \text{ with } \gamma_{-1} = 0 \\ \rho_t &= \frac{\rho_{t-1}}{\rho_{t-1} + (1-\rho_{t-1})\lambda_t} \text{ given } \rho_0 \\ \lambda_{t+1} &= \lambda(\pi_t, a_t, \alpha_t) \text{ with probability } g(\pi_t|a_t) \end{aligned}$$

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

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

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

A.9 State space reduction

For computational and analytical benefits, it is desirable to reduce the state space. We now show how to eliminate the likelihood ratio (λ) from the state vector so that we only need three state variables instead of four. Start by rewriting (21) as

$$\psi_t = \gamma_t e_t - \frac{\beta}{\beta_1(1-q)}\gamma_{t-1}\rho_{t-1}\{[(1-q)a_t + qz_t] + \frac{(1-\rho_{t-1})\lambda_t}{\rho_{t-1}}[(1-q)\alpha_t + qz_t]\} \tag{A36}$$

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

$$\psi_t = \gamma_t e_t - \frac{\beta}{\beta_1(1-q)}\gamma_{t-1}\rho_{t-1}\{[(1-q)a_t + qz_t] + \boxed{\frac{(1-\rho_t)}{\rho_t}}[(1-q)\alpha_t + qz_t]\} \tag{A37}$$

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:

$$\mu_{t+1} = \frac{\beta}{\beta_1(1-q)}\gamma_t\rho_t \text{ with } \mu_0 = 0 \tag{A38}$$

$$\rho_{t+1} = b(\pi_t, a_t, \alpha_t, \rho_t) \text{ with probability } g(\pi_t|a_t) \tag{A39}$$

A.10 Extended recursive program

The recursive optimization (A35) can now be written with only three state variables. While doing so, we extend the program to make it easy to shut down each of the two key mechanisms: endogenous reputation and optimizing behavior by the alternative type.

$$\begin{aligned} W(\varsigma_t, \rho_t^s, \mu_t) &= \min_{\gamma}\max_{a,\alpha,e}\{\underline{u}(a_t, e_t, \varsigma_t, \tau = 1) + \gamma_t e_t + \mu_t\omega_t \\ &\quad + \beta_1(1-q)\int\sum_{\varsigma_{t+1}}\varphi(\varsigma_{t+1};\varsigma_t)W(\varsigma_{t+1}, \rho_{t+1}^s, \mu_{t+1})g(\pi; a_t)d\pi\} \end{aligned} \tag{A40}$$

where

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

subject to the IC constraint

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

with state dynamics allowing exogenous reputation (y=0) or endogenous reputation (y=1)

$$
\begin{aligned}
\mu_{t+1} &= \frac{\beta}{\beta_1(1-q)}\gamma_t \rho_t \\
\rho_{t+1} &= y\rho^s_{t+1} + (1-y)\,\rho \\
\rho^s_{t+1} &= b(\pi_t, a_t, \alpha_t, \rho_t)
\end{aligned}
$$

The recursive program here is written in a general form that allows (i) optimizing or mechanical alternative type and (ii) endogenous or exogenous reputation. The program in Proposition 1 of the main text is a special form of (A40) where there is an optimizing alternative type and endogenous reputation. Hence, in that setting, there is no need to distinguish ρ^s from ρ.

A.11 The Fixed Point Requirement

At the start of a period, there is a reputation ρ_t of a policymaker that would continue its tenure. If there is a replacement, then we assume that this reputation is partly inherited by a new policymaker whose date clock is set to zero, which we write as $\rho_0 = l(\rho_{t+1})$.

Since we now know that optimal policies take the form $a^*(\varsigma, \rho, \mu)$ and $\alpha^*(\varsigma, \rho, \mu)$ given a particular function $z(\varsigma, \rho)$, rational expectations across regimes requires that the expected inflation conditional on a replacement $z(\varsigma, \rho)$ must satisfy the following fixed point:

$$
z^*(\varsigma, \rho) = \rho_0 a^*\left(\varsigma, \rho_0, 0; z^*(\varsigma, \rho)\right) + (1-\rho_0)\,\alpha^*\left(\varsigma, \rho_0, 0; z^*(\varsigma, \rho)\right) \tag{A41}
$$

where $\rho_0 = l(\rho)$ is the new policymaker's initial reputation as described above and we impose $\mu = 0$ as appropriate at the start of a regime.

