Monetary Policy & Anchored Expectations An Endogenous Gain Learning Model

Laura Gáti*

February 9, 2022

Abstract

This paper analyzes monetary policy in a model that accounts for the potential unanchoring of inflation expectations. The extent to which expectations are anchored is determined by an adaptive learning mechanism in which the sensitivity of the public's expectations varies endogenously. Within the context of an otherwise standard macro model with nominal rigidities and natural-rate and cost-push shocks, I find that optimal policy involves an aggressive interest rate response to any threat of expectations unanchoring. This stabilizes both inflation expectations and inflation, allowing the central bank to accommodate inflation fluctuations when expectations are well anchored. Furthermore, I estimate the model-implied relationship that determines the extent of unanchoring. The data suggest that the expectations process is nonlinear and asymmetric: expectations exhibit a higher sensitivity to large surprises than to smaller ones, and to downside surprises than to upside ones.

JEL codes: E52, E71, D84
Keywords: anchored expectations, adaptive learning, optimal monetary policy

*European Central Bank, Directorate General Research, Sonnemannstraße 20, 60314 Frankfurt am Main, Germany. Email: laura_veronika.gati@ecb.europa.eu. I would like to thank Ryan Chahrour, Susanto Basu and Peter Ireland, as well as Klaus Adam, Bob King, Sergey Slobodyan and Jenny Tang, and seminar participants at Brown University, the Federeal Reserve Bank of Boston, SED 2021, the Expectations in Dynamic Macro Models conference 2021 and the EEA-ES Summer Meeting 2021. The views expressed herein are my own and do not necessarily reflect those of the ECB or the Eurosystem.

1 Introduction

> Inflation that runs below its desired level can lead to an unwelcome fall in *longer-term inflation expectations*, which, in turn, can pull actual inflation even lower, resulting in an adverse cycle of ever-lower inflation and inflation expectations. [...] *Well-anchored inflation expectations* are critical[.]
> Jerome Powell, Chairman of the Federal Reserve[1]
> *(Emphases added.)*

There is broad consensus among policymakers that anchoring inflation expectations is central to the modern conduct of monetary policy. As Powell emphasizes, policymakers think of expectations as anchored when long-run expectations do not fluctuate systematically with short-run inflation surprises. If expectations were to unanchor, policymakers fear this would result in an adverse cycle of self-enforcing movements in expectations. And yet, this is a topic that until recently had received little attention in the literature. While there existed an empirical literature proposing measures of central bank credibility often associated with the idea of expectations anchoring, there were no models with a formal concept of anchoring and unanchoring.

This paper builds on Carvalho et al. (2021), the first paper with a formal notion of unanchoring, to introduce an anchoring mechanism based on a smooth, continuously variable sensitivity to inflation surprises. I embed the anchoring mechanism in an otherwise standard New Keynesian model to study how a concern to anchor inflation expectations affects the conduct of monetary policy. The main contribution of the paper is to study optimal monetary policy when the degree of expectations anchoring may vary. I take the model to the data to quantify my anchoring mechanism, and use it to characterize optimal policy analytically and numerically.

Fig. 1 provides some motivating evidence for continuously time-varying levels of unanchoring. It shows the regression coefficient of a rolling window regression of long-run inflation expectations on short-run forecast errors. Letting $\bar{\pi}_t^i$ denote the 10-year-ahead inflation expectation of forecaster i from the Survey of Professional Forecasters (SPF), and $f_{t|t-1}^i \equiv \pi_t - \mathbb{E}_{t-1}^i \pi_t$ denote forecaster i's one-year-ahead forecast error, I run

$$\Delta \bar{\pi}_t^i = \beta_0^w + \beta_1^w f_{t|t-1}^i + \epsilon_t^i, \tag{1}$$

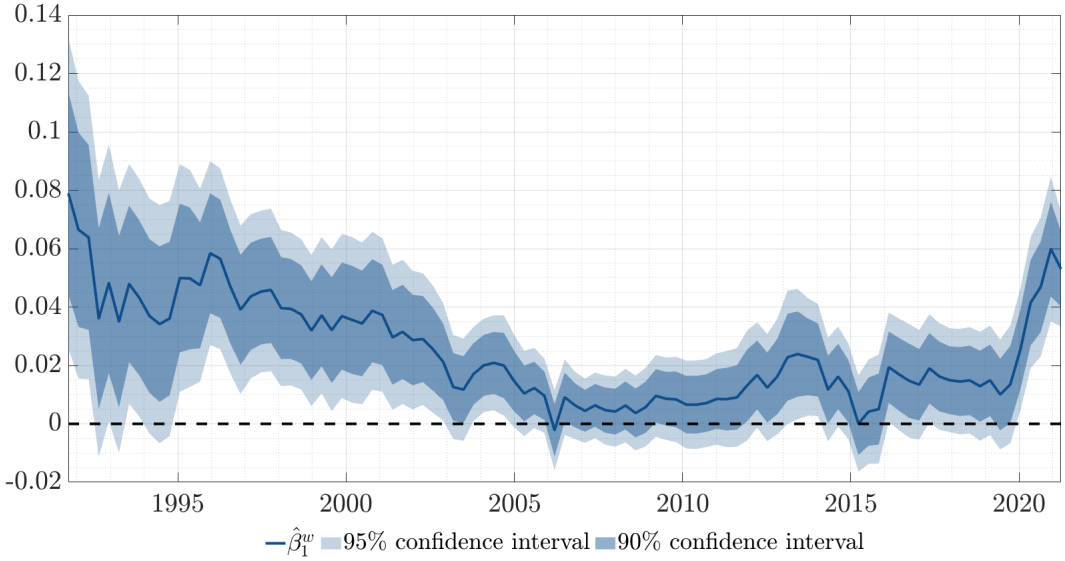
where w indexes windows of 20 quarters. Fig. 1 plots the time series of the estimated coefficient $\hat{\beta}_1^w$, along with 90% and 95% confidence intervals.[2]

As seen on Fig. 1, the relationship between 10-year inflation expectations and one-year surprises, $\hat{\beta}_1^w$, is time-varying. Up to about 2005, $\hat{\beta}_1^w$ is positive and significant, with varying magnitude. Between 2005 and 2010, by contrast, it is not significantly different from zero. In this period, then, long-run expectations were unresponsive to short-run conditions. Around 2010, this changed once more as $\hat{\beta}_1^w$ became slightly larger than in the preceeding years. Lastly, following the spike

[1] "New Economic Challenges and the Fed's Monetary Policy Review," August 27, 2020, Jackson Hole.

[2] Appendix A presents robustness checks. Appendix B discusses alternative measures of long-run inflation expectations.

Figure 1: Time series of responsiveness of long-run inflation expectations to inflation surprises



in inflation due to the Covid-19 crisis, it has become significantly positive; a worrying sign that short-run inflation surprises are passed through to long-run expectations once again.

In order to provide a convincing analysis of optimal policy in an environment with time-varying unanchoring, I need a qualitatively and quantitatively realistic model of expectation formation. To this end, I build on new work by Carvalho et al. (2021), who model the anchoring of expectations as a discrete sensitivity to surprises, in which agents either learn about long-run inflation using a small or large weight on past expectations errors. I extend this work along two dimensions. First, in order to be able to study optimal monetary policy analytically, I consider a smooth, continuous sensitivity, so that agents choose what weight to put on past expectations errors based on recent inflation surprises. Second, I embed the anchoring model of expectation formation in a general equilibrium New Keynesian model, allowing me to formally consider the monetary policy problem.

As an empirical contribution, I bring my smooth sensitivity specification to the data to explore the functional form describing the anchoring of expectations. Calibrating the parameters of the New Keynesian core of my model, I employ a simulated method of moments estimation strategy (Duffie and Singleton, 1990, Lee and Ingram, 1991, Smith, 1993) to back out a piecewise linear approximation to the true functional form in the data.

The estimation results show that my anchoring framework has several reasonable properties. First, I estimate an implied sensitivity time series with a median of 0.098, very close to the values obtained by Milani (2014) and Carvalho et al. (2021) (0.082 and 0.145 respectively). This number means that, on average, agents in the model put the bulk of weight on the 10 most recent quarters of data when updating their long-run expectations. Most importantly, the expectations process is nonlinear: larger forecast errors in absolute value lead to a higher sensitivity of long-run expectations. In line with other studies, like Hebden et al. (2020), I also find that inflation surprises on the

downside unanchor expectations more than same-sized surprises on the upside.

Having calibrated the model to the data, and establishing that it provides a realistic description of expectations, I perform my analysis of policy in the model in three stages. First, I present an analytical characterization of the Ramsey problem of the monetary authority. I demonstrate that the first-order condition of the Ramsey problem delivers a target criterion for the policymaker. The target criterion prescribes that the central bank should smooth out the effects of shocks by taking advantage of the current and future expected degree of unanchoring.

Second, I solve the optimality conditions of the Ramsey problem numerically. Because the endogenous anchoring function renders the model nonlinear, I rely on global methods to obtain the optimal policy function. Like the degree of unanchoring, the optimal interest-rate setting is time-varying. Since the interest rate responds to movements in long-run expectations, the monetary authority acts aggressively when expectations unanchor. When expectations are well-anchored, by contrast, the central bank accommodates fluctuations in inflation.

Third, I consider the class of simple, time-invariant Taylor rules that constitute the most common specification of monetary policy in practice. I solve for the optimal response coefficient of the interest rate to inflation numerically in the case of both the anchoring and the rational expectations versions of my model. It turns out that the central bank finds it optimal to respond much more to inflation under varying levels of anchoring than under rational expectations. The reason is that under the anchoring expectation formation, there is positive feedback between expectations and observables. This positive feedback simultaneously amplifies the shocks that hit the model, as well as introduces the need to stabilize expectations in order to stabilize inflation and output. This corroborates the findings of Orphanides and Williams (2004) that monetary policy needs to be more aggressive on inflation under learning than under rational expectations.

I also show that under an alternative specification for interest rate expectations, one obtains the opposite conclusion for the optimal Taylor rule. In particular, as Eusepi et al. (2018a) show, when interest rate expectations internalize drifting long-run inflation expectations, monetary policy can no longer manage interest rate expectations without introducing excess volatility into the model. In this case, then, the monetary authority should be less aggressive on inflation to reduce the volatility it inflicts on the economy. This bears out Eusepi et al. (2018a)'s point that a monetary authority that seeks to anchor inflation expectations also needs to pay attention to how interest rate expectations are formed because that matters for how effective the central bank's interest rate policy can be.

The paper is related to two main strands of literature. The model I use to study the interaction between monetary policy and anchoring is a behavioral version of the standard New Keynesian (NK) model of the type widely used for monetary policy analysis. Monetary policy in the rational expectations (RE) version of this model has been studied extensively, for example in Clarida et al. (1999) or Woodford (2003a). The formulation of a target criterion to implement optimal policy is in the tradition of Svensson (1999).

The behavioral part of the model is the departure from rational expectations on the part of the private sector. Instead, I allow the private sector to form expectations via an adaptive learning

scheme, where the learning gain - the parameter governing the sensitivity of expectations to inflation surprises - is endogenous. The learning framework represents an extension to the adaptive learning literature advocated by Evans and Honkapohja (2001). This literature replaces the rational expectations assumption by postulating an ad-hoc forecasting rule, the perceived law of motion (PLM), as the expectation-formation process. Agents use the PLM to form expectations and update it in every period using recursive estimation techniques. My contribution to this literature is to study optimal monetary policy in a learning model with an endogenous gain.

Within the adaptive learning literature, the paper touches base with three sets of papers. First, Molnár and Santoro (2014) and Mele et al. (2019) show that adaptive learning in general introduces an intertemporal tradeoff to monetary policy that is absent under rational expectations. This general insight, which also implies that optimal policy under commitment coincides what that under discretion, also carries over to my model because it comes from the fact that agents learn from past data.

Second, Marcet and Nicolini (2003), Cho and Kasa (2015), Kostyshyna (2012) Milani (2014) and Carvalho et al. (2021) study time-varying and state-dependent gain functions. In particular, Carvalho et al. (2021) use such a framework as a first model of anchoring, and validate the empirical performance of their anchoring model using inflation expectations surveys. My paper is intimately connected to these papers, in particular as my anchoring model is a version of Carvalho et al. (2021)'s. The main difference is that as opposed to the discrete gain specification in Carvalho et al. (2021), I consider a smooth and continuous gain function in order to make the optimal policy problem tractable. This way I can provide what I think is a novel contribution: the study of monetary policy for varying degrees of unanchoring.

Third, my results on the inflation coefficient in the Taylor rule speak to the literature on whether monetary policy should respond more or less to inflation under learning than under rational expectations. On the one hand, Orphanides and Williams (2004) obtain that in the presence of learning, monetary policy should respond more to inflation in order to subdue volatile expectations. In contrast, Eusepi et al. (2018a) and Eusepi and Preston (2018) find that an aggressive response to inflation leads to excess volatility when expectations of long-run means have a drift. As a result, monetary policy cannot and should not respond strongly to inflation fluctuations. While my baseline results here echo those of Orphanides and Williams (2004), the model can generate the Eusepi et al. (2018a) result if an alternative assumption on interest rate expectations is used. In particular, which of the two results one gets hinges on whether one assumes that interest rate expectations incorporate the fact that long-run inflation expectations are drifting. If yes, one obtains the Eusepi et al. (2018a) result, if no, the Orphanides and Williams (2004) one. For the monetary authority, this suggests that in order to be able anchor inflation expectations, it is first order to understand how the private sector forms its expectations of the interest rate.

There are several reasons for why adaptive learning is an attractive alternative to rational expectations in general, and a suitable framework to study anchoring in particular. First, many studies document the ability of adaptive learning models to match empirical properties of both expectations and macro aggregates. Adaptive learning models imply that forecast errors are corre-

lated with forecast revisions, a feature of expectations documented by Coibion and Gorodnichenko (2015). In fact, the prediction of adaptive learning models concerning the response of expectations to shocks exactly aligns with new evidence by Angeletos et al. (2021), suggesting that in response to shocks, expectations initially underreact, and then overshoot. As for the macro evidence, Milani (2007) demonstrates that estimated constant gain learning models match the persistence of inflation without recourse to backward-looking elements in the Phillips curve. Eusepi and Preston (2011) show that a calibrated adaptive learning version of the real business cycle (RBC) model outperforms the rational expectations version. In particular, even with a small gain, the learning model leads to persistent and hump-shaped responses to iid shocks, resolving the long-standing critique of RBC models of Cogley and Nason (1993).

Secondly, there is very strong empirical backing for state-dependent gains. Milani (2014) documents that endogenous gain models can generate endogenous time-varying volatility in models without any time-variation in the exogenous processes. Additionally, Carvalho et al. (2021) estimate the evolution of the endogenous gain for the last fifty years in the US. Not only does their model display excellent out-of-sample forecasting performance in terms of matching long-run expectations, but the estimated gain time series invites a reinterpretation of the Great Inflation as a period of unanchored expectations.

Thirdly, an extensive experimental literature in the spirit of Anufriev and Hommes (2012) demonstrates that simple, state-dependent adaptive learning rules provide the best fit among competing models to how individuals form expectations in controlled lab settings. A key finding in this literature is that experiment participants rely on simple heuristics to form expectations, and they do so in a state-dependent way. In particular, changing the economic environment of the experiment induces the forecasters to change the forecasting rule they use.

Finally, my work is also related to the literature attempting to explain features of expectations data using departures from the full information rational expectations (FIRE) paradigm. This literature consists of two main lines of attack. The first is relaxing the assumption of full information. In this body of work, information is either not fully or not symmetrically available, or information acquisition is costly (Mankiw and Reis, 2002, Sims, 2003, Mackowiak and Wiederholt, 2009, Angeletos and Pavan, 2009). The second strand of this literature, to which this paper belongs, instead emphasizes that expectation formation departs from rational expectations. One example of this literature is Sargent (1999), who explores the role of adaptive learning on the part of the Federal Reserve for the Great Inflation. A more recent example is Bordalo et al. (2018). These authors complement the empirical results of Coibion and Gorodnichenko (2015) with new evidence to suggest that the pattern of over- and underreaction to news is consistent with diagnostic expectations instead of dispersed information. My model of anchored expectations is also consistent with these data facts.

The paper is structured as follows. Section 2 introduces the model. Section 3 describes the learning framework and spells out the anchoring mechanism. Section 4 estimates the anchoring process. Section 5 presents the results in three parts. First, Section 5.1 discusses an analytical characterization of the Ramsey policy. Second, Section 5.2 solves for the interest rate sequence that

implements the optimal Ramsey allocation using global methods. Third, Section 5.3 investigates the optimal choice of response coefficients if monetary policy is restricted to follow a Taylor rule. Section 6 concludes.

2 The model

Apart from expectation formation, the model is a standard New Keynesian (NK) model with nominal frictions à la Calvo (1983). The advantage of having a standard NK backbone to the model is that one can neatly isolate the way the anchoring mechanism alters the behavior of the model. Since the mechanics of the rational expectations version of this model are well understood, I only lay out the model briefly and pinpoint the places where the assumption of non-rational expectations matters.[3]

2.1 Households

The representative household is infinitely-lived and maximizes expected discounted lifetime utility from consumption net of the disutility of supplying labor hours:

$$\hat{\mathbb{E}}_t^i \sum_{T=t}^{\infty} \beta^{T-t} \left[U(C_T^i) - \int_0^1 v(h_T^i(j)) dj \right]. \tag{2}$$

In the above, the only non-standard element is the expectations operator, $\hat{\mathbb{E}}^i$. As detailed in Section 3, this operator captures non-rational expectations and is assumed to satisfy the law of iterated expectations, so that $\hat{\mathbb{E}}_t^i \hat{\mathbb{E}}_{t+1}^i = \hat{\mathbb{E}}_t^i$. The remaining elements are familiar from the rational expectations version of the NK model. $U(\cdot)$ and $v(\cdot)$ denote the utility of consumption and disutility of labor, respectively, and β is the discount factor of the household. I am using standard CRRA utility of the form $U(c_t) = \frac{c_t^{1-\sigma}}{1-\sigma}$. $h_t^i(j)$ denotes the supply of labor hours of household i at time t in the production of good j, and the household participates in the production of all goods j. Similarly, household i's consumption bundle at time t, C_t^i, is a Dixit-Stiglitz composite of all goods in the economy:

$$C_t^i = \left[\int_0^1 c_t^i(j)^{\frac{\theta-1}{\theta}} dj \right]^{\frac{\theta}{\theta-1}}, \tag{3}$$

where $\theta > 1$ is the elasticity of substitution between the varieties of consumption goods. Denoting by $p_t(j)$ the time-t price of good j, the aggregate price level in the economy can then be written as

$$P_t = \left[\int_0^1 p_t(j)^{1-\theta} dj \right]^{\frac{1}{\theta-1}}. \tag{4}$$

The budget constraint of household i is given by

$$B_t^i \leq (1 + i_{t-1}) B_{t-1}^i + \int_0^1 w_t(j) h_t^i(j) + \Pi_t^i(j) dj - T_t - P_t C_t^i, \tag{5}$$

[3] For more details on the NK model, see Woodford (2003a).

where $\Pi_t^i(j)$ denotes profits from firm j remitted to household i, T_t taxes, and B_t^i the riskless bond purchases at time t.[4]

The only difference to the standard New Keynesian model is the expectations operator, $\hat{\mathbb{E}}^i$. This is the subjective expectations operator that differs from its rational expectations counterpart, $\mathbb{E}$, in that it does not encompass knowledge of the model. In particular, households have no knowledge of the fact that they are identical. By extension, they also do not internalize that they hold identical beliefs about the evolution of the economy. This implies that while the modeler can suppress the index i, understanding that $\hat{\mathbb{E}}^i = \hat{\mathbb{E}}^j = \hat{\mathbb{E}}$, households cannot do so. As we will see in Section 2.3, this has implications for their forecasting behavior and will result in decision rules that differ from those of the rational expectations version of the model.

