

Viana de Carvalho, Carlos; Eusepi, Stefano; Mönch, Emanuel; Preston, Bruce

Working Paper

Anchored inflation expectations

Texto para discussão, No. 689

Provided in Cooperation with:

Departamento de Economia, Pontifícia Universidade Católica do Rio de Janeiro

Suggested Citation: Viana de Carvalho, Carlos; Eusepi, Stefano; Mönch, Emanuel; Preston, Bruce (2021) : Anchored inflation expectations, Texto para discussão, No. 689, Pontifícia Universidade Católica do Rio de Janeiro (PUC-Rio), Departamento de Economia, Rio de Janeiro

This Version is available at:

<https://hdl.handle.net/10419/249737>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

TEXTO PARA DISCUSSÃO

No. 689

Anchored Inflation Expectations

Carlos Carvalho
Stefano Eusepi
Emanuel Moench
Bruce Preston



ANCHORED INFLATION EXPECTATIONS*

Carlos Carvalho[†]
PUC-Rio

Stefano Eusepi[‡]
University of Texas at Austin

Emanuel Moench[§]
Deutsche Bundesbank, Goethe University Frankfurt and CEPR

Bruce Preston[¶]
The University of Melbourne

ABSTRACT

We develop a theory of low-frequency movements in inflation expectations, and use it to interpret joint dynamics of inflation and inflation expectations for the United States and other countries over the post-war period. In our theory long-run inflation expectations are endogenous. They are driven by short-run inflation surprises, in a way that depends on recent forecasting performance and monetary policy. This distinguishes our theory from common explanations of low-frequency properties of inflation. The model, estimated using only inflation and short-term forecasts from professional surveys, accurately predicts observed measures of long-term inflation expectations and identifies episodes of unanchored expectations.

Keywords: Anchored expectations, inflation expectations, survey data

JEL Codes: E32, D83, D84

*April 5, 2021. First version August 9, 2015. The authors thank participants in numerous seminars and conferences for comments, especially at the NBER Summer Institute 2015 Behavioral Macroeconomics and 2016 Impulse and Propagation Mechanisms meetings. We particularly thank Sushant Acharya, Klaus Adam, Olivier Coibion, Pierre-Olivier Gourinchas, Greg Kaplan, Mariano Kulish, Albert Marcet, Giorgio Primiceri, Ricardo Reis, Andrea Tambalotti, Jesus Fernandez-Villaverde, Michael Woodford and our discussants Cars Hommes, Elmar Mertens, Christie Smith and Eric Swanson for useful exchange of ideas. The usual caveat applies. The views expressed in the paper are those of the authors and are not necessarily reflective of views at the Deutsche Bundesbank or the Eurosystem. Preston acknowledges research support from the Australian Research Council, under the grant FT130101599.

[†]PUC-Rio. E-mail: cvianac@econ.puc-rio.br.

[‡]University of Texas at Austin. E-mail: stefano.eusepi@austin.utexas.edu.

[§]Deutsche Bundesbank. E-mail: emanuel.moench@bundesbank.de.

[¶]The University of Melbourne. E-mail: bruce.preston@unimelb.edu.au

Long-run inflation expectations do vary over time. That is, they are not perfectly anchored in real economies; moreover, the extent to which they are anchored can change, depending on economic developments and (most important) the current and past conduct of monetary policy. In this context, I use the term “anchored” to mean relatively insensitive to incoming data. So, for example, if the public experiences a spell of inflation higher than their long-run expectation, but their long-run expectation of inflation changes little as a result, then inflation expectations are well anchored. If, on the other hand, the public reacts to a short period of higher-than-expected inflation by marking up their long-run expectation considerably, then expectations are poorly anchored — Bernanke (2007)

1 INTRODUCTION

How best to provide a nominal anchor is a central question in monetary economics. Inflation targeting regimes emphasize a credible commitment to a numerical objective for inflation in the medium to long term. By anchoring long-term expectations, central banks are free to pursue activist short-run stabilization policy. Yet despite the obvious importance of long-term inflation expectations, models for policy analysis provide little guidance on how market participants form these expectations, or how policy ensures expectations to be consistent with central bank objectives. Indeed, most models simply assume long-term expectations are consistent with the policy strategy of the central bank.¹

A complete description of inflation targeting as a framework for monetary policy must include a theory of how and when inflation targeting successfully anchors long-run inflation expectations. Only with a theory can we answer questions like: Are inflation expectations anchored? Will chronic undershooting of inflation targets lead to downward drift and unanchoring of long-term expectations? How can we reconcile large negative output gaps and stable inflation over the past decade, with positive output gaps and high inflation in the 1970s?

We study a New Keynesian model in which poorly anchored long-term inflation expectations are the source of low-frequency movements in inflation. The degree to which expectations are anchored depends on the *endogenous* link between long-term expectations and short-term forecast errors. The strength of this connection depends on historical forecasting performance. In our model the evolution of long-term beliefs depends on monetary policy

¹This is true in a wide class of New Keynesian models, including those with indeterminacy of equilibrium and sunspots [Clarida, Gali, and Gertler (2000), Albanesi, Chari, and Christiano (2003), Lubik and Schorfheide (2004)], regime switching [Davig and Leeper (2006), Bianchi and Melosi (2017)], and exogenous time-varying inflation targets [Cogley and Sbordone (2008)].

and structural shocks. The model, estimated using only inflation and short-term forecasts from professional surveys, accurately predicts observed measures of long-term inflation expectations. This is the first paper to directly test this mechanism, central to many models of imperfect information and learning.

The model has a central bank with a fixed inflation target. The policy regime is fixed for all time. Subject to nominal rigidities, monopolistically competitive firms set prices optimally as a function of expectations about the future path of marginal costs and inflation. As in the standard New Keynesian model, marginal costs are determined by monetary policy and aggregate demand. The key additional state variable in our model is firms' beliefs about long-run inflation. Firms are uncertain about the long-run mean of inflation: this could be because there is no publicly announced inflation target, as was the case for the United States over much of our sample, or that the central bank lacks the commitment, the tools or right incentives to implement the inflation target.

Price setters behave like econometricians. Following Marcet and Nicolini (2003) firms are unsure about the correct model to forecast inflation and must choose between two estimators for the long-run mean inflation rate: either i) a decreasing-gain algorithm; or ii) a constant-gain algorithm. The first estimator is ordinary least squares which is consistent with estimating a time-invariant inflation mean. The gain is the inverse of the sample size, so that accumulating evidence of a stationary mean leads to declining sensitivity to new information. The second estimator implies a constant and relatively high sensitivity to new information. Geometrically discounting older data permits tracking different forms of structural change such as a sudden shift in the inflation mean.

Firms select their estimator using a model selection criterion. Similar in spirit to Brown, Durbin, and Evans' (1975) CUSUM specification test and Cho and Kasa's (2015) model validation procedure, the criterion has the property that large persistent forecast errors lead agents to doubt a constant mean of inflation. When this weighted average is larger than a specific threshold, firms suspect the inflation target is shifting and use a constant gain forecasting model.

Collectively these assumptions give formal expression to the ideas of Bernanke (2007). How sensitive expectations are to forecast errors depends on historical forecasting performance. For a given forecast error, the size of long-run mean revisions will depend on which forecasting model is being used, and, for the decreasing-gain algorithm, how long that estimator has been used. Importantly, the framework permits a formal definition of anchored expectations. Following Preston (2006) and Eusepi and Preston (2010), we say long-term expectations are anchored when beliefs are consistent with the policy regime. That is, long-term expectations are determined under the correct assumption of a constant inflation mean

and firms update their estimate using the decreasing gain algorithm. Conversely, we say expectations are un-anchored when agents doubt a constant long-run inflation mean and switch to a constant gain algorithm. Expectations display high sensitivity to forecast errors.

The paper makes empirical, theoretical and methodological contributions. Empirically, we assess the link that the model forges between short-term inflation surprises and long-run beliefs. We use observed data on inflation and short-term inflation expectations to measure forecast errors. Conditional on observing forecast errors, the model makes sharp testable predictions about the evolution of agents' estimated inflation mean, the main driver of long-run inflation expectations.

Using Bayesian methods, we compare model-implied predictions of long-term inflation expectations with equivalent measures from survey data (which are not used in estimation). Over a long sample of US data, 1955 to 2015, the model explains well both short- and long-run historical inflation expectations from the Survey of Professional Forecasters and the Michigan Survey of US households. The model demonstrates expectations were poorly anchored before the late 1990s, with long-run expectations displaying high sensitivity to new information.² Formal predictive likelihood comparisons reveal the model provides a better fit of long-term expectations when compared to an otherwise equivalent model with a single constant gain; a model with rational expectations; and a model with an exogenous time-varying inflation target. Using the posterior parameter distribution from US data, we provide a further evaluation of the model, using it to explain professional survey expectations for a range of OECD countries. The model also fits these data and detects episodes of un-anchored expectations.

Large and persistent forecast errors lead firms to doubt a constant inflation target and adopt a constant gain forecasting technology. A firm's optimal price becomes more sensitive to new information because revisions to long-term inflation expectations are more sensitive to short-run forecast errors. In aggregate, this forges strong feedback between expected and realized inflation: beliefs are partially self-fulfilling. The model is self-referential in the language of Marcat and Sargent (1989) — beliefs determine inflation, which in turn determines beliefs. This propagation of forecast errors delivers trend inflation.

Absent self-referential behavior the model can explain neither the extent and pace of the rise of the Great Inflation, nor the gradual decline and ultimate stabilization of inflation expectations over the Great Moderation.³ That long-term expectations don't shift during

²Having a structural model provides a richer interpretation of what constitutes anchored expectations than approaches based on pass-through regressions of either macroeconomic news or movements in short-term expectations to long-term expectations as developed by Gurkaynak, Levin, and Swanson (2010) and Beechey, Johannsen, and Levin (2011). Low sensitivity of long-run beliefs need not imply consistency with a central bank's inflation target.

³For example, counterfactuals demonstrate that an identical model with the same beliefs but without

the Great Recession (despite shocks being estimated to have comparable magnitude to the 1970s) reflects the anchoring of expectations. Time variation in the the degree of expectations' anchoring explains the coexistence of large negative output gaps and missing deflation in the Great Recession, with the positive output gaps and high inflation of the Great Inflation. The standard New Keynesian Phillips curve is consistent with data once the nature of expectations formation and the degree of anchoring are properly accounted for.⁴

Theoretically, we show that the belief structure places lower bounds on rationality and that monetary policy plays a central role in shaping forecast errors and self-referentiality. We show firms will learn the inflation target given enough data. This distinguishes our paper from other studies on expectations anchoring which imply agents can never hold expectations consistent with monetary strategy. Following Honkapohja and Evans (1993) and Marcet and Nicolini (2003), we show that parameters governing firms' forecasting technology are optimal in the sense that given other firms' chosen parameters, it is optimal for any one firm to also adopt them. However, optimality is sustained by a high degree of self-referentiality. Stronger monetary policy responses reduce self-referentiality, leading firms to choose a forecasting model with lower sensitivity to new information (lower gain) given other firms' behavior.

Methodologically, we advance understanding of how to estimate models with learning. We identify and resolve a conceptual problem in extant work, by introducing techniques that are novel to the macroeconomics literature. Our model is an example of a state-dependent learning rule, which includes as a special case recursive least-squares estimation of statistical models, such as vector auto-regressions, common to the adaptive learning literature. Building on Primiceri (2005) and Sargent, Williams, and Zha (2006), the learning literature estimates this class of model using standard linear methods. This requires learning rules to depend only on observable data. But in most empirical applications, beliefs depend on unobserved states (from the perspective of the econometrician outside the model), which requires a non-linear estimation procedure.⁵ We make use of the marginalized particle filter of Schön, Gustafsson, and Nordlund (2005), to exploit the fact that the model is conditionally linear. This permits a highly efficient algorithm without requiring large numbers of particles. The approach can be applied more widely to models with state-dependent learning rules.

Related literature. Our paper is most closely related to Marcet and Nicolini (2003) and Milani (2014). We adapt and implement Marcet and Nicolini's (2003) learning mechanism in a New Keynesian model. While they study the determinants of recurring hyperinflation episodes, we study low-frequency properties of inflation. Milani (2014) estimates a New

self-referential feedback cannot generate the Great Inflation.

⁴See Del Negro, Giannoni, and Schorfheide (2015) for a discussion of relevant literature.

⁵Milani (2007, 2014) and Slobodyan and Wouters (2012) estimate such models within a linear framework. This requires specific assumptions on agents' information sets, which are discussed in section 3.

Keynesian model with closely related beliefs on US data. He shows that agents switch often between constant- and decreasing-gain estimators over the postwar sample. In these papers learning dynamics generate an endogenous trend in inflation with time-varying volatility. Our work also builds on Cho and Kasa (2015) by appealing to the idea that agents use econometric tests based on past forecast errors to select from a set of forecasting models.

Other approaches also generate an endogenous inflation trend. Using a New Keynesian Phillips curve and boundedly rational beliefs, Lansing (2009) demonstrates the existence of a self-confirming equilibrium in which agents' Kalman gain is uniquely determined. Introducing learning in this model induces time variation in the estimated gain, which is argued to generate time-varying volatility of the kind observed in US inflation data. Branch and Evans (2007) and Cornea-Madeira, Hommes, and Massaro (2019) instead emphasize belief heterogeneity and endogenous predictor selection based on evolutionary fitness. Empirical estimates show that time variation in the number of firms adopting misspecified random walk inflation beliefs can explain low-frequency properties of inflation. Primiceri (2005) and Sargent, Williams, and Zha (2006) emphasize central bank learning based on a mis-specified Phillips curve as the source of inflation trend. Estimated on US data, these models explain the rise and fall of inflation during the post war period.

We share with these papers the idea that imperfect information and the propagation of forecast errors over time is empirically and quantitatively relevant to low-frequency movements in US inflation data. What sets our study apart is the emphasis on long-term inflation expectations: to explain low-frequency properties of inflation and to provide a coherent policy-relevant definition of anchored expectations. Long-term expectations play no role in the above papers — decisions are based solely on short-term forecasts. Critically, none of these papers use survey measures of inflation expectations to empirically test the relationship between forecast errors and long-run expectations. We therefore provide the first paper to directly identify the learning mechanisms proposed in the literature. Lastly, none of these papers make predictions about long-term inflation expectations, which is one of our key results.

Finally, the endogeneity of the inflation trend distinguishes our results from Cogley and Sbordone (2008) and Erceg and Levin (2003) which assume low-frequency properties of inflation are given exogenously.⁶ The model provides an alternative interpretation of the driving forces behind inflation dynamics, which have distinct implications for policy. Indeed, the results suggest the Great Inflation occurred despite the Federal Reserve's commitment to

⁶Formally Erceg and Levin (2003) assume agents must learn about exogenous transitory and permanent shocks to the inflation target. But this is isomorphic to a model with full information for a certain process for the exogenous inflation target — see Hamilton (1994).

a constant inflation target. For the same reason, our work differs from a large literature which uses reduced-form statistical models to explain inflation dynamics. For example, building on Stock and Watson (2007), papers such as Kozicki and Tinsley (2012) and Chan, Clark, and Koop (2015) use expectations data to better fit low-frequency properties of inflation.

A road map. Section 2 develops a New Keynesian model with imperfect information. Section 3 provides estimates. Section 4 discusses predictions and the model’s ability to explain observed long-term expectations. Section 5 provides counterfactuals to understand how beliefs generate an endogenous inflation trend. Section 6 demonstrates beliefs satisfy lower bounds on rationality: firms can learn the inflation mean, and given other firms have the proposed beliefs, it is optimal to have those beliefs. However, under a more aggressive monetary policy it would be optimal to adopt a forecasting model with a lower gain. Section 7 uses formal model comparisons to demonstrate our model provides a superior account of long-term survey expectations data than do a range of alternative models. Section 8 provides an example of how to interpret the model selection criterion as the outcome of a formal sequential statistical test, as well as evidence that discrete shifts in forecasting model are required to explain observed data. Section 9 demonstrates the model can explain long-term forecast data from US household and professional forecasters from other countries. Section 10 concludes.

2 A MODEL WITH ENDOGENOUS INFLATION DRIFT

This section presents a New Keynesian model. The optimal pricing decisions of firms relate beliefs about the long-run inflation target to current inflation outcomes, with monetary policy determining the strength of this relationship. We show how these beliefs are updated over time, derive the state-space representation of the model, which includes an equation for the endogenous inflation trend, and relate it to other prominent models of inflation.

Theory of Price Setting. A continuum of monopolistically competitive firms face a price-setting problem subject to quadratic adjustment costs as in Rotemberg (1982). Given subjective beliefs \hat{E}_t^f , each firm $f \in [0, 1]$ maximizes the expected present discounted value of profits

$$\hat{E}_t^f \sum_{T=t}^{\infty} \Lambda_{t,T} \Gamma_T(f)$$

by choice of $P_t(f)$ subject to the demand and profit functions

$$Y_t(f) = \left(\frac{P_t(f)}{P_t} \right)^{-\psi_t} Y_t$$

$$\Gamma_t(f) = Y_t(f) \left(\frac{P_t(f)}{P_t} - S_t \right) - \frac{\Phi}{2} \left(\frac{P_t(f)}{P_{t-1}(f)} - e^{\gamma\pi_{t-1}} \right)^2$$

for all $T \geq t$, where P_t and Y_t give the aggregate level of prices and output in period t , and S_t a real marginal cost function. The exogenous time-varying elasticity of demand across differentiated goods satisfies $\psi_t > 1$, with mean $\bar{\psi}$. The quadratic costs of price adjustment are determined by price movements relative to an inflation index, $\gamma\pi_{t-1}$, a linear function of aggregate inflation, $\pi_{t-1} = \ln(P_{t-1}/P_{t-2})$.⁷ The constant $\Phi > 1$ scales the size of adjustment costs. The parameter $0 \leq \gamma \leq 1$ measures the degree of price indexation. When setting prices in period t , firms value future streams of income using the stochastic discount factor $\Lambda_{t,T}$ which in steady state takes the value β^{T-t} , for $0 < \beta < 1$ and all $T > t$.

Deriving the first-order condition and taking a log-linear approximation around a zero inflation steady state provides the optimal price

$$p_t^*(f) = \alpha p_{t-1}^*(f) - \alpha(\pi_t - \gamma\pi_{t-1})$$

$$+ \alpha \hat{E}_t^f \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} (\xi_p s_T + \beta(1-\alpha)(\pi_{T+1} - \gamma\pi_T)) + \alpha\mu_t$$

where

$$p_t^*(f) = \ln \left(\frac{P_t(f)}{P_t} \right); \quad s_t = \ln \left(\frac{S_t}{\bar{S}} \right); \quad \hat{\psi}_t = \ln \left(\frac{\psi_t}{\bar{\psi}} \right); \quad \mu_t = -\xi_p \frac{\hat{\psi}_t}{\bar{\psi} - 1}$$

and $\xi_p = \alpha^{-1}(1-\alpha)(1-\alpha\beta)$ with $0 < \alpha < 1$, the model's stable eigenvalue.⁸ The markup shock μ_t is a normally distributed and serially independent process. Optimal price setting requires firms to project the future evolution of marginal costs and aggregate inflation, adjusted for indexation.⁹ Aggregating over the continuum in a symmetric equilibrium, where

⁷To the first order this is equivalent to a model of Calvo price setting in which firms, when not optimally re-setting prices, adjust prices according to the index $\gamma\pi_{t-1}$.

⁸In a model of Calvo price setting α corresponds to the probability a firm cannot adjust their price in any given period. While it is well understood the predictions of Rotemberg and Calvo pricing are different at non-zero average rates of inflation, Cogley and Sbordone (2008) show in a closely related model with Calvo pricing it is the assumption of an exogenous inflation trend, rather than the point of approximation, that is most relevant to explaining US inflation data.