B Consolidation

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

B.1 Relationship between U and W

If $W(\cdot)$ in (A40) is differentiable, there are two notable implications of this structure.

The **envelope theorem implication** for μ is

$$
W_\mu(\varsigma_t, \rho^s_t, \mu_t) = -\{[(1-q)a_t + qz_t] + \boxed{\tfrac{(1-\rho^s_t)}{\rho^s_t}}[(1-q)\alpha_t + qz_t]\}
$$

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

$$\begin{aligned} 0 &= e_t - \beta_1(1-q)\sum_{\varsigma_{t+1}}\varphi(\varsigma_{t+1};\varsigma_t)\int W_\mu(\varsigma_{t+1},\rho_{t+1},\mu_{t+1})\frac{\partial\mu_{t+1}}{\partial\gamma_t}g(\pi_t|a_t)d\pi_t \\ &= e_t - \beta\sum_{\varsigma_{t+1}}\varphi(\varsigma_{t+1};\varsigma_t)\int W_\mu(\varsigma_{t+1},\rho_{t+1},\mu_{t+1})\rho_t g(\pi_t|a_t)d\pi_t \end{aligned}$$

where the state evolution equation (A38) implies $\partial\mu_{t+1}/\partial\gamma_t = \rho_t\beta/(\beta_1(1-q))$.

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

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

where

$$\frac{1-\rho^s_{t+1}}{\rho^s_{t+1}} = \frac{(1-\rho_t)\,\lambda_{t+1}}{\rho_t}.$$

Hence, in equilibrium where the rational expectation constraint must hold, we obtain the following relationship between the value function $W(\cdot)$ and the optimized objective $U^*(\cdot)$:

(B1) $$\begin{aligned} W(\varsigma_t,\rho^s_t,\mu_t) - \mu_t\omega^*_t &= U^*(\varsigma_t,\rho^s_t,\mu_t) \\ &= \underline{u}(a^*_t, e^*_t,\varsigma_t,\tau=1) + ... \\ &\quad + \beta_1(1-q)\int\sum_{\varsigma_{t+1}}\varphi(\varsigma_{t+1};\varsigma_t)U^*(\varsigma_{t+1},\rho^s_{t+1},\mu_{t+1})g(\pi_t|a_t)d\pi_t \end{aligned}$$

where $\omega^*_t = -\{[(1-q)a^*_t + qz^*_t] + \boxed{\frac{(1-\rho^s_t)}{\rho^s_t}}[(1-q)\alpha^*_t + qz^*_t]\}$.

B.2 Further consolidation

We now show that imposing the rational expectation constraint (A31) on the choice of e_t implies Lemma 1, which allows us to further reduce the recursive program in Proposition 1 to the one in Proposition 2. The key idea is that only the policy difference $\delta = a - \alpha$ matters rather than the levels of a and α.

Recall that (A31) comes from (A27) before undertaking a "change of measure". So the original form of the rational expectation constraint on e_t is:

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

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

Recall $\rho_{t+1} = b(\pi_t, a_t, \alpha_t, \rho_t)$ from (A39) and $b(\cdot)$ is the Bayes' learning rule specified in (A23). The inflation distribution is $\pi = a + \varepsilon_1$ under the committed type and $\pi = \alpha + \varepsilon_2$ under the opportunistic type, with ε_1 and ε_2 being zero mean random variables. We can therefore rewrite the Bayes' learning rule (A23) as

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

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

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

$$\rho' = b\left(\varepsilon_1, \varepsilon_1 + \delta, \rho\right) \text{ conditional on } \tau = 1 \tag{B4}$$

$$\rho' = b\left(\varepsilon_2 - \delta, \varepsilon_2, \rho\right) \text{ conditional on } \tau = 2 \tag{B5}$$

Replacing $g(\pi|a)$ and $g(\pi|\alpha)$ in (B2) with $\phi_1(\pi - a)$ and $\phi_2(\pi - \alpha)$, ρ_{t+1} with (B4) and (B5), and realizing choosing γ_t is equivalent to choosing μ_{t+1} due to $\mu_{t+1} = \frac{\beta}{\beta_1(1-q)}\gamma_t\rho_t$, we obtain Lemma 1 with the added details as follows:

LEMMA 1. Given (ς, ρ), and that future policymakers follow the equilibrium strategies $a^*(\varsigma', \rho', \mu')$, $\alpha^*(\varsigma', \rho', \mu')$ and $z^*(\varsigma', \rho')$, rationally expected inflation is uniquely determined by the contemporaneous policy difference $\delta = a - \alpha$, and the future pseudo-state variable μ'.

$$\begin{aligned} e &= e(\delta, \mu'; \varsigma, \rho) = \beta\rho \int \widehat{M_1}(\varsigma, b(\varepsilon_1, \varepsilon_1 + \delta, \rho), \mu')\phi_1(\varepsilon_1)\, d\varepsilon_1 + \\ & \quad \beta(1-\rho) \int \widehat{M_2}(\varsigma, b(\varepsilon_2 - \delta, \varepsilon_2, \rho), \mu')\phi_2(\varepsilon_2)\, d\varepsilon_2; \end{aligned}$$

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

$$\begin{aligned} \widehat{M_1}(\varsigma, \rho', \mu') &:= \sum_{\varsigma'} \varphi(\varsigma'; \varsigma)\left[(1-q)\, a^*(\varsigma', \rho', \mu') + qz^*(\varsigma', \rho')\right]; \\ \widehat{M_2}(\varsigma, \rho', \mu') &:= \sum_{\varsigma'} \varphi(\varsigma'; \varsigma)\left[(1-q)\, \alpha^*(\varsigma', \rho', \mu') + qz^*(\varsigma', \rho')\right]; \end{aligned}$$

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

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

$$\underline{u}(a, e, \varsigma, \tau = 1) + \mu\omega(a, \alpha) + \beta_1(1-q) \int \sum_{\varsigma'} \varphi(\varsigma'; \varsigma) U^*(\varsigma', \rho', \mu') g(\pi|a) d\pi$$

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

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

$$\text{(B6)} \qquad \underline{\underline{u}}(\delta, \mu') := \underline{u}(Ae + B(\varsigma), e, \varsigma, \tau = 1)$$

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

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

C Forecasting Functions and Matching the SPF

C.1 SPF Data

This project is our first use of Survey of Professional Forecasters data. Many researchers employ the summary data files from the Federal Reserve Bank of Philadelphia, particularly the "annualized percent change of median responses" file, available for the GDP deflator at https://www.philadelphiafed.org/surveys-and-data/pgdp. This file includes an inflation "nowcast" and forecasts at the 1,2,3, and 4 quarter horizons.

In the middle of each quarter, each survey participant submits a forecast for the price level in that quarter and the next four. The FRBP first constructs a median price level for each horizon, say $P_{t+k|t}$ for k=0,1,...4. It then constructs an annualized percentage growth rate using the formula $100 * ([P_{t+k|t}/P_{t+k-1|t}]^4 - 1)$.

The series used in our research differ in two ways. First, we compute annualized percent growth rates as $400 * log(P_{t+k|t}/P_{t+k-1|t})$. Second, we start by calculating these growth rates for each forecaster at each date. We then take take the median of these inflation rates.

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

At some point, we plan to investigate these differences in more detail, as well as looking into the behavior of mean and trimmed mean inflation rates, but the time series employed seemed to us to be the best combination of conventional practice and attention to the underlying survey data. As our theory does not start with microfoundations, it is silent on the best manner to undertake such constructions.

C.2 Recursive forecasting in our theory

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

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

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

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

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

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

C.2.1 Forecasting inflation

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

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

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

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

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

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

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

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

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

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

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

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

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

With probability q there is a regime change, in which case $\mu_{t+1} = 0$ and ρ_{t+1} evolves according to an inheritance mechanism that relates the new policymaker's initial reputation ρ_0 to what it would have been if there was no replacement, i.e., ρ_{t+1}, which we write as $\rho_0 = l(\rho_{t+1})$.