2.2 Firms

Firms are monopolistically competitive producers of the differentiated varieties $y_t(j)$. The production technology of firm j is $y_t(j) = A_t f(h_t(j))$, whose inverse, $f^{-1}(\cdot)$, signifies the amount of labor input. Noting that A_t is the level of technology and that $w_t(j)$ is the wage per labor hour, firm j profits at time t can be written as

$$\Pi_t^j = p_t(j)y_t(j) - w_t(j)f^{-1}(y_t(j)/A_t). \tag{6}$$

Firm j's problem then is to set the price of the variety it produces, $p_t(j)$, to maximize the present discounted value of profit streams

$$\hat{\mathbb{E}}_t^j \sum_{T=t}^{\infty} \alpha^{T-t} Q_{t,T} \left[\Pi_t^j(p_t(j)) \right], \tag{7}$$

subject to the downward-sloping demand curve

$$y_t(j) = Y_t \left(\frac{p_t(j)}{P_t} \right)^{-\theta}, \tag{8}$$

where

$$Q_{t,T} = \beta^{T-t} \frac{P_t U_c(C_T)}{P_T U_c(C_t)} \tag{9}$$

is the stochastic discount factor from households. Nominal frictions enter the model through the parameter α in Equation (7). This is the Calvo probability that firm j is not able to adjust its price in a given period.

Analogously to households, the setup of the production side of the economy is standard up to the expectation operator. Also here the model-consistent expectations operator $\mathbb{E}$ has been replaced by the subjective expectations operator $\hat{\mathbb{E}}^j$. This implies that firms, like households, do not know the model equations, nor do they internalize that they are identical. Thus their decision rules, just like those of the households, will be distinct from their rational expectations counterparts.

[4] For ease of exposition I have suppressed potential money assets here. This has no bearing on the model implications since it represents the cashless limit of an economy with explicit money balances.

2.3 Aggregate laws of motion

The model solution procedure entails deriving first-order conditions, taking a loglinear approximation around the nonstochastic steady state and imposing market clearing conditions to reduce the system to two equations, the New Keynesian Phillips curve and IS curve. The presence of subjective expectations, however, implies that firms and households are not aware of the fact that they are identical. Thus, as Preston (2005) points out, imposing market clearing conditions in the expectations of agents is inconsistent with the assumed information structure.[5]

Instead, I prevent firms and households from internalizing market clearing conditions. As Preston (2005) demonstrates, this leads to long-horizon forecasts showing up in firms' and households' first-order conditions. As a consequence, instead of the familiar expressions, the IS and Phillips curves take the following form:

$$x_t = \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t}((1-\beta)x_{T+1} - \sigma(\beta i_T - \pi_{T+1}) + \sigma r_T^n), \tag{10}$$

$$\pi_t = \kappa x_t + \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t}(\kappa\alpha\beta x_{T+1} + (1-\alpha)\beta\pi_{T+1} + u_T). \tag{11}$$

Here x_t, π_t and i_t are the log-deviations of the output gap, inflation and the nominal interest rate from their steady state values, and σ is the intertemporal elasticity of substitution. The variables r_t^n and u_t are exogenous disturbances representing a natural rate shock and a cost-push shock respectively.

The laws of motion (10) and (11) are obtained by deriving individual firms' and households' decision rules, which involve long-horizon expectations, and aggregating across the cross-section. Importantly, agents in the economy have no knowledge of these relations since they do not know that they are identical and thus are not able to impose market clearing conditions required to arrive at (10) and (11). Thus, although the evolution of the observables (π, x) is governed by the exogenous state variables (r^n, u) and long-horizon expectations via these two equations, agents in the economy are unaware of this. As I will spell out more formally in Section 3, it is indeed the equilibrium mapping between states and jump variables the agents are attempting to learn.

To simplify notation, I gather the exogenous state variables in the vector s_t and observables in the vector z_t as

$$s_t = \begin{bmatrix} r_t^n \\ \bar{i}_t \\ u_t \end{bmatrix} \qquad \text{and} \qquad z_t = \begin{bmatrix} \pi_t \\ x_t \\ i_t \end{bmatrix}, \tag{12}$$

[5]Two ways have been proposed in the adaptive learning literature to deal with this issue. The first is the so-called Euler-equation approach, see for example by Bullard and Mitra (2002). This approach involves writing down the log-linearized first-order conditions of the model, and simply replacing the rational expectations operators with subjective ones. The other approach, advocated by Preston (2005) and adopted in this paper, instead involves rederiving the aggregate laws of motion from the individual agents' optimization problems that incorporate subjective expectations from the start. Lastly, a related but distinct alternative to Preston (2005)'s long-horizon approach is the shadow price learning framework advocated by Evans and McGough (2009).

where $\bar{i}_t$ is a shock to the interest rate that only shows up in the model for particular specifications of monetary policy.[6] This allows me to denote long-horizon expectations by

$$f_{a,t} \equiv \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} z_{T+1} \qquad \text{and} \qquad f_{b,t} \equiv \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} (\beta)^{T-t} z_{T+1}. \tag{13}$$

As detailed in Appendix D, one can use this notation to reformulate the laws of motion of jump variables (Equations (10) and (11)) and a given monetary policy rule compactly as

$$z_t = A_a f_{a,t} + A_b f_{b,t} + A_s s_t, \tag{14}$$

where the matrices $A_i,\ i = \{a, b, s\}$ gather coefficients and are given in Appendix D. Assuming that exogenous variables evolve according to independent AR(1) processes, I write the state transition matrix equation as

$$s_t = h s_{t-1} + \epsilon_t \qquad \text{with} \qquad \epsilon_t \sim \mathcal{N}(\mathbf{0}, \Sigma), \tag{15}$$

where h gathers the autoregressive coefficients ρ_j, ϵ_t the Gaussian innovations ε_t^j, and η the standard deviations σ_t^j, for $j = \{r, i, u\}$. $\Sigma = \eta\eta'$ is the variance-covariance matrix of disturbances.[7]

3 Learning with an anchoring mechanism

The informational assumption of the model is that agents do not know the equilibrium mapping between states and jumps in the model. Without knowing the form of the observation equation, they are not able to form rational expectations forecasts. Instead, agents postulate an ad-hoc forecasting relationship between states and jumps and seek to refine it in light of incoming data. In other words, they act like an econometrician: they estimate a simple statistical model and constantly attempt to improve their model as new data arrive.

3.1 Perceived law of motion

I assume agents consider a forecasting model for the endogenous variables of the form

$$\hat{\mathbb{E}}_t z_{t+1} = a_{t-1} + b_{t-1} s_t, \tag{16}$$

where a and b are estimated coefficients of dimensions 3×1 and 3×3 respectively. This perceived law of motion (PLM) reflects the assumption that agents forecast jumps using a linear function of current states and a constant, with last period's estimated coefficients. Note that a can be interpreted as long-run expectations of the observables z. Since inflation is the first element of z, the first element of a corresponds to long-run inflation expectations. This object, which I will denote by $\bar{\pi}$, will be the main focus of the paper in what follows.

[6] For generality, I treat the exogenous state vector as three-dimensional throughout the paper, even when the monetary policy shock is absent and the second element of the state vector is zero.

[7] For the sake of conciseness, I have suppressed the expressions for these in the main text. See Appendix D.

Summarizing the estimated coefficients as $\phi_{t-1} \equiv \begin{bmatrix} a_{t-1} & b_{t-1} \end{bmatrix}$, here 3×4, I can rewrite Equation (16) as

$$\hat{\mathbb{E}}_t z_{t+1} = \phi_{t-1} \begin{bmatrix} 1 \\ s_t \end{bmatrix}. \tag{17}$$

I also assume that

$$\hat{\mathbb{E}}_t \phi_{t+k} = \phi_t \quad \forall\, k \geq 0. \tag{18}$$

This assumption, known in the learning literature as anticipated utility (Kreps, 1998), means that agents do not internalize that they will update the forecasting rule in the future.[8] Since the states s_t are exogenous, I assume that agents know Equation (15), the equation governing the evolution of s_t.[9] Then, the PLM together with anticipated utility implies that k-period-ahead forecasts in the beginning of period t are constructed as

$$\hat{\mathbb{E}}_t z_{t+k} = a_{t-1} + b_{t-1} h^{k-1} s_t \quad \forall k \geq 1. \tag{19}$$

The timing assumptions of the model are as follows. In the beginning of period t, the current state s_t is realized. Agents then form expectations according to (16) using last period's estimate ϕ_{t-1} and the current state s_t. Given exogenous states and expectations, today's jump vector z_t is realized. This allows agents to evaluate the forecast error made at the end of last period

$$f_{t|t-1} \equiv z_t - \phi_{t-1} \begin{bmatrix} 1 \\ s_{t-1} \end{bmatrix} \tag{20}$$

to update their forecasting rule.[10] The estimate is updated according to the following recursive least-squares algorithm:

$$\phi_t = \left(\phi'_{t-1} + k_t R_t^{-1} \begin{bmatrix} 1 \\ s_{t-1} \end{bmatrix} f'_{t|t-1} \right)', \tag{21}$$

$$R_t = R_{t-1} + k_t \left(\begin{bmatrix} 1 \\ s_{t-1} \end{bmatrix} \begin{bmatrix} 1 & s_{t-1} \end{bmatrix} - R_{t-1} \right). \tag{22}$$

Here R_t is the 4×4 variance-covariance matrix of the regressors and k_t is the learning gain, specifying to what extent the updated estimate loads on the forecast error. Clearly, a high gain implies high loadings and thus strong changes in the estimated coefficients ϕ_t. A low gain, by contrast, means that a given forecast error only has a small effect on ϕ_t. Thus, one can interpret the gain as the sensitivity of the expectations process to short-run surprises.

[8] Clearly, anticipated utility poses a higher level of irrationality than not knowing the model and using statistical techniques to learn it. Nevertheless, because Sargent (1999) demonstrates that it does not alter the dynamics of the linearized model, anticipated utility has become a standard simplifying assumption in the adaptive learning literature.

[9] This is another common simplifying assumption in the adaptive learning literature. In an extension, I relax this assumption and find that when agents have to learn the evolution of state variables, they do so very quickly. Therefore the dynamics in that case resemble the dynamics here, except that responses to shocks are muted as long as agents are learning about the state variables.

[10] Alternatively, one could assume that agents use the forecast they made in the morning of the current period. This does not change the dynamics of the model in any way.

3.2 Endogenous gain as anchoring mechanism

The vast majority of the learning literature specifies the gain either as a constant, $\bar{g}$, or decreasing with time, so that $k_t = t^{-1}$, where t indexes time. Instead, I follow Carvalho et al. (2021) in allowing the gain to fluctuate in a time-varying way in response to short-run forecast errors. I assume the gain evolves as

$$k_t = \mathbf{g}(k_{t-1}, f_{t|t-1}), \tag{23}$$

where $\mathbf{g}(\cdot)$ is a smooth, continuous function. I refer to $\mathbf{g}(\cdot)$ as the anchoring function since it specifies how the sensitivity of the expectations process is determined. As Carvalho et al. (2021) point out, one can think of the gain as a formal notion of the degree of unanchoring, because a high gain implies a high sensitivity to surprises, while a low gain a low one.

To understand the intuition, consider the learning vector a, the intercept in the private sector's forecasting rule in Equation (16). Recall that the first element of this vector, $\bar{\pi}$, corresponds to the expectation of long-run inflation. When the gain is high, agents update their expectations process strongly in light of surprises, thus also updating $\bar{\pi}$. What this means is that episodes of high gains correspond to episodes when the private sector changes its expectation of long-run inflation very strongly in response to current realized inflation. So if the central bank misses its inflation target, this surprise inflation will show up in long-run inflation expectations only if the gain is large enough. In other words, one can interpret the gain as a metric of unanchoring because it specifies how strongly current realized inflation will cause the public to revise its view of the inflation regime it is in.

Notice that the anchoring function is an extension to the conventional decreasing or constant gain specifications, as it nests both as special cases. A decreasing gain implies $\mathbf{g}(k_{t-1}, f_{t|t-1}) = t^{-1}$, while a constant gain $\mathbf{g}(k_{t-1}, f_{t|t-1}) = \bar{g}$. Furthermore, in both special cases, $\mathbf{g}_k = \mathbf{g}_f = 0$.[11] The smooth gain function is also an extension to Carvalho et al. (2021)'s framework, as I explain below.

Apart from the smoothness and continuity assumptions, I make the following assumptions on the anchoring function $\mathbf{g}(\cdot)$:

$$\mathbf{g}_k \geq 0, \tag{24}$$

$$\mathbf{g}_{ff} \geq 0. \tag{25}$$

Equation (24) says that the current value of the gain depends positively on its past value. Equation (25) states that $\mathbf{g}(\cdot)$ is convex in its second argument, meaning that the gain is increasing in the absolute value of forecast errors.

To interpret these assumptions, consider what time series of inflation surprises imply for the evolution of the gain. During the Great Inflation of the 1970s in the US, for example, when inflation kept exceeding the public's expectations, large positive forecast errors resulted in high gains through the convexity of $\mathbf{g}(\cdot)$. In other words, a volatile environment with lots of large inflation surprises drives up the gain, rendering long-run inflation expectations sensitive to short-run fluctuations, unanchoring them. The convexity makes sure that both upside and downside surprises

[11] To be precise, for a decreasing gain this statement only holds in the limit.

raise the sensitivity of long-run expectations, so that what matters for the unanchoring of expectations is the absolute size of the surprise, not its sign.

Once long-run inflation expectations were already very sensitive to short-run surprises in the early 1980s, the private sector was tracking inflation developments very closely, despite the fact that aggressive monetary policy under Chairman Paul Volcker resulted in smaller inflation surprises. It is this feature that the positive derivative in Equation (24) captures. It says that a high gain in the previous period pushes up today's gain. What this means is that the public remains alert to inflation developments if it has been surprised in the past, even if it may not be surprised today.

The usefulness of specifying the evolution of the gain as endogenous is that it offers a time- and state-dependent sensitivity of the expectations process to surprises. In both of the standard, exogenous gain schemes the gain is divorced from the current environment. It either decreases deterministically ($k_t = t^{-1}$), or is a constant ($k_t = \bar{g}$). This means that the level of unanchoring is either deterministic or constant.

An endogenous gain, instead, generates varying degrees of unanchoring. When high volatility results in large forecast errors in absolute value, the anchoring function outputs a larger gain than when forecast errors are close to zero. Thus the learning coefficients a_t and b_t respond more to forecast errors in volatile states of the world than in stable ones. The time-varying nature of the gain therefore generates periods of well-anchored expectations (low gain) as well as highly unanchored episodes (high gain), and everything in-between. The endogenous gain can thus be interpreted as a metric of the varying degrees of unanchoring.

Time- and state-dependent learning has significant empirical backing. For a start, the fact that the gain is a convex function of forecast errors allows the model to match the time-varying sensitivity of long-term inflation expectations documented in Fig. 1. The model also implies a pattern of delayed reaction and overshooting of expectations in response to incoming information, in line with recent evidence by Angeletos et al. (2021). Furthermore, as demonstrated by Milani (2014), endogenous gain learning models can match time-varying volatility in the data, a feature that constant gain learning or rational expectations models cannot account for. Lastly, an expectation formation that allows individuals to regulate the extent to which they track incoming information also aligns with micro-level evidence on state-dependent expectation formation in lab settings, such as in Anufriev and Hommes (2012).

To be clear, I am not the first to use an endogenous gain learning model in the macro context. That innovation is due to Marcet and Nicolini (2003), and the papers in their wake: Kostyshyna (2012), Milani (2014) and Carvalho et al. (2021). In these papers, agents switch between a constant and a decreasing gain based on a switching criterion. Furthermore, Carvalho et al. (2021) are the first to use an endogenous gain learning scheme as a model of anchoring.[12] Their focus is to show that an anchoring expectation formation can both match untargeted moments of long-run inflation expectations from surveys, as well as rationalize the historical dynamics of US inflation through a novel narrative of transitioning from an unanchored phase to an anchored one. To do so, they

[12]See Appendix F for a detailed description of Carvalho et al. (2021)'s anchoring function.

estimate a partial equilibrium model which incorporates their novel anchoring theory.

My focus here instead is to analyze the interaction between expectations unanchoring and monetary policy. I thus embed an endogenous gain model of anchoring in a general equilibrium context to understand how it affects the conduct of monetary policy. The modification of Carvalho et al. (2021)'s anchoring theory using a smooth gain function serves the purpose of rendering the optimal monetary policy problem tractable. With this modification, the derivatives of $\mathbf{g}(\cdot)$ exist, allowing me to obtain analytical expressions for the Ramsey problem of the monetary authority. Thus the paper can shed light on what varying degrees of expectations unanchoring means for monetary policy.

3.3 Actual law of motion

To complete the model, I now use the specifics of the anchoring expectation formation to characterize the evolution of the jump variables under learning. Using the PLM from Equation (16), I write the long-horizon expectations in (13) as

$$f_{a,t} \equiv \frac{1}{1-\alpha\beta} a_{t-1} + b_{t-1}(I_3 - \alpha\beta h)^{-1} s_t \qquad \text{and} \qquad f_{b,t} \equiv \frac{1}{1-\beta} a_{t-1} + b_{t-1}(I_3 - \beta h)^{-1} s_t. \quad (26)$$

Substituting these into the law of motion of observables (Equation (14)) yields the actual law of motion (ALM):

$$z_t = g^l_{t-1} \begin{bmatrix} 1 \\ s_t \end{bmatrix}, \quad (27)$$

where g^l is a 3×4 matrix given in Appendix E. Thus, instead of the state-space solution of the RE version of the model (Equations (15) and (E.1)), the state-space solution for the learning model is characterized by the pair of equations (15) and (27), together with the PLM (19), the learning equations (21) and (22), as well as the anchoring function (23).

3.4 Simplifying assumptions

To simplify the analytical work in Section 5.1, I make two assumptions that I maintain for the rest of the paper.

Assumption 1. $a_t = \begin{pmatrix} \bar{\pi}_t \\ 0 \\ 0 \end{pmatrix}$, *and* $b_t = g\ h, \quad \forall t.$

Assumption 1 amounts to restricting the intercepts in the forecasts of the output gap and the interest rate, as well as the slope coefficients of all forecasts to what they would be under rational expectations. This means that instead of learning the intercept and slope parameters for all three endogenous variables, the private sector only learns the intercept of the inflation process. Thus the single learning parameter is the long-run inflation expectation, $\bar{\pi}$.

The rationale behind this assumption is that it is the smallest possible deviation from rational expectations that at the same time has enough flexibility to study the unanchoring of inflation

expectations. It renders the Ramsey problem very tractable and makes the comparison with the rational expectations benchmark more transparent. In particular, since the inflation intercept is learned using my anchoring mechanism, the formulation is able to capture the time-varying sensitivity of long-run inflation expectations to short-run forecast errors in inflation.

Relaxing the assumption that slope coefficients are not learned does not change the implications of the model because it only makes impulse responses to shocks more bumpy after impact. As I explain in Section 5.3, the assumption on whether all variables are learnt or not has stronger implications. I maintain this assumption for the sake of tractability of the Ramsey problem, and for ease of comparison with rational expectations.

With Assumption 1, the updating equations (21) and (22) simplify to a single equation:

$$\bar{\pi}_t = \bar{\pi}_{t-1} + k_t f_{t|t-1}, \tag{28}$$

where, using the notation that b_1 is the first element of the vector b, the k-period-ahead inflation forecast in the beginning of period t is given by

$$\hat{\mathbb{E}}_t \pi_{t+k} = \bar{\pi}_{t-1} + b_1^k s_t, \quad \forall k \geq 1. \tag{29}$$

Lastly, the one-period-ahead forecast error in inflation from the end of last period simplifies to

$$f_{t|t-1} = \pi_t - (\bar{\pi}_{t-1} + b_1 s_{t-1}). \tag{30}$$

Assumption 2. $\mathbf{g}_k = 0$.

The interpretation of this assumption is that the current gain depends only on forecast errors, not on its own past level. The motivation is again to make the analytical solution of the Ramsey problem in Section 5.1 tractable. Appendix I provides a treatment of the Ramsey problem with the general specification of the anchoring function, relaxing Assumption 2. As seen in the Appendix, this does not change the mechanics nor the intuition behind the solution of the Ramsey problem, it only makes the algebra more involved.