⁹See Preston (2005) and Eusepi and Preston (2018b) for discussions of optimal price setting under arbi-

firms set identical prices, $P_t(f) = P_t$, and hold the same subjective beliefs, $\hat{E}_t^f = \hat{E}_t$, provides the aggregate supply curve

$$\pi_t - \gamma\pi_{t-1} = \mu_t + \hat{E}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} [\xi_p s_T + (1 - \alpha)\beta(\pi_{T+1} - \gamma\pi_T)]. \quad (1)$$

Monetary Policy and Demand. To close the model we assume the central bank implements monetary policy using the targeting rule

$$\pi_t - \gamma\pi_{t-1} + \lambda_x x_t = \varphi_t, \quad (2)$$

where x_t is the output gap; λ_x a policy parameter indexing the relative weight given to stabilizing the output gap; and

$$\varphi_t = \rho\varphi_{t-1} + \epsilon_t \quad (3)$$

an exogenous monetary policy shock. From Giannoni and Woodford (2002) this target criterion is the optimal policy under discretion given our assumption that the central bank has a zero inflation target. As frequently assumed in the New Keynesian literature the output gap is the instrument of policy.¹⁰ We interpret monetary policy shocks as being due to policy mistakes or to mis-measurement of underlying inflation and the output gap.¹¹ Marginal costs are assumed to be proportional to the output gap so that

$$s_t = \phi x_t, \quad (4)$$

where the constant of proportionality is determined by preferences and technology. Without loss of generality, we assume $\phi = 1$. For a detailed account of this relationship and structural interpretation see Woodford (2003).

Rational Expectations Equilibrium. Under rational expectations inflation satisfies the first-order difference equation

$$\pi_t - \gamma\pi_{t-1} = \xi_p s_T + \beta\mathbb{E}_t(\pi_{t+1} - \gamma\pi_t) + \mu_t, \quad (5)$$

where \mathbb{E}_t denotes mathematical expectations. Combining with (2) and (4) and solving gives

trary beliefs.

¹⁰Because we don't exploit data on interest rates and output in the empirical work, we don't model the transmission mechanism of monetary policy. But standard preference and technology assumptions give the stated policy.

¹¹See, for example, Orphanides (2001) and Primiceri (2005).

equilibrium inflation

$$\pi_t = \gamma\pi_{t-1} + \rho\bar{\omega}\varphi_{t-1} + \eta_t \quad (6)$$

where

$$\bar{\omega} = (1 + (1 - \beta\rho)\lambda_x)^{-1} \text{ and } \eta_t = \bar{\omega}\epsilon_t + (1 + \xi_p\lambda_x^{-1})^{-1}\mu_t = \tilde{\epsilon}_t + \tilde{\mu}_t.$$

See the appendix for details. Inflation is a stationary process with zero mean, consistent with the central bank objective.

Imperfect Knowledge. The optimal price decision rule implies firms must forecast future marginal costs and inflation. To simplify the analysis, we assume firms learn only about the long-run mean of inflation. They are endowed with the correct forecasting model of marginal costs so that forecasts for s_{T+1} for $T \geq t$ are computed as

$$\hat{E}_t^f s_{T+1} = \frac{1}{\lambda_x} \hat{E}_t^f [\varphi_{T+1} - (\pi_{T+1} - \bar{\pi}_t) + \gamma(\pi_T - \bar{\pi}_t)] \quad (7)$$

$$\hat{E}_t^f \varphi_{T+1} = \rho^{T-t} \varphi_t. \quad (8)$$

The term $\bar{\pi}_t$ reflects the firm's belief that the inflation target is non-zero. Inflation is forecast using the statistical model

$$\pi_t = \gamma\pi_{t-1} + (1 - \gamma)\bar{\pi}_t + \rho\bar{\omega}\varphi_{t-1} + f_t, \quad (9)$$

where f_t is the forecast error, given an estimate of the inflation target $\bar{\pi}_t$. This model nests rational expectations — compare to (6). It assumes firms have perfect information about the *short-run* dynamics of the economy, governed by monetary policy shocks and lagged inflation. We study the anticipated utility solution of the model — see Kreps (1998), Sargent (1999) and Eusepi and Preston (2018b) — so that beliefs about the inflation target satisfy

$$\hat{E}_{t-1} \bar{\pi}_T = \bar{\pi}_t. \quad (10)$$

From (9), long-run conditional inflation expectations must satisfy

$$\lim_{T \rightarrow \infty} \hat{E}_t \pi_T = \bar{\pi}_t$$

which Kozicki and Tinsley (2001) call a shifting end-point model. This expression makes clear that beliefs about the inflation target are equivalent to long-term inflation expectations. Movements in beliefs are then informative about the anchoring of expectations.

Evaluating subjective expectations using (7)-(10) in the aggregate supply curve (1) provides the true data-generating process

$$\pi_t = \gamma\pi_{t-1} + (1 - \gamma)\Gamma\bar{\pi}_t + \rho\bar{\omega}\varphi_{t-1} + \eta_t, \quad (11)$$

where

$$\Gamma = \frac{1}{1 + \xi_p\lambda_x^{-1}} \frac{(1 - \alpha)\beta}{1 - \alpha\beta}.$$

The subjective and objective probability models, (9) and (11), differ when the coefficient Γ differs from unity. This coefficient measures the degree of self-referentiality. Firms raise their optimal price if they believe the inflation target has risen, so that aggregate inflation also rises. The stance of monetary policy regulates the strength of this connection. When $\lambda_x \rightarrow 0$ the targeting criterion is equivalent to strict inflation targeting which eliminates self-referentiality. When $\lambda_x \rightarrow \infty$ the targeting criterion is equivalent to output gap targeting and, for standard parameters values, implies values of Γ near unity.¹²

Estimating the inflation mean. We assume firms behave like econometricians who are unsure about the correct model to forecast inflation. To confront this uncertainty, they choose between two estimation algorithms: a decreasing-gain algorithm that is appropriate when the inflation target is constant; and a constant-gain algorithm that, by discounting past observations, is appropriate when either there is a sudden but infrequent shift in the mean or the inflation mean drifts over time. The two estimators lead to the following recursive updating

$$\bar{\pi}_t = \bar{\pi}_{t-1} + \mathbf{k}_t^{-1} \times f_{t-1}, \quad (12)$$

linking revisions in the current estimate of the inflation mean to the last one-step-ahead forecast error, $f_{t-1} = \pi_{t-1} - \hat{E}_{t-2}\pi_{t-1}$.¹³ Let \mathbf{k}_t^{-1} denote the gain. Recursive least square learning implies

$$\mathbf{k}_t = \mathbf{k}_{t-1} + 1.$$

Because the gain is the inverse of the sample size, accumulating evidence of a stationary mean leads to declining sensitivity to new information. Discounting past information leads to a constant gain

$$\mathbf{k}_t^{-1} = \bar{g} > 0$$

¹²See the appendix for derivations. For example, taking uncontroversial values $\beta = 0.99$ and α in the range $0.5 - 0.85$, implies Γ falls in the range $0.94 - 0.98$.

¹³Here $\hat{E}_{t-2}\pi_{t-1} = \gamma\pi_{t-2} + (1 - \gamma)\Gamma\bar{\pi}_{t-1} + \rho\bar{\omega}\varphi_{t-2}$. A complicated simultaneity is resolved by assuming the estimate, $\bar{\pi}_t$, depends on the previous-period's forecast error so that $\bar{\pi}_{t-1}$ is determined by information available at $t - 2$. This is standard in the learning literature.

and therefore a constant and relatively high sensitivity to new information.

The final component of the model specifies the conditions under which firms adopt each estimator. Being alert to possible changes in the inflation mean, firms have explicit concern for model specification. In each period their model is tested: if the hypothesis of a time-invariant mean is rejected agents switch to a constant-gain algorithm. In our baseline specification this process of model selection and validation is not explicitly modeled. We assume that firms are able to detect model miss-specification when it is sufficiently large.

To capture this in a parsimonious way, the evolution of the learning gain is directly tied to the relative distance of forecasts under the subjective ($\hat{E}_{t-1}\pi_t$) and objective ($\mathbb{E}_{t-1}\pi_t$) probability distributions:

$$\mathbf{k}_{t+1} = I_{(\theta_t \leq \bar{\theta})} \times (\mathbf{k}_t + 1) + \left(1 - I_{(\theta_t \leq \bar{\theta})}\right) \times \bar{g}^{-1}, \quad (13)$$

where $I_{(\cdot)}$ is an indicator function which depends on the magnitude

$$\theta_t = \left| \hat{E}_{t-1}\pi_t - \mathbb{E}_{t-1}\pi_t \right| / \sigma_\eta, \quad (14)$$

relative to a threshold value $\bar{\theta}$. While the criterion (14) involves model-consistent expectations (which they are assumed not to know), it is a function of *observable forecast errors*:

$$\begin{aligned} \left| \hat{E}_{t-1}\pi_t - \mathbb{E}_{t-1}\pi_t \right| &= |(1 - \gamma)(\Gamma - 1)\bar{\pi}_t| \\ &= \left| (1 - \gamma)(\Gamma - 1) \left[\pi_0 + \sum_{\tau=0}^{t-1} \mathbf{k}_\tau^{-1} f_\tau \right] \right| \end{aligned} \quad (15)$$

given some initial conditions $\bar{\pi}_0$, f_0 and \mathbf{k}_0 .¹⁴ It is equivalent to a weighted average of the entire history of past forecast errors with weights given by the gain that applied in the period of each forecast error. The distance tends to be large when forecast errors happen to be of the same sign for several periods. This signals subjective beliefs about the inflation target are far from the true value of zero and therefore a reasonable basis for specification tests. These calculations are not meant to suggest that firms would naturally arrive at this criterion as the basis for model selection. Rather the point is that a criterion based on past forecast errors is natural. Indeed, the dependency on the cumulative history of past forecast errors has much in common with the classic CUSUM test of Brown, Durbin, and Evans (1975).

¹⁴It is worth noting that the distance criterion takes the value of zero when either $\gamma = 1$ or $\Gamma = 1$. With complete indexation inflation dynamics are independent of long-run expectations. In a self confirming equilibrium subjective and objectives probability models are identical, though different to the full information rational expectations equilibrium.

Expectations anchoring. The model permits a formal definition of anchored expectations which is directly tied to agents’ beliefs about the inflation mean. Following Eusepi (2005), Preston (2006) and Eusepi and Preston (2010), we define expectations as anchored when beliefs are consistent with the policy regime in place: a fixed inflation target. Expectations are then anchored when firms fail to reject a time-invariant mean: long-term expectations display decreasing sensitivity to forecast errors and, if they remain anchored, eventually converge to the correct mean. However, large and persistence forecast errors can lead firms to abandon the hypothesis of a constant mean. Long-run beliefs become un-anchored and tightly connected to short-term forecast errors.

Bounds to rationality. Is the estimator optimal given the information available to agents? Firms are boundedly rational because they do not fully understand the economic environment. This friction is central to modeling anchored inflation expectations. By assumption, agents cannot write down a law of motion for the mean for inflation and then adopt an optimal estimator. Firms do not know *all possible future policy regimes* and therefore can not assign probabilities to these regimes. We believe it is implausible to suppose that firms, during the 1960s, could accurately assign probabilities to the likelihood of the Great Inflation, or, standing at the peak of the Great Inflation, to the likelihood of the Federal Reserve adopting inflation targeting some three decades later. However, following Marcet and Nicolini (2003), we require beliefs to satisfy lower bounds to rationality. Specifically, we impose that: i. agents eventually learn the truth—that is beliefs will not be anchored asymptotically to the wrong mean;¹⁵ ii. belief parameters are such that agents will stick to the estimator given the option to re-adjust; and iii. these ‘optimal’ belief parameters are not invariant to monetary policy. Details about these bounds are discussed in section 6.

Model validation: discussion. In our baseline specification we do not explicitly model agents’ recursive testing of their statistical model. This is done for two reasons. First, simple validation rules are likely to violate our imposed rationality bounds. Second, a more sophisticated testing procedure would make estimation unfeasible.

A simple criterion to test for model mis-specification is to evaluate current and recent forecast errors against a threshold. For example, Marcet and Nicolini (2003) and Milani (2014) present criteria based on recent forecast errors. However, recent forecast errors need not be informative about low-frequency developments in a volatile process such as inflation.

¹⁵Cho and Kasa (2015) provide a detailed discussion of why statistical inference in this class of models is complicated. Even though the decreasing-gain algorithm spans the truth, during the learning process, both estimators are misspecified. Indeed, Evans, Honkapohja, Sargent, and Williams (2012) and Cho and Kasa (2017) provide examples of models with our beliefs in which Bayesian model averaging leads decision makers to adopt the incorrect model with probability 1. Our approach, therefore, follows Cho and Kasa’s (2015) model validation procedure, which selects the correct model with probability 1.

To be concrete, forecast errors in our model can be written

$$f_t = \eta_t + (1 - \gamma) (\Gamma - 1) \bar{\pi}_t. \quad (16)$$

The first component is short-term noise, η_t , from marginal cost and mark-up innovations. This is the dominant source of volatility in the forecast error. The second component arises from model mis-specification. The magnitude of this term depends both on structural factors affecting the degree of self-referentiality through Γ and on the size of subjective estimate $\bar{\pi}_t$ relative to its true time-invariant value of zero. The contribution of this term to the volatility of the short-term forecast error is negligible because of a high degree of self-referentiality (empirically Γ takes values near unity) and the fact that the inflation trend is slow moving (it has small variance). For this reason, criteria based on current and recent forecast errors lack power to detect shifts in the long-run mean of inflation. They would lead to frequent switches in forecasting model, even when the the learning gain is arbitrarily small and model mis-specification small. Firms' would fail to converge to the true model asymptotically, violating the lower bound to rationality discussed above.

The principal advantage of our baseline specification is it renders estimation of a non-linear model tractable. It introduces no additional state variables so that the model is conditionally linear given the gain. This allows us to exploit techniques from the statistics literature in the form of the marginalized particle filter. Beliefs also introduce only two parameters: the value of the constant gain, \bar{g} , and the threshold statistic, $\bar{\theta}$. However, in section 8, we present a fully specified model where agents test their current model using a recursive LM test like in Cho and Kasa (2015). While this model is calibrated, it delivers very similar results.

Model Summary. The evolution of the gain in (13) can be written in compact form

$$\mathbf{k}_{t+1} = \mathbf{f}_{\mathbf{k}}(\mathbf{k}_t, \bar{\pi}_t). \quad (17)$$

Given the sequence of exogenous processes $\{\varphi_t, \eta_t\}$, model dynamics are described by the pair of equations

$$\pi_t = \gamma\pi_{t-1} + (1 - \gamma) \Gamma \bar{\pi}_t + \rho\bar{\omega}\varphi_{t-1} + \eta_t \quad (18)$$

$$\bar{\pi}_{t+1} = [1 + \mathbf{f}_{\mathbf{k}}^{-1}(\mathbf{k}_t, \bar{\pi}_t) (1 - \gamma) (\Gamma - 1)] \bar{\pi}_t + \mathbf{f}_{\mathbf{k}}^{-1}(\mathbf{k}_t, \bar{\pi}_t) \eta_t. \quad (19)$$

The estimated mean inflation rate is a first-order auto-regressive process with time-varying persistence and volatility. The persistence depends on the stance of policy. When Γ takes

values near unity, beliefs about the inflation target are close to a unit root. As policy responds more aggressively to inflation, Γ declines and beliefs are less persistent. The volatility also depends on the estimated inflation target. When the estimate is close to the central bank’s target, the learning gain is declining over time, lowering volatility in the inflation trend.

Because the model endogenously generates low-frequency properties of inflation, the inflation trend is a function of all model disturbances, which in this simple environment are monetary policy and markup shocks. This contrasts with Cogley and Sbordone (2008) which captures low-frequency movements in inflation with an exogenous process. Similarly, Erceg and Levin’s (2003) imperfect information model assumes agents observe a noisy signal of the true inflation target, itself a function of exogenous permanent and transitory shocks. Again, trend inflation is exogenous.¹⁶

A further implication of this endogeneity is the inflation trend — equivalently, long-run expectations — displays time-varying sensitivity to inflation forecast errors. Imperfect information rational expectations models imply a constant sensitivity of long-run beliefs to current forecast errors. A central contribution of this paper is to document substantial and significant changes in this sensitivity over time. This contribution is independent of the assumptions we make on the model selection criterion. Finally, the time-varying persistence and volatility of the inflation trend provides a structural interpretation of Stock and Watson’s (2007) reduced-form inflation model. Time variation in the extent to which expectations are anchored explain changes in the persistence and volatility of inflation.

3 ESTIMATION

The data and observation equation. We estimate the model with Bayesian methods using US data on both inflation and survey measures of *short-term* inflation expectations from professional forecasters. The estimation strategy employs only these data for inference on model parameters. The use of short-term forecasts directly identifies the forecast errors that are central to the mechanism. Conditional on observing short-term forecasts, the updating algorithm in (12) and (13) implies tight predictions for the evolution of the estimated inflation mean. Having tied our hands in this fashion, model success is evaluated by a comparison of model-implied predictions of long-term inflation expectations with equivalent measures available from survey data.

For data comparability, we use the log-difference in the CPI as the measure of inflation. Four survey measures of CPI inflation forecasts are used: one- and two-quarter-ahead

¹⁶Kozicki and Tinsley (2005) and Ireland (2007) also study models in which structural shocks affect the central bank’s inflation target in an exogenously determined way. Models of exogenous drift are discussed further in Section 7.

forecasts from the Survey of Professional Forecasters (SPF); and two measures of six-month-ahead forecasts from the Livingston Survey. The first Livingston Survey measure is computed as the growth rate between the forecast of the CPI level six-months ahead and the last monthly price level available to forecasters at the time of the survey. The second measure is the growth rate between the forecast of the CPI level six-months ahead and the forecast of the current CPI level. The latter is a more accurate measure of inflation expectations but it is available only for a short sample.

The sample spans 1955Q1-2015Q4. The survey data are available over different sample sizes and at different frequencies. The SPF measures are available starting in 1981Q3, at a quarterly frequency. The Livingston survey is available only at a bi-annual frequency, but its first measure of six-months-ahead forecasts is available since the beginning of the sample, while the second starts only in 1992Q2. Reflecting the structure of these data, the model observation equation is

$$\begin{bmatrix} \pi_t \\ E_t^{SPF} \pi_{t+1} \\ E_t^{SPF} \pi_{t+2} \\ E_t^{LIV_1} \left(\frac{1}{2} \sum_{i=1}^2 \pi_{t+i} \right) \\ E_t^{LIV_2} \left(\frac{1}{2} \sum_{i=1}^2 \pi_{t+i} \right) \end{bmatrix} = \pi^* + H_t' \begin{bmatrix} \bar{\pi}_t \\ \xi_t \end{bmatrix} + R_t o_t,$$

where $\xi_t = (\eta_t, \tilde{\varphi}_t, \pi_t)'$; we estimate $\tilde{\varphi}_t = \bar{\omega} \varphi_t$ which rescales the monetary shock. The rescaled shock's process is then $\tilde{\varphi}_t = \rho \tilde{\varphi}_{t-1} + \tilde{\varepsilon}_t$. The structural parameter $\bar{\omega}$ is not estimated. The variable π^* is the mean inflation rate; and o_t is measurement error attached to both the survey data and CPI inflation. The observation matrix H_t captures the true data generating process of inflation and the model-implied firm expectations of inflation. The measurement error on inflation captures the fact the CPI measure of inflation exhibits substantial quarter-to-quarter volatility that is not incorporated in short-term forecasts and is understood to be temporary. For example, the technical appendix shows that while CPI inflation is substantially more volatile than the GDP deflator, the survey-based forecasts of these two variables are very similar. Finally, the matrices H_t and R_t are time varying because of missing observations.

