Boundedly rational expectations and the optimality of flexible average inflation targeting[*]

Anthony Brassil[†] Christopher G. Gibbs[‡] Callum Ryan[†]

23 January 2026

Abstract

Expectations play a central role in the transmission of monetary policy, but how people form expectations is widely debated. We study optimal monetary policy design in a model with behavioral expectations that nests rational and adaptive learning beliefs as special cases, and approximates the aggregate implications of several bounded-rationality theories. We show that optimal policy is robustly characterized by a single policy framework across expectation theories: Flexible Average Inflation Targeting. Distinct from existing characterizations of such policies, we make precise what *flexible* and *average* mean, and how they depend on expectation formation and constraints faced by the central bank (imperfect information and the zero lower bound).

JEL Classifications: E31; E32; E52; E71; D83; D84.

Key Words: Optimal policy; behavioral macroeconomics; average inflation targeting; Zero lower bound.

---

[*] We thank Bruce Preston, Ippei Fujiwara, Nigel McClung, Kevin Lansing, Matthew Read, and Jonathan Hambur for helpful comments and suggestions. We thank seminar participants at the Bank of Canada, George Washington University, Federal Reserve Board of Governors, Federal Reserve Bank of San Francisco, Reserve Bank of New Zealand, BIS Inflation Expectations Network, University of Oregon, the Bank of Finland, and the Reserve Bank of Australia. We thank the organisers of the Workshop of the Australasian Economic Society 2022, Econometric Society Australasian Meetings 2023, the 2023 Reserve Bank of Australia's Quantitative Macroeconomics Workshop, and the 2025 e61 UNSW Macroeconomics Workshop for allowing us to present.

[†] Reserve Bank of Australia, Sydney 2000, Australia. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Reserve Bank of Australia.

[‡] School of Economics, The University of Sydney, Camperdown, NSW 2006. Email: christopher.gibbs@sydney.edu.au. All errors are the fault of the authors.

1 Introduction

There is agreement among economists that expectations are central to monetary policy design. However, there is disagreement over how to model expectations, and prescriptions about the optimal conduct of monetary policy can be sensitive to assumptions about expectation formation. We formulate a general framework for expectation formation that nests or approximates the aggregate implications of alternative expectation theories. We use this general environment to derive optimal target criteria under commitment (from a timeless perspective) for a central bank that seeks to minimize inflation and output volatility given different assumptions about the constraints that it faces – such as an inability to observe contemporaneous economic conditions or the zero lower bound (ZLB) on nominal interest rates. We characterize the dependence of optimal policy prescriptions on features of the aggregate expectation formation process and provide recommendations that are independent of how expectations are formed.

We find that the optimal target criterion in our general boundedly rational New Keynesian environment is always a form of Flexible Average Inflation Targeting (FAIT) where we make precise what *average* and *flexible* mean. Importantly, our characterization of FAIT departs in several key ways from the policy pursued by the Federal Reserve from 2020 until 2025.

We show that the *average* in FAIT should be a weighted average of current and past inflation with declining weights over time. Policy should react most strongly to current inflation and continue to partially "make-up" past inflation gaps. Adherence to the weighted-average target, however, must also be *flexible.* When a policymaker's ability to achieve the weighted-average target is constrained, it should seek to influence expectations more aggressively than usual.

This optimal policy flexibility has two components: additional make-up and preemption. A policymaker should only commit to stronger-than-usual make-up policy (the emphasis of the Federal Reserve's former FAIT policy) when they have actually been unable to implement their desired policy, such as when the ZLB is currently binding or when they are uncertain about the current state of the economy. In these cases, policy should seek to make up for past output gaps as well as past inflation misses.

When constraints are anticipated but not yet binding, or shocks are anticipated, optimal policy calls for "preemptive" deviations from the weighted-average target.[1] Especially with respect to an anticipated ZLB constraint, policy is a "use it or lose it" proposition for all the cases that we study.

Central to understanding the advantage of FAIT's focus on a weighted-average inflation target across differing aggregate expectation environments is recognizing a parallel

---

[1] Powell (2025) Jackson hole speech specifically highlights "preemptive" actions as a missing element in the Federal Reserve's FAIT framework, which was too focused on make-up policy and complicated communication about preemptive tightening.

between the history dependence optimal policy requires when there is some forward-looking expectation formation (such as with rational expectations) and the history dependence embedded within the aggregate expectations if beliefs are partially formed by adaptively learning from past observations.

If aggregate expectations contain forward-looking elements, the weighted-average inflation target commits the policymaker to continue responding to current shocks into the future. This commitment changes expectations today in a way that optimally spreads the welfare costs of shocks over time. The more aggregate expectations deviate from the full-information rational expectations (FIRE) benchmark, the more the optimal weights decline over time.[2]

If aggregate expectations contain elements of learning from past observations, expectations themselves become tied to weighted averages of past data and evolve as belief states that policy must steer. Importantly, we show that this creates a distinct rationale for targeting a weighted average of past inflation outcomes. Hence, despite different underlying mechanisms, optimal policy across these environments naturally features targets for weighted averages of past outcomes, which explain why we find FAIT is a robust policy.