4 Quantification of learning channel

The numerical analysis of monetary policy requires a functional specification for the anchoring function $\mathbf{g}(\cdot)$ of Equation (23). For this reason, I back out the functional form of $\mathbf{g}(\cdot)$ from data. Apart from its usefulness for numerical analysis, the form of the anchoring function is interesting in its own right because it describes empirical properties of the anchoring of expectations. In particular, central bankers may want to know how much forecast errors of a particular sign and magnitude unanchor expectations.

I carry out the estimation in two steps. I first calibrate the parameters of the underlying New Keynesian model. Conditional on these parameter values, I estimate the anchoring function by simulated method of moments à la Lee and Ingram (1991), Duffie and Singleton (1990) and Smith (1993). I target the autocovariance structure of the Baxter-King filtered observables of the model

and expectations. The observables are CPI inflation from the Bureau of Labor Statistics (BEA), the output gap and the federal funds rate from the Board of Governors of the Federal Reserve System.[13] For expectations, I rely on 12-month-ahead CPI inflation forecasts from the Survey of Professional Forecasters (SPF). The dataset is quarterly and ranges from 1981-Q3 to 2020-Q1. Appendix G contains a detailed description of the estimation methodology.

4.1 Calibration

For the calibration of the New Keynesian backbone of the model, I split the parameters into two subsets. The first is calibrated using values from the literature, while the second is calibrated to match the moments outlined above, for an initial set of expectations parameters. Tables 1 and 2 show the two subsets of calibrated parameters respectively. As Table 1 depicts, I adopt standard parameters from the literature where possible. In particular, for β, σ and other parameters underlying κ, the slope of the Phillips curve, I rely on the parameterization of Chari et al. (2000), advocated in Woodford (2003a).

The composite parameter κ is given by $\kappa = \frac{(1-\alpha)(1-\alpha\beta)}{\alpha}\zeta$, where ζ is a measure of strategic complementarity in price setting. Assuming specific factor markets, constant desired markups with respect to output levels and no intermediate inputs, $\zeta = \frac{\omega+\sigma^{-1}}{1+\omega\theta}$. Here θ is the price elasticity of demand and ω is the elasticity of the marginal cost function with respect to output. Chari et al. (2000)'s calibration involves $\theta = 10, \sigma = 1, \omega = 1.25, \beta = 0.99$, so that together with my choice of α detailed below, κ is pinned down. Note that I lower β slightly (0.98 instead of Chari et al. (2000)'s 0.99). This allows the model to better match the autocovariance structure of the output gap because it lowers the pass-through of long-horizon expectations in the IS curve.

The probability of not adjusting prices, α, is set to match an average price duration of 2 quarters, which is slightly below the average number found in empirical studies of around 7 months.[14] I choose a slightly lower number in order to allow the learning mechanism, and not price stickiness, to drive the bulk of the model's dynamics.

To simplify the numerical analysis as well as the interpretation, I follow Molnár and Santoro (2014) in restricting the shocks to be iid. As those authors show, while this is a useful assumption to make sure that the optimal monetary policy problem has an analytical solution, it is actually not a restrictive assumption at all. The reason is that because learning models introduce endogenous persistence into inflation dynamics, autocorrelation in innovations becomes superfluous.

The volatilities of the disturbances and, where applicable, the output-coefficient of the Taylor rule, are set to match the above-mentioned moments for an initial set of expectations parameters. As shown in Table 2, this implies standard deviations of 0.01 for the natural rate and monetary policy shock and 0.5 for the cost-push shock. In sections of the paper where I assume a Taylor rule,

[13] The output gap measure is constructed as the difference between real GDP from the Bureau of Economic Analysis (BEA) and the Congressional Budget Office's (CBO) estimate of real potential output.

[14] On the lower end of the empirical values for α, Bils and Klenow (2004) find a mean duration of 4.3 months and Klenow and Malin (2010) 6.9 months. Klenow and Kryvtsov (2008) and Nakamura and Steinsson (2008) agree on between 7-9 months, while Eichenbaum et al. (2011)'s number is 10.6 months.

Table 1: Parameters calibrated from literature

β	0.98	stochastic discount factor
σ	1	intertemporal elasticity of substitution
α	0.5	Calvo probability of not adjusting prices
κ	0.0842	slope of the Phillips curve
ψ_π	1.5	coefficient of inflation in Taylor rule*
$\bar{g}$	0.145	initial value of the gain

Table 2: Parameters set to match data moments

ψ_x	0.3	coefficient of the output gap in Taylor rule*
σ_r	0.01	standard deviation, natural rate shock
σ_i	0.01	standard deviation, monetary policy shock*
σ_u	0.5	standard deviation, cost-push shock

*Parameters with an asterisk refer to sections of the paper where a Taylor rule is in effect.

the moment-matching exercise results in a 0.3 coefficient on the output gap. The last parameter of the New Keynesian core is the inflation coefficient of the Taylor rule, ψ_π. Unless otherwise specified, I set ψ_π to 1.5, the value recommended by Taylor (1993). Note that the asterisks in Tables 1 and 2 demarcate parameters that pertain to the Taylor rule and thus only to sections of the paper which assume that a Taylor rule is in effect.

I initialize the value of the gain at $\bar{g} = 0.145$. This is Carvalho et al. (2021)'s estimate for the case of unanchored expectations in their discrete anchoring function model. In the case of a discrete anchoring function, as in Carvalho et al. (2021), this parameter has important implications for model dynamics because by construction, the gain takes on this value very frequently. However, since my specification for $\mathbf{g}(\cdot)$ is smooth, $\bar{g}$ does not have a strong bearing on model dynamics.

4.2 Estimation

In inferring the functional form of the anchoring function $\mathbf{g}(\cdot)$, I proceed as follows. First, since the analysis in Section 5 relies on Assumptions 1 and 2, I similarly impose these on the estimation. In other words, I seek to back out the functional form of the relationship between the size of the gain and forecast errors in inflation, $k_t = \mathbf{g}(f_{t|t-1})$.

While it would be very interesting to estimate the general gain function, $k_t = \mathbf{g}(k_{t-1}, f_{t|t-1})$, the number of moments available constrains the amount of parameters I can estimate with precision. The assumption of the simple gain function that does not depend on past gains is thus also helpful

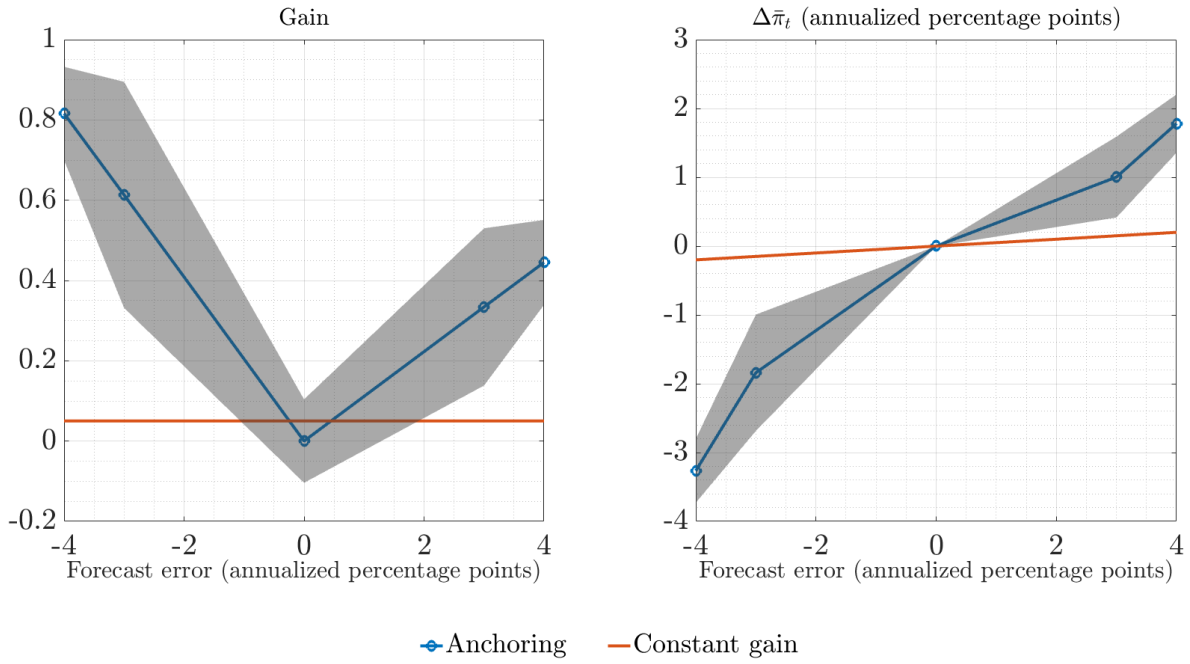
to reduce the number of estimated parameters below the number of moments.[15]

To be as close as possible to a non-parametric estimate of $\mathbf{g}(\cdot)$, while at the same time preserve the shape of the function, I employ a piecewise linear approximation of the form:

$$\mathbf{g}(f_{t|t-1}) = \sum_i \gamma_i b_i(f_{t|t-1}). \tag{31}$$

Here $b_i(\cdot)$ is a piecewise linear basis and γ is a vector of approximating coefficients. The index i refers to the breakpoints of the piecewise linear approximation. As explained in Appendix G, I estimate γ by simulated method of moments, targeting the autocovariance structure of the observables of the model and expectations.

Figure 2: The estimated gain function



Estimates for 5 knots (indicated by circles), cross-section of size $N = 1000$. The shaded gray areas are 95% confidence intervals constructed using Wald's method from 100 bootstrap samples. The constant gain takes the consensus value of 0.05 for the gain.

The left panel of Fig. 2 presents the estimated coefficients: $\hat{\gamma} = (0.82; 0.61; 0; 0.33; 0.45)$. For comparison, the red line plots a constant gain model with the gain calibrated to the consensus

[15] In an attempt to overcome this, I have experimented with ways of estimating the anchoring function directly using data on short-run and long-run inflation expectations. In particular, an alternative is to select the parameters of the gain function so as to make the updating equation in (28) hold in every period for every forecaster in the SPF. The advantage of this approach is that a much larger number of moments allows many more parameters to be estimated. In theory, this should make it possible to estimate the general gain function $k_t = \mathbf{g}(k_{t-1}, f_{t|t-1})$. The issue with this approach is that long-run inflation expectations data (5 or 10-year ahead) are only available starting in the early 1990s. This means that one misses the time period where most of the action in long-run expectations is. This is what lead me to opt for the approach presented here instead.

value from the empirical literature (see below). The circles on the blue line indicate the nodes of the piecewise linear approximation. The elements of $\hat{\gamma}$ represent the value of the gain the private sector chooses when it observes a forecast error of a particular magnitude. For example, a forecast error of -4 pp in inflation is associated with a gain of 0.82. To make these numbers easier to interpret, the right panel of Fig. 2 depicts what the elements of $\hat{\gamma}$ imply for the updating of long-run inflation expectations. For instance, a gain of 0.82 that arises when the private sector observes a -4 pp forecast error implies a downward revision in long-run inflation expectations by about 3.2 pp. Also here, the red line corresponds to a constant gain case. For the constant gain model, the -4 pp forecast error implies a 0.1 pp downward revision in long-run expectations.

Before interpreting the magnitudes of the estimated anchoring function, it is helpful to work out what amounts to a large gain. In particular, since the gain is the model's metric of the extent of unanchoring, from what size of the gain should one consider expectations to be significantly unanchored? The consensus in the literature on estimating learning gains is that if the true model is one with constant gain learning, then the gain lies between 0.01-0.05. On the higher end of the empirical estimates, Branch and Evans (2006) obtain a constant gain on inflation of 0.062. Milani (2007) finds 0.0183. The estimates of the maximal value for endogenous gains lie somewhat higher, at 0.082 in Milani (2014) and at 0.145 in Carvalho et al. (2021). While Eusepi and Preston (2011) find that the value of 0.002 is sufficient to significantly alter the dynamic behavior of a standard RBC model, calibrated models tend to use values between 0.01-0.05. Since the estimation and calibration literature has settled on the number 0.05 in particular as the benchmark, I calibrate the constant gain model I use for comparison throughout the paper using this value. Accordingly, the constant gain model on Fig. 2 also features a 0.05 gain.

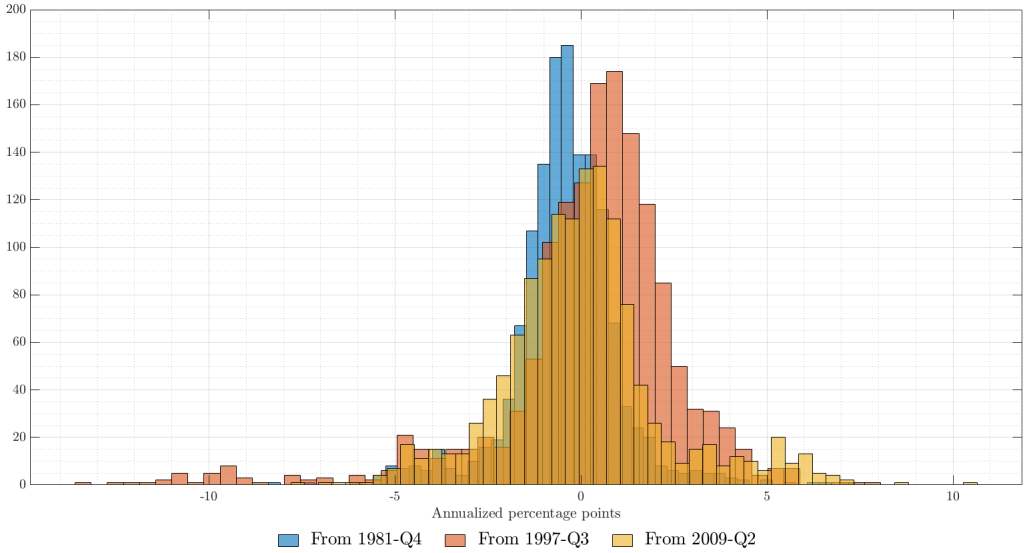
An intuitive interpretation of the gain is that its inverse gives the number of past observations the private sector uses to form its current forecasts. Eusepi and Preston (2011)'s number, 0.002, thus implies that firms and households rely on the last 125 years of data. By contrast, the consensus number of 0.05 translates to using 5 years of data.

Seen from both of these perspectives, the consensus gain value of 0.05 corresponds to a sizable, but not excessive sensitivity to surprises. With this benchmark number in mind, the message of Fig. 2 seems stark: the estimated coefficients are very large. The highest gain value in my estimation, 0.82, can be interpreted as seeing forecast errors of 4 pp in absolute value prompting the private sector to discount any observations older than about 5 months. That is a very short time.

Presumably, such large forecast errors are rare, however. To confirm this, Fig. 3 plots individual-level forecast errors in the SPF, split into three distinct time periods so as to have roughly the same number of forecast errors in each bin. As the figure shows, the bulk of surprises are approximately in the $[-3, 3]$ percentage point interval. What is also noticable is that forecast errors changed pattern over time. In the first time period, inflation tended to surprise forecasters on the downside, while in the second period (roughly between 1997 and 2010) it instead mostly surprised on the upside.[16] In the last decade, the tails of the distribution shrank, reflecting more stable inflation, less inflation surprises and thus a smaller need for forecasters to reconsider their view on the long-run

[16] In this period, there are some outliers with extreme downward surprises, which may be due to the Great Recession.

Figure 3: Forecast errors in the SPF data



1-quarter-ahead forecast errors of individual forecasters in the SPF, 1981-Q4 - 2021-Q4, annualized percentage points. The forecast errors are split into three non-overlapping bins with roughly the same number of observations.

mean of inflation.

A changing pattern in inflation surprises is thus a feature of the data. The anchoring framework allows forecasters to adjust their forecasting behavior according to the volatility of their environment, which is something they would not be able to do if they used a constant gain forecasting model.[17] One sees this by looking at the right panel of Fig. 2. Whereas the anchoring framework allows forecasters to change their long-run inflation expectations more sensitively in volatile environments, the constant gain implies the same sensitivity across the board.

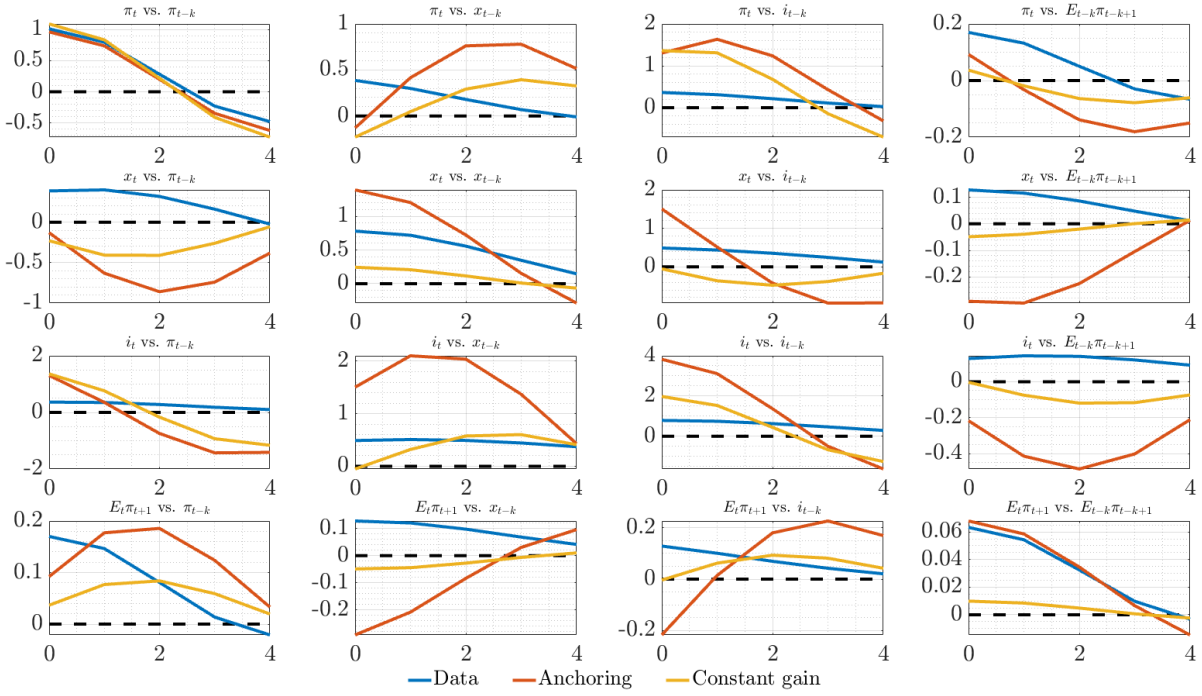
That time-varying anchoring is a feature consistent with data can be also seen on the autocovariogram, shown in Fig. 4, which plots the moments of the model (red) against those in the data (blue). For comparison, I also plot moments generated by a constant gain learning model with the consensus value for the gain of 0.05 (yellow).

While a constant gain can match inflation dynamics (top left panel), its fit deteriorates compared to the anchoring model when it comes to the output gap (second diagonal panel) and, importantly, inflation expectations (bottom right panel). The reason the anchoring model fits the expectations data better than a constant gain alternative is because it does not constrain the private sector to always respond to surprises in the same way. In a volatile environment, such as the 1980s, high degrees of unanchoring allow the private sector to update its view very strongly on what kind of inflation regime it is in. With low and stable inflation instead, such as in the 2000s, occasional small inflation surprises do not incite the private sector to change its mind about average inflation. With a constant gain, instead, the private sector would respond as sensitively during

[17]Note that forecasting using Kalman filtering corresponds to a constant gain model with the gain given by the underlying signal-to-noise ratio of the true data-generating process.

the 2000s as during the 1980s. That is at odds with the data.

Figure 4: Autocovariogram



Estimates for a cross-section of $N = 1000$. Moments are computed for the best estimate for the sample length of 151 with 200 additional burn-in periods. The constant gain case shows the moments of the model for the consensus value of the gain of 0.05 from the literature.