Marginalized particle filter. Because of non-linearities, we cannot estimate our model using techniques based on the Kalman filter. However, our baseline model permits an efficient estimation procedure. To proceed, partition the model states into a subset of non-linear variables $(\bar{\pi}_t, \mathbf{k}_t)'$ and a subset of linear variables, ξ_t , and write the state-space representation

as

$$\begin{aligned}
 \mathbf{k}_t &= \mathbf{f}_k(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) \\
 \bar{\pi}_t &= \mathbf{f}_{\bar{\pi}}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) + A_{\bar{\pi}}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) \xi_{t-1} \\
 \xi_t &= \mathbf{f}_{\xi}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) + A_{\xi}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) \xi_{t-1} + S_{\xi} \epsilon_t,
 \end{aligned}
 \tag{20}$$

where $\epsilon_t = (\tilde{\epsilon}_t, \tilde{\mu}_t)'$ are normally distributed innovations, and remaining coefficients are defined in the appendix. As our state-space system is linear conditional on the states $(\bar{\pi}_t, \mathbf{k}_t)$, we employ the Marginalized Particle Filter of Schön, Gustafsson, and Nordlund (2005): using Bayes rule, the linear state variables ξ_t are marginalized out and estimated using the linear Kalman filter. The nonlinear state variables are estimated using the particle filter — see, for example, Kitagawa (1996).¹⁷ Because there is an explicit dependence between the linear and nonlinear state variables the prediction of the nonlinear state variables can be used to improve the estimates of the linear state variables. This feature distinguishes our application from Rao-Blackwellization techniques commonly used in economics and finance.

Despite these methods being new to the macroeconomics literature, we are not the first to estimate models with learning dynamics. Earlier contributions by Primiceri (2005) and Sargent, Williams, and Zha (2006) also estimate models with state-dependent learning gains, which include as a special case recursive least-squares estimation of statistical models, such as vector auto-regressions. However, these models study beliefs which are solely a function of variables that are *observable* to the econometrician. This special class of learning models can be expressed in linear state-space form. In contrast, and in common with most dynamic stochastic general equilibrium model applications, our framework assumes firm beliefs are updated using information that is not directly observable to the econometrician, namely the monetary policy shocks $\tilde{\varphi}_t$. Correctly specified inference requires non-linear methods.¹⁸

Estimated parameters. Table 1 shows prior and posterior distributions for each parameter. The posterior distribution is obtained by first computing the mode of the distribution. In a second step we use the Metropolis-Hastings algorithm to compute the full distribution.¹⁹

¹⁷See Fernandez-Villaverde and Rubio-Ramirez (2007) for a discussion of theory and application of the particle filter to macroeconomic models without a conditionally linear structure.

¹⁸Milani (2007, 2014) and Slobodyan and Wouters (2012) use linear Kalman filter techniques to estimate models of this kind. This requires arbitrary assumptions on agents' information sets. In particular, to compute forecast errors, agents are assumed to form predictions using the *econometrician's mean estimate* of processes such as the marginal cost. This is inconsistent with model micro-foundations which assume agents observe all exogenous processes.

¹⁹We started the Metropolis-Hastings step using a diagonal matrix with prior variances on the diagonal.

		Prior distribution		Posterior distribution			
	Distr.	Mean	St. Dev.	Mode	Mean	5 percent	95 percent
π^*	Normal	2.250	0.400	2.461	2.472	2.102	2.880
$\bar{\theta}$	Gamma	0.050	0.040	0.020	0.029	0.012	0.052
\bar{g}	Gamma	0.100	0.090	0.125	0.145	0.103	0.203
γ	Beta	0.500	0.265	0.115	0.128	0.090	0.173
Γ	Beta	0.500	0.265	0.926	0.891	0.813	0.948
ρ	Beta	0.500	0.200	0.874	0.877	0.832	0.917
$\sigma_{\bar{\varepsilon}}$	Inv.-Gamma	0.100	2.000	0.086	0.084	0.071	0.098
$\sigma_{\bar{\mu}}$	Inv.-Gamma	0.100	1.000	0.371	0.359	0.301	0.416
$\sigma_{o,1}$	Inv.-Gamma	0.100	1.000	0.264	0.277	0.217	0.335
$\sigma_{o,2}$	Inv.-Gamma	0.100	1.000	0.043	0.042	0.035	0.050
$\sigma_{o,3}$	Inv.-Gamma	0.100	1.000	0.020	0.021	0.014	0.028
$\sigma_{o,4}$	Inv.-Gamma	0.100	1.000	0.071	0.073	0.063	0.084
$\sigma_{o,5}$	Inv.-Gamma	0.100	1.000	0.048	0.049	0.041	0.059

Note: The posterior distribution is obtained using the Metropolis-Hastings algorithm.

Table 1: **Prior and Posterior Distribution of Structural Parameters**

The data permit fairly tight identification of key model parameters. Measurement errors on the survey forecasts have small variance. This provides confidence we identify the core mechanism of the model. The measurement error on inflation implies price changes in the model are somewhat less volatile than CPI inflation, but track its movements fairly closely.²⁰ The sample average inflation rate, π^* , has a posterior mean of about 2.5%, in annualized terms, and the 90% posterior interval ranges from 2.1% to 2.9%. The parameter γ is tightly distributed around 0.1, suggesting a small degree of price indexation. This is consistent with Cogley and Sbordone (2008), showing that once low-frequency movements in inflation are properly accounted for, there is little evidence of price indexation. Concomitantly, the feedback effects from drifting beliefs, measured by Γ , are substantial, with the 90% interval between 0.8 and 0.95. Later counterfactual analysis demonstrates self-confirming beliefs are the driving force behind the rise and fall of inflation over the sample period, and central to long-term expectations being well anchored or not.

Turning to the parameters defining the learning algorithm, $\bar{\theta}$ has a posterior mean of

The transition probability function was iteratively updated using short chains of between 20,000 and 40,000 draws. This process is repeated until the resulting variance-covariance matrix is stable. The variance-covariance matrix so obtained is used to generate 5 samples of 200,000 draws. A step size of 0.2 gave a rejection rate of 0.64 in each sample. Convergence is evaluated using the Gelman and Rubin potential-scale-reduction factor, which was well below 1.01 for all estimated parameters.

²⁰In particular, model-implied inflation remains significantly more volatile than the GDP deflator. See the additional Appendix for details.

0.03 with 90% credible interval spanning 0.01 and 0.05. To evaluate these magnitudes, for each draw from the posterior distribution, use (15) to measure the minimum absolute distance between $\bar{\pi}_t$ and its true value of zero required for firms to switch to a constant gain. Using the 90% posterior interval, this threshold ranges between 0.3% and 0.8% in annual terms. More concretely, the mean threshold implies that as $\bar{\pi}_t$ drifts outside the interval $2\% < \pi^* + \bar{\pi}_t < 3\%$ firms switch to the constant-gain regime. Firms are responsive to small deviations from the true inflation mean. However, we will show it can take a sequence of persistent and potentially large shocks to shift the estimated inflation drift significantly.

The constant gain has mean posterior estimate 0.14. Values of the gain within the 90% interval imply firms give minimal weight to observations older than five years. For example, at the boundaries of the 90% posterior interval, corresponding to $\bar{g} = 0.1$ and 0.2, the weight on observations two-years old is 0.2 and 0.4 respectively, with five-year weights being negligible. When inflation expectations are un-anchored, long-term inflation expectations are quite sensitive to short-term forecast errors.

4 MODEL PREDICTIONS

Figure 1 displays the model’s fit of short-term inflation forecasts. In each panel, the dashed gray line denotes CPI inflation; the solid black line displays the median prediction, while the gray shaded area shows the 95% credible set. Consistent with the small size of the observation errors in Table 1, the model-implied short-term forecasts correspond closely with the survey forecasts represented in red and blue: forecast surprises are well disciplined by the data.²¹ This permits testing the link between short-term surprises and long-term expectations.

To this end, we compare model-based predictions with available survey-based long-term forecasts. Because many surveys do not cover the entire sample, we use a number of measures from various surveys to build a comprehensive picture of long-term inflation expectations. Prior to the late 1970s there are no professional forecast data. We therefore use the five-to-ten-year-ahead forecasts from the Michigan Survey, restricted to the period 1974Q2-1977Q2.²² After this time the following professional forecast data are available. For one-to-ten-year-ahead average inflation forecasts: the Decision Makers Poll Survey (1978Q3-1980Q4); Livingston Survey (1990Q2-2015Q4); Blue Chip Economic Forecasts (1979Q4-1991Q1); and the Survey of Professional Forecasters (1991Q1-2015Q4). For five-to-ten-year-ahead average inflation forecasts: Blue Chip Economic Forecasts (1984Q1-2015Q4); Blue Chip Financial Forecasts (1986Q1-2015Q4); Survey of Professional Forecasters (2005Q3-

²¹As mentioned above model-implied inflation, while somewhat less volatile, tracks closely observed CPI inflation.

²²We study the behavior of the Michigan survey in section 8.

ANCHORED INFLATION EXPECTATIONS




Figure 1: Short-Term Forecasts.

The top and middle panels show the evolution of one- and two-quarters-ahead forecasts from SPF; The bottom panel shows two measures of the six-months-ahead forecasts from Livingston. The black lines show median predictions, while the gray area measures the 95% credible interval; the red and blue dots denote survey-based forecasts. Finally, the thin gray line measures CPI inflation.

2015Q4); Consensus Economics (1989Q4-2105Q4).

Long-term inflation expectations. Figure 2 shows model-predicted long-term inflation expectations with corresponding survey data. The survey data are given by the red and blue dots in the upper panel, which correspond to the five-to-ten- and one-to-ten-year-ahead forecasts. When multiple surveys are available for the same forecast the figure reports the average forecast across surveys. The solid black line in the top panel measures the median prediction for the five-to-ten-year-ahead expectation, while the gray areas denote the 70% and 95% credible sets; the dashed black line shows the median prediction for the one-to-ten-year-ahead forecast. The proximity of the two model-generated long-term forecasts underscores the low-frequency movement in inflation is for the most part driven by the drift $\bar{\pi}_t$. The model captures well the evolution of long-term forecasts. The credible sets are fairly tight, especially beyond the 1970s, and the survey-based forecasts largely lie within the 70% credible set. More broadly, the model captures the mild increase in expectations in the mid-seventies; the surge up to the early 1980s; the gradual decline through to the 1990s; and the stabilization in the 2000s, which persists beyond the crisis period.

The results are consistent with contemporary narratives of US monetary history. The pattern of short-term forecasts, in the middle panel of Figure 1, shows that starting in the early 1970s firms faced persistent positive inflation surprises, leading to poorly anchored inflation expectations, and a switch to a constant gain illustrated in the bottom panel. The constant-gain regime lasts until the mid-1990s, because of persistent negative surprises, initially during the Volcker disinflation, and, subsequently, the disinflation in the early 1990s under Greenspan. Since the late 1990s, conditional on observed inflation and short-term forecasts, the 95% credible set excludes a switch to a constant-gain regime, despite occasionally large forecast errors. Indeed, from Figure 1, these forecast errors are of a similar magnitude to those observed over the 1970s and 1980s. However, the pattern of forecast errors is crucially different: they were not as persistent and, therefore, did not generate large enough deviations in (15) to lead to a drift in inflation beliefs. The model therefore provides a rationale for the observed stability of expectations during this period which include sizable shocks like the 2008 recession.²³

Endogenous inflation trend and monetary policy. The relative stability of the persistent component of inflation has been the basis of much recent skepticism about models

²³The result that long-term expectations have been stable over the past decade does not appeal to any informational friction predicting inertia in belief updating. The results show that reported survey inflation expectations represent a coherent view of actual inflation developments — not simply mechanical reportage of central bank inflation objectives. Long-term forecasts adjust significantly only in response to systematic and sufficiently large forecast errors. Once again: the fact that expectations are anchored in itself affects the size of forecast errors — the stability of expectations arises endogenously in this model. The counterfactuals in Section 6 further clarify this point.

ANCHORED INFLATION EXPECTATIONS

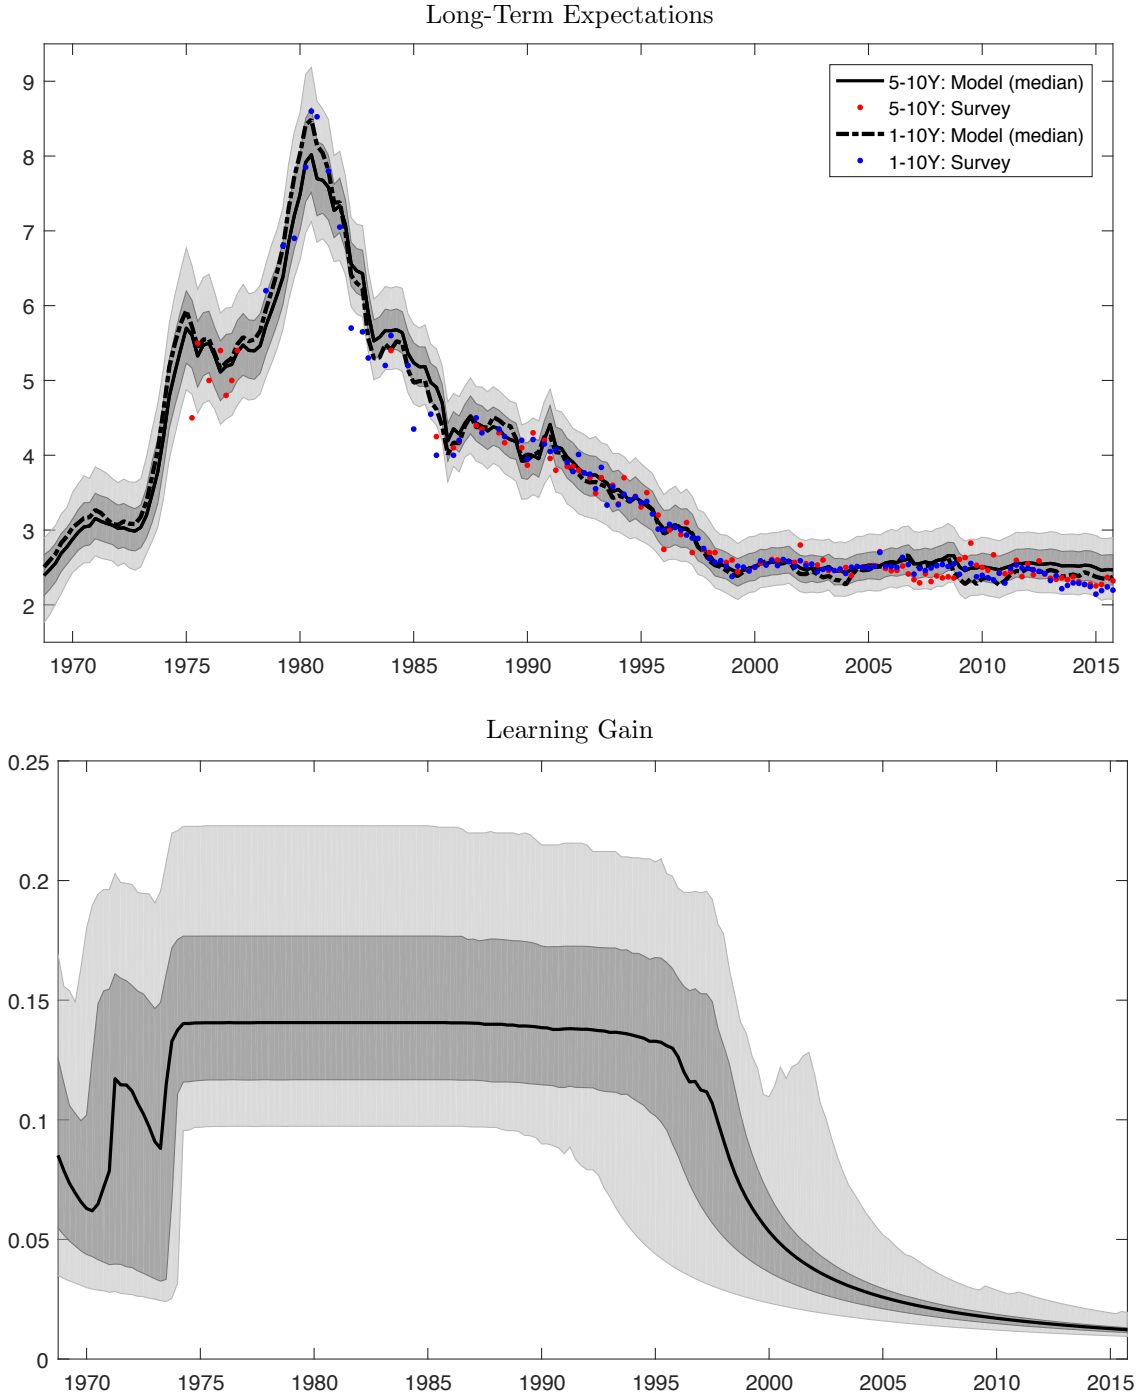


Figure 2: **Predictions: Professional Forecasters.**

The top panel shows the evolution long-term expectations at five-to-ten-year (5-10Y) and one-to-ten-year (1-10Y) horizons. The black lines show median predictions, while the gray area measures the 70% and 95% credible intervals; the red and blue dots denotes survey-based forecasts. The bottom panel shows the evolution of the learning gain.

relating inflation to measures of economic slack, such as the Phillips curve in its various forms.²⁴ But our analysis is entirely consistent with time variation in the sensitivity of inflation to economic activity. The weak link between shocks and inflation dynamics during the financial crisis period is because inflation expectations are anchored. Conversely, the large and persistent shocks realized over the 1970s and 1980s produced sizable effects on inflation because they led to unanchored expectations. Beliefs transmit these shocks to the persistent component of inflation, creating a strong link between real activity and prices.²⁵

Prior accounts of the low-frequency movements in inflation, typified by the Great Inflation, rely on purely exogenous specifications of the inflation drift — see, for example, Smets and Wouters (2003), Cogley and Sbordone (2008) and Cogley, Primiceri, and Sargent (2010). While Kozicki and Tinsley (2005) and Ireland (2007) permit more general specifications in which the drift is determined in part by identified disturbances, such as supply shocks, the specification remains exogenous. These papers interpret low-frequency movements as the result of variation in preferences over inflation outcomes. Higher inflation in the 1970s reflects the Federal Reserve’s rising tolerance for inflation. In contrast, because the drift is endogenous in our account, the Great Inflation arises despite the Federal Reserve’s commitment to price stability.²⁶ In this way the paper has much in common with Sargent (1999) and Primiceri (2005), which emphasize learning by policy makers rather than price setters.

More generally, the endogeneity of the inflation trend fundamentally alters impulse and propagation mechanisms, with implications for policy design — see Eusepi and Preston (2018b) for a recent survey. To give one pertinent example, Eusepi, Giannoni, and Preston (2018) show in a New Keynesian model that optimal policy under un-anchored expectations (large constant gain) exhibits history dependence and resembles optimal commitment under rational expectations. But anchored inflation expectations (small constant gain) induce no history dependence, as in the case of discretion under rational expectations.

5 UNDERSTANDING THE ENDOGENOUS TREND

The two defining features of the model are that beliefs and optimal pricing decisions combine to generate self-referential dynamics, and the learning gain is state dependent. The following counterfactuals isolate the contributions of each of these properties. All counterfactuals are generated by simulating the economy using the posterior distribution of the model parameters, and the smoothed distribution of initial states and structural innovations.

²⁴See for example Hall (2011).

²⁵Del Negro, Giannoni, and Schorfheide (2015) provides a similar argument in a medium-size DSGE model, where long-term inflation expectations evolve exogenously.

²⁶Eusepi and Preston (2018a) provides evidence for a related belief structure, in an empirical medium-scale general equilibrium model of the US economy.

The role of self-referential beliefs. Long-term inflation expectations are a state variable that engender low-frequency movements in inflation: through optimal price setting, shifts in beliefs affect inflation which, in turn, shapes beliefs. This produces substantial variation in beliefs for certain sequences of forecast errors. Importantly, low-frequency drift in beliefs is directly related to episodes of poorly anchored inflation expectations.

To illustrate this connection, we run the following counterfactual simulation. Assume strict inflation targeting, $\Gamma = 0$, which eliminates the feedback from beliefs to actual inflation. Furthermore, assume firms update their estimates of $\bar{\pi}_t$ using a constant-gain algorithm in all periods, so that $\bar{\theta} = 0$. Long-run beliefs are therefore quite sensitive to inflation surprises, even though we show in section 6 that this would not have been optimal. The top panel in Figure 3 shows the median prediction for the agents' long-term expectations (solid black line) together with the 95% credible interval in this counterfactual economy. Also shown are the long-term survey forecasts, together with the 95% credible set from the predicted long-term inflation expectations of the baseline model (light green). Even with the assumed high responsiveness to inflation surprises, inflation expectations exhibit limited drift. Absent feedback effects, inflation expectations fluctuate very little. This illustrates how the sensitivity of inflation to disturbances can vary over time, depending on general equilibrium effects from belief updating, price setting and the policy regime. Drift in inflation is not the mechanical outcome of a specific choice of learning algorithm.