Showing the similarity of optimal policy with different underlying mechanisms allows us to reconcile two distinct strands of literature on optimal policy under bounded rationality. Eusepi and Preston (2018) and Eusepi, Giannoni, and Preston (2018) show that optimal monetary policy under adaptive learning shares many features with the optimal commitment policy under FIRE, despite little ability for policy commitments to shape beliefs. In contrast, Gabaix (2020) argues that when agents are forward-looking but myopic or engage in low levels of reasoning, optimal policy must depart from the FIRE commitment benchmark because these commitments have a weaker effect on expectations. We show that these different insights actually point to the same policy framework: FAIT. Moreover, many other varieties of bounded rationality, which we can approximate in the aggregate dynamics in our model, point to the same optimal policy.

To establish these general results, we take a reduced-form approach to expectations. Rather than committing to a particular microfoundation, we leverage the growing recognition that the aggregate implications of several different bounded rationality models are similar, even though the microeconomic information structures and predictions for individual level outcomes differ. An example of this growing recognition is Angeletos, Guerreiro, and Zhang (2025), who prove observational equivalence for an aggregate IS relationship under: noisy information, overconfidence and ambiguity, level-k thinking, representative agent inattention, sparsity, cognitive discounting, and confusion between

---

[2] While we focus on how expectation formation influences optimal policy, the weighted-average inflation target aspect of our optimal policy prescription could also be derived in a setup with FIRE but where the private sector has a different discount factor to the policymaker or their expectations are attenuated for some other structural feature. See Section 2 for further discussion.

aggregate and idiosyncratic shocks. Given we study a central bank with a mandate just to stabilise aggregate inflation and output outcomes, the similarity in the aggregate implications of these theories lets us characterise optimal policy under a broad range of microfoundations.

Specifically, we propose an expectations operator for inflation and output gaps that combines a rational forward-looking component with an adaptive-learning component.[3] By adjusting the relative weights on the forward-looking and adaptive learning components, and by adjusting the parameters that govern learning, we are able to capture the aggregate implications of:

- models with bounded rationality, such as the cognitive discounting model of Gabaix (2020) and the level-k rasoning models of Farhi and Werning (2019) and Evans, Gibbs, and McGough (2025);
- models with information frictions, such as the imperfect common knowledge models of Angeletos and Lian (2018) and Angeletos and Huo (2021); and
- adaptive learning models in the tradition of Evans and Honkapohja (2001) and Preston (2005), in which agents estimate and update reduced-form forecasting rules over time.

At an aggregate level, each of these microfoundations generates some combination of myopia (i.e. over-discounting of rational expectations of the future) and/or anchoring to past outcomes (i.e. persistent dependence on lagged outcomes) (see e.g. Angeletos and Huo, 2021 and Angeletos, Guerreiro, and Zhang, 2025).

In addition to capturing the aggregate implications of several bounded-rationality theories, related papers (Beckers and Brassil, 2022 and Brassil, Haidari, Hambur, Nolan, and Ryan, 2024) show that our general form of expectations can accurately capture the surveyed average expectations of households and unions (with parameter estimates that significantly differ from FIRE, myopia, or adaptive learning benchmarks).

By changing the parameters that govern this general expectations operator, we can trace how optimal policy varies with the features of expectation formation. We use this capability to demonstrate the robustness of FAIT policy by running two horse races comparing a range of different inflation and output gap policy targets as the expectation formation process is varied.

We first show that when the expectation formation process is known, policy targeting a combination of a weighted average of past inflation gaps and a weighted average of past output gaps performs similarly to the optimal FAIT policy. And that this simple combination of weighted averages outperforms price-level targeting, inflation targeting, and arithmetic average inflation targeting using a fixed window.

---

[3] We assume the intended path of the policy rate is credible and commonly understood.

We then compare how these competing criteria perform when policymakers do not know how expectations are formed. Specifically, we choose a single calibration for each different target criteria and then compare performance as we change the underlying expectation formation process. We again find that target criterion combining weighted averages of current and past inflation and output gaps are robust, generating similar losses to the fully optimal FAIT policy regardless of how expectations are formed. We discuss how this finding stems from the ability of policy to capture some aspects of the optimal FAIT flexibility by additionally targeting a weighted average of output gaps.