We also notice that unanchoring introduces another flexibility to expectations formation compared to constant gains: asymmetry. As seen on Fig. 2, negative inflation surprises raise the gain about twice as much as positive ones do, and lead to almost twice as large updates in long-run inflation expectations. Consistent with the findings of Hebden et al. (2020), this indicates that long-run expectations are more sensitive to negative than to positive surprises. In other words, surprises on the downside are almost twice as likely to unanchor expectations as those on the upside. This has some interesting implications for the current time period, where rising inflation in the wake of the Covid-19 crisis has resulted in a series of very large inflation surprises. On the one hand, this is clearly putting upward pressure on the expectations anchor. On the other, the asymmetry of unanchoring implies that this situation would be more concerning if inflation was surprising on the downside rather than on the upside.

5 Monetary policy and anchoring

This section sets up and solves the optimal monetary policy problem in the model with the anchoring expectation formation. In Section 5.1, I begin by analyzing the Ramsey problem of determining optimal paths for the endogenous variables that policy seeks to bring about. While the anchoring

mechanism introduces substantial nonlinearity into the model, it is possible to derive analytically an optimal target criterion for the policymaker to follow. As we shall see, the optimal rule prescribes for monetary policy to act conditionally on the stance of expectations, and will thus be time-varying and state-dependent. In particular, whether expectations are anchored or not matters for the extent to which there is a tradeoff between inflation and output gap stabilization, and also for the volatility cost of getting expectations anchored.

I then turn to the question of how to implement optimal policy. Section 5.2 uses global methods to solve for the interest rate sequence that implements the target criterion. I then discuss the properties of the optimal interest rate policy and why it is successful in stabilizing both inflation expectations and inflation.

The optimal interest rate policy will thus be a nonlinear function of all the states in the model - a complicated object to compute. In practice, however, monetary policy is most commonly modeled using time-invariant rules like the Taylor rule that are both simple to compute and to communicate to the public. In Section 5.3, I therefore restrict attention to Taylor-type feedback rules for the interest rate. I solve for the optimal Taylor-rule coefficient on inflation numerically and investigate how this choice affects the anchoring mechanism.

A Taylor rule with given coefficients involves higher fluctuations in the anchoring model than under rational expectations because it does not allow the central bank to respond to long-run expectations directly. As we will see, whether an aggressive interest-rate response to inflation subdues this volatility or amplifies it depends crucially on the specification of interest-rate expectations. For the benchmark assumption that constrains interest-rate expectations to be formed exactly as under the Ramsey policy, the optimal Taylor-rule inflation-coefficient is higher than under rational expectations, as in Orphanides and Williams (2004). For an alternative assumption that allows expected interest rates to have a drift, however, one obtains Eusepi et al. (2018a)'s conclusion that monetary policy just inflicts too much excess volatility, and should therefore be less aggressive than under rational expectations.

5.1 The Ramsey policy under anchoring

I assume the monetary authority seeks to maximize welfare of the representative household under commitment. As shown in Woodford (2003a), a second-order Taylor approximation of household utility delivers a central bank loss function of the form

$$L^{CB} = \mathbb{E}_t \sum_{T=t}^{\infty} \{(\pi_T - \pi^*)^2 + \lambda_x (x_T - x^*)^2\}, \tag{32}$$

where λ_x is the weight the central bank assigns to stabilizing the output gap. In the rational expectations New Keynesian model, λ_x is a function of deep parameters: $\lambda_x = \frac{\kappa}{\theta}$. Just as under rational expectations, one can show that it is optimal to set the central bank's targets, π^* and x^*, to zero. The central bank's problem, then, is to determine paths for inflation, the output gap and the interest rate that minimize the loss in Equation (32), subject to the model equations (10) and (11),

as well as the PLM (19), the learning equations (21) and (22), and the anchoring function (23). The full statement of the Ramsey problem under Assumptions 1 and 2 is as follows:

$$
\begin{aligned}
&\min_{\{\pi_t, x_t, i_t, \bar{\pi}_{t-1}, k_t\}_{t=t_0}^{\infty}} \mathbb{E}_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0}(\pi_t^2 + \lambda_x x_t^2) \qquad \text{s.t.} \\
&x_t = \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} \beta^{T-t}((1-\beta)x_{T+1} - \sigma(\beta i_T - \pi_{T+1}) + \sigma r_T^n), \\
&\pi_t = \kappa x_t + \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t}(\kappa\alpha\beta x_{T+1} + (1-\alpha)\beta\pi_{T+1} + u_T), \\
&\hat{\mathbb{E}}_t \pi_{t+k} = \bar{\pi}_{t-1} + b_1^k s_t \quad \forall k \geq 1, \\
&\bar{\pi}_t = \bar{\pi}_{t-1} + k_t f_{t|t-1}, \\
&f_{t|t-1} = \pi_t - \hat{\mathbb{E}}_{t-1}\pi_t, \\
&k_t = \mathbf{g}(f_{t|t-1}),
\end{aligned} \tag{33}
$$

where $\mathbb{E}$ is the central bank's expectation. The use of the mathematical expectation $\mathbb{E}$ reflects the assumption that the central bank has model-consistent expectations and observes the private sector's expectations. This assumption, which Gaspar et al. (2010) refer to as "sophisticated central banking," is quite strong. In practice, it is likely that the central bank's measure of private sector expectations is noisy at best. Nevertheless, it is a useful case that can serve as a benchmark for future research.

A last thing to note about the Ramsey problem in the anchoring model is that while it is set up under commitment, the optimal plans under commitment and discretion will coincide, and will thus not be subject to the time-inconsistency problem of Kydland and Prescott (1977). As shown by Mele et al. (2019) and Molnár and Santoro (2014), this is a general feature of adaptive learning models and comes from the fact that adaptive expectations cannot incorporate promises of the policymaker.

5.1.1 Optimal Ramsey policy as a target criterion

As foreshadowed above, the nonlinearity of the model due to the anchoring function prevents a full analytical solution to the Ramsey problem. Therefore I now characterize the first-order conditions of the problem analytically, and proceed in Section 5.2 to solve the full problem numerically. The details of the derivations are given in Appendix H, which also illustrates how the endogeneity of the gain introduces nonlinearity into the model.

The solution of the Ramsey problem is stated in the following proposition.

Proposition 1. *Target criterion in the anchoring model*

The targeting rule in the simplified learning model with anchoring is given by

$$\pi_t + \frac{\lambda_x}{\kappa} x_t = c\left(k_t + f_{t|t-1}\mathbf{g}'_t\right) \mathbb{E}_t \left((1-\beta)x_{t+1} + \beta \sum_{i=1}^{\infty} x_{t+i} \prod_{j=0}^{i-1} 1 - k_{t+j} - f_{t+j|t-1+j}\mathbf{g}'_{t-1+j}\right), \quad (34)$$

$$c \equiv \frac{\lambda_x}{\kappa} \frac{(1-\alpha)\beta^2}{1-\alpha\beta}. \quad (35)$$

Proved in Appendix H. For a general target criterion without Assumption 2, see Appendix I.

Here $\mathbf{g}'_t$ denotes the derivative of $\mathbf{g}(\cdot)$ with respect to the forecast error $f_{t|t-1}$, evaluated at time t. Note that $\mathbf{g}' = \mathbf{g}_\pi = -\mathbf{g}_{\bar{\pi}}$.

The interpretation of Equation (34) is that the *intra*temporal tradeoff between inflation and the output gap due to cost-push shocks is complemented by two *intertemporal* tradeoffs. One is due to learning in general and corresponds to Molnár and Santoro (2014)'s result, and one is due to anchoring in particular, and is thus novel here. The first intertemporal effect comes from the current level of the gain, k_t, which captures how far learning is from converging to rational expectations. The second, novel intertemporal tradeoff is manifest in the derivative of the anchoring function today, $\mathbf{g}'_t$, as well as in all expected levels and changes in the gain in the future in the expression $(1 - k_{t+j} - f_{t+j|t-1+j}\mathbf{g}'_{t-1+j})$ in the second bracket on the right-hand side. These expressions say that the presence of anchoring qualifies the first intertemporal tradeoff because now the degree and direction in which the gain changes today and is expected to change in the future matter too. In other words, the central bank needs to consider whether its chosen interest rate sequence contributes to anchoring expectations in future periods, or whether it actually serves to unanchor them.

Let me investigate these channels in isolation. To see exactly what the role of anchoring is in the target criterion, consider first the special case of exogenous gain adaptive learning, for simplicity with a constant gain specification. In this case the anchoring function and the forecast error are irrelevant (since $\mathbf{g}'_t = 0, \ \forall t$) and (34) boils down to

$$\pi_t = -\frac{\lambda_x}{\kappa} x_t + \frac{\lambda_x}{\kappa} \frac{(1-\alpha)\beta^2}{1-\alpha\beta} k\, \mathbb{E}_t \left((1-\beta)x_{t+1} + \sum_{i=0}^{\infty} x_{t+i+1}(1-k)^i\right), \quad (36)$$

which replicates Molnár and Santoro (2014)'s expression, with the exceptions that the parameters $\frac{(1-\alpha)\beta^2}{1-\alpha\beta}$ and an extra $(1-\beta)x_{t+1}$ appear.[18] These differences come from the fact that in my model the entire term structure of expectations enters the evolution equations of inflation and output, while in Molnár and Santoro (2014), only one-period-ahead expectations do.

As Molnár and Santoro (2014) emphasize, this result suggests that already the presence of learning by itself is responsible for the first intertemporal tradeoff between inflation and output gap stabilization. The fact that the central bank now has future output gaps as a margin of adjustment means that it does not have to face the full tradeoff in the current period. Learning allows

[18] Since Molnár and Santoro (2014) has no equation number at the relevant expression, I refer the reader to Equation (24) of Gaspar et al. (2010), who provide a parsimonious summary of Molnár and Santoro (2014).

the central bank to improve the current output gap without sacrificing inflation stability today; however, this results in a worsened tradeoff in the future. In other words, adaptive learning by itself allows the central bank to postpone the current tradeoff to later periods.

Intuitively, this happens because adaptive expectations are slow in converging to rational expectations. In the transition, the private sector's expectations do not adjust to fully internalize the intratemporal tradeoff. This gives the monetary authority room to transfer the tradeoff to the future.

Contrasting Equations (36) and (34) highlights the novel role of the anchoring channel. With anchoring, the extent to which policy can transfer the intratemporal tradeoff to future periods depends not only on the stance of the learning process, as in (36), but also on whether expectations are anchored, and in which direction they are moving. In fact, not only the current stance and change of anchoring matters, but also all expected future levels and changes.

Anchoring, however, complicates the possibility of transferring today's tradeoff to the future. One can see this on the fact that forecast errors and the derivatives of the anchoring function are able to flip the sign of the second term in (34). This means that anchoring can alleviate or worsen the intertemporal tradeoff. To see the intuition, consider the equation system of first-order conditions from solving the Ramsey problem. While the full system is presented in Appendix H, let us focus solely on Equation (H.2), the equation governing the dynamics of observables in the model:

$$2\pi_t = -2\frac{\lambda_x}{\kappa}x_t + \varphi_{5,t}k_t + \varphi_{6,t}\mathbf{g}'_t. \tag{37}$$

The Lagrange multipliers $\varphi_5 \geq 0$ and $\varphi_6 \geq 0$ are the multipliers of the updating equation (21) and the anchoring function respectively. This equation, upon substitution of the solutions for the two multipliers, yields the target criterion. It is therefore easy to read off the intuition at a glance. First, since $\varphi_{5,t}k_t > 0$, one immediately obtains the above-discussed conclusion that as long as the adaptive learning equation is a constraint to the policymaker ($\varphi_{5,t} > 0$), the central bank has more room to transfer the contemporaneous tradeoff between inflation and the output gap to the future.[19]

However, whether the anchoring equation alleviates or exacerbates the inflation-output gap tradeoff depends on the sign of $\mathbf{g}'_t$. If the derivative is positive, the effect is the same as above, and the central bank has more leeway to postpone the tradeoff to the future. By contrast, if the derivative is negative, that is expectations are becoming anchored, the intratemporal tradeoff is worsened.

Why do unanchored expectations give the central bank the possibility to postpone its current inflation-output gap tradeoff? The reason is that when expectations become unanchored, the learning process is restarted. A not-yet converged learning process implies, as discussed above,

[19] φ_5 and φ_6 are zero if the learning process has converged. In absence of shocks, this eventually happens because decreasing forecast errors lead to decreasing gains. But with exogenous disturbances, forecast errors may increase throughout the lifetime of the economy, raising the gain and restarting the learning process. The model thus exhibits a weak form of weak expectational stability.

that postponing the tradeoff is possible. Restarting the convergence process thus unlocks this possibility.

This seems to suggest that from a smoothing standpoint, the central bank should prefer to have unanchored expectations. As will be shown in Sections 5.2-5.3, volatility considerations will suggest otherwise. But in fact, even the smoothing viewpoint involves some ambiguity on whether expectations should be anchored from the perspective of the central bank. Clearly, the central bank prefers to face a learning process that on the one hand has not yet converged, and on the other is converging only slowly. A high gain under unanchored expectations implies both a sizable distance from convergence as well as faster learning and thus faster convergence. Therefore, ideally the central bank would like to have expectations anchored but the gain far from zero; a contradiction. Once the gain approaches zero, only unanchored expectations can raise it again to restart the learning process. But once the gain is large, the only way to slow down learning is to anchor expectations, that is, to lower the gain.

5.2 Implementing the Ramsey policy: the optimal interest rate sequence

Having a characterization of optimal policy in the anchoring model as a first-order condition, the next relevant question is how the central bank should set its interest rate tool in order to implement the target criterion in (34). In other words, we would like to know what time-path of interest rates implements the optimal sequence of inflation and output gaps. As emphasized in Section 5.1, the nonlinearity of the model does not admit an analytical answer to this question. I therefore solve for the optimal interest rate policy numerically using global methods. I rely on the calibration presented in Tables 1 and 2 and the estimated parameters of the expectations process in Section 4. Furthermore, I set the central bank's weight on output gap fluctuations, λ_x, to 0.05, the value estimated by Rotemberg and Woodford (1997).

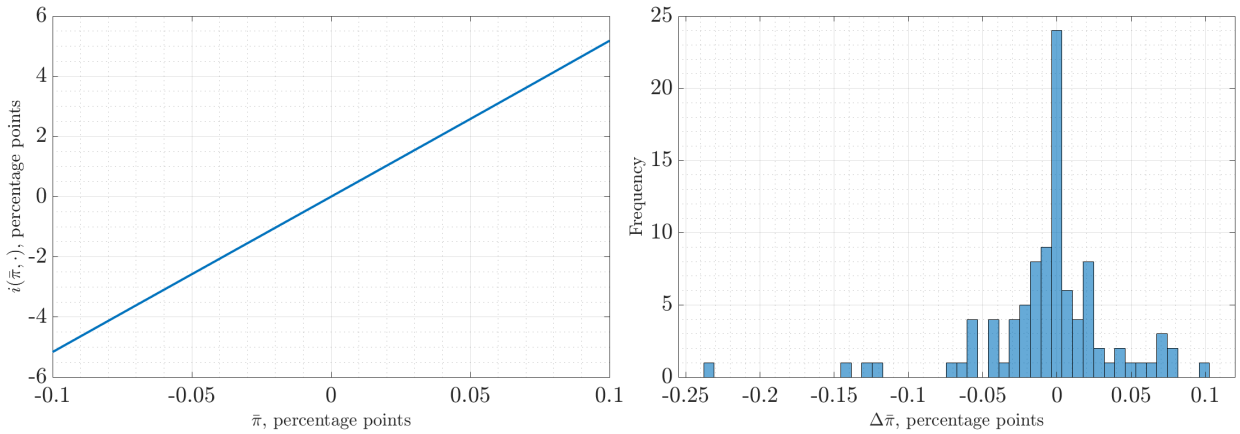
Appendix J outlines my preferred solution procedure, the parameterized expectations approach, while Appendix K gives the details of the parametric value function iteration approach I implement as a robustness check. The main output of this procedure is an approximation of the optimal interest rate policy as a function of the vector of state variables. Due to Assumption 2, the relevant state variables are expected mean inflation and the exogenous states at time t and $t-1$, rendering the state vector five-dimensional:

$$X_t = (\bar{\pi}_{t-1}, r_t^n, u_t, r_{t-1}^n, u_{t-1}). \tag{38}$$

As a first step, I plot how the approximated policy function depends on $\bar{\pi}_{t-1}$, while keeping all the other states at their mean. The result, depicted on Panel (a) of Fig. 5, suggests that optimal interest-rate setting responds linearly and very sensitively to the stance of expectations, $\bar{\pi}_{t-1}$. If expected mean inflation decreases by 5 basis points, the interest rate drops by about 250 basis points.[20]

[20]This result is qualified if one relaxes Assumption 1. If the private sector learns about all observables, not just inflation, interest-rate expectations play a major role in determining forecasts of future inflation and thus add a stabilizing

Figure 5: Policy function and implied volatility in long-run expectations



(a) $i(\bar{\pi}, \text{all other states at their means})$

(b) Histogram of changes in $\bar{\pi}$

This is a large response. Clearly, optimal policy involves subduing unanchored expectations by injecting massive negative feedback to the system. One may then wonder why optimal policy is so aggressive on unanchored expectations when the analysis of the target criterion in Section 5.1 suggested that learning can alleviate the stabilization tradeoff between output and inflation.

The reason is that the anchoring expectation formation introduces another intertemporal tradeoff to monetary policy: a volatility tradeoff. One can see this on Fig. 6, portraying the dynamics of the system following a two-standard-deviation inflationary cost-push shock, conditional on a Taylor rule with baseline parameters. The figure contrasts the rational expectations version of the model with a well-anchored, weakly anchored and strongly unanchored scenario in the anchoring model, as well as with a constant gain model.

The same shock that under rational expectations completely vanishes by the second period, triggers a large, persistent and oscillatory response in the anchoring model.[21] Clearly, the size and persistence of the shock, as well as the magnitude of oscillations increase the more unanchored expectations are. This comes from the fact that if expectations are anchored, stable expectations lower the pass-through between shocks and observables. Instead, if expectations are unanchored, they become volatile, passing through the shocks and amplifying them. Having unanchored expectations, then, comes at a volatility cost in the central bank's target variables.

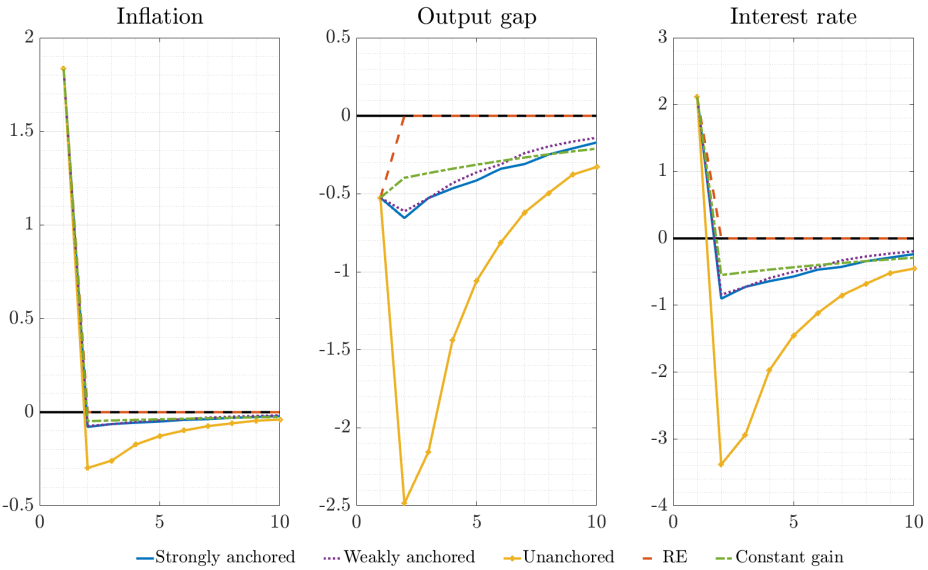
This volatility cost is also present in the constant gain model, and its size depends on the size of the gain. But as the literature on estimating constant gain models shows (Branch and Evans, 2006, Milani, 2007), a too large constant gain is inconsistent with data, thus putting an upper bound on

channel that is absent from the current specification. In the more general case, then, an order of magnitude smaller responses in the current interest rate are sufficient. For more on the role of interest-rate expectations, see Section 5.3.