The role of timing of switching beliefs. How important is the timing of switches in the forecasting algorithm? Can beliefs that are relatively insensitive to new information be counterproductive from the perspective of a central bank trying to engineer a disinflation? Long-term inflation expectations plateau briefly at around 4% in the late 1980s, before continuing to decline. Suppose at this time agents become convinced the observed decline in inflation represents the end of the Federal Reserve's commitment to disinflation. Believing inflation to be close to its time-invariant mean, firms switch to a decreasing-gain algorithm to construct forecasts. The middle panel in Figure 3 shows this counterfactual. In this scenario, long-term inflation expectations decline very slowly and remain above 3% at the end of 2015. Again, self-referential dynamics explain slow convergence. Because beliefs display less sensitivity to new information, negative inflation surprises lead to smaller downward revisions in long-term inflation expectations, which results in actual inflation declining more gradually. This hinders the ability of firms to learn the true underlying inflation mean.

To gauge the importance of these self referential effects we run an additional experiment. Assume now inflation expectations for all firms are updated using a constant gain, so that inflation evolves exactly as in the baseline model. Now suppose a measure-zero firm updates their beliefs using a decreasing-gain algorithm. In contrast to the previous simulation, the

measure-zero agent faces a data-generating process which has inflation declining much faster. Moreover, their own expectations do not feed back into actual inflation outcomes. The dashed black line demonstrates the evolution of this firm’s inflation expectations. Despite using a decreasing-gain algorithm, when there is no self-referentiality inflation expectations converge noticeably faster, falling below 3% by 2005 as compared to 2015. The stickiness of inflation expectations is a general equilibrium outcome.

Stable expectations: “well anchored” or “good luck”? From the early 2000s long-term inflation expectations are very stable. This stability reflects neither inattention nor inertial adjustment on the part of professional forecasters. Similarly, it does not reflect good luck from favorable disturbances. Indeed, large forecast errors arising from substantial energy and commodity price changes, and the events of the financial crisis underscore anchored expectations are not simply the result of a stable underlying inflation process. Rather, the stability of long-term expectations is due to the relatively small, and decreasing, learning gain, a dividend of well-anchored inflation expectations. What if firms remained skeptical of the Federal Reserve’s commitment to price stability, and never switched to a decreasing gain in the mid-to-late 1990s? With the US economy experiencing the same shocks as in the baseline model, the bottom panel of Figure 3 demonstrates inflation under constant-gain learning would have exhibited substantial volatility. In contrast to the baseline prediction, long-term inflation expectations would have been very uncertain, with the credible interval ranging from 0.5% to 4%. In addition, toward the end of the sample, the model’s predictions imply some downside risk to inflation expectations, accompanied by median predictions trending below 2%.

6 MONETARY POLICY AND BOUNDS ON IRRATIONALITY


The learning algorithm depends on two parameters. This section demonstrates these parameters satisfy a lower bound on rationality — agents cannot hold beliefs that are implausible. As a corollary, it shows that more aggressive monetary policy toward inflation weakens the link between short-term inflation surprises and long-term inflation expectations, delivering more strongly anchored beliefs.

Asymptotic convergence. A reasonable requirement of a theory of long-term expectations is that firms can learn the inflation target in a stationary environment. The estimate $\bar{\pi}_t$ should converge to zero, the rational expectations equilibrium, with probability one.²⁷ Suppose firms construct forecasts using a recursive least-squares algorithm. Using standard results in Marcet and Sargent (1989) and Evans and Honkapohja (2001) we have the following


²⁷Formally, we consider a log-linear approximation in the neighborhood of price stability (zero inflation).

ANCHORED INFLATION EXPECTATIONS

No Feed-Backs ($\Gamma = 0$).



Early Switch to Decreasing Gain



Stay with Constant Gain




Figure 3: Counterfactuals: Endogenous Inflation Trend

The three panels show the predicted behavior of long-term expectations under alternative counterfactual simulations described in the main text. The black line denotes median predictions; the gray area shows the 95% credible interval in the counterfactual; the green lines show the 95% credible interval under the baseline model; the blue and red dots measure long-term forecasts from surveys.

result.

Result 1. *Consider the equilibrium evolution of $\bar{\pi}_t$ under least-squares learning. Provided $\Gamma < 1$, then $P(\bar{\pi}_t \rightarrow 0) = 1$.*

A formal proof can be found in Evans and Honkapohja (2001). Here we offer a heuristic argument. Recall from (19)

$$\bar{\pi}_{t+1} = \bar{\pi}_t + \mathbf{t}^{-1} \times [(1 - \gamma)(\Gamma - 1)\bar{\pi}_t + \eta_t].$$

The mean dynamics are then approximated by the ordinary differential equation

$$\dot{\bar{\pi}} = (1 - \gamma)(\Gamma - 1)\bar{\pi}$$

which is (globally) stable, provided $\Gamma < 1$.

This asymptotic result establishes the local stability of the inflation target: so long as initial beliefs are sufficiently close to the true long-run mean of inflation, given enough data firms can learn the objective of the central bank. Now consider the properties of beliefs under a constant-gain algorithm.

Result 2. *Consider the evolution of $\bar{\pi}_t$ under constant-gain learning. Provided*

$$|1 + \bar{g}(1 - \gamma)(\Gamma - 1)| < 1$$

then $\bar{\pi}_t$ evolves according to an auto-regressive process of order one, with mean 0.

To see this, from (19), the dynamics of the inflation drift satisfy

$$\bar{\pi}_{t+1} = [1 + \bar{g}(1 - \gamma)(\Gamma - 1)]\bar{\pi}_t + \bar{g}\eta_t$$

giving the statistical properties we desire if the condition of the result is met. Beliefs about trend inflation are then ergodically distributed around the true long-run mean of inflation.

While we do not offer a formal proof of convergence, these two results are useful to build intuition about convergence in our model. To frame the discussion, consider the following experiment. Generate 50,000 simulations with each lasting 10,000 periods. In each simulation, draw both initial conditions for agents' estimate of the inflation mean and the entire sequence of exogenous shocks. Figure 4 displays two statistics that establish firms eventually learn the true model — that is, beliefs are anchored at the true inflation target. The top panel shows in each period the fraction of simulations in which the gain is larger than the minimum constant gain estimate from the parameter distribution in the empirical model.

That is, the fraction of simulations in which agents still have a constant gain. This fraction decreases monotonically over time and reaches a number close to zero by the end of the simulation. The bottom panel shows the fraction of simulations where $|\bar{\pi}_t| > 0.05$. Again this fraction decreases monotonically over time, implying long-run expectations converge to the true inflation target.

To gain intuition, consider the economy when agents update using a decreasing gain. From Result 1, for beliefs within the domain of attraction agents never switch forecasting models and convergence occurs. However, sufficiently large shocks can induce switching. To see this, equations (14) and (19) imply firms will use a decreasing gain algorithm if

$$|[1 + \mathbf{t}^{-1}(1 - \gamma)(\Gamma - 1)]\bar{\pi}_t + \mathbf{t}^{-1}\eta_t| < \bar{\theta}\tau \tag{21}$$

for some threshold $\bar{\theta}$ and defining $\tau = \sigma_\eta / (1 - \gamma)(1 - \Gamma)$. For small values of \mathbf{t} sufficiently large shocks can lead to a switch to the constant gain algorithm, preventing convergence. Under the constant gain the economy behaves as described in Result 2. Given that $\bar{\pi}_t$ fluctuates around zero (and so the criterion in 14 is zero), the mean of this process, firms will switch to the decreasing gain model again with probability one. Importantly, the longer agents adopt a decreasing gain the smaller the probability of escaping. From equation (21) it is immediate that as the gain gets smaller the probability of receiving a shock large enough to produce a switch decreases and vanishes as $\mathbf{t} \rightarrow \infty$.²⁸

Given sufficient data, firms will always learn the long-run mean of inflation.²⁹ This property distinguishes our model from various other contributions in the literature which seek to study questions of central bank credibility, and, specifically, whether inflation expectations are anchored or not. For example, Orphanides and Williams (2005) and Kozicki and Tinsley (2005) develop models of central bank credibility in which agents must estimate the inflation target using a constant-gain algorithm.³⁰ While imperfect knowledge has implications for monetary policy design, both models have the property that agents cannot ever learn the time-invariant inflation target. More recently, papers such as Hommes and Lustenhouwer

²⁸This is easier to see if we are willing to assume a finite bound, no matter how large, on the support of the shock.

²⁹Much of the literature on learning and monetary policy can be interpreted as models of central bank credibility, and as answering the question of whether expectations are well anchored or not. As one example, Eusepi and Preston (2010) provide theoretical results on central bank communication, showing certain types of information ensure consistency of beliefs with monetary policy strategy — providing a well-defined notion of anchored expectations — which improves stabilization outcomes of a given monetary policy framework. However, such theoretical analyses, relying on asymptotic convergence results, do not provide an account of dynamics when beliefs are poorly anchored.

³⁰See also Lansing (2009); Hommes and Lustenhouwer (2015) for a model of central bank credibility with heterogeneous expectations; and Bomfim and Rudebusch (2000) and Gibbs and Kulish (2015) for explorations of the role of anchored expectations in the costs of disinflation.




Figure 4: **Convergence.**

Statistics on convergence properties of the model from 50,000 simulations over 10,000 periods. For each simulation, we draw an initial estimate of the inflation target and the full sequence of exogenous shocks.

(2015) explore similar questions in a model of heterogeneous beliefs and predictor selection. While the composition of predictors is endogenous to the environment, the model nonetheless has the property that the ergodic distribution of beliefs never converges to the central bank's inflation target.

Optimality. The second lower bound on rationality is given by a criterion proposed by Marcet and Nicolini (2003). We show that within the adopted class of learning algorithm beliefs are in a certain sense optimal. Specifically, conditional on a data-generating process where all firms learn using parameters $(\bar{g}, \bar{\theta})$, an individual choosing a potentially different parameter pair $(\bar{g}', \bar{\theta}')$ should not produce much better forecasts than the parameters $(\bar{g}, \bar{\theta})$

according to a mean-squared error criterion. Formally, the choice of beliefs must satisfy

$$\bar{E} [f_t(\bar{g}, \bar{\theta})]^2 \leq \min_{\bar{g}', \bar{\theta}'} \bar{E} [f_t(\bar{g}, \bar{\theta}, \bar{g}', \bar{\theta}')]^2 + \epsilon$$

for some small $\epsilon > 0$. We approximate the above moments using the mean-squared error from a sample of size T , averaged over S replications, where each replication represents a simulation of the model at the modal estimates. Consistent with the sample size in estimation, take $T = 264$, and $S = 200,000$. This is done for all candidate beliefs on a grid covering the interval $\bar{g} \in [0.001, 0.3]$ and $\bar{\theta} \in [0.001, 0.1]$.

Figure 5 shows the ratio of the mean-squared errors of the counterfactual beliefs relative to the baseline model, as a function of $(\bar{g}', \bar{\theta}')$ in two different economies.³¹ Numbers greater than unity imply the baseline parameters are optimal. The top panel provides the results for the estimated model. The estimated parameters $(\bar{g}, \bar{\theta})$ are optimal, with other forecasting models exhibiting a deterioration in performance. There is no incentive to deviate from these beliefs. This is explained by the high degree of self-referentiality in the system. Because movements in beliefs are in large part reflected in the data, it is optimal to maintain those beliefs. Indeed, the bottom panel shows the same exercise in a counterfactual economy exhibiting substantially less feedback than at the mode: $\Gamma = 0.4$. As discussed in section 2, a more aggressive policy towards inflation lowers Γ . As a result, an individual firm would find it desirable to adopt a lower value of the gain and a higher switching threshold. This shows how the expectation formation mechanism is not independent of policy. A more aggressive policy toward inflation endogenously lowers the sensitivity of long-term inflation expectations to short-term surprises and, with that, limits the drift observed in inflation.


7 MODEL COMPARISON

We now offer a formal assessment of model fit, comparing its predictive likelihood against three alternatives: the rational expectations version of our baseline model where $\bar{\pi}_t^* = 0$ in every period; a model where beliefs are updated using a constant-gain algorithm — equivalent to setting $\bar{\theta} = 0$ in our baseline model; and a model where $\bar{\pi}_t^*$ evolves exogenously according to the auto-regressive process

$$\bar{\pi}_{t+1} = \rho_{\bar{\pi}} \bar{\pi}_t + \varepsilon_t^{\bar{\pi}}$$

³¹To better visualize the result, the Figure shows a smaller set of parameter values relative to the size of the grid used in the simulation. The result, however, holds for all parameters in the grid.

High Feedback (Mode).



Low Feedback




Figure 5: **Optimal choice of $\bar{\theta}$ and \bar{g}**

These panels show the ratio between the mean-squared error from predictions of a measure-zero agent using alternative values of $\bar{\theta}$ and \bar{g} and that from the representative agent using modal estimates. The true data-generating process is obtained assuming agents use the modal estimates. The top panel assumes the feedback parameter Γ at the mode; the bottom panel assumes $\Gamma = 0.4$.

where the innovation is independent of the structural shocks in η_t . As a statistical description of inflation, this final class of models has been studied in a range of contexts, including reduced-form, partial equilibrium and general equilibrium analyses [Smets and Wouters (2003), Stock and Watson (2007), Cogley and Sbordone (2008), Cogley, Primiceri, and Sargent (2010), Del Negro, Giannoni, and Schorfheide (2015) and Mertens (2016)]. These papers have in common a highly persistent inflation trend, with auto-regressive coefficient in the neighborhood of, or equal to, unity. In the subsequent exercise we set $\rho_{\bar{\pi}} = 0.99$.³² Because the process for the inflation trend is exogenous, we set $\Gamma = 1$ in (19).

Following Del Negro and Eusepi (2011), we evaluate how these models, estimated to fit inflation and short-term survey forecasts, perform in predicting long-term inflation expectations. To achieve this we compute the predictive likelihood

$$\begin{aligned} p(y^{LT}|y^{ST}, M_j) &= \int p(y^{LT}|y^{ST}, \Theta, M_j) p(\Theta|y^{ST}, M_j) d\Theta \\ &= p(y^{LT}, y^{ST}|M_j) / p(y^{ST}|M_j) \end{aligned}$$

where y^{LT} includes the two series measuring long-term inflation expectations shown in Figure 2: the five-to-ten- and one-to-ten-year inflation forecasts. The vector y^{ST} includes CPI inflation and all the short-term survey forecasts used in the estimation of the baseline model. Finally Θ contains the estimated parameters in each model. This vector includes two additional measurement errors associated with the long-term forecasts.³³ The predictive likelihood measures the fit of long-term inflation expectations conditional on the parameter distribution delivering the best fit of both short-term inflation expectations and inflation. In other words, we evaluate each model's fit of long-term expectations by using as a prior the posterior parameter distribution obtained using only short-term forecasts and inflation. As shown above, it can be computed by taking the ratio of the marginal likelihoods resulting from estimating the model using two different data sets: one including (y^{ST}) , and the second additionally including long-term forecasts, (y^{ST}, y^{LT}) .

Table 2 compares the marginal and predictive likelihood for the baseline model and the three alternative models. The baseline model performs best. Figure 6 gives intuition. The three panels show the predictive density for each alternative model (estimated using only short-term forecasts), compared to the survey-based long-term expectations and the median predictions from the baseline model. The rational expectations model gives the worst predictive likelihood, fails to deliver the large increase in long-term expectations during the

³²This choice is consistent with a range of papers in the literature.

³³We set the standard deviation of both measurement error innovations to 0.15.

	$\ln p(Y_{1,T}^{ST})$ Dataset without LT Expectations (1)	$\ln p(Y_{1,T}^{ST}, Y_{1,T}^{LT})$ Dataset with LT Expectations (2)	$\ln p(Y_{1,T}^{LT} Y_{1,T}^{ST})$ (2) - (1)
Baseline	-567.68	-677.97	-110.29
Rational Expectations	-588.51	-727.34	-138.83
Constant Gain	-571.43	-686.66	-115.23
Exogenous π_t^*	-579.44	-698.55	-119.11

Note: The table shows the the log-marginal likelihood for the four alternative models. The first column corresponds to the case where long-term expectations are not used in the estimation, while the second column shows the marginal likelihood when long-term expectations are added to the data set. Finally, the third column shows the predictive likelihood for each model.

Table 2: **Model Comparison**

1970s and 1980s, and predicts excessively volatile expectations in the 2000s. The constant-gain model performs well until the mid-2000s and then produces a countefactual run-up in inflation expectations in the mid-2000s. Finally, the model with exogenous inflation mean under-predicts long-term inflation expectations in both the 1980s and 1990s, and produces excessively volatile expectations in the 2000s. The estimated volatility of the innovation to the inflation drift is the result of a trade-off between fitting the large variations in the inflation trend during the first part of the sample with the stability in the latter part. As shown by Mertens (2016), for example, even higher-dimensional models with an exogenous inflation trend which use several measures of inflation and economic activity in the estimation present similar shortcomings in fitting long-term inflation forecasts.


8 LEARNING AND MODEL UNCERTAINTY IN REAL TIME

To allay concern about the use of the true data generating process in our learning algorithm, we provide two robustness exercises. The first shows that our criterion is consistent with the implications of a recursive statistical test. The second shows a learning algorithm which continuously updates the learning gain fails to capture key features of the data.


Detecting structural change. This section provides a structural interpretation of how agents update their model in real time. We assume firms assess the stability of the long-run inflation mean by conducting a sequential test on their model’s forecast errors, f_t . This

ANCHORED INFLATION EXPECTATIONS

Rational Expectations



Constant Gain



Exogenous π_t^*




Figure 6: Model Comparison

The three panels show three alternative models estimated on the same observables as our baseline model.

corresponds to the sequential LM test in Cho and Kasa (2015). The test is based on the recursive estimate of the regression constant, similar to the well-known CUSUM of squares test of Brown, Durbin and Evans (1975). The aim is to detect a shift in the mean of inflation from a volatile signal, as described in equation (16). To this end, agents employ a test based on the statistic f_t^2/ω_t , where ω_t is a recursive estimate of the variance of the forecast error implied by firms' forecasting model

$$\omega_t = \omega_{t-1} + \kappa \mathbf{k}_{t-1}^{-1} (f_t^2 - \omega_{t-1}).$$

Allowing for discounting of past data, agents evaluate the model using a recursively estimated statistic

$$\theta_t = \theta_{t-1} + \kappa \mathbf{k}_{t-1}^{-1} (f_t^2/\omega_t - \theta_{t-1})$$

where $0 < \kappa < 1$.³⁴ The introduction of the parameter κ permits the test statistic to be revised at a different rate than firms' estimates of the long-run inflation rate. The null hypothesis of a constant inflation mean is rejected if the statistic θ_t crosses a threshold $\tilde{\theta}$. When firms cannot reject a constant mean, they forecast inflation by using a recursive least-squares algorithm (decreasing gain). When the test indicates the mean has shifted, they switch to a constant-gain algorithm, permitting tracking of structural change. The evolution of the time-varying learning gain can then be written

$$\mathbf{k}_{t+1} = I_{(\theta_t \leq \tilde{\theta})} \times (\mathbf{k}_{t-1} + 1) + \left(1 - I_{(\theta_t \leq \tilde{\theta})}\right) \times \bar{g}^{-1}$$

where I is an indicator function.

Cho and Kasa (2015, 2017) propose an alternative model-validation procedure. Suppose there are two forecasting models, as in our application, one with a time-invariant mean, and one with drifting, random-walk, coefficients. The mean in the former is estimated using a Kalman filter with decreasing gain, and in the latter using a constant gain. Both models are revised in real time using using all available data. However, at any point in time, only one model is used for forecasting, and this same model is used until rejected by a statistical test of the kind developed above. When rejected, the alternative model is adopted with some probability.