Then, we can determine

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

C.2.2 Forecasting output

We now turn to forecasting output, determined by

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

so that a "nowcast" of output is

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

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

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

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

C.3 Matching the SPF: motivation and mechanics

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

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

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

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

C.3.1 The strategy for extracting states

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

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

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

With the date t extracted states $\widehat{\varsigma}_t$ and $\widehat{\rho}_t$ and the predetermined state $\widehat{\mu}_t$ in hand, we

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

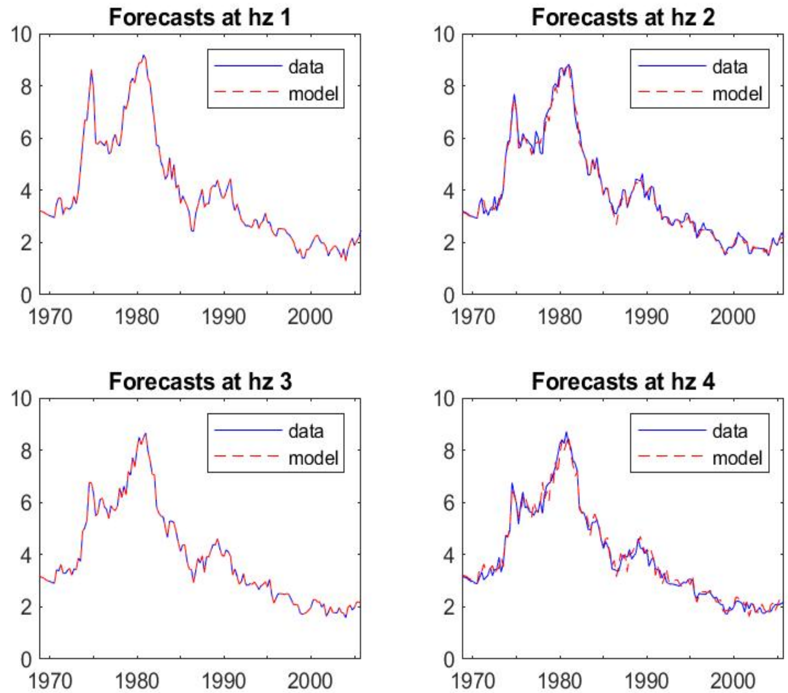

Figure 11: Model-implied and SPF forecasts of inflation

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

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

C.3.2 Application and fitting performance

As discussed in section 6.1 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 11 shows our match is nearly perfect given that we choose two state variables ς and ρ each period to match two data points SPF1Q and SPF3Q. Using the extracted states, we can also compute model-implied inflation forecasts at horizons 2 and 4, and compare them with SPF two-quarter-ahead and four-quarter-ahead forecasts. The comparison is shown in the right panels of Figure 11. It is notable that our model-implied forecasts lie almost entirely on top of the SPF data for both forecasting horizons, which are not explicitly targeted. We view this figure as evidence in support of our state extraction approach.

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

The results in the main text are based on using a decision rule $\widehat{\mu}' = \mu'^*(0, \widehat{\rho}, \widehat{\mu})$ rather than $\widehat{\mu}' = \mu'^*(\widehat{\varsigma}, \widehat{\rho}, \widehat{\mu})$. That is, we do not allow the extracted shock to influence the dynamics of the pseudo state variable. Figure 12 displays the results when we alternatively allow this influence. The main messages from the text are maintained.