[21] As periodically noted in the literature, adaptive learning models tend to produce impulse responses that exhibit damped oscillations. Authors making explicit note of this phenomenon include Evans and Honkapohja (2001), Evans et al. (2013) and Anufriev and Hommes (2012). The reason is that under an adaptive learning framework, forecast errors following an impulse are oscillatory. In fact, the higher the learning gain, the higher the amplitude of forecast error oscillations. Appendix L presents a simple illustration for why this is the case.

how much excess volatility can come from a constant gain model. Also, the constant gain model implies the same volatility cost to a particular shock in any state of the world. A state-dependent gain, instead, introduces the distinction between the effects of the same shock when expectations are unanchored as opposed to anchored. As seen on the difference between the blue and yellow lines, the same shock leads to much more volatility if expectations are unanchored when the shock hits than if they are strongly anchored.

Figure 6: Impulse responses after a cost-push shock



Shock imposed at $t = 25$ of a sample length of $T = 400$ (with 200 additional burn-in periods), cross-sectional average with a cross-section size of $N = 100$. The remark on whether expectations are anchored or not refers to whether the gain is below the 10^{th}, below the 50^{th}, or above the 90^{th} percentile of simulated gains at the time the shock hits. For the constant gain model, the gain is set to the consensus value of 0.05.

The fact that the volatility cost rises the more unanchored expectations are dictates that, in the long run, the central bank wishes to have expectations anchored. In the constant gain model, this would not be an option because the central bank would not be able to influence the size of the constant gain. In the anchoring model, instead, the central bank can affect the expectation formation process.

What complicates the central bank's life is that anchoring expectations itself comes at a convex volatility cost. Anchoring expectations requires an aggressive interest response because by these means the central bank can introduce negative feedback to the system. But changes in the interest rate surprise the private sector, raising forecast errors. The more unanchored expectations are, the more volatility the interest rate movement inflicts on the economy.

We can see from the policy function how optimal policy resolves this tradeoff: it reacts extremely aggressively to movements in long-run expectations. This way, the central bank hopes to avoid even larger interventions that would become necessary were expectations to unanchor further. To avoid having to pay so high a price, the central bank is extra aggressive in the short run to

prevent massive unanchoring from ever materializing. Thus the optimal response to the volatility tradeoff is to temporarily increase volatility in order to reduce it in the long-run. In this way, the central bank's aggressiveness in the model is driven by the desire to prevent upward (downward) drifting long-run expectations from becoming a self-fulfilling inflationary (deflationary) spiral.

The large interest-rate responses resemble the idea advocated by Goodfriend (1993) that the central bank moves to offset "inflation scare" (or "deflation scare") episodes. As Goodfriend shows, it was historically not uncommon to move the interest rate by hundreds of basis points to subdue inflation scares. For example, in March 1980, the Fed raised the interest rate by 230 bp to convince the public that it would not tolerate high inflation. The optimal policy function prescribes that this is exactly what the policymaker should do to fight unanchored expectations.

What Fig. 5 also suggests is that acting aggressively in the short run indeed delivers the long-run benefits of stabilizing economic fluctuations. As Panel (b) depicts, employing the optimal policy implies that realized changes in long-run inflation expectations are very small. As seen on the histogram of $\Delta\bar{\pi}$, with the optimal policy in place, the model spends most of its time in the region of minuscule fluctuations in $\bar{\pi}$. The mode of the distribution is a change of 0.3 basis points in absolute value, implying that in normal times, the central bank only needs to raise or lower the interest rate by 15 basis points. In other words, the aggressive nature of optimal policy allows the central bank to keep expectations anchored or quickly reanchor them following shocks. In this way, the monetary authority eliminates as much volatility stemming from unanchored expectations as it possibly can.

5.3 Optimal Taylor rule under anchoring

Monetary policy is often formulated using a Taylor rule. Proponents of such a characterization, like Taylor (1993) himself, emphasize the benefits of having a simple, time-invariant and easily verifiable rule. Also in the anchoring model, a policymaker may thus be interested in using a Taylor-type approximation to optimal policy in order to combine the benefits of having a simple, yet near-optimal rule.[22] Therefore I now consider the restricted set of Taylor-type policy rules and ask what value of the time-invariant Taylor-rule coefficient on inflation is optimal in the case of the anchoring model.

In this section, I thus restrict attention to a standard Taylor rule:

$$i_t = \psi_\pi(\pi_t - \pi^*) + \psi_x(x_t - x^*) + \bar{i}_t, \tag{39}$$

where ψ_π and ψ_x represent the responsiveness of monetary policy to inflation and the output gap respectively. Lastly, $\bar{i}_t$ is a monetary policy shock. I also assume that when the Taylor rule is in effect, the central bank publicly announces this. In order to avoid comparing apples with oranges, I start by assuming that, exactly as in the Ramsey policy, interest-rate expectations are formed using the Taylor rule with rational expectations of inflation. This is a crucial assumption which

[22] Recall from Woodford (2003a) that even under rational expectations, a standard Taylor rule is not fully optimal because its purely forward-looking nature precludes the use of promises of future policy.

will determine whether the results corroborate the findings of Orphanides and Williams (2004) or Eusepi et al. (2018a). I relax the assumption and discuss its implications at the end of the section.

I compute the optimal Taylor rule coefficient on inflation numerically by minimizing the central bank's expected loss in a cross-section of $N = 100$ simulations of the rational expectations, constant gain learning and anchoring models. I continue to use the calibration of Tables 1 and 2 and to parameterize the anchoring function using the estimated approximating coefficients of the gain function ($\hat{\gamma}$ in Section 4). For the constant gain model, I continue using the consensus gain value of 0.05.

Table 3 presents the optimal Taylor rule coefficient ψ_π for the three models. The message is clear: in the anchoring model, the central bank needs to work much harder than in the other models and push strongly against inflation fluctuations. This is necessary because mistakes by the central bank in achieving the inflation target unanchor expectations, but unanchored expectations feed back into inflation, leading to further unanchoring. To get out of this unpleasant feedback loop, the central bank moves the interest rate exceedingly aggressively in response to movements in inflation. Corroborating the conventional wisdom in the adaptive learning literature, for example in Orphanides and Williams (2004), I thus find that monetary policy specified as a Taylor rule should be more aggressive on inflation than what would be optimal under rational expectations. As I explain below, however, one can overturn this result by maintaining alternative assumptions on interest rate expectations.

Table 3: Optimal coefficient on inflation, RE against anchoring

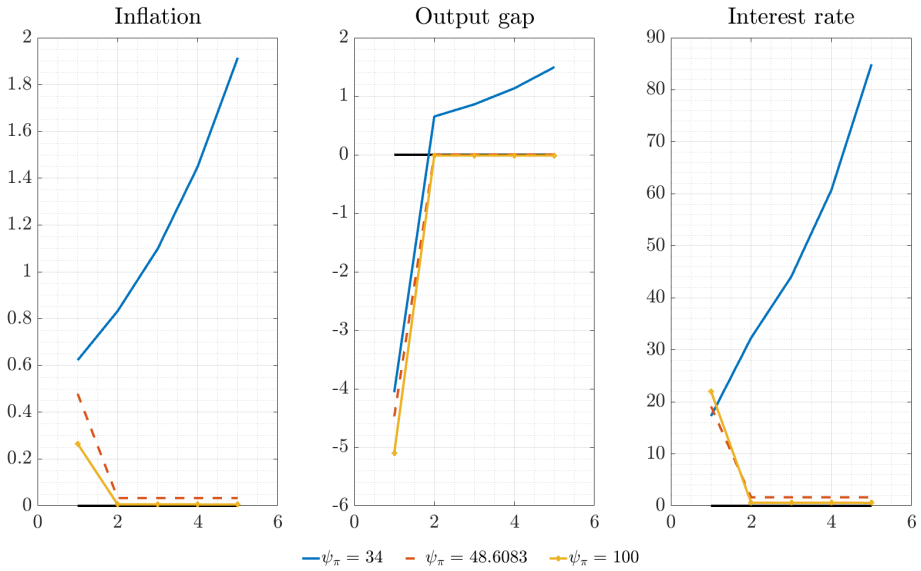
$\psi_\pi^{*,RE}$	$\psi_\pi^{*,Anchoring}$	$\psi_\pi^{*,Constant\ gain}$
2.2079	48.6083	45.6499

Sample length is $T = 400$ with a cross-section of $N = 100$ and 200 additional burn-in periods. For the constant gain model, the gain is set to the consensus value of 0.05.

The reason behind this high aggressiveness is that when the central bank's interest rate rule is not state-dependent, it moves the interest rate too little when expectations are unanchored, and too much when they are anchored. We can see this on Fig. 7, which depicts impulse responses of the model for an inflationary cost-push shock for various levels of inflation aggressiveness ψ_π. A too small value of $\psi_\pi = 34$ relatively to the optimal value imposes too little negative feedback to get inflation back to its pre-shock value. A too high value of ψ_π instead is effective in stabilizing inflation, but also involves costs in terms of making interest rates and output very volatile. But as the figure suggests, it is relatively less costly to err on the upside, i.e. to choose a too high ψ_π, than on the downside, since a too low inflation coefficient in the Taylor rule is not sufficient to stabilize the economy at all.

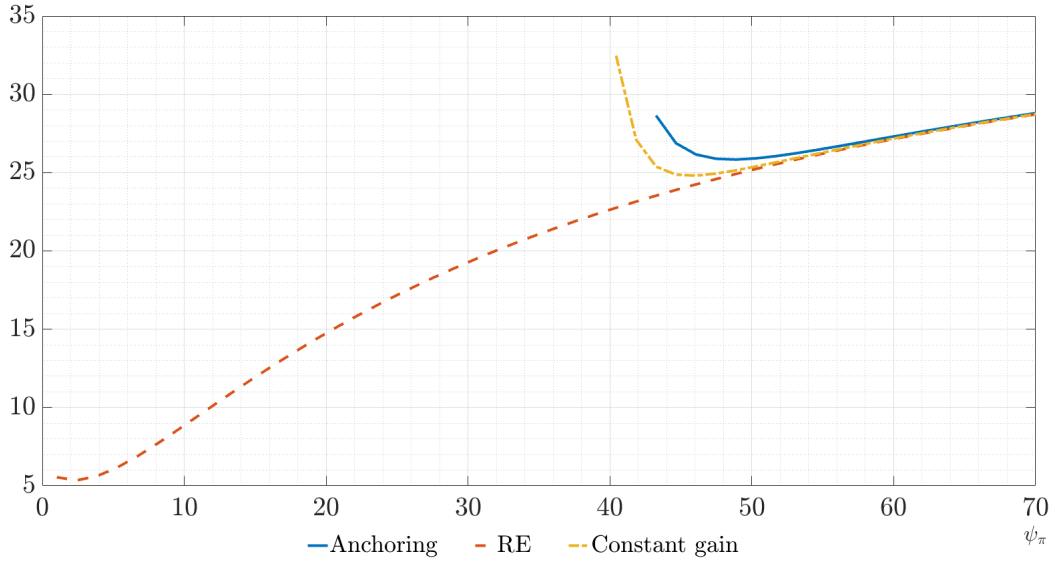
This intuition is reinforced in Fig. 8, which depicts the central bank's loss as a function of the inflation coefficient ψ_π. One notices that for values of ψ_π below about 42, the loss in the anchoring

Figure 7: Impulse responses for various values of ψ_π



Cost-push shock imposed at $t = 25$ of a sample length of $T = 400$ (with 200 additional burn-in periods), cross-sectional average with a cross-section size of $N = 100$.

Figure 8: Central bank loss as a function of ψ_π



Sample length is $T = 400$ with a cross-section of $N = 100$ and 200 additional burn-in periods. Note that for the anchoring and constant gain models, I only plot the loss function if it takes on sufficiently small values. For the constant gain model, the gain is set to the consensus value of 0.05.

model goes to infinity so fast that it exits the figure. Intuitively, this comes from the fact that since ψ_π is too low in this case to stabilize the model, the explosive paths for inflation, output gaps and interest rates lead to infinite losses for the central bank. A too high Taylor coefficient, instead, while increasing the loss somewhat by virtue of making interest rates more volatile, is not nearly as

costly. In fact, in this range of ψ_π-values (about > 50), the loss in the anchoring model approaches that in the rational expectations version of the model. This makes sense because since a sufficiently high inflation coefficient stabilizes inflation expectations, fluctuations in expectations cease being the main driver of losses. Thus the wedge between RE and anchoring disappears, and in both cases the central bank just incurs losses from volatility coming from volatile interest rates.

A similar intuition applies to the constant gain model as to the anchoring model: the central bank needs to lean against fluctuating expectations to stabilize the economy. But there is a crucial difference. In the constant gain world, the central bank does not need to work as hard as in the anchoring model, and its loss is also always weakly below that in the anchoring model. This is because expectations are always (un)anchored to the same degree in a constant gain model. The central bank can do nothing to influence anchoring, and thus there are no benefits of acting extra aggressively today to anchor expectations tomorrow. Thus a lower inflation coefficient is sufficient to do the job.

The anchoring and the constant gain models look similar here in the sense that the anchoring model involves a higher optimal aggressiveness to subdue higher losses, but the order of the magnitude of losses and optimal coefficients is similar. This is because the fact that the Taylor rule policy is not state-dependent constrains the policymaker to act *as if* the anchoring model was a constant gain model. There is simply not much the policymaker can do to take state-dependent anchoring into account apart from being a little more aggressive.

This highlights that the time-varying nature of Ramsey-optimal policy is its key characteristic. The ability to take the stance of anchoring into account is what enables the central bank to be exceedingly aggressive if and only if expectations are about to unanchor, and maintain a dovish stance otherwise. Thus it seems reasonable to expect that if the central bank were allowed to select state-dependent Taylor-rule coefficients, it would choose low coefficients when expectations are anchored and high ones when expectations show signs of unanchoring. Such a central bank, as well as one following the optimal policy, would appear to the econometrician as time-varying, echoing the findings of Lubik and Matthes (2016).

Lastly, the optimal Taylor rule exercise entails an interesting lesson for optimal policy when the public is learning. I noted at the beginning of the section that interest-rate expectations are formed using a Taylor rule with rational inflation expectations. This assumption served to make the comparison between the Ramsey and the Taylor-rule policies fair. However, if one recomputes the inflation coefficient under the alternative assumption that the private sector forms interest-rate expectations with a Taylor rule with the anchoring inflation expectation, one obtains an optimal coefficient of 1.1, about half of the rational expectations value. In other words, in this case, optimal Taylor-rule policy should be *less* aggressive on inflation, not more.

In this case, then, one obtains Eusepi et al. (2018a)'s result that when there is a drift in long-run interest-rate expectations, there is an upper limit to how aggressive the central bank can be. In such a case, interest-rate expectations inherit the volatility of inflation expectations induced by an aggressive interest rate response. This leads to a lot of additional volatility on top of the volatility coming from inflation expectations. This excess volatility outweighs the benefits of stabilization.

In fact, as I remark in the beginning of Section 5.2, a similar conclusion holds for the Ramsey policy as well. If one introduces a drift in interest-rate expectations, for example by allowing the private sector to form interest-rate expectations using a Taylor rule with the anchoring expectation of inflation in it, the aggressiveness of the optimal Ramsey policy is scaled down. This points to that when central bankers conduct policy to take the anchoring and unanchoring of inflation expectations into account, it is instrumental for them to understand not only how inflation expectations, but also how interest-rate expectations are formed.

6 Conclusion

Central bankers frequently voice a concern to anchor expectations, that is, to render expectations of long-run inflation unresponsive to short-run economic conditions. The contribution of this paper is to investigate how this affects the conduct of monetary policy. I use a simple behavioral model in the spirit of Carvalho et al. (2021) to capture the notion of anchoring as the time-varying sensitivity of long-run expectations to short-run conditions. I quantify my anchoring mechanism by estimating the form of the function that determines the degree of unanchoring. I use the model to characterize monetary policy both analytically and numerically.

The simulated method of moments estimation establishes that the anchoring function provides a realistic description of expectations. On the one hand, expectations become more sensitive to forecast errors when the private sector made larger mistakes in predicting inflation in the past. On the other hand, like in Hebden et al. (2020), negative mistakes unanchor expectations more than positive ones of the same magnitude do. Moreover, I use the estimated anchoring function to back out an implied time series for the gain, the model's metric for the degree of unanchoring. The median gain, 0.098, is in line with estimates in the literature such as Milani (2014) and Carvalho et al. (2021), and implies that agents in the model discount observations older than 10 quarters when updating their long-run expectations.

Using the thus quantified model, I provide three sets of results on monetary policy. I first consider the Ramsey policy of the central bank, deriving an analytical target criterion that prescribes how the monetary authority should respond to shocks. I show that the presence of my novel anchoring channel makes it desirable and feasible to smooth out shocks over time. However, the extent this is feasible varies over time in tandem with the current and expected future degree of unanchoring.

Second, I use global methods to solve the nonlinear system of first-order conditions of the Ramsey problem numerically. I thus obtain an approximation to the optimal policy function, providing the fully optimal path of interest rates conditional on the sequence of exogenous disturbances. Like in Goodfriend (1993)'s account of US monetary policy, the optimal policy involves responding aggressively when expectations unanchor in order to suppress the volatility that high degrees of unanchoring cause. By contrast, the central bank need not intervene when expectations are anchored, resulting in accommodating inflation fluctuations in this case.

Lastly, I explore the implications of the model for the most common specification of monetary

policy, the Taylor rule. In the benchmark case with rational inflation expectations in the Taylor rule, the central bank should be much more aggressive on inflation in my model than under rational expectations, as in Orphanides and Williams (2004). This is because a time-invariant Taylor rule involves the same response to a given movement in inflation regardless of the degree of unanchoring. Therefore the central bank needs to be unconditionally aggressive on inflation fluctuations to sufficiently stabilize and anchor expectations.

This result crucially hinges on how interest-rate expectations are formed, however. In an alternative with anchoring inflation expectations in the Taylor rule, one obtains the opposite result of Eusepi et al. (2018a) that monetary policy should be less aggressive. This highlights that for the proper anchoring of inflation expectations, central bankers need to understand not only how inflation expectations are formed, but also the formation of interest-rate expectations.

References

Andrade, P., O. Coibion, E. Gautier, and Y. Gorodnichenko (2020, June). No firm is an island? how industry conditions shape firms' aggregate expectations. Working Paper 27317, National Bureau of Economic Research.

Andre, P., C. Pizzinelli, C. Roth, and J. Wohlfart (2022). Subjective models of the macroeconomy: Evidence from experts and a representative sample. *Review of Economic Studies*.

Andreasen, M. M., J. H. Christensen, and S. Riddell (2017). Extracting information from inflation markets: The role of tips liquidity.

Andreasen, M. M., J. H. Christensen, and S. Riddell (2018). The TIPS liquidity premium. Federal Reserve Bank of San Francisco.

Angeletos, G.-M., Z. Huo, and K. A. Sastry (2021). Imperfect macroeconomic expectations: Evidence and theory. *NBER Macroeconomics Annual 35*(1), 1–86.

Angeletos, G.-M. and A. Pavan (2009). Policy with dispersed information. *Journal of the European Economic Association 7*(1), 11–60.

Anufriev, M. and C. Hommes (2012). Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments. *American Economic Journal: Microeconomics 4*(4), 35–64.

Assenza, T., P. Heemeijer, C. Hommes, and D. Massaro (2019). Managing self-organization of expectations through monetary policy: a macro experiment. *Journal of Monetary Economics*.

Bai, J., R. L. Lumsdaine, and J. H. Stock (1998). Testing for and dating common breaks in multivariate time series. *The Review of Economic Studies 65*(3), 395–432.

Ball, L. (1994). Credible disinflation with staggered price-setting. *American Economic Review 84*(1), 282–289.

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

Barro, R. J. and D. B. Gordon (1983). Rules, discretion and reputation in a model of monetary policy. *Journal of monetary economics 12*(1), 101–121.

Baxter, M. and R. G. King (1999). Measuring business cycles: approximate band-pass filters for economic time series. *Review of Economics and Statistics 81*(4), 575–593.

Beaudry, P. and F. Portier (2006). Stock prices, news, and economic fluctuations. *American Economic Review 96*(4), 1293–1307.