The testing strategy used here is closer to Marcet and Nicolini (2003) and Milani (2014) and differs from Cho and Kasa (2015, 2017) in an important way. In our model agents entertain the possibility that the policy regime might experience a sudden change. Their forecast selection mechanism assumes that the policy environment might be changing over

³⁴See Cho and Kasa (2015) for a detailed discussion of these recursive tests.

time. When agents switch back to a decreasing-gain algorithm, the initial learning gain is reset to a value equal to the constant gain. This choice reflects the fact that, while agents believe the inflation mean constant under the current regime, it might not have been in the past. Consequently, they discount older observations. In Cho and Kasa (2017) agents estimate the model with constant parameters using the full sample, whether it is adopted or not. For example, under the null of a constant mean the model gives equal weight to all data. This reflects the assumption that agents believe there exists only one fixed regime.

We don't estimate the model because this implementation of sequential testing breaks conditional linearity and because the number of nonlinear state variables increases by two. We instead calibrate the parameters $(\kappa, \tilde{\theta})$ and initial conditions (θ_0, ω_0) by simulating the model using the distribution of smoothed states obtained in the baseline estimation. For each path of the smoothed shocks, the parameters are chosen to minimize the distance between the sequential testing model and our baseline predictions for the learning gain and the estimate of the long-run mean of inflation. Using the resulting distribution of parameters (one for each of the paths drawn), if we can replicate our benchmark results we have a proof of concept: there exists a sequential statistical test consistent with our criterion.

Consistent with agents detecting low-frequency developments in volatile environments, the calibration reveals statistics are computed with very little discounting of old data. The gain κ is small, in the range 0.01 – 0.06. The threshold $\tilde{\theta}$ ranges from 0.6 (5th percentile) to 0.90 (95th percentile). The top panel of Figure 7 shows the model predicts the evolution of inflation expectations is nearly identical to our baseline results. The middle panel shows the statistic θ_t increases in the 1970s until the mid 1990s and then declines. Interestingly, the bottom panel shows the estimated volatility of forecast errors is relatively flat until the 1990s and then declines markedly — not too different from the volatility estimates of Cogley and Sargent (2005).³⁵ We interpret these results as providing evidence that our criterion provides meaningful detection of model mis-specification.

Continuous updating of the learning gain. A key feature of the learning algorithm we assume is the discontinuous nature of the changes in the learning gain. If firms doubt the forecasting performance of their model, they switch to a constant-gain algorithm, potentially leading to a substantial change in the degree of sensitivity of long-run beliefs to new information. Such adjustments afford flexibility to adapt to shifting economic circumstances. To evaluate the role of the discontinuity, we consider an alternative real-time estimator of the long-run mean of inflation, where the learning gain is also state dependent. The adaptive-step-size algorithm proposed by Kushner and Yin (2003) permits the learning gain to be adjusted continuously, leading to a more gradual adjustment in the sensitivity of

³⁵See also Lansing (2009).

ANCHORED INFLATION EXPECTATIONS




Figure 7: **Model selection in real time.**

The top panel shows the predicted behavior of long-term expectations. The black line denotes median predictions; the gray area shows the 95% credible interval; the green lines show the 95% credible interval under the baseline model; the blue and red dots measure survey forecasts. The remaining two panels show the median and interquartile range for the evolution of the test statistic θ_t and the recursive estimate of the forecast errors' variance ω_t .




Figure 8: **Adaptive-step-size-algorithm.**

The top panel shows the predicted behavior of long-term expectations under the alternative learning model. The black line denotes median predictions; the gray area shows the 95% credible interval; the green lines show the 95% credible interval under the baseline model; the blue and red dots measure long-term forecasts from surveys. The bottom panel shows the median (black line) and the interquartile range for the evolution of the learning gain.

long-run beliefs to short-term surprises. This algorithm was first used in macroeconomics by Kostyshyna (2012) in the context of the hyperinflation model of Marcet and Nicolini (2003).

The algorithm for the learning gain g_t is

$$g_t = G(g_{t-1} + \nu f_{t-1} V_{t-1})$$

$$V_t = (1 - g_{t-1}) V_{t-1} + f_{t-1}, \quad \text{given } V_0.$$

The variable V_t measures discounted past forecast errors, f_t ; and the function $G(\cdot)$ satisfying

$$G(g_{t-1} + \nu f_{t-1} V_{t-1}) = \begin{cases} g_{t-1} + \nu f_{t-1} V_{t-1}, & \text{if } g_- < g_{t-1} + \nu f_{t-1} V_{t-1} < g^+ \\ g^+, & \text{if } g^+ < g_{t-1} + \nu f_{t-1} V_{t-1} \\ g_-, & \text{if } g_- > g_{t-1} + \nu f_{t-1} V_{t-1}. \end{cases}$$

The parameter ν now captures how the learning gain is adjusted in response to past forecast errors. The bounds g^+ and g_- ensure that the algorithm is well behaved.³⁶ When the current forecast error, f_t , has the same sign as discounted past errors, V_t , the gain increases. This has similar intuition to our baseline algorithm: the gain changes in response to persistent forecast errors of the same sign. Figure 8 shows outcomes of the same calibration as described above, where the parameter ν and initial conditions (g_0, V_0) are calibrated to minimize the distance between this model’s predictions for $\bar{\pi}_t$ and its evolution in our baseline.

The top panel demonstrates expectations fail to track the survey data, and are substantially more volatile in the latter part of the sample. The bottom panel reveals the source of difficulty is that the gain coefficient rises in the 1970s, but fails to decline much at all after 1980. This counterfactual adduces evidence that our threshold algorithm is better able to capture the relation between long-term expectations and short-term surprises, relative to algorithms which continuously update the gain. A primary reason for this shortcoming is that negative forecast errors after 1980 display smaller size compared to the positive errors in the 1970s — recall figure 1. As a result, the increase in the gain cannot be fully reversed.

9 PREDICTING OTHER LONG-TERM FORECASTS

This final section demonstrates the basic belief structure provides a good characterization of other survey data of inflation expectations. The parameter distribution is chosen to match the joint behavior of US inflation and short-term professional forecasts. We now document evidence supporting the proposed theory of belief formation for different types of economic agents (households rather than professional forecasters), and for a range of different countries.

9.1 HOUSEHOLD EXPECTATIONS

The Michigan Survey has collected data on short-term inflation forecasts since nearly the beginning of our sample, and on long-term forecasts since 1975. Despite the long sample,

³⁶In the numerical analysis that follows they are set to $g^+ = 0.6$ and $g_- = 0.01$ and the estimate for \bar{g} never reaches these boundaries.

the survey presents some challenges. First, in contrast with professional surveys, short-term forecasts are measured as one-year-ahead forecasts, so we can measure quarterly inflation surprises only indirectly. Second, survey participants are not asked to forecast a specific price index. In what follows we therefore continue to use CPI as our measure of inflation, acknowledging this is only a noisy measure of the underlying inflation rate about which households form expectations.³⁷ Third, short-term inflation forecasts display a substantial upward bias in the last fifteen years when compared to CPI inflation. To avoid modeling this bias directly, we mitigate the problem by using median forecasts, instead of the mean. Furthermore, following the finding of Coibion and Gorodnichenko (2012) that the difference between professional forecasters and Michigan forecasts is related to oil prices, we purge the bias by regressing the difference between short-term professional and household forecasts on oil prices. Our measure of household expectations is then given by the sum of the survey of professional forecasts data and the residual from the regression.³⁸

With these limitations in mind, the observation equation for household expectations is

$$\left[\begin{array}{c} \pi_t \\ E_t^{Mich} \left(\frac{1}{4} \sum_{i=1}^4 \pi_{t+i} \right) \end{array} \right] = \pi^* + H_t' \xi_t + R_t o_t.$$

The observation error on inflation now additionally captures the discrepancy between the CPI and the survey participants' notion of the inflation measure being predicted. We then use the posterior distribution from the baseline model using professional forecasts to infer an estimate of household beliefs about trend inflation. That is, conditioning on CPI inflation and the Michigan Survey short-term inflation expectations, we predict household long-term expectations. We assume the measurement error on inflation and the short-term forecast correspond to those estimated for inflation and the one-period-ahead professional forecast in the benchmark estimation. These assumptions are inconsequential given the small size of estimated measurement errors.

Figure 9 shows the model predictions for the five-to-ten-year-ahead expectations, along with the corresponding Michigan survey forecasts. The pattern matches broadly that for professional forecasters, adducing complementary evidence to Malmendier and Nagel (2016) that adaptive learning structures provide a reasonable description of household forecasts.³⁹ However, there are differences. First, the middle panel shows the evolution of one-year-ahead forecasts for both households (red dots) and professional forecasters (blue dots).⁴⁰

³⁷Note that measurement error on CPI inflation can at least partially address this data limitation.

³⁸The transformed and un-transformed measures are shown in the technical appendix.

³⁹Malmendier and Nagel (2016) exploit the cross-section providing evidence that different age-cohorts use different constant gains to form inferences about future inflation.

⁴⁰These are computed using the median response from the Livingston survey. A similar pattern obtains if

ANCHORED INFLATION EXPECTATIONS

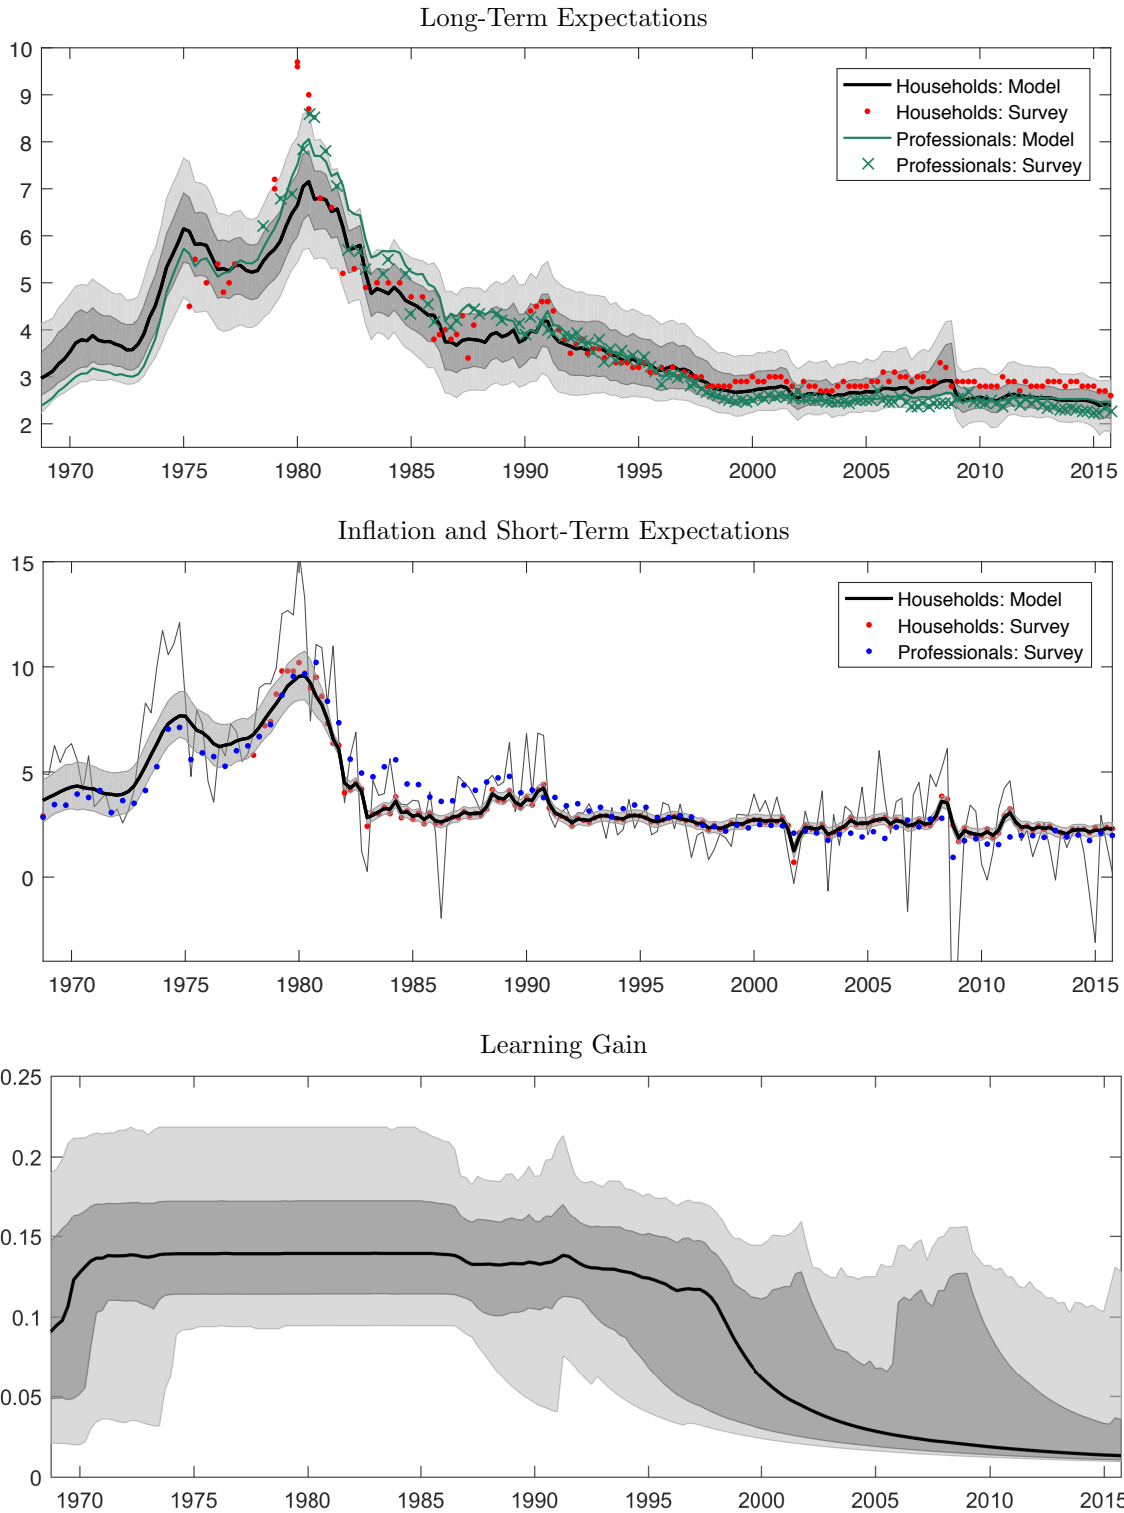


Figure 9: **Predictions: Households Forecasts.**

The panels show model predictions for long-term forecasts (top), short-term forecasts (middle) and the learning gain (bottom). The black solid (dashed) line denotes the median prediction; the gray area measures the 70% and 95% credible intervals; in the top panel, the red (green) dots indicate long-term survey forecasts from the Households Michigan survey (Professional forecasters); Finally the green solid line in the top panel shows the mean prediction from the baseline model

Interestingly, household forecasts decline more sharply than professional forecasts from the early 1980s to the early 1990s. As a result, the model on average predicts a lower path for long-term inflation expectations. In fact, over the same period long-term household forecasts from the Michigan survey (red dots) are consistently below professional forecasts (green crosses). In line with the pattern of short-term forecasts, household and professional long-term expectations move in sync during the 1990s. Second, while the median prediction is that expectations remained anchored in the aftermath of the financial crises, model predictions are more uncertain than for professionals: the model attributes a substantial probability to long-term expectations not being fully anchored during that period. This is consistent with the temporary increase in household long-term forecasts during the crisis.

More generally, the model predictions depart from observed expectations in two dimensions. First, the model falls short of capturing the sharp increase in expectations in the late 1970s, suggesting other factors beyond observed inflation surprises triggered such a large spike. Second, while the predicted long-term forecasts mimic reasonably closely the variation in long-term survey forecasts, by the end of the sample the model predicts long-term expectations to be somewhat below realized forecasts on average. This discrepancy can be traced to the role of oil prices in affecting Michigan Survey forecasts since the late 1990s. Indeed, the model can be used to gauge to what extent changes in energy prices have shifted long-term inflation expectations. Recall the short-term inflation forecasts have been corrected for the influence of oil price movements, but not the long-term forecasts. Conditional on the model being right, the difference between model predictions and the corresponding survey data therefore measures the effects of oil on long-term survey forecasts. We can then conclude that the effects of oil prices on average long-term forecasts over the 2000s are not particularly large.

9.2 PROFESSIONAL FORECASTS: OTHER COUNTRIES

The sample of international data comprises the following countries: France, Germany, Italy, Spain, Japan, Canada, Switzerland and Sweden. Inflation expectations are measured using data available from Consensus Economics. While the Consensus Economics data set includes short-term and long-term professional survey forecasts for a wide set of countries, it presents two challenges for estimation. First, forecasts in this data set are made on a year-over-year basis. As for the Michigan Survey, this formulation prevents a clean identification of the mechanism of the model, which links one-step-ahead forecast errors to the beliefs about long-term inflation, $\bar{\pi}_t$. Second, in contrast with the US forecast data, for most countries expectations data are only available from 1991, providing a limited time series for estimation.

the SPF are used, but the sample is shorter.

The following discussion details how each of these complications is handled.

Mapping model concepts to the data. Estimation employs available Consensus Economics inflation forecasts for the one- and two-years-ahead horizons. In contrast with SPF forecasts which have a constant horizon, independently of when they are polled, in the Consensus Economics data set the forecast horizon is shrinking in each quarter, and the target forecast becomes progressively less uncertain towards the end of the year. Since forecast horizons are different in each quarter, we have only one observation for each particular forecast per year.

To map the data concept into the model concept, year-over-year forecasts can be approximated as a weighted average of quarterly forecasts at different horizons, with tent-shaped weights.⁴¹ For the estimation we use six sets of forecasts. The first two are forecasts for the current year, made in the first and second quarter of the year. The remaining four are forecasts for the next year made in each quarter of the current year. Common to all these forecasts is most of the weight is given to quarterly forecasts ranging from one-to-four-quarters ahead. Details on the observation equation can be found in the Appendix. Similarly to the Michigan survey, these are reasonably interpreted as short-term forecasts. Figure 10 plots short-term forecasts with a horizon below one year and above one year respectively, along with the CPI for each country. These forecasts are fairly close and respond to current inflation developments.

Confronting a short sample. To handle the short available sample for the international data (survey data are available only from 1991), which prohibits sharp identification of the inflation trend using only CPI data (available from the mid-to-late 1950s), we employ posterior information from the US model to shape the priors adopted in estimation for these countries. In particular, the US posterior is used as a prior for all parameters except for the steady-state inflation rate and all observation errors. For these parameters, we use the same prior distributions as specified for the US. One final assumption is made. When simulating the posterior distribution of foreign parameters, the foreign likelihood function is scaled by the parameter λ^F so that

$$P^F(\Theta^F | Y_t^F, Y_t^{US}, \Theta^{US}) = L(Y_t^* | \Theta^{US}, \Theta^F)^{\lambda^F} L(Y_t^{US} | \Theta^{US}) p(\Theta^{US}) p(\Theta^F).$$

Smaller values of the parameter λ^F imply model predictions are more closely tied to the US posterior distribution. The presented results consider the case $\lambda^F = 0.05$. This weight delivers a posterior distribution of the common parameters that is extremely close to the distribution for the US model, while gaining some information about country-specific mean

⁴¹See Crump, Eusepi, and Moench (2015) for a discussion.

ANCHORED INFLATION EXPECTATIONS




Figure 10: Short-Term Surveys: OECD countries

These panels show the evolution of CPI inflation and short-term survey forecasts. The green line and the blue dots denote short-term survey forecasts of different horizons, from Consensus Economics.

inflation rates and observation errors.⁴² This provides a clear test of the out-of-sample forecasting power of the model estimated on US data.

Model predictions. Figure 11 presents model predictions for long-term inflation expectations.⁴³ For each country the top panel provides the model-implied long-term inflation forecast and 95% credible interval (the solid black line and gray band), along with the corresponding survey data (red dots). The dashed line provides the CPI inflation rate. The bottom panel displays the estimated gain, along with 50%, 70% and 95% credible intervals.