Figure 12: Inflation history and model-implied policies

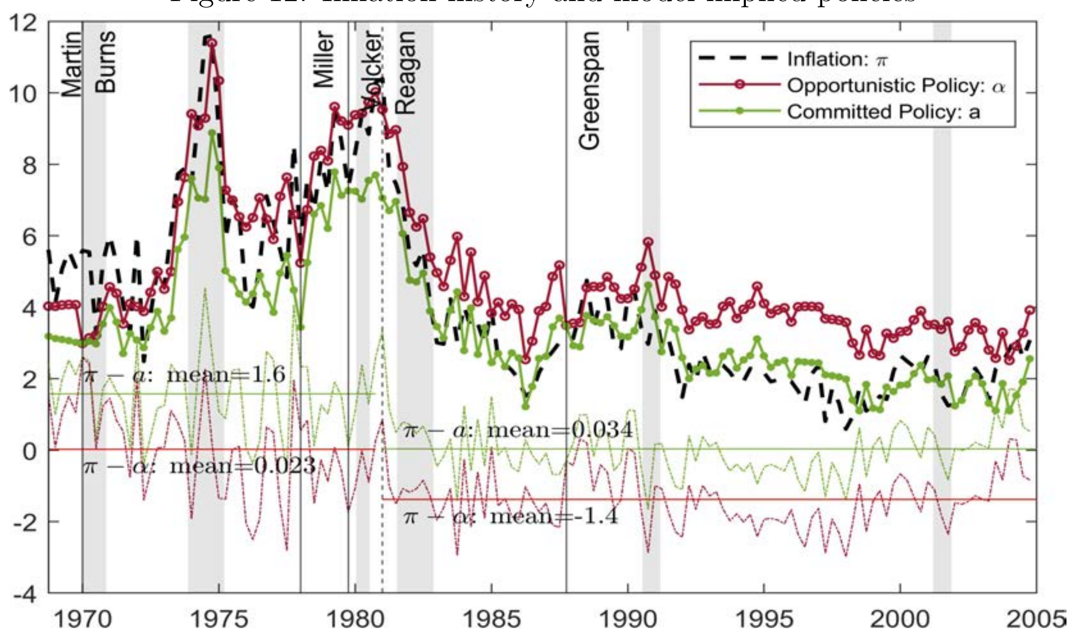

Note: Model-implied policies are based on extracted states produced using $\widehat{\mu}' = \mu'^*(\widehat{\varsigma}, \widehat{\rho}, \widehat{\mu})$ recursively.

C.5 Shock Comparisons

We have explored four indicators of the cost-push shock. First, there is a food and and energy shock constructed along the lines of Watson (2014). Second, there is the SPF spread. Third, there is the extracted shock series from the main text. Fourth, there is the extracted shock using the procedure that we just discussed. Figure 13 displays these alternative series. Note first that all measures rise dramatically during the famous "oil price shock" of late 1973 and early 1974 and also during the late 1970s interval that preceded Volcker's appointment.

Note next that the extracted shocks and the SPF spread are more persistent during the earlier episode. Contemporary sources, such as the January 1975 Economic Report of the President prepared by Alan Greenspan and his CEA colleagues, point to other price shocks in addition to oil during the preceding year. Econometric studies such as those of R.J. Gorden (2013) and Watson (2014) estimate price shocks, including those from price decontrols in the 1970s, of more lasting form. So, on this basis, we are led to prefer extracted shocks as a parsimonious approach. Note further that the extracted shocks depart from each other toward the end of the period, which is the motivation for us to adopt the extraction strategy employed in the main text rather than that discussed in the prior section. Our extraction procedure is a straightforward and transparent way to induce the extracted price shocks to be mean-reverting, but does not explicitly impose the requirement that extracted price shocks are stationary. More sophisticated methods, applicable to hidden Markov models such as ours, would impose that requirement.

Figure 13: Various indicators for cost-push shocks

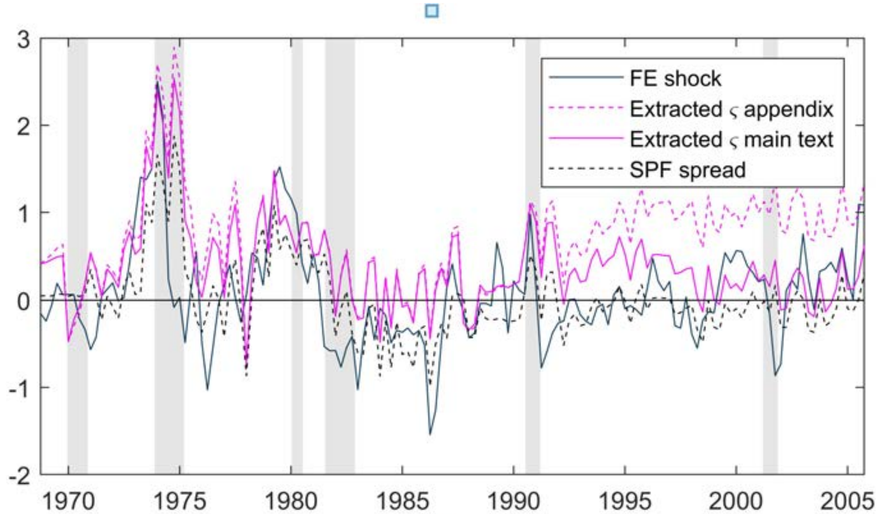

Note: Comparing $\widehat{\varsigma}$ extracted using $\widehat{\mu}' = \mu'^{*}(\widehat{\varsigma}, \widehat{\rho}, \widehat{\mu})$ to its counterpart in main text using $\widehat{\mu}' = \mu'^{*}(0, \widehat{\rho}, \widehat{\mu})$, the two series behave similarly before 1990, but the former is higher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