Beck, G. W. and V. Wieland (2002). Learning and control in a changing economic environment. *Journal of Economic Dynamics and Control 26*(9), 1359 – 1377. Special issue in honour of David Kendrick.

Benhabib, J., S. Schmitt-Grohé, and M. Uribe (2001). The perils of Taylor rules. *Journal of Economic Theory 96*(1-2), 40–69.

Bernanke, B. (2007). Inflation expectations and inflation forecasting. Technical report, Board of Governors of the Federal Reserve System (US).

Bernanke, B. S. and M. Woodford (1997). Inflation forecasts and monetary policy. *Journal of Money, Credit and Banking 29*(4), 653–684.

Bianchi, F., S. C. Ludvigson, and S. Ma (2020, June). Belief distortions and macroeconomic fluctuations. Working Paper 27406, National Bureau of Economic Research.

Bianchi, F., L. Melosi, and M. Rottner (2019, September). Hitting the elusive inflation target. Working Paper 26279, National Bureau of Economic Research.

Bils, M. and P. J. Klenow (2004). Some evidence on the importance of sticky prices. *Journal of Political Economy 112*(5), 947–985.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2018). Over-reaction in macroeconomic expectations. Technical report, National Bureau of Economic Research.

Bordalo, P., N. Gennaioli, and A. Shleifer (2018). Diagnostic expectations and credit cycles. *The Journal of Finance 73*(1), 199–227.

Bottan, N. L. and R. Perez-Truglia (2020, June). Betting on the house: Subjective expectations and market choices. Working Paper 27412, National Bureau of Economic Research.

Branch, W. A. (2004). Restricted perceptions equilibria and learning in macroeconomics. *Post walrasian macroeconomics: Beyond the dynamic stochastic general equilibrium model*, 135–160.

Branch, W. A. and G. W. Evans (2006). A simple recursive forecasting model. *Economics Letters 91*(2), 158–166.

Branch, W. A. and G. W. Evans (2011a, July). Learning about risk and return: A simple model of bubbles and crashes. *American Economic Journal: Macroeconomics 3*(3), 159–91.

Branch, W. A. and G. W. Evans (2011b, Jun). Monetary policy and heterogeneous expectations. *Economic Theory 47*(2), 365–393.

Branch, W. A., G. W. Evans, and B. McGough (2014). Perpetual learning and stability in macroeconomic models. *Unpublished manuscript*.

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), 149–163.

Bullard, J. and K. Mitra (2002). Learning about monetary policy rules. *Journal of Monetary Economics 49*(6), 1105–1129.

Caballero, R. J. and A. Simsek (2020, June). Monetary policy with opinionated markets. Working Paper 27313, National Bureau of Economic Research.

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics 12*(3), 383–398.

Campbell, J. R., C. L. Evans, J. D. M. Fisher, and A. Justiniano (2012). Macroeconomic effects of Federal Reserve forward guidance [with comments and discussion]. *Brookings Papers on Economic Activity*, 1–80.

Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics letters 91*(3), 312–320.

Carvalho, C., S. Eusepi, E. Moench, and B. Preston (2021). Anchored inflation expectations. *Available at SSRN 3018198*.

Chari, V. V., P. J. Kehoe, and E. R. McGrattan (2000). Sticky price models of the business cycle: can the contract multiplier solve the persistence problem? *Econometrica 68*(5), 1151–1179.

Cho, I.-K. and K. Kasa (2015). Learning and model validation. *The Review of Economic Studies 82*(1), 45–82.

Cho, I.-K. and K. Kasa (2017). Gresham's law of model averaging. *American Economic Review 107*(11), 3589–3616.

Cho, I.-K. and A. Matsui (1995). Induction and the Ramsey policy. *Journal of Economic Dynamics and Control 19*(5-7), 1113–1140.

Christiano, L. J. and J. D. Fisher (2000). Algorithms for solving dynamic models with occasionally binding constraints. *Journal of Economic Dynamics and Control 24*(8), 1179 – 1232.

Clarida, R., J. Gali, and M. Gertler (1999). The science of monetary policy: a new Keynesian perspective. *Journal of Economic Literature 37*(4), 1661–1707.

Cogley, T. and J. M. Nason (1993). Impulse dynamics and propagation mechanisms in a real business cycle model. *Economics Letters 43*(1), 77–81.

Cogley, T. and T. J. Sargent (2008, 2). Anticipated utility and rational expectations as approximations of Bayesian decision making. *International Economic Review 49*(1), 185–221.

Coibion, O. and Y. Gorodnichenko (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review 105*(8), 2644–78.

Coibion, O., Y. Gorodnichenko, S. Kumar, and M. Pedemonte (2018, June). Inflation expectations as a policy tool? Working Paper 24788, National Bureau of Economic Research.

Coibion, O., Y. Gorodnichenko, and M. Weber (2020, June). Does policy communication during Covid work? Working Paper 27384, National Bureau of Economic Research.

Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer (2016, March). Parameter learning in general equilibrium: The asset pricing implications. *American Economic Review 106*(3), 664–98.

Cukierman, A. and A. H. Meltzer (1986). A theory of ambiguity, credibility, and inflation under discretion and asymmetric information. *Econometrica 54*(5), 1099–1128.

D'Acunto, F., D. Hoang, and M. Weber (2020, June). Managing households' expectations with unconventional policies. Working Paper 27399, National Bureau of Economic Research.

D'Acunto, F., U. Malmendier, J. Ospina, and M. Weber (2019, September). Exposure to daily price changes and inflation expectations. Working Paper 26237, National Bureau of Economic Research.

Davig, T. and E. M. Leeper (2007). Generalizing the Taylor principle. *American Economic Review 97*(3), 607–635.

Del Negro, M., M. Lenza, G. E. Primiceri, and A. Tambalotti (2020, April). What's up with the Phillips curve? Working Paper 27003, National Bureau of Economic Research.

Dombeck, B. (2017). The effects of news shocks and bounded rationality on macroeconomic volatility.

Dosi, G., M. Napoletano, A. Roventini, J. E. Stiglitz, and T. Treibich (2020, April). Rational heuristics? expectations and behaviors in evolving economies with heterogeneous interacting agents. Working Paper 26922, National Bureau of Economic Research.

Duffie, D. and K. J. Singleton (1990). Simulated moments estimation of Markov models of asset prices. Technical report, National Bureau of Economic Research.

Durbin, J. (1969). Tests for serial correlation in regression analysis based on the periodogram of least-squares residuals. *Biometrika 56*(1), 1–15.

Eichenbaum, M., N. Jaimovich, and S. Rebelo (2011). Reference prices, costs, and nominal rigidities. *American Economic Review 101*(1), 234–62.

Erceg, C. J. and A. T. Levin (2003). Imperfect credibility and inflation persistence. *Journal of monetary economics 50*(4), 915–944.

Eusepi, S., M. Giannoni, and B. Preston (2018a). On the limits of monetary policy.

Eusepi, S., M. P. Giannoni, and B. Preston (2018b). Some implications of learning for price stability. *European Economic Review 106*, 1–20.

Eusepi, S. and B. Preston (2010, July). Central bank communication and expectations stabilization. *American Economic Journal: Macroeconomics 2*(3), 235–71.

Eusepi, S. and B. Preston (2011). Expectations, learning, and business cycle fluctuations. *American Economic Review 101*(6), 2844–72.

Eusepi, S. and B. Preston (2018). 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 Macroeconomics*. Princeton University Press.

Evans, G. W. and S. Honkapohja (2003). Expectations and the stability problem for optimal monetary policies. *The Review of Economic Studies 70*(4), 807–824.

Evans, G. W. and S. Honkapohja (2006). Monetary policy, expectations and commitment. *Scandinavian Journal of Economics 108*(1), 15–38.

Evans, G. W. and S. Honkapohja (2009). Learning and macroeconomics. *Annu. Rev. Econ. 1*(1), 421–449.

Evans, G. W., S. Honkapohja, et al. (2009). Expectations, learning and monetary policy: an overview of recent research. *Monetary policy under uncertainty and learning 13*, 27–76.

Evans, G. W., S. Honkapohja, and K. Mitra (2009). Anticipated fiscal policy and adaptive learning. *Journal of Monetary Economics 56*(7), 930–953.

Evans, G. W., S. Honkapohja, T. J. Sargent, and N. Williams (2013). *Bayesian model averaging, learning and model selection*, Chapter 6, pp. 99–119. Oxford University Press Oxford.

Evans, G. W. and B. McGough (2009). Learning to optimize. Technical report, Citeseer.

Fernández-Villaverde, J., P. Guerrón-Quintana, and J. F. Rubio-Ramírez (2015). Estimating dynamic equilibrium models with stochastic volatility. *Journal of Econometrics 185*(1), 216–229.

Fernández-Villaverde, J., J. F. Rubio-Ramírez, and F. Schorfheide (2016). Solution and estimation methods for dsge models. Volume 2, Chapter Chapter 9, pp. 527–724. Elsevier.

Ferrero, G. (2007). Monetary policy, learning and the speed of convergence. *Journal of Economic Dynamics and Control 31*(9), 3006–3041.

Frankel, A. and N. Kartik (2016). What kind of central bank competence? Technical report, Working Paper, Columbia.

Fuster, A., R. Perez-Truglia, and B. Zafar (2018, June). Expectations with endogenous information acquisition: An experimental investigation. Working Paper 24767, National Bureau of Economic Research.

Gaspar, V., F. Smets, and D. Vestin (2006). Adaptive learning, persistence, and optimal monetary policy. *Journal of the European Economic Association 4*(2-3), 376–385.

Gaspar, V., F. Smets, and D. Vestin (2010). Inflation expectations, adaptive learning and optimal monetary policy. In *Handbook of Monetary Economics*, Volume 3, pp. 1055–1095. Elsevier.

Gaus, E. and S. Ramamurthy (2019). A new approach to modeling endogenous gain learning. *Topics in Identification, Limited Dependent Variables, Partial Observability, Experimentation, and Flexible Modeling: Part A (Advances in Econometrics 40*, 203–227.

Giannoni, M. P. (2010). Optimal interest-rate rules in a forward-looking model, and inflation stabilization versus price-level stabilization. Technical report, National Bureau of Economic Research.

Giannoni, M. P. and M. Woodford (2003). Optimal interest-rate rules: I. general theory. Technical report, National Bureau of Economic Research.

Gobbi, L., R. Mazzocchi, and R. Tamborini (2019). Monetary policy, de-anchoring of inflation expectations, and the "new normal". *Journal of Macroeconomics 61*, 103070.

Goodfriend, M. (1993). Interest rate policy and the inflation scare problem: 1979-1992. *FRB Richmond Economic Quarterly 79*(1), 1–23.

Goodfriend, M. and R. G. King (2005). The incredible Volcker disinflation. *Journal of Monetary Economics 52*(5), 981–1015.

Gürkaynak, R. S., B. Sack, and E. Swanson (2005). The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models. *American Economic Review 95*(1), 425–436.

Gust, C., E. Herbst, and D. López-Salido (2020). Short-term planning, monetary policy, and macroeconomic persistence.

Hamilton, J. D. (2018). Why you should never use the hodrick-prescott filter. *Review of Economics and Statistics 100*(5), 831–843.

Haubrich, J., G. Pennacchi, and P. Ritchken (2012). Inflation expectations, real rates, and risk premia: Evidence from inflation swaps. *The Review of Financial Studies 25*(5), 1588–1629.

Hazell, J., J. Herreño, E. Nakamura, and J. Steinsson (2020, October). The slope of the Phillips curve: Evidence from U.S. states. Working Paper 28005, National Bureau of Economic Research.

Hebden, J., E. P. Herbst, J. Tang, G. Topa, and F. Winkler (2020). How robust are makeup strategies to key alternative assumptions?

Hodrick, R. J. and E. C. Prescott (1997). Postwar US business cycles: an empirical investigation. *Journal of Money, credit, and Banking*, 1–16.

Holbrook, R. S. (1972). Optimal economic policy and the problem of instrument instability. *The American Economic Review 62*(1/2), 57–65.

Ireland, P. N. (1997). Stopping inflations, big and small. *Journal of Money, Credit, and Banking*, 759–775.

Ireland, P. N. (2000). Expectations, credibility, and time-consistent monetary policy. *Macroeconomic Dynamics 4*(4), 448–466.

Judd, K. L. (1998). *Numerical Methods in Economics*. Number 0262100711 in MIT Press Books. The MIT Press.

Kilian, L. and H. Lütkepohl (2017). *Structural vector autoregressive analysis*. Cambridge University Press.

Klenow, P. J. and O. Kryvtsov (2008). State-dependent or time-dependent pricing: Does it matter for recent US inflation? *The Quarterly Journal of Economics 123*(3), 863–904.

Klenow, P. J. and B. A. Malin (2010). Microeconomic evidence on price-setting. In *Handbook of Monetary Economics*, Volume 3, pp. 231–284. Elsevier.

Kostyshyna, O. (2012). Application of an adaptive step-size algorithm in models of hyperinflation. *Macroeconomic Dynamics 16*(S3), 355–375.

Kozlowski, J., L. Veldkamp, and V. Venkateswaran (2015, November). The tail that wags the economy: Beliefs and persistent stagnation. Working Paper 21719, National Bureau of Economic Research.

Kreps, D. M. (1998). Anticipated utility and dynamic choice. *Econometric Society Monographs 29*, 242–274.

Kydland, F. E. and E. C. Prescott (1977). Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy 85*(3), 473–491.

Lane, T. D. (1984). Instrument instability and short-term monetary control. *Journal of Monetary Economics 14*(2), 209 – 224.

Lee, B.-S. and B. F. Ingram (1991). Simulation estimation of time-series models. *Journal of Econometrics 47*(2-3), 197–205.

Lubik, T. A. and C. Matthes (2016). Indeterminacy and learning: An analysis of monetary policy in the Great Inflation. *Journal of Monetary Economics 82*, 85 – 106.

Lucas, R. E. (1976). Econometric policy evaluation: A critique. In *Carnegie-Rochester conference series on public policy*, Volume 1, pp. 19–46.

Lucas, R. E. (1987). *Models of business cycles*, Volume 26. Basil Blackwell Oxford.

Lütkepohl, H. (2013). *Introduction to multiple time series analysis*. Springer Science & Business Media.

Mackowiak, B. and M. Wiederholt (2009). Optimal sticky prices under rational inattention. *American Economic Review 99*(3), 769–803.

Malmendier, U., S. Nagel, and Z. Yan (2016). The making of hawks and doves: Inflation experiences and voting on the fomc. Technical report, Working paper.

Mankiw, N. G. and R. Reis (2002). Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve. *The Quarterly Journal of Economics 117*(4), 1295–1328.

Marcet, A. and D. A. Marshall (1994). Solving nonlinear rational expectations models by parameterized expectations: convergence to stationary solutions.

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 mechanisms in self-referential linear stochastic models. *Journal of Economic Theory 48*(2), 337–368.

McKay, A., E. Nakamura, and J. Steinsson (2017). The discounted Euler equation: A note. *Economica 84*(336), 820–831.

Mele, A., K. Molnár, and S. Santoro (2019). On the perils of stabilizing prices when agents are learning. *Journal of Monetary Economics*.

Milani, F. (2007). Expectations, learning and macroeconomic persistence. *Journal of Monetary Economics 54*(7), 2065–2082.

Milani, F. (2014). Learning and time-varying macroeconomic volatility. *Journal of Economic Dynamics and Control 47*, 94–114.

Molavi, P. (2019). Macroeconomics with learning and misspecification: A general theory and applications. 2019 Meeting Papers 1584, Society for Economic Dynamics.

Molnár, K. and S. Santoro (2014). Optimal monetary policy when agents are learning. *European Economic Review 66*, 39–62.

Moscarini, G. (2007). Competence implies credibility. *The American Economic Review 97*(1), 37–63.

Nakamura, E. and J. Steinsson (2008). Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics 123*(4), 1415–1464.

Orlik, A. and L. Veldkamp (2014). Understanding uncertainty shocks and the role of black swans. Technical report, National Bureau of Economic Research.

Orphanides, A. and J. Williams (2004). Imperfect knowledge, inflation expectations, and monetary policy. In *The inflation-targeting debate*, pp. 201–246. University of Chicago Press.

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

Orphanides, A. and J. C. Williams (2007). Robust monetary policy with imperfect knowledge. *Journal of monetary Economics 54*(5), 1406–1435.

Preston, B. (2005). Learning about monetary policy rules when long-horizon expectations matter. MPRA paper, University Library of Munich, Germany.

Preston, B. (2008). Adaptive learning and the use of forecasts in monetary policy. *Journal of Economic Dynamics and Control 32*(11), 3661 – 3681.

Primiceri, G. E. (2006, 08). Why Inflation Rose and Fell: Policy-Makers' Beliefs and U. S. Postwar Stabilization Policy*. *The Quarterly Journal of Economics 121*(3), 867–901.

Reis, R. (2020). Keynote speech - the anchoring of long-run inflation expectations today. In B. f. I. Settlements (Ed.), *Inflation dynamics in Asia and the Pacific*, Volume 111, pp. 11–20. Bank for International Settlements.

Rotemberg, J. J. and M. Woodford (1997). An optimization-based econometric framework for the evaluation of monetary policy. *NBER Macroeconomics Annual 12*, 297–346.

Sargent, T. J. (1999). *The Conquest of American Inflation*. Princeton University Press.

Schmitt-Grohé, S. and M. Uribe (2004). Solving Dynamic General Equilibrium Models Using a Second-Order Approximation to the Policy Function. *Journal of Economic Dynamics and Control 28*(4), 755 – 775.

Schumaker, L. L. (1983). On shape preserving quadratic spline interpolation. *SIAM Journal on Numerical Analysis 20*(4), 854–864.

Sims, C. A. (1974). Optimal stable policies for unstable instruments. In *Annals of Economic and Social Measurement, Volume 3, number 1*, pp. 257–266. NBER.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics 50*(3), 665 – 690. Swiss National Bank/Study Center Gerzensee Conference on Monetary Policy under Incomplete Information.

Slobodyan, S. and R. Wouters (2012). Learning in an estimated medium-scale DSGE model. *Journal of Economic Dynamics and Control 36*(1), 26–46.

Smith, A. A. (1993). Estimating nonlinear time-series models using simulated vector autoregressions. *Journal of Applied Econometrics 8*, S63–S84.

Svensson, L. E. (1999). Inflation targeting as a monetary policy rule. *Journal of Monetary Economics 43*(3), 607–654.

Svensson, L. E. (2003). What is wrong with Taylor rules? using judgment in monetary policy through targeting rules. *Journal of Economic Literature 41*(2), 426–477.

Taylor, J. B. (1993). Discretion versus policy rules in practice. In *Carnegie-Rochester conference series on public policy*, Volume 39, pp. 195–214. Elsevier.

Townsend, R. M. (1983). Forecasting the forecasts of others. *Journal of Political Economy 91*(4), 546–588.

Tversky, A. and D. Kahneman (1974). Judgment under uncertainty: Heuristics and biases. *Science 185*(4157), 1124–1131.

Weiser, A. and S. E. Zarantonello (1988). A note on piecewise linear and multilinear table interpolation in many dimensions. *Mathematics of Computation 50*(181), 189–196.

Whittle, P. (1983). Prediction and regulation by linear least-square methods.

Williams, J. C. (2006). The Phillips curve in an era of well-anchored inflation expectations. In *Federal Reserve Bank of San Francisco*. Citeseer.

Williams, N. (2003). Adaptive learning and business cycles. *Manuscript, Princeton University*.

Winkler, F. (2019). The role of learning for asset prices and business cycles. *Journal of Monetary Economics*.

Woodford, M. (1999). Optimal monetary policy inertia. *The Manchester School 67*, 1–35.

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

Woodford, M. (2003b). Optimal interest-rate smoothing. *The Review of Economic Studies 70*(4), 861–886.

Woodford, M. (2013). Macroeconomic analysis without the rational expectations hypothesis. *Annu. Rev. Econ. 5*(1), 303–346.

Woodford, M. (2019). Monetary policy analysis when planning horizons are finite. *NBER Macroeconomics Annual 33*(1), 1–50.