We offer the following observations. First, the model characterizes the evolution of long-term forecasts surprisingly well, despite model parameters being largely inferred from US data. For most of the sample, the survey-based forecasts are inside the 95% credible set. Second, while the precise timing differs, in general long-term expectations are more stable from the early-to-mid 1990s, with the estimated median gain declining for all countries relative to the 1980s. Perhaps not surprisingly, long-term inflation expectations in Canada follow similar dynamics to the US: the model places high probability on anchored expectations in recent years. However, several countries exhibit episodes in which expectations are poorly anchored. For example, Germany and Spain experience a temporary rise in the estimated gain after German reunification and during the Global Financial Crisis, respectively. Also, Italy, Germany and France display a high degree of uncertainty about the size of the gain during the crisis period, despite the median estimate being consistent with anchored expectations. At the end of the sample, all of these countries show some downside risk to long-term inflation expectations according to the 95% credible set.

Japan displays more prolonged periods of heightened sensitivity of long-term inflation expectations to short-term forecast errors, especially in the latter part of the sample. This is reflected in both the median estimate and also the 95% credible set which provides a metric of the risks that long-term expectations are not well anchored. Despite several years of deflation, both predicted and actual long-term inflation expectations remain above zero. Moreover, they appear to revert to about 2% over the latter part of the sample. Similarly to the US, this suggests the observed behavior of long-run expectations does not simply reflect inertia in survey responses by forecasters; rather, it is consistent with the observed pattern of surprises. Finally, Switzerland provides a second example of an economy for which our model implies unanchored expectations. However, here the model predictions and observed forecast data are in stark contrast, as survey-based long-term forecasts have remained fairly stable despite persistent forecast errors.

⁴²Details can be found in the appendix.

⁴³The appendix documents the evolution of short-term expectations used for the estimation and prediction, together with the results for two additional countries: Netherlands and Norway.

This serves to introduce our final observation. The comparison between the model-predicted paths and survey-based forecasts suggests that in some countries short-term forecast errors might not have been the only determinant of long-term expectations. For example, for much of the sample Japan had higher reported long-term inflation expectations than predicted by the model. This possibly reflects various fiscal and monetary interventions implemented to combat deflation over the period, or, more recently, concerns about long-term fiscal sustainability. Sweden displays a faster decline in survey-based long-term forecasts compared to model predictions in the early 1990s. This coincides with the announcement and adoption of an inflation targeting regime in 1992 – 1995. And in the case of Switzerland, which displays the largest discrepancy over the final years of the sample, the Swiss National Bank operated a formal exchange rate policy to limit appreciation of the Swiss Franc. This might have granted the policy framework an independent source of credibility despite short-term inflation behavior.

In each of these cases, announcement effects and other policy-related factors might have plausibly shifted expectations beyond what is justified by short-term inflation forecast errors alone. And in practice, it will almost certainly be the case that firms and other market participants will condition long-term inflation forecasts on a range of information, not least other indicators of the state of the macro-economy, such as short-term interest rates and output. What is remarkable about our results, for both the US and other countries, is that measures of short-term inflation expectations go a long way in explaining the dynamics of long-term inflation expectations. They don't explain everything, but they clearly have strong predictive content for longer-term developments.


10 CONCLUSIONS

The introduction motivated our theory of anchored inflation expectations with three questions: Are inflation expectations anchored? Will chronic undershooting of inflation targets lead to downward drift and un-anchoring of long-term expectations? How can we reconcile large negative output gaps and stable inflation over the past decade, with positive output gaps and high inflation in the 1970s? We now have answers. Expectations are anchored when agents hold beliefs that are consistent with the policy strategy of the central bank. Our empirical work adduces clear evidence that expectations were un-anchored prior to the late 1990s. When agents doubt there is a constant inflation target, beliefs are un-anchored and highly sensitive to forecast errors. This is the source of trend inflation. Doubt arises from systematic forecast errors of either sign. In same way the chronic overshooting of inflation targets led to the unstable inflation expectations and the Great Inflation, chronic undershooting may lead to un-anchoring and downward movement in long-term inflation ex-

pectations. Finally, an implication of anchored long-term expectations is that large forecast errors will not lead to mark-ups or mark-downs in long-term inflation expectations. The different experiences of the 1970s and 2000s reflects the success of the Federal Reserve in anchoring expectations, so that large output gaps and large inflation forecast errors had little effects on inflation expectations.

Figure 11: Selected OECD countries

These panels show model predictions for long-term inflation forecasts (top) and the learning gain (bottom). Black solid lines denote median; gray areas measure 50th, 70th and 95th credible intervals; red dots denote five-to-ten years inflation forecasts from Consensus Economics.



REFERENCES

- ALBANESI, S., V. V. CHARI, AND L. J. CHRISTIANO (2003): “Expectation Traps and Monetary Policy,” *Review of Economic Studies*, 70(4), 715–741.
- BEECHEY, M. J., B. K. JOHANNSEN, AND A. T. LEVIN (2011): “Are Long-Run Inflation Expectations Anchored More Firmly in the Euro Area Than in the United States?,” *American Economic Journal: Macroeconomics*, 3(2), 104–129.
- BIANCHI, F., AND L. MELOSI (2017): “Escaping the Great Recession,” *American Economic Review*, 17(4), 1030–58.
- BOMFIM, A. N., AND G. D. RUDEBUSCH (2000): “Opportunistic and Deliberate Disinflation under Imperfect Credibility,” *Journal of Money, Credit and Banking*, 32(4), 707–721.
- BRANCH, W. A., AND G. W. EVANS (2007): “Model uncertainty and endogenous volatility,” *Review of Economic Dynamics*, 10(2), 207–237.
- BROWN, R. L., J. DURBIN, AND J. M. EVANS (1975): “Techniques for Testing the Constancy of Regression Relationships over Time,” *Journal of the Royal Statistical Society. Series B (Methodological)*, 37(2), 149–192.
- CHAN, J. C. C., T. E. CLARK, AND G. KOOP (2015): “A New Model of Inflation, Trend Inflation, and Long-Run Inflation Expectations,” Working Paper 1520, Federal Reserve Bank of Cleveland.
- CHO, I.-K., AND K. KASA (2015): “Learning and Model Validation,” *Review of Economic Studies*, 82(1), 45–82.
- (2017): “Gresham’s Law of Model Averaging,” *American Economic Review*, 107(11), 3589–3616.
- CLARIDA, R., J. GALI, AND M. GERTLER (2000): “Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory,” *Quarterly Journal of Economics*, 115, 147–180.
- COGLEY, T., G. E. PRIMICERI, AND T. J. SARGENT (2010): “Inflation-Gap Persistence in the US,” *American Economic Journal: Macroeconomics*, 2(1), 43–69.
- COGLEY, T., AND T. J. SARGENT (2005): “Drifts and volatilities: monetary policies and outcomes in the post WWII US,” *Review of Economic Dynamics*, 8(2), 262 – 302.

- COGLEY, T., AND A. M. SBORDONE (2008): “Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve,” *American Economic Review*, 98(5), 2101–26.
- COIBION, O., AND Y. GORODNICHENKO (2012): “What Can Survey Forecasts Tell Us about Information Rigidities?,” *Journal of Political Economy*, 120(1), 116–159.
- CORNEA-MADEIRA, A., C. HOMMES, AND D. MASSARO (2019): “Behavioral Heterogeneity in U.S. Inflation Dynamics,” *Journal of Business & Economic Statistics*, 37(2), 288–300.
- CRUMP, R., S. EUSEPI, AND E. MOENCH (2015): “The Term Structure of Expectations and Bond Yields,” unpublished manuscript, Federal Reserve Bank of New York.
- DAVIG, T., AND E. LEEPER (2006): “Fluctuating Macro Policies and the Fiscal Theory,” in *NBER Macroeconomics Annual*, ed. by D. Acemoglu, K. Rogoff, and M. Woodford, vol. 21.
- DEL NEGRO, M., AND S. EUSEPI (2011): “Fitting Observed Inflation Expectations,” *Journal of Economic Dynamics and Control*, 35(12), 2105–2131.
- DEL NEGRO, M., M. P. GIANNONI, AND F. SCHORFHEIDE (2015): “Inflation in the Great Recession and New Keynesian Models,” *American Economic Journal: Macroeconomics*, 7(1), 168–196.
- ERCEG, C. J., AND A. T. LEVIN (2003): “Imperfect Credibility and Inflation Persistence,” *Journal of Monetary Economics*, 50(4), 915–944.
- EUSEPI, S. (2005): “Central Bank Transparency under Model Uncertainty,” FRBNY Working Paper.
- EUSEPI, S., M. P. GIANNONI, AND B. PRESTON (2018): “Some implications of learning for price stability,” *European Economic Review*, 106(C), 1–20.
- EUSEPI, S., AND B. PRESTON (2010): “Central Bank Communication and Macroeconomic Stabilization,” *American Economic Journal: Macroeconomics*, 2, 235–271.
- (2018a): “Fiscal Foundations of Inflation: Imperfect Knowledge,” *American Economic Review*, 108(9), 2551–2589.
- (2018b): “The Science of Monetary Policy: An Imperfect Knowledge Perspective,” *Journal of Economic Literature*, 56(1), 3–59.

- EVANS, G. W., AND S. HONKAPOHJA (2001): *Learning and Expectations in Economics*. Princeton University Press.
- EVANS, G. W., S. HONKAPOHJA, T. SARGENT, AND N. WILLIAMS (2012): “Bayesian Model Averaging, Learning and Model Selection,” CDMA Working Paper Series 201203, Centre for Dynamic Macroeconomic Analysis.
- FERNANDEZ-VILLAVARDE, J., AND J. F. RUBIO-RAMIREZ (2007): “Estimating Macroeconomic Models: A Likelihood Approach,” *Review of Economic Studies*, 74(4), 1059–1087.
- GIANNONI, M. P., AND M. WOODFORD (2002): “Optimal Interest-Rate Rules: I. General Theory,” NBER Working Paper no. 9419.
- GIBBS, C., AND M. KULISH (2015): “Disinflations in a Model of Imperfectly Anchored Expectations,” unpublished manuscript, University of New South Wales.
- GODSILL, S. J., A. DOUCET, AND M. WEST (2004): “Monte Carlo Smoothing for Nonlinear Time Series,” *Journal of the American Statistical Association*, 99(465), 156–168.
- GURKAYNAK, R. S., A. LEVIN, AND E. SWANSON (2010): “Does Inflation Targeting Anchor Long-Run Inflation Expectations? Evidence from the U.S., UK, and Sweden,” *Journal of the European Economic Association*, 8(6), 1208–1242.
- HALL, R. E. (2011): “The Long Slump,” *American Economic Review*, 101(2), 431–469.
- HAMILTON, J. D. (1994): *Time Series Analysis*. Princeton University Press.
- HOL, J. D., T. B. SCHÖN, AND F. GUSTAFSSON (2006): “On Resampling Algorithms for Particle Filters,” in *Nonlinear Statistical Signal Processing Workshop, 2006 IEEE*, pp. 79–82. IEEE.
- HOMMES, C. H., AND J. LUSTENHOUWER (2015): “Inflation Targeting and Liquidity Traps under Endogenous Credibility,” unpublished, University of Amsterdam.
- HONKAPOHJA, S., AND G. W. EVANS (1993): “Adaptive forecasts, hysteresis, and endogenous fluctuations,” *Economic Review*, pp. 3–13.
- IRELAND, P. N. (2007): “Changes in the Federal Reserve’s Inflation Target: Causes and Consequences,” *Journal of Money, Credit and Banking*, 39(8), 1851–1882.
- KITAGAWA, G. (1996): “Monte Carlo Filter and Smoother for Non-Gaussian Nonlinear State Space Models,” *Journal of Computational and Graphical Statistics*, 5, 1–25.

- KOSTYSHYNA, O. (2012): “Application of an Adaptive Step-Size Algorithm in Models of Hyperinflation,” *Macroeconomic Dynamics*, 16, 355–375.
- KOZICKI, S., AND P. A. TINSLEY (2001): “Shifting Endpoints in the Term Structure of Interest Rates,” *Journal of Monetary Economics*, 47(3), 613–652.
- (2005): “Permanent and Transitory Policy Shocks in an Empirical Macro Model with Asymmetric Information,” *Journal of Economic Dynamics and Control*, 29(11), 1985–2015.
- (2012): “Effective Use of Survey Information in Estimating the Evolution of Expected Inflation,” *Journal of Money, Credit and Banking*, 44(1), 145–169.
- KREPS, D. (1998): “Anticipated Utility and Dynamic Choice,” in *Frontiers of Research in Economic Theory*, ed. by D. Jacobs, E. Kalai, and M. Kamien, pp. 242–274. Cambridge: Cambridge University Press.
- KUSHNER, H. J., AND G. G. YIN (2003): *Stochastic Approximation and Recursive Algorithms and Applications*. New York: Springer-Verlag, 2nd edn.
- LANSING, K. (2009): “Time Varying U.S. Inflation Dynamics and the New Keynesian Phillips Curve,” *Review of Economic Dynamics*, 12(2), 304–326.
- LINDSTEN, F., AND T. B. SCHÖN (2013): “Backward Simulation Methods for Monte Carlo Statistical Inference,” *Foundations and Trends in Machine Learning*, 6(1), 1–143.
- LUBIK, T. A., AND F. SCHORFHEIDE (2004): “Testing for Indeterminacy: An Application to U.S. Monetary Policy,” *American Economic Review*, 94, 190–217.
- MALMENDIER, U., AND S. NAGEL (2016): “Learning from Inflation Experiences,” *The Quarterly Journal of Economics*, 131(1), 53–87.
- MARCET, A., AND J. P. NICOLINI (2003): “Recurrent Hyperinflations and Learning,” *American Economic Review*, 93(5), 1476–1498.
- MARCET, A., AND T. J. SARGENT (1989): “Convergence of Least-Squares Learning in Environments with Hidden State Variables and Private Information,” *Journal of Political Economy*, pp. 1306–1322.
- MERTENS, E. (2016): “Measuring the Level and Uncertainty of Trend Inflation,” *The Review of Economics and Statistics*, 98(5), 950–967.

- MILANI, F. (2007): “Expectations, Learning and Macroeconomic Persistence,” *Journal of Monetary Economics*, 54, 2065–2082.
- (2014): “Learning and Time-varying Macroeconomic Volatility,” *Journal of Economic Dynamics and Control*, 47(C), 94–114.
- ORPHANIDES, A. (2001): “Monetary Policy Rules Based on Real-Time Data,” *American Economic Review*, 91(4), 964–985.
- ORPHANIDES, A., AND J. C. WILLIAMS (2005): “Imperfect Knowledge, Inflation Expectations, and Monetary Policy,” in *The Inflation Targeting Debate*, ed. by B. S. Bernanke, and M. Woodford. University of Chicago Press.
- PRESTON, B. (2005): “Learning About Monetary Policy Rules when Long-Horizon Expectations Matter,” *International Journal of Central Banking*, 1(2), 81–126.
- (2006): “Adaptive Learning, Forecast-Based Instrument Rules and Monetary Policy,” *Journal of Monetary Economics*, 53, 507–535.
- PRIMICERI, G. E. (2005): “Why Inflation Rose and Fell: Policymakers’ Beliefs and US Postwar Stabilization Policy,” *The Quarterly Journal of Economics*, 121, 867–901.
- ROTEMBERG, J. J. (1982): “Monopolistic Price Adjustment and Aggregate Output,” *Review of Economic Studies*, 64, 517–31.
- SARGENT, T., N. WILLIAMS, AND T. ZHA (2006): “Shocks and Government Beliefs: The Rise and Fall of American Inflation,” *American Economic Review*, 96(4), 1193–1224.
- SARGENT, T. J. (1999): *The Conquest of American Inflation*. Princeton University Press.
- SCHÖN, T., F. GUSTAFSSON, AND P.-J. NORDLUND (2005): “Marginalized Particle Filters for Mixed Linear/Non-linear State-Space Models,” *IEEE Transactions on Signal Processing*, 53(7), 2279–2289.
- SLOBODYAN, S., AND R. WOUTERS (2012): “Learning in an Estimated Medium-scale DSGE Model,” *Journal of Economic Dynamics and Control*, 36(1), 26–46.
- SMETS, F., AND R. WOUTERS (2003): “An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area,” *Journal of the European Economic Association*, 1(5), 1123–1175.
- STOCK, J. H., AND M. W. WATSON (2007): “Why Has U.S. Inflation Become Harder to Forecast?,” *Journal of Money, Credit and Banking*, 39(s1), 3–33.

WOODFORD, M. (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*.
Princeton University Press.

A APPENDIX (NOT FOR PUBLICATION)

Here we present the marginalized particle filter and smoother. For ease of notation note that we use φ_t for $\tilde{\varphi}_t$, and (ε_t, μ_t) for $(\tilde{\varepsilon}_t, \tilde{\mu}_t)$ in the main text. Recall the model is summarized by the following equations

$$\begin{aligned}\pi_t &= (1 - \gamma) \Gamma \bar{\pi}_t + \gamma \pi_{t-1} + \varphi_t + \mu_t \\ \bar{\pi}_t &= \bar{\pi}_{t-1} + \mathbf{k}_t^{-1} \times f_{t-1} \\ \mathbf{k}_t &= I(\bar{\pi}_{t-1}) \times (\mathbf{k}_{t-1} + 1) + (1 - I(\bar{\pi}_{t-1})) \times \bar{g}^{-1} \\ f_t &= (1 - \gamma) (\Gamma - 1) \bar{\pi}_t + \mu_t + \epsilon_t \\ \varphi_t &= \rho_\varphi \varphi_{t-1} + \epsilon_t,\end{aligned}$$

where the function $I(\bar{\pi}_t)$ is described as

$$I(\bar{\pi}) = \begin{cases} 1, & \text{if } |(1 - \gamma) (\Gamma - 1) \bar{\pi}| \leq \bar{\theta} \sigma^\eta \\ 0, & \text{otherwise.} \end{cases}$$

The model can then be re-written as

$$\begin{aligned}\mathbf{k}_t &= f_{\mathbf{k}}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) \\ \bar{\pi}_t &= f_{\bar{\pi}}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) + f_{\mathbf{k}}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1})^{-1} \times \eta_{t-1} \\ \eta_t &= \mu_t + \epsilon_t \\ \varphi_t &= \rho_\varphi \varphi_{t-1} + \epsilon_t \\ \pi_t &= (1 - \gamma_p) \Gamma f_{\bar{\pi}}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) + (1 - \gamma_p) \Gamma f_{\mathbf{k}}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1})^{-1} \eta_{t-1} \\ &\quad + \gamma \pi_{t-1} + \rho_\varphi \varphi_{t-1} + \epsilon_t + \mu_t.\end{aligned}$$

where

$$f_{\mathbf{k}}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) = I(\bar{\pi}_{t-1}) \times (\mathbf{k}_{t-1} + 1) + (1 - I(\bar{\pi}_{t-1})) \times \bar{g}^{-1},$$

$$f_{\bar{\pi}}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) = [1 - (1 - \Gamma)(1 - \gamma) f_{\mathbf{k}}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1})^{-1}] \bar{\pi}_{t-1}.$$

We can also re-write the system in matrix notation. One way to write it is by separating linear and nonlinear states. For the linear variables we have:

$$\xi_t = f_{\xi}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) + A_{\xi}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) \xi_{t-1} + S_{\xi} \begin{bmatrix} \epsilon_t \\ \mu_t \end{bmatrix},$$

where

$$\xi_t = \begin{bmatrix} \eta_t \\ s_t \\ \pi_t \end{bmatrix};$$

$$f_{\xi}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) = \begin{bmatrix} \mathbf{0}_{2 \times 1} \\ (1 - \gamma) \Gamma f_{\bar{\pi}}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) \end{bmatrix};$$

$$A_{\xi}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & & \rho_{\varphi} & 0 \\ (1 - \gamma) \Gamma f_{\mathbf{k}}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1})^{-1} & & \rho_{\varphi} & \gamma \end{bmatrix};$$

$$S_{\xi} = \begin{bmatrix} 1 & 1 \\ 1 & 0 \\ 1 & 1 \end{bmatrix}.$$

For the nonlinear variables we can express

$$\mathbf{k}_t = f_{\mathbf{k}}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1})$$

and

$$\bar{\pi}_t = f_{\bar{\pi}}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) + A_{\bar{\pi}}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) \xi_{t-1}.$$

where

$$A_{\bar{\pi}}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) = \begin{bmatrix} f_{\mathbf{k}}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1})^{-1} \\ \mathbf{0}_{2 \times 1} \end{bmatrix}'.$$

Notice that \mathbf{k}_t does not depend on the linear state. In yet another formulation we can express the system in more compact notation:

$$\mathbf{k}_t = f_{\mathbf{k}}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1})$$

$$\begin{bmatrix} \bar{\pi}_t \\ \xi_t \end{bmatrix} = f(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) + A(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) \xi_{t-1} + \begin{bmatrix} 0 \\ S_\xi \end{bmatrix} \begin{bmatrix} \epsilon_t \\ \mu_t \end{bmatrix}$$

where

$$f(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) = \begin{bmatrix} f_{\bar{\pi}}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) \\ f_{\xi}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) \end{bmatrix}$$

$$A(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) = \begin{bmatrix} A_{\bar{\pi}}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) \\ A_{\xi}(\bar{\pi}_{t-1}, \mathbf{k}_{t-1}) \end{bmatrix}$$

and

$$\Sigma = E \left(\begin{bmatrix} \epsilon_t \\ \mu_t \end{bmatrix} \begin{bmatrix} \epsilon_t \\ \mu_t \end{bmatrix}' \right)$$

is the variance covariance of the innovations.