A Robustness of rolling regression results

I investigate the following alternative specifications. As a first check, I replace the forecast error with realized CPI inflation. Thus Equation (1) becomes

$$\Delta\bar{\pi}_t^i = \beta_0 + \beta_1^w \pi_t + \epsilon_t^i. \tag{A.1}$$

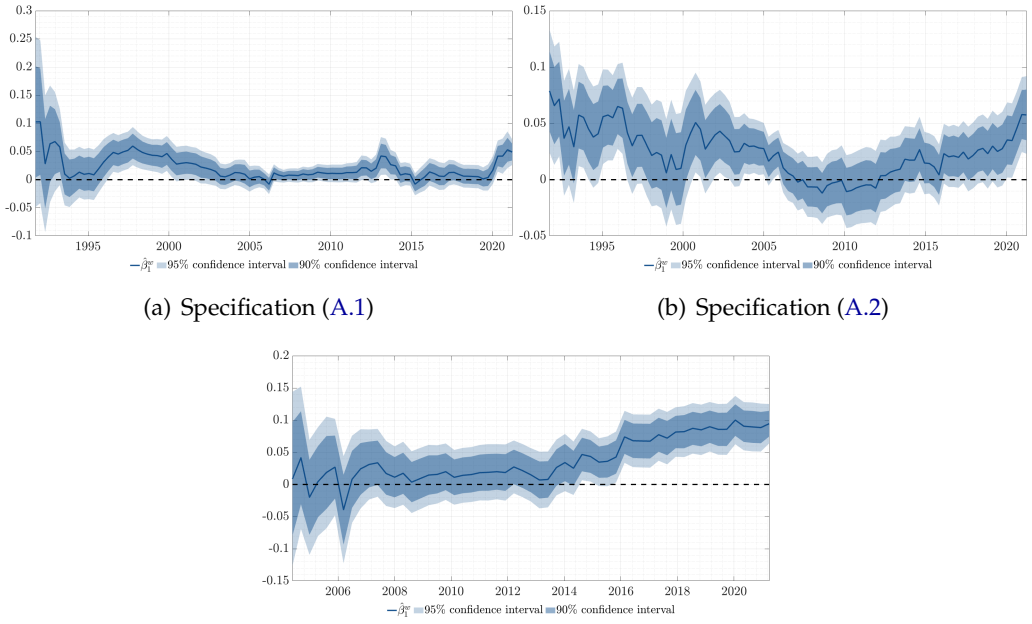
I then consider the original specification of Equation (1), adding realized CPI inflation to control for inflation levels.[23] This results in the following regression:

$$\Delta\bar{\pi}_t^i = \beta_0^w + \beta_1^w f_{t|t-1}^i + \beta_2^w \pi_t + \epsilon_t^i. \tag{A.2}$$

In the last alternative specification, I replace both the ten and the one-year ahead CPI inflation expectation with PCE core. The regression equation thus looks identical to Equation (1) but cleans out the effects of movements in oil prices:

$$\Delta\bar{\pi}_t^i = \beta_0^w + \beta_1^w f_{t|t-1}^i + \epsilon_t^i. \tag{A.3}$$

Figure 9: Time series of responsiveness of long-run inflation expectations to inflation



(a) Specification (A.1)

(b) Specification (A.2)

(c) Specification (A.3)

The three alternative time series of the estimate $\hat{\beta}_1^w$ are depicted in Fig. 9. Qualitatively, the same conclusion obtains as in the specification in the main text: the coefficients are time-varying

[23] In a separate specification, I have similarly added uncertainty as a control, but it is so similar that I suppress it here.

in size and in statistical significance. In particular, large sensitivity in the 1990s gives way to an insignificant coefficient in the early 2000s, to drift back up shortly after 2010. The data is much shorter for the PCE core specification, but also in this case one observes the upward drift in the estimated coefficient starting around 2010.

B Alternative measures of long-run inflation expectations

Measuring long-run expectations is challenging. In particular, two dimensions of expectations are hard to line up between models and data: the identity of the economic agent forming expectations, and the forecast horizon. To start with the former, one would ideally wish to measure the expectations of the private sector, firms and households, since their expectations correspond to the ones in economic models. As for the forecast horizon, no single horizon is an exact counterpart to the concept of expectations of average inflation. Arguably, the longer the forecast horizon, the closer the measured expectation is to expectations of the average.

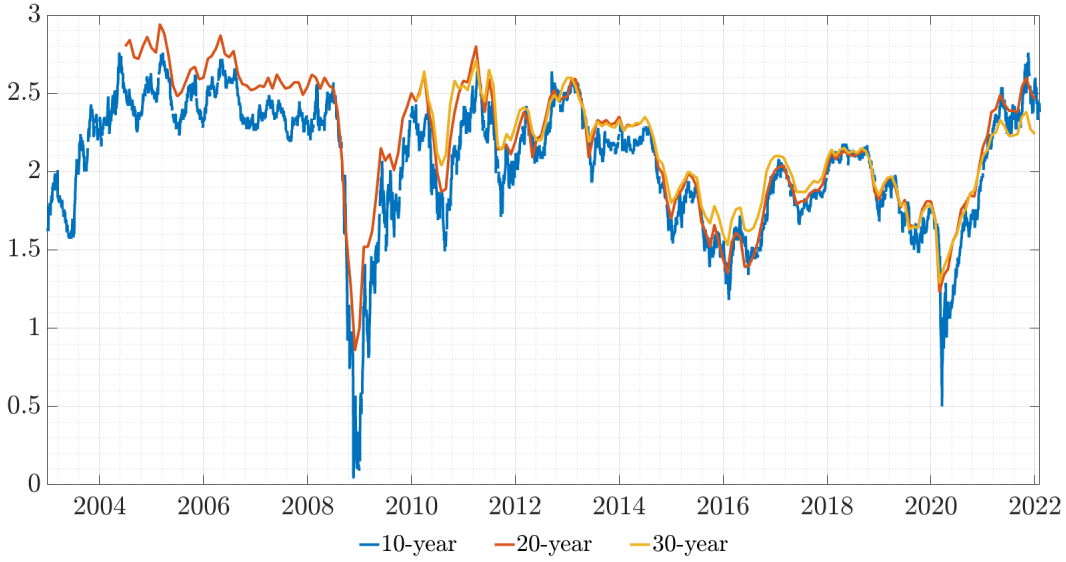
These considerations can lead the researcher to utilize breakeven inflation, that is, the difference between Treasury yields and Treasury Inflation Protected Securities (TIPS) (top panel of Fig. 10). There are two main advantages of constructing inflation expectation measures from TIPS. First, since this measure is based on trades that happened in the Treasury market, it pertains to agents who are active in a market where inflation expectations matter. Such expectations are therefore a good proxy for the private sector's expectations in economic models. Second, the fact that breakeven inflation is not elicited from surveys allows the econometrician to bypass many of the challenges that survey data involve. It has been widely documented that survey participants may have poor understanding of the economic concept elicited, may misunderstand the survey questions or be unduly influenced by the wording. Most troubling is the fact that there is no way the econometrician could control for the noise thus introduced in survey data.

Breakeven inflation measures are not a panacea, however. The illiquidity of TIPS markets introduces a distortion into the measured expectations in the form of a time-varying liquidity premium (Appendix C corrects for this by filtering out an approximate liquidity premium series from the 10-year expectation series.) One could also contend that many firms and the majority of households are not active on the TIPS market, and therefore that breakeven inflation is not representative of the economy-wide expectation. For this reason, I turn to alternative long-run inflation measures.

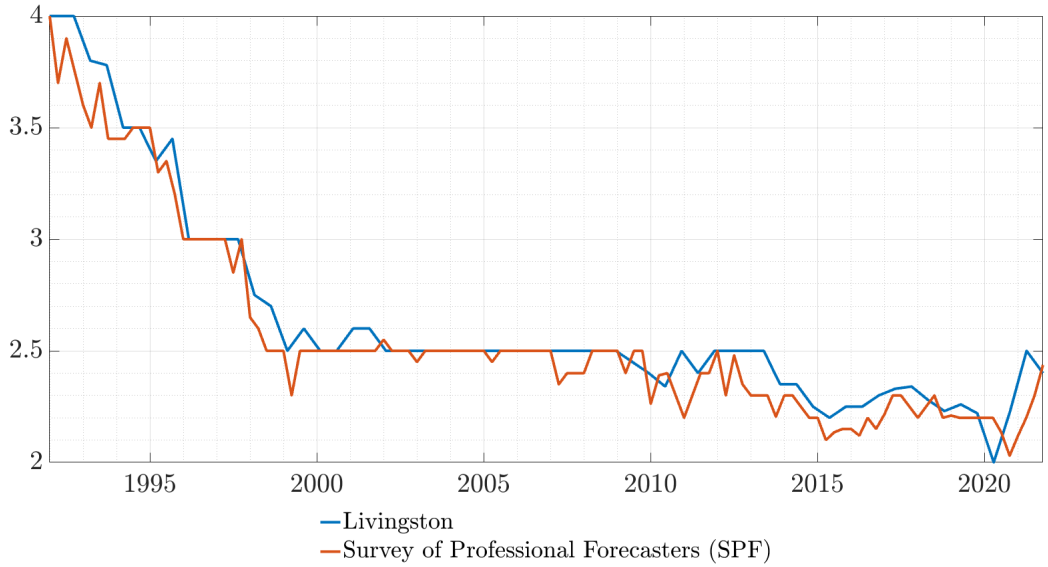
The forecast horizon is a binding constraint: consumer expectations data have a maximal forecast horizon of three (New York Fed Survey of Consumer Expectations (SCE)) or five years (University of Michigan Survey of Consumers). To have at least a forecast horizon of ten years, I thus resort to the Livingston Survey of the Philadelphia Fed and the Survey of Professional Forecasters (SPF). Another advantage of the SPF is that comprehensive individual-level data is available, a feature which I exploit for the regressions in the Introduction and in Appendix A.

Fig. 10 plots the breakeven inflation series for a horizon of 10, 20 and 30 years (top panel) alongside Livingston and SPF (bottom panel). The breakeven inflation measures are clearly and consistently below the Livingston and SPF measures. Part of this may be driven by liquidity premia,

Figure 10: Alternative long-run inflation expectations measures



(a) Market-based inflation expectations from TIPS, various horizons



(b) 10-year ahead inflation expectations of firms (Livingston) and professional forecasters (SPF)

All inflation rates are annualized percentages. The series are at the following frequency: daily (10-year TIPS), monthly (20- and 30-year TIPS), quarterly (SPF), twice per year (Livingston).

which I control for in Appendix C. At the same time, both the Livingston and the SPF expectations may potentially suffer from representativeness issues. The fact that the two expectations series are so closely aligned indicates that the Livingston survey may only be reaching a subset of firms.[24]

[24]The category codes of individual survey participates in the Livingston survey seems to underscore this. According to the documentation, the surveyed businesses fall into the following categories: academic institution, commercial banking, consulting, Federal Reserve, government, industry trade group, insurance company, investment banking,

This subset seems to be highly correlated with the set of professional forecasters filling out the SPF. This is a concern because it suggests that the elicited expectations may not be representative of the economy-wide firm expectations. Moreover, if professional forecasters respond to the survey by running a (by assumption stationary) econometric model, then the survey responses may not even accurately reflect their held beliefs.

But all these caveats notwithstanding, there is a single element that is consistent across all of these measures of long-run inflation expectations: following 2010, they all shift downward. In other words, they all show responsiveness to the fact that the Federal Reserve has undershot the 2% inflation target. Given that the Livingston and SPF measures are initially above the 2% target, one may argue that expectations are only adjusting following the 2012 announcement of an official target on the part of the Fed.

However, the 2% target was already understood prior to the 2012 announcement. Moreover, the Fed's target is specified in terms of PCE inflation, while the Livingston and SPF forecasts pertain to CPI inflation. The fact that the latter tends to be higher than the former also helps to explain why expectations of CPI inflation are above the Fed's PCE inflation target.

C Filtering out liquidity risk from TIPS

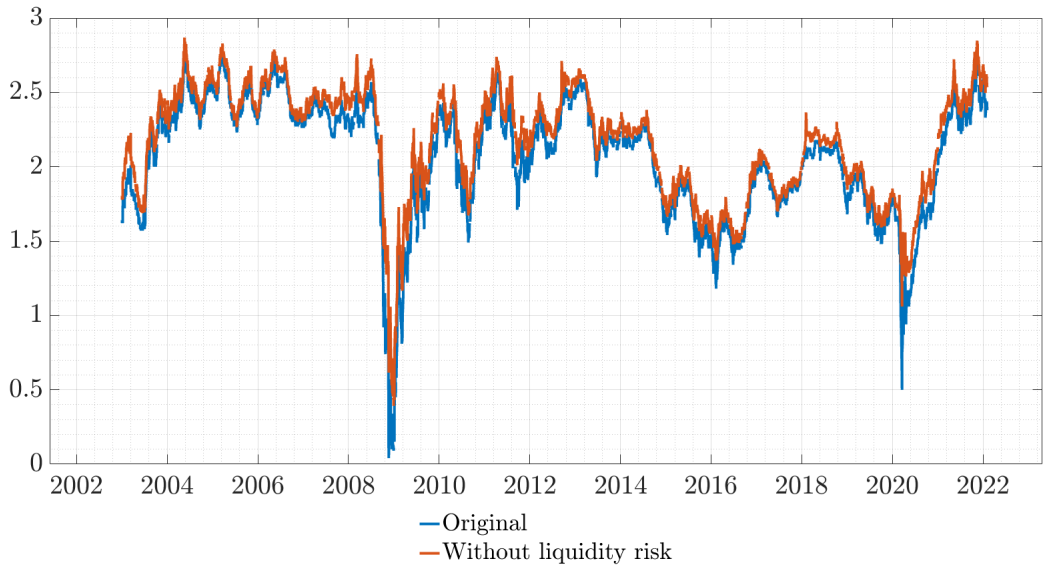
Some caution is needed when inferring inflation expectations from TIPS. The underlying idea is that the difference between real (indexed to inflation) and nominal yields should be a good metric for the market's expectations of inflation on average for the duration of the particular maturity. But since TIPS markets face liquidity issues, especially for seasoned securities, the TIPS yield also incorporates a liquidity premium. The positive bias in the TIPS yield thus leads to a negative bias in expected inflation.

To gauge the presence of liquidity risk in TIPS, I rely on Andreasen et al. (2018)'s estimation of the liquidity premium. Since their series only covers the period between July 11, 1997 - Dec 27, 2013, I make use of the fact that they demonstrate a high correlation between liquidity risk and uncertainty. In particular, a regression of their average TIPS liquidity premium measure on the VIX index and controls yields an estimated coefficient of 0.85 (significant at the 1 percent level), with a constant of -5.21. Thus I can use the VIX to back out a fitted Andreasen et al. (2018) estimate of the TIPS liquidity premium after 2013. Doing so, I subtract this fitted liquidity premium series from the TIPS yields, allowing me to construct an estimate of breakeven inflation corrected for the liquidity risk bias. Fig. 11 presents the original breakeven inflation series, along with the bias-corrected estimate.

In line with Andreasen et al. (2018)'s findings, I obtain a negative bias in the inflation expectations series throughout the sample. This bias averages -0.0966 pp, which is sizable, but not as significant as one might have suspected. Also in analogy with Andreasen et al. (2018)'s results, I find that liquidity issues in the TIPS market pose a bigger problem in recessions, when the market worries more about future TIPS becoming illiquid. In particular once the Covid-19-shock raises

labor, non-financial business and other.

Figure 11: Market-based inflation expectations, 10 year, average, %



Breakeven inflation, constructed as the difference between the yields of 10-year Treasuries and 10-year TIPS (blue line), difference between 10-year Treasury and 10-year TIPS, the latter cleaned from liquidity risk (red line).

worries at the end of the sample, the estimated liquidity premium hits its highest value of 0.6508 pp. Even this large upward correction in breakeven inflation does not change the overall picture, however. The conclusion that long-run inflation expectations trend downward in the second half of the decade remains solid.

D Compact model notation

The A matrices are given by

$$A_a = \begin{pmatrix} g_{\pi a} \\ g_{xa} \\ \psi_\pi g_{\pi a} + \psi_x g_{xa} \end{pmatrix}, \quad A_b = \begin{pmatrix} g_{\pi b} \\ g_{xb} \\ \psi_\pi g_{\pi b} + \psi_x g_{xb} \end{pmatrix}, \quad A_s = \begin{pmatrix} g_{\pi s} \\ g_{xs} \\ \psi_\pi g_{\pi s} + \psi_x g_{xs} + \begin{bmatrix} 0 & 1 & 0 \end{bmatrix} \end{pmatrix}, \tag{D.1}$$

$$g_{\pi a} = (1 - \frac{\kappa\sigma\psi_\pi}{w}) \left[(1-\alpha)\beta, \kappa\alpha\beta, 0\right], \tag{D.2}$$

$$g_{xa} = \frac{-\sigma\psi_\pi}{w} \left[(1-\alpha)\beta, \kappa\alpha\beta, 0\right], \tag{D.3}$$

$$g_{\pi b} = \frac{\kappa}{w} \left[\sigma(1-\beta\psi_\pi), (1-\beta-\beta\sigma\psi_x, 0\right], \tag{D.4}$$

$$g_{xb} = \frac{1}{w} \left[\sigma(1-\beta\psi_\pi), (1-\beta-\beta\sigma\psi_x, 0\right], \tag{D.5}$$

$$g_{\pi s} = (1 - \frac{\kappa\sigma\psi_\pi}{w}) \begin{bmatrix} 0 & 0 & 1 \end{bmatrix} (I_3 - \alpha\beta h)^{-1} - \frac{\kappa\sigma}{w} \begin{bmatrix} -1 & 1 & 0 \end{bmatrix} (I_3 - \beta h)^{-1}, \tag{D.6}$$

$$g_{xs} = \frac{-\sigma\psi_\pi}{w} \begin{bmatrix} 0 & 0 & 1 \end{bmatrix} (I_3 - \alpha\beta h)^{-1} - \frac{\sigma}{w} \begin{bmatrix} -1 & 1 & 0 \end{bmatrix} (I_3 - \beta h)^{-1}, \tag{D.7}$$

$$w = 1 + \sigma\psi_x + \kappa\sigma\psi_\pi. \tag{D.8}$$

The matrices of the state transition equation (15) are

$$h \equiv \begin{pmatrix} \rho_r & 0 & 0 \\ 0 & \rho_i & 0 \\ 0 & 0 & \rho_u \end{pmatrix}, \quad \epsilon_t \equiv \begin{pmatrix} \varepsilon_t^r \\ \varepsilon_t^i \\ \varepsilon_t^u \end{pmatrix}, \quad \text{and} \quad \eta \equiv \begin{pmatrix} \sigma_r & 0 & 0 \\ 0 & \sigma_i & 0 \\ 0 & 0 & \sigma_u \end{pmatrix}. \tag{D.9}$$

Note that this is the formulation for the case where a Taylor rule is in effect and is known by the private sector. It is straightforward to remove any of these two assumptions.

E The observation matrix for learning

Instead of the matrix g in the rational expectations observation equation

$$z_t = g s_t, \tag{E.1}$$

agents in the anchoring model use the estimated matrix g^l

$$g_{t-1}^l = \begin{bmatrix} F_{t-1} & G_{t-1} \end{bmatrix}, \tag{E.2}$$

with

$$F_{t-1} = \left(A_a \frac{1}{1-\alpha\beta} + A_b \frac{1}{1-\beta}\right) a_{t-1}, \tag{E.3}$$

$$G_{t-1} = A_a b_{t-1} \Big(I_3 - \alpha\beta h\Big)^{-1} + A_b b_{t-1} \Big(I_3 - \beta h\Big)^{-1} + A_s. \tag{E.4}$$

F The discrete gain function as a model of anchoring

Carvalho et al. (2021) define their discrete choice anchoring function as follows. Let θ_t be a criterion to be defined. Then, for a threshold value $\tilde{\theta}$, the gain evolves according to

$$k_t = \begin{cases} (k_{t-1}+1)^{-1} & \text{if} \quad \theta_t < \tilde{\theta}, \\ \bar{g} & \text{otherwise.} \end{cases} \tag{F.1}$$

In other words, agents choose a decreasing gain when the criterion θ_t is lower than the threshold $\tilde{\theta}$; otherwise they choose a constant gain. The criterion employed by Carvalho et al. (2021) is computed as the absolute difference between subjective and model-consistent expectations, scaled by the variance of shocks:

$$\theta_t = \max|\Sigma^{-1}(\phi_{t-1} - \begin{bmatrix} F_{t-1} & G_{t-1} \end{bmatrix})|, \tag{F.2}$$

where Σ is the VC matrix of shocks, ϕ_{t-1} is the estimated matrix and $[F, G]$ is the ALM (see Appendix E).