This notation is used below when computing the smoothed states. Finally, given the data $Y_T = y_1 \dots y_T$, the model observation equation is

$$y_t = h_{0,t} + h_{\bar{\pi},t} \bar{\pi}_t + H'_t \xi_t + R_t^{1/2} e_t^o$$

where the vectors and matrices h_0 , $h_{\bar{\pi}}$, H'_t and R_t are defined to be consistent with the timing of available data, and e_t denotes observation errors.

A.1 ALGORITHM FOR THE MARGINALIZED PARTICLE FILTER

This follows Schön, Gustafsson, and Nordlund (2005). For details of the particle filter we use Kitagawa (1996). We are looking for the following distributions:

$$p(\xi_t, [\bar{\pi}_t, \mathbf{k}_t] | Y_t) = p(\xi_t | [\bar{\pi}_t, \mathbf{k}_t], Y_t) \times p([\bar{\pi}_t, \mathbf{k}_t] | Y_t).$$

The following describes the algorithm. Discussion and proofs are given below.

Algorithm:

1. Initialization. Choose $\bar{\pi}_{1|0}^{(i)}$, $\mathbf{k}_{1|0}^{(i)}$ from some distributions (drawing from normal for $\bar{\pi}$ and $\mathbf{k}_{0|0}^{(i)} = \bar{\mathbf{k}}_0$ (or draw from $U(0, \bar{g}^{-1})$), and $\xi_{1|0}^{(i)}$, $P_{1|0}^{(i)} = [\xi_0, P_{1|0}]$, where $P_{1|0}$ denotes the

initial precision matrix in the linear Kalman filter.

2. For each $t = 1 \dots T$, compute

$$\Omega_t = H_t' P_{t|t-1} H_t + R_t$$

and its inverse. For $i = 1, \dots, N$, evaluate the importance weights

$$q_t^{(i)} = p\left(y_t | \bar{\pi}_{t|t-1}^{(i)}, \xi_{t|t-1}^{(i)}\right).$$

In order to do this, use

$$p\left(y_t | \bar{\pi}_{t|t-1}^{(i)}, \xi_{t|t-1}^{(i)}\right) = N\left(h_{0,t} + h_{\bar{\pi},t} \bar{\pi}_{t|t-1}^{(i)} + H_t' \xi_{t|t-1}^{(i)}, H_t' P_{t|t-1} H_t + R_t\right)$$

so that

$$q_t^{(i)} = w_{t-1}^{(i)} \times |\Omega_t|^{-1/2} \times$$

$$\exp\left\{-\frac{1}{2} \left(y_t - h_{0,t} - h_{\bar{\pi},t} \bar{\pi}_{t|t-1}^{(i)} - H_t' \xi_{t|t-1}^{(i)}\right)' \times \Omega_t^{-1} \times \left(y_t - h_{0,t} - h_{\bar{\pi},t} \bar{\pi}_{t|t-1}^{(i)} - H_t' \xi_{t|t-1}^{(i)}\right)\right\}.$$

where $w_{t-1}^{(i)}$ denotes the particle weight from the previous period. (In the expression above we eliminate the constant coefficient that is independent of (i) and the model parameters.)

3. Re-sampling.⁴⁴ Provided the number of effective particles (effective sample size), computed as

$$ESS = \frac{1}{\sum \left(w_t^{(j)}\right)^2},$$

falls below the threshold ($ESS < 0.75 * N$) we re-sample such that

$$p\left(\left[\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)}\right] = \left[\bar{\pi}_{t|t-1}^{(j)}, \mathbf{k}_{t|t-1}^{(j)}\right]\right) = \frac{q_t^{(j)}}{\sum q_t^{(j)}}.$$

Here we use **systematic resampling**: see Kitagawa (1996), Hol, Schön, and Gustafsson (2006) for a discussion of resampling and different methods. The outcome of systematic resampling is a discrete distribution with particles $\left\{\bar{\pi}_{t|t}^{(k)}, \mathbf{k}_{t|t}^{(k)}\right\}_{k=1}^N$ and corresponding weights $w_t(i) = 1/N$ for $i = 1, \dots, N$. In case of not resampling the weights are $w_t(i) = q_t^{(j)} / \sum q_t^{(j)}$.

⁴⁴This means (roughly speaking) increasing the number of particles receiving high weight $\left(\frac{q_t^{(j)}}{\sum q_t^{(j)}}\right)$ and eliminating particles with very low weight, while keeping the number of particles equal to N .

4. Linear measurement equation: for $i = 1, \dots, N$, evaluate

$$\xi_{t|t}^{(i)} = \xi_{t|t-1}^{(i)} + K_t \left(y_t - h_{0,t} - h_{\bar{\pi},t} \bar{\pi}_{t|t}^{(i)} - H_t' \xi_{t|t-1}^{(i)} \right)$$

$$K_t = P_{t|t-1} H_t \Omega_t^{-1}$$

$$P_{t|t} = P_{t|t-1} - K_t H_t' P_{t|t-1}$$

5. Particle filter prediction. For $i = 1, \dots, N$, compute

$$\mathbf{k}_{t+1|t}^{(i)} = f_{\mathbf{k}}(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)})$$

and then draw $\bar{\pi}_{t+1|t}^{(i)}$ from distribution

$$\begin{aligned} & p \left(\bar{\pi}_{t+1|t} | Y_t, \bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)} \right) \\ &= N \left(f_{\bar{\pi}}(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)}) + f_{\mathbf{k}}(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)})^{-1} z_{t|t}^{(i)}, f_{\mathbf{k}}(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)})^{-2} P_{t|t}^{[\eta,\eta]} \right) \end{aligned}$$

where we use the notation: $P_{t|t}^{[x,z]} = P_{t|t}(x, z)$.

6. Linear model prediction

$$\tilde{\xi}_{t|t}^{(i)} = \xi_{t|t}^{(i)} + \tilde{K}_t^{(i)} \left(\bar{\pi}_{t+1|t}^{(i)} - f_{\bar{\pi}}(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)}) - f_{\mathbf{k}}(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)})^{-1} z_{t|t}^{(i)} \right)$$

$$\xi_{t+1|t}^{(i)} = f_{\xi}(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)}) + A_{\xi}(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)}) \tilde{\xi}_{t|t}^{(i)}$$

$$P_{t+1|t} = Q_{\xi} + \tilde{P}_{t|t}; \quad Q_{\xi} = S_{\xi} \Sigma S_{\xi}';$$

where

$$\begin{aligned} \tilde{K}_t^{(i)} &= P_{t|t} A_{\bar{\pi}}' \left(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)} \right) \left(A_{\bar{\pi}} \left(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)} \right) P_{t|t} A_{\bar{\pi}}' \left(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)} \right) \right)^{-1} \\ &= f_{\mathbf{k}}(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)}) \begin{bmatrix} 1 \\ P_{t|t}^{[\eta,\varphi]} / P_{t|t}^{[\eta,\eta]} \\ P_{t|t}^{[\eta,\pi]} / P_{t|t}^{[\eta,\eta]} \end{bmatrix}; \end{aligned}$$

and

$$A_l(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)}) \tilde{K}_t^{(i)} = f_{\mathbf{k}}(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)}) \begin{bmatrix} 0 \\ \frac{P_{t|t}^{[\eta,s]}}{P_{t|t}^{[\eta,\eta]}} \rho_\varphi \\ \frac{P_{t|t}^{[\eta,\pi]}}{P_{t|t}^{[\eta,\eta]}} \gamma + \frac{P_{t|t}^{[\eta,s]}}{P_{t|t}^{[\eta,\eta]}} \rho_\varphi + f_{\mathbf{k}}(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)})^{-1} (1 - \gamma) \Gamma \end{bmatrix}$$

$$\tilde{P}_{t|t} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & \iota_1 & \iota_2 + \iota_1 \\ 0 & \iota_2 + \iota_1 & 2\iota_2 + \iota_1 + \left(P_{t|t}^{[\pi,\pi]} - \frac{(P_{t|t}^{[\eta,\pi]})^2}{P_{t|t}^{[\eta,\eta]}} \right) \gamma^2 \end{bmatrix}$$

where

$$\iota_1 = \left(P_{t|t}^{[\varphi,\varphi]} - \frac{(P_{t|t}^{[\eta,\varphi]})^2}{P_{t|t}^{[\eta,\eta]}} \right) \rho_\varphi^2;$$

$$\iota_2 = \left(-\frac{P_{t|t}^{[\eta,\varphi]} P_{t|t}^{[\eta,\pi]}}{P_{t|t}^{[\eta,\eta]}} + P_{t|t}^{[s,\pi]} \right) \rho_\varphi \gamma.$$

Notice that, importantly, $P_{t+1|t}$ is independent of particles. This is key for a fast evaluation of the likelihood.

Finally, the log-Likelihood is approximated by

$$L(\cdot) = \sum_{t=1}^T \log p(y_t | Y_{t-1})$$

where

$$\begin{aligned} p(y_t | Y_{t-1}) &= p(y_t | \xi_t, [\bar{\pi}_t, \mathbf{k}_t]) p(\xi_t, [\bar{\pi}_t, \mathbf{k}_t] | Y_{t-1}) \\ &= p(y_t | \xi_t, [\bar{\pi}_t, \mathbf{k}_t]) p(\xi_t | [\bar{\pi}_t, \mathbf{k}_t], Y_{t-1}) p([\bar{\pi}_t, \mathbf{k}_t] | Y_{t-1}), \\ &\rightarrow \\ L(\cdot) &\simeq \sum_{t=1}^T \log \left(\sum_{i=1}^N q_t^{(i)} \right). \end{aligned}$$

Algorithm ends.

In the estimation performed in this paper we set the number of particles $N = 2500$; To

avoid injecting randomness in the calculation of the likelihood, the ‘‘chatter’’ of changing random numbers, we keep the simulated (standardized) innovations constant as we evaluate different parameters—see the discussion in Fernandez-Villaverde and Rubio-Ramirez (2007). In detail, we fix the following innovations: random initial conditions for the nonlinear state variables; random draws to compute shocks in the nonlinear prediction step; and random draws in the resampling step.

B MARGINALIZED SMOOTHER

We follow Lindsten and Schön (2013) using the ‘joint backward simulation’ Rao-Blackwellised particle smoother. See also Godsill, Doucet, and West (2004). We compute a smoothed path for the states, conditional on a parameter draw, for the sample $t = 1 \dots T$. The algorithm, in conjunction with the forward filter above, allows producing the full distribution of state and parameters using Carter and Khon (1994).

The objective is to draw $j = 1 \dots M$ trajectories of the model variables $\left\{ \tilde{\pi}_{t|T}^{(j)}, \tilde{\mathbf{k}}_{t|T}^{(j)}, \tilde{\xi}_{t|T}^{(j)} \right\}_{t=1}^T$. The forward filter allows drawing $\tilde{\pi}_{T|T}^{(j)}, \tilde{\mathbf{k}}_{T|T}^{(j)}$ from the empirical distribution of $\left\{ \bar{\pi}_{t|t}^{(k)}, \mathbf{k}_{t|t}^{(k)} \right\}_{k=1}^N$ where each particle has weight $w_t^{(k)}$. Moreover, conditional on the draw $\tilde{\pi}_{t|T}^{(j)}, \tilde{\mathbf{k}}_{T|T}^{(j)}$, it allows drawing the linear state $\tilde{\xi}_{T|T}^{(j)}$ from the normal distribution $N\left(\xi_{T|T}^{(j)}, P_{T|T}\right)$. Given this we compute M paths as follows.

Algorithm:

For $t = T - 1 : -1 : 1$

 For each $j = 1 \dots M$

 For each $i = 1 \dots N$, compute

$$w_{t|t+1}^j(i) = \frac{w_t(i) p\left(\tilde{\pi}_{t+1|t+1}^{(j)}, \tilde{\mathbf{k}}_{t+1|t+1}^{(j)}, \tilde{\xi}_{t+1|t+1}^{(j)} \mid \bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)}, Y_t\right)}{\sum_k^N w_t(k) p\left(\tilde{\pi}_{t+1|t+1}^{(j)}, \tilde{\mathbf{k}}_{t+1|t+1}^{(j)}, \tilde{\xi}_{t+1|t+1}^{(j)} \mid \bar{\pi}_{t|t}^{(k)}, \mathbf{k}_{t|t}^{(k)}, Y_t\right)}$$

where the last line makes use of the fact that $w_t(i) = 1/N$ because of resampling in the forward filter. The probability distribution above can be expressed as

$$p\left(\bar{\pi}_{t+1|T}, \mathbf{k}_{t+1|T}, \xi_{t+1|T} \mid \bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)}, Y_t\right) =$$

$$p\left(\bar{\pi}_{t+1|T}, \xi_{t+1|T} \mid \bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)}, Y_t\right) \times p\left(\mathbf{k}_{t+1|T} \mid \bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)}, Y_t\right),$$

which uses $\mathbf{k}_{t+1|t} = f_{\mathbf{k}}(\bar{\pi}_t, \mathbf{k}_t)$. We can then evaluate

$$p \left(\tilde{\pi}_{t+1|T}^{(j)}, \tilde{\xi}_{t+1|T}^{(j)} \mid \bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)}, Y_t \right) \times \mathbf{1}_{\tilde{\mathbf{k}}_{t+1}^{(j)} = f_{\mathbf{k}}(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)})} =$$

$$\begin{cases} p_{t|t+1}^{j,i}, & \text{if } \tilde{\mathbf{k}}_{t+1|T}^{(j)} = f_{\mathbf{k}} \left(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)} \right) \\ \mathbf{0}, & \text{otherwise} \end{cases}$$

where

$$p_{t|t+1}^{j,i} = \propto \exp \left\{ -\frac{1}{2} \ln \left(\left| \tilde{\Omega}_t^{(i)} \right| \right) - \frac{1}{2} \left(\eta_{t+1}^{j,i} \right)' \times \left(\tilde{\Omega}_t^{(i)} \right)^{-1} \left(\eta_{t+1}^{j,i} \right) \right\}$$

$$\eta_{t+1}^{j,i} = \begin{bmatrix} \tilde{\pi}_{t+1|t+1}^{(j)} \\ \tilde{\xi}_{t+1|t+1}^{(j)} \end{bmatrix} - \left[f \left(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)} \right) + A \left(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)} \right) \xi_{t|t}^{(i)} \right].$$

$$\tilde{\Omega}_t^{(i)} = Q + A \left(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)} \right) P_{t|t} A \left(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)} \right)'$$

$$Q = \begin{bmatrix} 0 & 0 \\ 0 & S_{\xi} \Sigma S_{\xi}' \end{bmatrix}.$$

and where we can use

$$\mathbf{1}_{\tilde{\mathbf{k}}_{t+1}^{(j)} = f_{\mathbf{k}}(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)})} =$$

$$\mathbf{1}_{\tilde{\mathbf{k}}_{t+1}^{(j)} = \bar{g}} \times \mathbf{1}_{\left| (1-\gamma)(\Gamma-1)\bar{\pi}_{t|t}^{(i)} \right| > \bar{\theta}\sigma^n} +$$

$$+ \left(1 - \mathbf{1}_{\tilde{\mathbf{k}}_{t+1}^{(j)} = \bar{g}} \right) \times \left(\mathbf{1}_{\left| (1-\gamma)(\Gamma-1)\bar{\pi}_{t|t}^{(i)} \right| \leq \bar{\theta}\sigma^n} \right) \times \mathbf{1}_{\mathbf{k}_{t|t}^{(i)} = \mathbf{k}_{t+1|t}^{(j)} - 1}.$$

Moving to the linear variables ξ_t , for each $j = 1 \dots M$, draw the nonlinear variables $\tilde{\pi}_{t|T}^{(j)}, \tilde{\mathbf{k}}_{t|T}^{(j)}$, from $\left\{ \bar{\pi}_{t|t}^{(k)}, \mathbf{k}_{t|t}^{(k)} \right\}_{k=1}^N$ using the new set of weights $\left\{ w_{t|t+1}^j(k) \right\}_{k=1}^N$. Conditional on the draw, sample from

$$p \left(\xi_t \mid \bar{\pi}_{1:t}^j, \mathbf{k}_{1:t}^j, \xi_{t+1}, \bar{\pi}_{t+1}, Y_t \right).$$

In particular we draw $\tilde{\xi}_{t|T}^{(j)}$ from the distribution

$$N \left(\xi_{t|t}^{(j)} + \Delta_t^{(j)} \left(\left[\bar{\pi}_{t+1|T}^{(j)}, \tilde{\xi}_{t+1|T}^{(j)} \right]' - f \left(\bar{\pi}_{t|T}^{(j)}, \tilde{\xi}_{t|T}^{(j)} \right) - A \left(\bar{\pi}_{t|T}^{(j)}, \tilde{\xi}_{t|T}^{(j)} \right) \xi_{t|t}^{(j)} \right), \Lambda_{t|t}^{(j)} \right)$$

where $\xi_{t|t}^{(j)}$ is the element in $\left\{ \tilde{\pi}_{t|t}^{(k)}, \tilde{\mathbf{k}}_{t|t}^{(k)} \right\}_{k=1}^N$ that corresponds to the same draw j from which the particles $\tilde{\pi}_{t|T}^{(j)}, \tilde{\mathbf{k}}_{t|T}^{(j)}$ are obtained, and where

$$\begin{aligned} \Delta_t^{(j)} &= P_{t|t} A \left(\tilde{\pi}_{t|T}^{(j)}, \tilde{\mathbf{k}}_{t|T}^{(j)} \right)' \left(Q + A \left(\tilde{\pi}_{t|T}^{(j)}, \tilde{\mathbf{k}}_{t|T}^{(j)} \right) P_{t|t} A \left(\tilde{\pi}_{t|T}^{(j)}, \tilde{\mathbf{k}}_{t|T}^{(j)} \right)' \right)^{-1} \\ \Lambda_{t|t}^{(j)} &= P_{t|t} - \Delta_t^{(j)} A \left(\tilde{\pi}_{t|T}^{(j)}, \tilde{\mathbf{k}}_{t|T}^{(j)} \right) P_{t|t}. \end{aligned}$$

Algorithm ends.