As a robustness check, Carvalho et al. (2021) also compute an alternative criterion.[25] Let ω_t denote agents' time t estimate of the forecast error variance and θ_t be a statistic evaluated by agents in every period as

$$\omega_t = \omega_{t-1} + \tilde{\kappa} k_{t-1}(f_{t|t-1} f'_{t|t-1} - \omega_{t-1}), \tag{F.3}$$

$$\theta_t = \theta_{t-1} + \tilde{\kappa} k_{t-1}(f'_{t|t-1} \omega_t^{-1} f_{t|t-1} - \theta_{t-1}), \tag{F.4}$$

where $\tilde{\kappa}$ is a parameter that allows agents to scale the gain compared to the previous estimation and $f_{t|t-1}$ is the most recent forecast error, realized at time t. Indeed, this is a multivariate time series version of the squared CUSUM test.[26]

G Estimation procedure

The estimation of Section 4 is a simulated method of moments (SMM) exercise. As elaborated in the main text, I target the autocovariances of CPI inflation, the output gap, the federal funds rate and the 12-months ahead inflation forecasts coming from the Survey of Professional Forecasters. For the autocovariances, I consider lags $0, \ldots, 4$. The target moment vector, Ω, is the vectorized autocovariance matrices for the lags considered, 80×1.

Let Θ denote the set of parameters in the New Keynesian model that I calibrate. Then, for each proposed coefficient vector γ, the estimation procedure consists of simulating the model conditional on γ, Θ and N different sequences of disturbances, then computing model-implied moments for each simulation, and lastly choosing γ such that the squared distance between the data- and model-implied mean moments is minimized. Thus

[25] Note that for both criteria, I present the matrix generalizations of the scalar versions considered by Carvalho et al. (2021).

[26] See Brown et al. (1975) and Lütkepohl (2013) for details.

$$\hat{\gamma} = \left(\Omega^{data} - \frac{1}{N}\sum_{n=1}^{N}\Omega^{model}(\gamma, \Theta, \{e_t^n\}_{t=1}^T)\right)' W^{-1} \left(\Omega^{data} - \frac{1}{N}\sum_{n=1}^{N}\Omega^{model}(\gamma, \Theta, \{e_t^n\}_{t=1}^T)\right), \quad \text{(G.1)}$$

where the observed data consist of $T = 151$ quarters. Here $\{e_t^n\}_{t=1}^T$ is a sequence of disturbances of the same length as the data; note that I use a cross-section of N such sequences and take average moments across the cross-section to wash out the effects of particular disturbances. Experimentation with the number N led me to choose $N = 1000$, as estimates no longer change upon selecting larger N.

Before computing moments, I filter both the observed and model-generated data using the Baxter and King (1999) filter, with thresholds at 6 and 32 quarters and truncation at 12 lags, the recommended values of the authors. I then compute the moments by fitting a reduced-form VAR to the filtered series and using the estimated coefficients to back out autocovariances. Because there are four observables to three structural shocks and occasionally low volatility in the expectation series, I estimate the VAR coefficients by ridge regression with a tuning parameter of 0.001. This is to ensure that the VAR coefficients are estimated with a lower standard error, so that estimated variances of the moments are more accurate. As the weighting matrix of the quadratic form in the moments, I use the inverse of the estimated variances of the target moments, W^{-1}, computed from 10000 bootstrapped samples.

To improve identification, I also impose restrictions on the estimates. First, I require that the γ-coefficients be convex, that is, that larger forecast errors in absolute value be associated with higher gains. Second, since forecast errors close to zero render the size of the gain irrelevant (cf. the learning equation (21)), I impose that the coefficient associated with a zero forecast error should be zero. Both restrictions are implemented with weights penalizing the loss function, and the weights are selected by experimentation.

Both additional assumptions reflect properties that the anchoring function should reasonably exhibit. The convexity assumption captures the very notion that larger forecast errors in absolute value suggest bigger changes to the forecasting procedure are necessary. This is thus a very natural requirement. As for the zero gain for zero forecast error assumption, the idea here is to supply the estimation with information where it is lacking. Since the updating of learning coefficients corresponds to gain times forecast error, as Equation (21) recalls, a zero forecast error supplies no information for the value of the gain. To impose a zero value here also seems natural, given that since forecast errors switch sign at zero, one would expect the zero forecast error point to be an inflection point in the anchoring function. By the same token, Gobbi et al. (2019) also impose a related restriction when they require that their deanchoring function should yield a zero value at the zero input. Furthermore, the objective function does not deteriorate upon imposing either assumption, suggesting that they are not at odds with the data. One does observe however that these additional moments do have bite in estimates the nodes at zero and at the edges, reflected in slightly tighter confidence intervals than at the other two nodes on Figure 2. Lastly, I supply 100 different initial points and select the estimate involving the lowest value for the loss (G.1).

H The policy problem in the simplified baseline model

Denote by $\mathbf{g}'_t = \mathbf{g}_{\pi,t} = -\mathbf{g}_{\bar{\pi},t}$ the derivative of the anchoring function $\mathbf{g}(\cdot)$ with respect to the forecast error, evaluated at time t. In this simplified setting, $\bar{\pi}_t = e_1 a_t$, the estimated constant for the inflation process. As in the main text, e_i is a selector vector, selecting row i of the subsequent matrix. I also use the notation $b_i \equiv e_i b$. The planner chooses $\{\pi_t, x_t, f_{a,t}, f_{b,t}, \bar{\pi}_t, k_t\}_{t=t_0}^{\infty}$ to minimize

$$
\begin{aligned}
\mathcal{L} = \mathbb{E}_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} \Big\{ & (\pi_t^2 + \lambda_x x_t^2) \\
& + \varphi_{1,t} \Big(\pi_t - \kappa x_t - (1-\alpha)\beta f_{a,t} - \kappa\alpha\beta b_2 (I_3 - \alpha\beta h)^{-1} s_t - e_3 (I_3 - \alpha\beta h)^{-1} s_t \Big) \\
& + \varphi_{2,t} \Big(x_t + \sigma i_t - \sigma f_{b,t} - (1-\beta) b_2 (I_3 - \beta h)^{-1} s_t + \sigma\beta b_3 (I_3 - \beta h)^{-1} s_t - \sigma e_1 (I_3 - \beta h)^{-1} s_t \Big) \\
& + \varphi_{3,t} \Big(f_{a,t} - \frac{1}{1-\alpha\beta} \bar{\pi}_{t-1} - b_1 (I_3 - \alpha\beta h)^{-1} s_t \Big) \\
& + \varphi_{4,t} \Big(f_{b,t} - \frac{1}{1-\beta} \bar{\pi}_{t-1} - b_1 (I_3 - \beta h)^{-1} s_t \Big) \\
& + \varphi_{5,t} \Big(\bar{\pi}_t - \bar{\pi}_{t-1} - k_t (\pi_t - (\bar{\pi}_{t-1} + b_1 s_{t-1})) \Big) \\
& + \varphi_{6,t} \Big(k_t - \mathbf{g}(\pi_t - \bar{\pi}_{t-1} - b_1 s_{t-1}) \Big) \Big\},
\end{aligned} \tag{H.1}
$$

where ρ_i are Lagrange-multipliers on the constraints. After a little bit of simplifying, the first-order conditions boil down to the following three equations:

$$
2\pi_t + 2\frac{\lambda_x}{\kappa} x_t - \varphi_{5,t} k_t - \varphi_{6,t} \mathbf{g}'_t = 0, \tag{H.2}
$$

$$
- \frac{2(1-\alpha)\beta^2}{1-\alpha\beta} \frac{\lambda_x}{\kappa} \mathbb{E}_t\, x_{t+1} + \varphi_{5,t} - \beta\, \mathbb{E}_t (1 - k_{t+1}) \varphi_{5,t+1} + \beta \mathbf{g}'_t\, \mathbb{E}_t\, \varphi_{6,t+1} = 0, \tag{H.3}
$$

$$
\varphi_{6,t} = (\pi_t - \bar{\pi}_{t-1} - b_1 s_{t-1}) \varphi_{5,t}. \tag{H.4}
$$

Note that Equation (H.2) is the analogue of Gaspar et al. (2010)'s Equation (22) (or, equivalently, of Molnár and Santoro (2014)'s (16)), except that there is an additional multiplier, φ_6. This multiplier reflects the fact that in addition to the constraint coming from the expectation process itself, with shadow value φ_5, learning involves the gain equation as a constraint as well. One can also clearly read off that when the learning process has converged, so that neither expectations nor the gain process are constraints ($\varphi_5 = \varphi_6 = 0$), the discretionary inflation-output gap tradeoff familiar from Clarida et al. (1999) obtains. Combining the above three equations and solving for $\varphi_{5,t}$, using the notation that $\prod_0^0 \equiv 1$, one obtains the target criterion (34).

The system of first-order conditions and model equations for this simplified system also reveals how the endogenous gain introduces nonlinearity to the equation system. In particular, notice how the gain k_t shows up multiplicatively with the Lagrange multiplier, φ_5. In fact, the origin of the

problem is the recursive least squares learning equation for the learning coefficient $\bar{\pi}_t$

$$\bar{\pi}_t = \bar{\pi}_{t-1} + k_t f_{t|t-1} \tag{H.5}$$

where the first interaction term between the gain and other endogenous variables show up. This results in an equation system of nonlinear difference equations that does not admit an analytical solution.

Considering equation (H.5) is instructive to see how it is indeed the endogeneity of the gain that causes these troubles. Were we to specify a constant gain setup, k_t would merely equal the constant $\bar{g}$ and the anchoring function $\mathbf{g}(\cdot)$ would trivially reduce to zero as well. In such a case, all interaction terms would reduce to multiplication between endogenous variables and parameters; linearity would be restored and a solution for the optimal time paths of endogenous variables would be obtainable. Similarly, a decreasing gain specification would also be manageable since for all t, the gain would simply be given by t^{-1}, and the anchoring function would also be deterministic and exogenous.

I Relaxing Assumption 2

Consider the general anchoring mechanism of Equation (23):

$$k_t = \mathbf{g}(k_{t-1}, f_{t|t-1}). \tag{I.1}$$

Using the notation $\mathbf{g}_{k,t}$ and $\mathbf{g}_{f,t}$ for the derivates of the anchoring function with respect to k and f evaluated at time t, the FOCs of the Ramsey problem are

$$2\pi_t + 2\frac{\lambda_x}{\kappa} x_t - \varphi_{5,t} k_t - \varphi_{6,t} \mathbf{g}_{f,t} = 0, \tag{I.2}$$

$$-\frac{2(1-\alpha)\beta^2}{1-\alpha\beta}\frac{\lambda_x}{\kappa}\,\mathbb{E}_t\, x_{t+1} + \varphi_{5,t} - \beta\,\mathbb{E}_t(1-k_{t+1})\varphi_{5,t+1} + \beta \mathbf{g}_{f,t}\,\mathbb{E}_t\,\varphi_{6,t+1} = 0, \tag{I.3}$$

$$\varphi_{6,t} - \mathbb{E}_t\, \mathbf{g}_{k,t+1}\varphi_{6,t+1} = (\pi_t - \bar{\pi}_{t-1} - b_1 s_{t-1})\varphi_{5,t}, \tag{I.4}$$

where the red term is the new element vis-à-vis the case where the anchoring function is specified in levels ($k_t = \mathbf{g}(f_{t|t-1})$). Thus a central bank that follows the target criterion in this general case has to compute the solution to this three-dimensional difference equation system, where the more involved algebra contains the same economic intuition.

J Parameterized expectations algorithm (PEA)

The objective of the parameterized expectations algorithm is to solve for the sequence of interest rates that solves the model equations including the target criterion, representing the first-order condition of the Ramsey problem. For convenience, I list the model equations:

$$x_t = -\sigma i_t + \begin{bmatrix}\sigma & 1-\beta & -\sigma\beta\end{bmatrix} f_{b,t} + \sigma \begin{bmatrix}1 & 0 & 0\end{bmatrix} (I_3 - \beta h)^{-1} s_t, \tag{J.1}$$

$$\pi_t = \kappa x_t + \begin{bmatrix}(1-\alpha)\beta & \kappa\alpha\beta & 0\end{bmatrix} f_{a,t} + \begin{bmatrix}0 & 0 & 1\end{bmatrix} (I_3 - \alpha\beta h)^{-1} s_t, \tag{J.2}$$

$$f_{a,t} = \frac{1}{1-\alpha\beta}\bar{\pi}_{t-1} + b(I_3 - \alpha\beta h)^{-1} s_t, \tag{J.3}$$

$$f_{b,t} = \frac{1}{1-\beta}\bar{\pi}_{t-1} + b(I_3 - \beta h)^{-1} s_t, \tag{J.4}$$

$$f_{t|t-1} = \pi_t - (\bar{\pi}_{t-1} + b_1 s_{t-1}), \tag{J.5}$$

$$k_t = \sum_i \gamma_i b_i (f_{t|t-1}), \tag{J.6}$$

$$\bar{\pi}_t = \bar{\pi}_{t-1} + k_t f_{t|t-1}, \tag{J.7}$$

$$\pi_t + \frac{\lambda_x}{\kappa} x_t = c\left(k_t + f_{t|t-1}\mathbf{g}_t'\right)\left((1-\beta)\,\mathbb{E}_t\, x_{t+1} + \beta\,\mathbb{E}_t \sum_{i=1}^{\infty} x_{t+i} \prod_{j=0}^{i-1} 1 - k_{t+j} - f_{t+j|t-1+j}\mathbf{g}_{t-1+j}'\right), \tag{J.8}$$

with $c \equiv \frac{\lambda_x}{\kappa}\frac{(1-\alpha)\beta^2}{1-\alpha\beta}$ as in the main text. Denote the expectation on the right hand side of (J.8) as E_t. The idea of the PEA is to approximate this expectation and to solve model equations given the approximation $\hat{E}_t$. The algorithm is as follows:[27]

Objective: Obtain the sequence $\{i_t\}_{t=1}^T$ that solves Equations (J.1) - (J.8) for a history of exogenous shocks $\{s_t\}_{t=1}^T$ of length T.

1. Conjecture an initial expectation $\hat{E}_t = \beta^0 s(X_t)$.
   The expectation is approximated as a projection on a basis, $s(X_t)$, where β^0 are initial projection coefficients, and $X_t = (k_t, \bar{\pi}_{t-1}, r_t^n, u_t)$ is the state vector. I use a monomial basis consisting of the first, second and third powers of X_t.

2. Solve model equations given conjectured $\hat{E}_t$ for a given sequence of shocks $\{s_t\}_{t=1}^T$.
   Compute residuals to the model equations (J.1) - (J.8) given $\{s_t\}_{t=1}^T$ and $\{\hat{E}_t\}_{t=1}^T$. Obtain a sequence $\{i_t\}_{t=1}^T$ that sets the residuals to zero. The output of this step is $\{v_t\}_{t=1}^T$, the simulated history of endogenous variables (Christiano and Fisher (2000) refer to this as a "synthetic time series").

3. Compute realized analogues of $\{E_t\}_{t=1}^T$ given $\{v_t\}_{t=1}^T$.

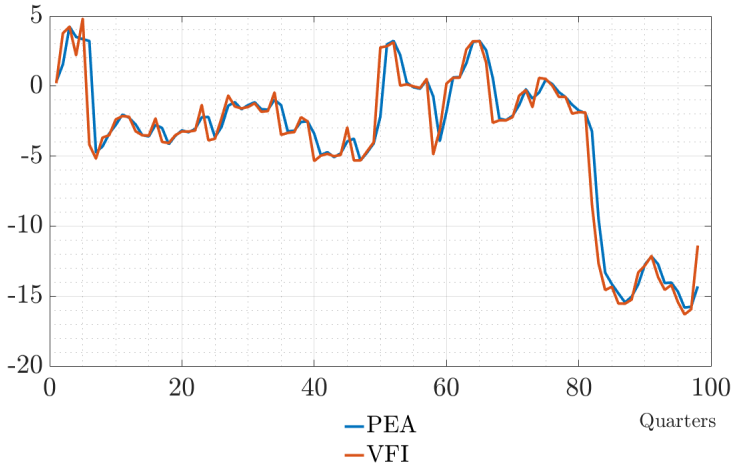
4. Update β regressing the synthetic E_t on $s(X_t)$.
   The coefficient update is $\beta^{i+1} = (s(X_t)'s(X_t))^{-1}s(X_t)'E_t$. Then iterate until convergence by evaluating at every step $||\beta^i - \beta^{i+1}||$.

[27] For a thorough treatment of the PEA, see Christiano and Fisher (2000).

K Parametric value function iteration

This is an alternative approach I implement as a robustness check to the PEA. The objective is thus the same: to obtain the interest rate sequence that solves the model equations. The general value function iteration (VFI) approach is fairly standard, for which reason I refer to the Judd (1998) textbook for details. Specific to my application is that the state vector is five-dimensional, $X_t = (\bar{\pi}_{t-1}, r^n_t, u_t, r^n_{t-1}, u_{t-1})$, and that I approximate the value function using a cubic spline. Thus the output of the algorithm is a cubic spline approximation of the value function and a policy function for each node on the grid of states. Next, I interpolate the policy function using a cubic spline as well. As a last step I pass the state vector from the PEA simulation, obtaining an interest rate sequence conditional on the history of states. Fig. 12 shows the resulting interest rate sequence, obtained through the two approaches, conditional on a simulated sequence for the exogenous states.

Figure 12: Policy function for a particular history of states, PEA against VFI



L Oscillatory dynamics in adaptive learning models

Here I present an illustration for why adaptive learning models produce oscillatory impulse responses if the gain is high enough. Consider a stylized adaptive learning model in two equations:

$$\pi_t = \beta e_t + u_t, \tag{L.1}$$

$$e_t = e_{t-1} + k(\pi_t - e_{t-1}). \tag{L.2}$$

The reader can recognize in (L.1) a simplified Phillips curve in which I am abstracting from output gaps to keep the presentation as clear as possible. Like in the simple model of Section 5.1 in the main text, e_t represents the one-period inflation expectation $\hat{\mathbb{E}}_t \pi_{t+1}$. (L.2), then, represents the simplest possible recursive updating of the expectations e_t. My notation of the gain as k indicates a constant gain specification, but the intuition remains unchanged for decreasing (or endogenous) gains.

Combining the two equations allows one to solve for the time series of expectations

$$e_t = \frac{1-k}{1-k\beta} e_{t-1} + \frac{k}{1-k\beta} u_t, \tag{L.3}$$

which, for β close but smaller than 1, is a near-unit-root process. (In fact, if the gain were to go to zero, this would be a unit root process.) Defining the forecast error as $f_{t|t-1} \equiv \pi_t - e_{t-1}$, one obtains

$$f_{t|t-1} = -\frac{1-\beta}{1-k\beta} e_{t-1} + \frac{1}{1-k\beta} u_t. \tag{L.4}$$

Equation (L.4) shows that in this simple model, the forecast error loads on a near-unit-root process with a coefficient that is negative and less than one in absolute value. Damped oscillations are the result.

Note that even if the gain would converge to zero, the coefficient on e_{t-1} would be negative and less than one in absolute value. Thus even for decreasing gain learning, one obtains oscillations, but the lower the gain, the more damped the oscillations become. This corroborates my findings in the impulse responses of Fig. 6. But importantly, the opposite extreme, when $k \to 1$, results in a coefficient of exactly -1, giving perpetual oscillations. This clearly illustrates how the oscillatory behavior of impulse responses comes from the oscillations in the forecast error that obtain when the gain is sufficiently large.