C ESTIMATION IN OTHER COUNTRIES

Observation equation. The Consensus forecasts can be expressed as

$$\begin{aligned} E_t^{Cons} \pi_{Y1,Q2} &= \sum_{j=1}^4 w(j) \pi_{t-j} + \hat{E}_t \sum_{i=0}^2 w(5+i) \pi_{t+i} \\ E_t^{Cons} \pi_{Y1,Q1} &= \sum_{j=1}^3 w(j) \pi_{t-j} + \hat{E}_t \sum_{i=0}^3 w(4+i) \pi_{t+i} \end{aligned}$$

where the vector

$$w = \left(\frac{1}{16} \quad \frac{2}{16} \quad \frac{3}{16} \quad \frac{4}{16} \quad \frac{3}{16} \quad \frac{2}{16} \quad \frac{1}{16} \right)$$

defines the appropriate weights, and the notation $\pi_{i,j}$ denotes the inflation forecast of year i inflation, taken in the current year in quarter j . The remaining four forecasts concern expectations for inflation in the next calendar year, taken in each quarter of the current year. Similarly they can be expressed as

$$\begin{aligned} E_t^{Cons} \pi_{Y2,Q4} &= \sum_{j=1}^2 w(j) \pi_{t-j} + \hat{E}_t \sum_{i=0}^4 w(3+i) \pi_{t+i} \\ E_t^{Cons} \pi_{Y2,Q3} &= \sum_{j=1}^1 w(j) \pi_{t-j} + \hat{E}_t \sum_{i=0}^5 w(2+i) \pi_{t+i} \\ E_t^{Cons} \pi_{Y2,Q2} &= \hat{E}_t \sum_{i=0}^6 w(1+i) \pi_{t+i} \\ E_t^{Cons} \pi_{Y2,Q1} &= \hat{E}_t \sum_{i=1}^7 w(i) \pi_{t+i} \end{aligned}$$

where the last two forecasts are purely forward looking. The observation equation can then be written

$$\begin{bmatrix} \pi_t \\ \hat{E}_t \sum_{i=1}^2 w(5+i) \pi_{t+i} \\ \hat{E}_t \sum_{i=1}^3 w(4+i) \pi_{t+i} \\ \hat{E}_t \sum_{i=1}^4 w(3+i) \pi_{t+i} \\ \hat{E}_t \sum_{i=1}^5 w(2+i) \pi_{t+i} \\ \hat{E}_t \sum_{i=1}^6 w(1+i) \pi_{t+i} \\ \hat{E}_t \sum_{i=1}^7 w(i) \pi_{t+i} \end{bmatrix} = \pi^{*,F} + H_t' \xi_t + R_t o_t^{Cons},$$

where $\pi^{*,F}$ denotes the country-specific sample mean of inflation.

As discussed in the main text, the aim is to evaluate model predictions under the posterior distribution obtained with US data on inflation and forecasts by professional forecasters. However, there are few parameters that we choose to estimate independently. In particular, we estimate the inflation mean and the standard deviation of measurement error on survey-forecasts. These parameters are necessarily country-specific and can impact significantly the model's predictions.

For the US, the Metropolis-Hasting algorithm is used to simulate the posterior distribution

$$P(\Theta^{US} | Y_t^{US}) = L(Y_t^{US} | \Theta^{US}) P(\Theta^{US})$$

where $L(Y_t^{US} | \Theta^{US})$ the model likelihood. For the other countries we use the US posterior distribution as prior for the common parameters. We can then simulate the posterior distribution

$$P^F(\Theta^F | Y_t^F, Y_t^{US}, \Theta^{US}) = \tilde{L}(Y_t^F | \Theta^{US}, \Theta^F)^{\lambda^F} L(Y_t^{US} | \Theta^{US}) p(\Theta^{US}) p(\Theta^F)$$

where the parameter λ^F is the weight that is given to the likelihood of the model for other countries. Notice that the case of $\lambda^F = 0$ corresponds to evaluating the parameters that are common to the US, Θ^{US} , at the posterior distribution for the US. The remaining parameters, Θ^F , are instead evaluated at their prior. In our estimation we set $\lambda^F = 0.05$ implying a very low weight on the foreign model likelihood, $\tilde{L}(Y_t^F | \Theta^{US}, \Theta^F)$. As a result, the posterior distribution of the common parameters with the US is essentially the same as for the US, while the likelihood informs about the country-specific parameters. Tables in the additional technical appendix give the parameter estimates for all other countries. They are obtained using 200000 draws from the simulated posterior distribution.

D MARGINAL LIKELIHOOD

To compute the marginal likelihood for the US baseline model we use the Geweke harmonic mean estimator. For each draw Θ_i we compute

$$p(y) = \left\{ \frac{1}{D} \sum_{i=1}^D \frac{f(\Theta_i)}{p(y|\Theta_i)p(\Theta_i)} \right\}^{-1}$$

where the function $f(\cdot)$ is the density of a Normal distribution with mean and variance corresponding to the mean and variance of the posterior draws sample. Moreover the distribution is truncated so that

$$\begin{aligned} f(\Theta_i) &= \tau^{-1} (2\pi)^{-d/2} |V_\Theta| \exp \left[-0.5 (\Theta_i - \bar{\Theta})' V_\Theta^{-1} (\Theta_i - \bar{\Theta}) \right] \\ &\quad \times \left\{ \Theta_i : (\Theta_i - \bar{\Theta})' V_\Theta^{-1} (\Theta_i - \bar{\Theta}) < \chi_{\tau,d}^2 \right\} \end{aligned}$$

where $\chi_{\tau,d}^2$ is the $(1 - \tau)$ quantile of the χ_d^2 distribution and d is the dimension of the parameters' vector. In order to compute the marginal likelihood we set $\tau = 0.5$.

APPENDIX: ADDITIONAL MATERIAL (NOT FOR PUBLICATION)

D.0.1 DERIVATION OF Γ . Substituting for the marginal cost (7) into the aggregate supply equation (1) gives

$$\begin{aligned} \pi_t - \gamma\pi_{t-1} &= \mu_t + \xi_p s_t + E_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} [\alpha\beta\xi_p s_{T+1} + (1 - \alpha)\beta(\pi_{T+1} - \gamma\pi_T)] \\ &= \tilde{\mu}_t + \tilde{\kappa}\varphi_t + E_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} [\alpha\beta\tilde{\kappa}[\varphi_{T+1} - (\pi_{T+1} - \bar{\pi}_t) + \gamma(\pi_T - \bar{\pi}_t)]] \\ &\quad + E_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} [(1 - \alpha)\tilde{\beta}(\pi_{T+1} - \gamma\pi_T)] \end{aligned}$$

where

$$\tilde{\kappa} = \left(1 + \frac{\xi_p}{\lambda_x \phi} \right)^{-1} \frac{\xi_p}{\lambda_x \phi} = \frac{\xi_p}{\xi_p + \lambda_x \phi} \quad \text{and} \quad \tilde{\beta} = \frac{\beta \lambda_x \phi}{\xi_p + \lambda_x \phi} = \frac{\beta}{1 + \lambda_x^{-1} \phi^{-1}}.$$

Rearranging gives

$$\begin{aligned}
 \left(1 + (1 - \alpha)\tilde{\beta} - \alpha\beta\tilde{\kappa}\right)\pi_t - \gamma\pi_{t-1} &= \tilde{\mu}_t + \tilde{\kappa}\varphi_t + \\
 &+ E_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \left[\alpha\beta\tilde{\kappa}\varphi_{T+1} - (1 - \alpha\beta\gamma) \left(\alpha\beta\tilde{\kappa} - (1 - \alpha)\tilde{\beta} \right) \pi_{T+1} \right] \\
 &+ E_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} [\alpha\beta\tilde{\kappa}(1 - \gamma)\bar{\pi}_t]
 \end{aligned}$$

The rational expectations equilibrium is computed from

$$\pi_t - \gamma\pi_{t-1} = \tilde{\mu}_t + \tilde{\kappa}E_t \sum_{T=t}^{\infty} \tilde{\beta}^{T-t}\varphi_T$$

giving

$$\pi_t - \gamma\pi_{t-1} = \tilde{\mu}_t + \omega_\varphi\varphi_t$$

where

$$\omega_\varphi = \frac{\tilde{\kappa}}{1 - \tilde{\beta}\rho}.$$

It can be further simplified as

$$\begin{aligned}
 \omega_\varphi &= \frac{\frac{\xi_p}{\xi_p + \lambda_x \phi}}{1 - \frac{\beta\lambda_x \phi}{\xi_p + \lambda_x \phi}\rho} \\
 &= \frac{1}{1 + (1 - \beta\rho)\lambda_x \xi_p^{-1}\phi}
 \end{aligned}$$

Substituting the discounted forecast for inflation

$$E_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \pi_{T+1} = \left(\frac{1}{1 - \alpha\beta} - \frac{\gamma}{1 - \alpha\beta\gamma} \right) \bar{\pi}_t + \frac{\gamma}{1 - \alpha\beta\gamma} \pi_t + \frac{\omega_\varphi \rho}{(1 - \alpha\beta\rho)(1 - \alpha\beta\gamma)} \varphi_t.$$

into the aggregate supply curve gives

$$\begin{aligned}\pi_t - \gamma\pi_{t-1} &= \tilde{\mu}_t + \left[\tilde{\kappa} + \frac{\alpha\beta\tilde{\kappa}\rho}{1-\alpha\beta\rho} - \frac{(1-\alpha\beta\gamma)(\alpha\beta\tilde{\kappa} - (1-\alpha)\tilde{\beta})}{(1-\alpha\beta\gamma)(1-\alpha\beta\rho)}\omega_\varphi\rho \right] \varphi_t + \\ &+ \left[\alpha\beta\tilde{\kappa}\frac{1-\gamma}{1-\alpha\beta} + \left((1-\alpha)\tilde{\beta} - \alpha\beta\tilde{\kappa} \right) \left(\frac{1-\alpha\beta\gamma}{1-\alpha\beta} - \frac{(1-\alpha\beta\gamma)\gamma}{1-\alpha\beta\gamma} \right) \right] \bar{\pi}_t.\end{aligned}$$

Using rational expectations about transitional dynamics and the fact that

$$\tilde{\kappa} + \frac{\alpha\beta\tilde{\kappa}\rho}{1-\alpha\beta\rho} - \frac{(1-\alpha\beta\gamma)(\alpha\beta\tilde{\kappa} - (1-\alpha)\tilde{\beta})}{(1-\alpha\beta\gamma)(1-\alpha\beta\rho)}\omega_\varphi\rho = \omega_\varphi$$

permits

$$\begin{aligned}\pi_t - \gamma\pi_{t-1} &= \tilde{\mu}_t + \omega_\varphi\varphi_t + \\ &+ \left[\alpha\beta\tilde{\kappa}\frac{1-\gamma}{1-\alpha\beta} + \left((1-\alpha)\tilde{\beta} - \alpha\beta\tilde{\kappa} \right) \left(\frac{1-\alpha\beta\gamma}{1-\alpha\beta} - \frac{(1-\alpha\beta\gamma)\gamma}{1-\alpha\beta\gamma} \right) \right] \bar{\pi}_t\end{aligned}$$

Simplifying then provides

$$\begin{aligned}\pi_t - \gamma\pi_{t-1} &= \tilde{\mu}_t + \omega_\varphi\varphi_t + \\ &+ \frac{1}{\xi_p + \lambda_x\phi} \left[\alpha\beta\xi_p\frac{1-\gamma}{1-\alpha\beta} + ((1-\alpha)\beta\lambda_x\phi - \alpha\beta\xi_p) \left(\frac{1-\alpha\beta\gamma}{1-\alpha\beta} - \frac{(1-\alpha\beta\gamma)\gamma}{1-\alpha\beta\gamma} \right) \right] \bar{\pi}_t\end{aligned}$$

or

$$\begin{aligned}\pi_t - \gamma\pi_{t-1} &= \tilde{\mu}_t + \omega_\varphi\varphi_t + \\ &+ \frac{1}{\xi_p + \lambda_x\phi} \left[\alpha\beta\xi_p \left(\frac{1-\gamma}{1-\alpha\beta} - \frac{1-\alpha\beta\gamma}{1-\alpha\beta} \right) \right] \bar{\pi}_t \\ &+ \frac{1}{\xi_p + \lambda_x\phi} \left[(1-\alpha)\beta\lambda_x\phi \left(\frac{1-\alpha\beta\gamma}{1-\alpha\beta} - \frac{(1-\alpha\beta\gamma)\gamma}{1-\alpha\beta\gamma} \right) + \alpha\beta\xi_p\gamma \right] \bar{\pi}_t\end{aligned}$$

or

$$\pi_t - \gamma\pi_{t-1} = \tilde{\mu}_t + \omega_\varphi\varphi_{t-1} + (1-\gamma)\Gamma\bar{\pi}_t + \omega_\varphi\varepsilon_t$$

where

$$\Gamma = \frac{1}{1 + \lambda_x^{-1} \phi^{-1} \xi_p} \frac{(1 - \alpha) \beta}{1 - \alpha \beta}.$$

Using $\phi = 1$ we have

$$\Gamma = \frac{1}{1 + \xi_p \lambda_x^{-1}} \frac{(1 - \alpha) \beta}{1 - \alpha \beta}$$

as in the main text.

D.0.2 DERIVATION AND PROOFS. The crucial step is the derivation of the prediction for the linear state (step 6). Notice first that given the link between $\bar{\pi}_t$ and the linear state we can use

$$N \left(\begin{bmatrix} \xi_{t|t}^{(i)} \\ A_{\bar{\pi},t}^{(i)} \xi_{t|t}^{(i)} \end{bmatrix}, \begin{bmatrix} P_{t|t}^{(i)} & A_{\bar{\pi},t}^{(i)} P_{t|t}^{(i)} \\ P_{t|t}^{(i)} A_{\bar{\pi}}^{(i)'} & A_{\bar{\pi},t}^{(i)} P_{t|t}^{(i)} A_{\bar{\pi}}^{(i)'} \end{bmatrix} \right) \sim N \left(\begin{bmatrix} \xi_t | Y_t, \bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)} \\ \bar{\pi}_{t+1|t} - f_{\bar{\pi}}(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)}) | Y_t, \bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)} \end{bmatrix} \right)$$

Using properties of the normal distribution, we can now get the conditional distribution

$$\xi_{t|t}^{*(i)} = \xi_{t|t}^{(i)} + P_{t|t}^{(i)} A_{\bar{\pi},t}^{(i)'} \left(A_{\bar{\pi},t}^{(i)} P_{t|t}^{(i)} A_{\bar{\pi},t}^{(i)'} \right)^{-1} \left(\bar{\pi}_{t+1|t}^{(i)} - f_{\bar{\pi}}(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)}) - f_{\mathbf{k},t}^{(i)-1} z_{t|t}^{(i)} \right)$$

$$P_{t|t}^{*(i)} = P_{t|t}^{(i)} - P_{t|t}^{(i)} A_{\bar{\pi},t}^{(i)'} \left(A_{\bar{\pi},t}^{(i)} P_{t|t}^{(i)} A_{\bar{\pi},t}^{(i)'} \right)^{-1} A_{\bar{\pi},t}^{(i)} P_{t|t}^{(i)}$$

where

$$A_{\bar{\pi},t}^{(i)} = A_{\bar{\pi}} \left(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)} \right);$$

$$f_{\mathbf{k},t}^{(i)-1} = f_{\mathbf{k}}(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)})^{-1}.$$

The predictions for the linear state are then

$$\xi_{t+1|t}^{(i)} = f_{\xi,t}^{(i)} + A_{\xi,t}^{(i)} \xi_{t|t}^{*(i)}$$

where

$$f_{\xi,t}^{(i)} = f_{\xi}(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)});$$

$$A_{\xi,t}^{(i)} = A_{\xi}(\bar{\pi}_{t|t}^{(i)}, \mathbf{k}_{t|t}^{(i)}),$$

and

$$P_{t+1|t}^{(i)} = A_{\xi,t}^{(i)} P_{t|t} A_{\xi,t}^{(i)'} +$$

$$- A_{\xi,t}^{(i)} \left[P_{t|t} A_{\bar{\pi},t}^{(i)'} \left(A_{\bar{\pi},t}^{(i)} P_{t|t} A_{\bar{\pi},t}^{(i)'} \right)^{-1} A_{\bar{\pi},t}^{(i)} P_{t|t} \right] A_{\xi,t}^{(i)'} + Q_{\xi}.$$

Here we show that $P_{t+1|t}^{(i)} = P_{t+1|t}$ for every (i) . For given initial $P_{t|t}$, it is straightforward to show that

$$\left(A_{\bar{\pi},t}^{(i)} P_{t|t} A_{\bar{\pi},t}^{(i)'} \right)^{-1} = \frac{1}{f_{\mathbf{k},t}^{(i)-2} P_{t|t}^{[\eta,\eta]}}.$$

Then a little algebra leads to the following:

$$\begin{aligned} \bar{P}_{t|t}^{(i)} &= P_{t|t} A_{\bar{\pi},t}^{(i)'} \left(\frac{1}{f_{\mathbf{k},t}^{(i)-2} P_{t|t}^{[\eta,\eta]}} \right) A_{\bar{\pi},t}^{(i)} P_{t|t} \\ &= \begin{bmatrix} P_{t|t}^{[\eta,\eta]} & P_{t|t}^{[\eta,\varphi]} & P_{t|t}^{[\eta,\pi]} \\ P_{t|t}^{[\eta,\varphi]} & \frac{\left(P_{t|t}^{[\eta,\varphi]} \right)^2}{P_{t|t}^{[\eta,\eta]}} & P_{t|t}^{[\eta,\varphi]} \frac{P_{t|t}^{[\eta,\pi]}}{P_{t|t}^{[\eta,\eta]}} \\ P_{t|t}^{[\eta,\pi]} & P_{t|t}^{[\eta,\varphi]} \frac{P_{t|t}^{[\eta,\pi]}}{P_{t|t}^{[\eta,\eta]}} & \frac{\left(P_{t|t}^{[\eta,\pi]} \right)^2}{P_{t|t}^{[\eta,\eta]}} \end{bmatrix} \\ &= \bar{P}_{t|t}. \end{aligned}$$

Next, evaluate

$$\tilde{P}_{t|t} = A_{\xi,t}^{(i)} P_{t|t} A_{\xi,t}^{(i)'} - A_{\xi,t}^{(i)} \bar{P}_{t|t} A_{\xi,t}^{(i)'} = A_{\xi,t}^{(i)} (P_{t|t} - \bar{P}_{t|t}) A_{\xi,t}^{(i)'}$$

where

$$(P_{t|t} - \bar{P}_{t|t}) = \begin{bmatrix} 0 & 0 & 0 \\ 0 & P_{t|t}^{[\varphi,\varphi]} - \frac{\left(P_{t|t}^{[\eta,\varphi]} \right)^2}{P_{t|t}^{[\eta,\eta]}} & P_{t|t}^{[\varphi,\pi]} - \frac{\left(P_{t|t}^{[\eta,\varphi]} \right)^2 P_{t|t}^{[\eta,\pi]}}{P_{t|t}^{[\eta,\eta]}} \\ 0 & P_{t|t}^{[\eta,\pi]} - \frac{\left(P_{t|t}^{[\eta,\varphi]} \right)^2 P_{t|t}^{[\eta,\pi]}}{P_{t|t}^{[\eta,\eta]}} & P_{t|t}^{[\pi,\pi]} - \frac{\left(P_{t|t}^{[\eta,\pi]} \right)^2}{P_{t|t}^{[\eta,\eta]}} \end{bmatrix}.$$

Finally,

$$\tilde{P}_{t|t} = A_{\xi,t}^{(i)} (P_{t|t} - \bar{P}_{t|t}) A_{\xi,t}^{(i)'} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & \iota_1 & \iota_2 + \iota_1 \\ 0 & \iota_2 + \iota_1 & 2\iota_2 + \iota_1 + \left(P_{t|t}^{[\pi,\pi]} - \frac{(P_{t|t}^{[\eta,\pi]})^2}{P_{t|t}^{[\eta,\eta]}} \right) \gamma^2 \end{bmatrix}$$

where

$$\begin{aligned} \iota_1 &= \rho_\varphi^2 \left(P_{t|t}^{[\varphi,\varphi]} - \frac{(P_{t|t}^{[\eta,\varphi]})^2}{P_{t|t}^{[\eta,\eta]}} \right) \\ \iota_2 &= \left(-\frac{P_{t|t}^{[\eta,\varphi]} P_{t|t}^{[\eta,\pi]}}{P_{t|t}^{[\eta,\eta]}} + P_{t|t}^{[\varphi,\pi]} \right) \rho_\varphi \gamma. \end{aligned}$$

So we can express

$$P_{t+1|t}^{(i)} = P_{t+1|t} = Q_\xi + \tilde{P}_{t|t}.$$