Literature review. We take a less used approach to the study of FAIT by using target criteria. Much of the policy work on this topic has focused on the study of modified interest rate rules that include averages of past inflation as arguments. For example, papers cited in strategy reviews by the Federal Reserve (e.g. Arias, Bodenstein, Chung, Drautzburg, and Raffo, 2020), European Central Bank (e.g. Cecioni, Coenen, Motto, Le Bihan, Ajevskis, Albertazzi, Gilbert, Al-Haschimi, Gomes, Bornemann, et al., 2021), and the Bank of Canada (e.g. Dorich, Mendes, and Zhang, 2021) have taken this approach.

There are several drawbacks to interest rate rule based analysis. Svensson (2003) and Svensson and Woodford (2005) argue that interest-rate rules are a fragile and non-transparent way of specifying a policy framework.[4] Interest rate rules are shock-specific and the mapping between policy framework and an interest-rate rule is not clear cut. In some cases, different rules can generate the same equilibrium outcomes (see, for example, Eskelinen, Gibbs, and McClung, 2024) or are very sensitive to the exact lag specification (see, for example, Honkapohja and McClung, 2024 or Jia and Wu, 2023). It is therefore difficult to draw robust conclusions about policy frameworks from a comparison of interest rate rules.

In contrast, target criteria are general. They do not depend on the statistical properties of or the number of economy-wide shocks that the central bank faces. They depend only on how beliefs are formed and on the constraints the central bank encounters when implementing its desired policy. In addition, policy is specified as a target for endogenous variables, i.e., a set objective for the evolution of economic outcomes policymakers hope to achieve. Target criterion are inherently less constraining than interest rate rules because they allow (subject to the ZLB) any choice for the path of the policy rate necessary to achieve their goals. Communication of policy in such a framework is, therefore, forward-

---

[4] There has been a gradual evolution of the language describing different characterizations of monetary policy. Svensson and Woodford define a monetary policy rule broadly as a prescribed guide for monetary policy. Target criterion and interest rate rules are two different ways of specifying a policy rule in their framework. In more recent treatments in the literature, there is no distinction made between interest rate rules and policy rules. The two are synonymous. We adopt the term policy framework to capture the broader notion of a general policy that may be implemented with a target criterion or an interest rate rule.

looking and data dependent, which reflects the way policymakers actually communicate.

Optimal target criteria for both unconstrained and constrained policymakers have been widely studied under FIRE. The theoretical justification for inflation targeting (e.g. Giannoni and Woodford, 2005), price level targeting (e.g. Giannoni, 2014), and inflation-forecast targeting (e.g. Svensson and Woodford, 2005) policy frameworks rest on this work. In addition, Eggertsson and Woodford (2003) extends optimal target criteria to the case of optimal policy at the ZLB, and many have studied imperfect central bank information such as discussed in Clarida, Gali, and Gertler (1999) or Woodford (2010). We extend these works by nesting it within a more general model of expectations. We show that some conclusions from the FIRE analysis are knife edge. Small departures from the FIRE assumption imply a new optimal policy framework: FAIT, which prescribes a clear role for preemptive policy missing from the FIRE conclusions.

Optimal target criteria for an unconstrained policymaker have been studied in the adaptive learning literature, such as in Molnár and Santoro (2014) and Eusepi, Giannoni, and Preston (2018, forthcoming). These papers establish that optimal policy shares some features with the optimal target criterion from the FIRE analysis in which it is optimal to engineer periods of price level overshooting or make-up policy in certain circumstances. We extend these works by approximating them within our general model of expectations and by deriving optimal target criteria in settings with imperfect information and the ZLB. We show that the similarity in optimal policy under both adaptive learning and FIRE carries over to other general forms of bounded rationality, and can be captured by a FAIT policy framework.

A related set of studies derive optimal target criteria in the unconstrained case under other kinds of deviations from FIRE. Gabaix (2020) and Benchimol and Bounader (2023) derive optimal policy when agents have myopic expectations and over-discount the future. In Dupraz and Marx (2023), agents have finite planning horizons and are similarly myopic, but also learn about long-run outcomes, which adds a backward-looking component to expectations. Gasteiger (2014, 2021) analyses optimal policy in a heterogeneous expectations setting, where some agents have rational expectations and others are adaptive. Di Bartolomeo, Di Pietro, and Giannini (2016), and Hagenhoff (2021) also consider a heterogeneous-expectations framework, but their focus is on how heterogeneity in expectations across agents affects welfare via price and consumption dispersion. We capture much of the aggregate implications for inflation-output dynamics of these theories but abstract from the issue of price and consumption dispersion that such heterogeneity may generate. We are instead interested in characterising the policy implications of a broad set of theories of expectation formation for a central bank with a mandate just to stabilise aggregate outcomes. We also extend the optimal policy analysis beyond previous work by deriving optimal target criteria in the constrained case, when the central bank has imperfect information or the ZLB can bind.

Outside of the unconstrained case, there are several papers that explore optimal policy at the ZLB in non-FIRE settings. Eusepi, Gibbs, and Preston (2024) and Evans, Gibbs, and McGough (2025) both study optimal forward-guidance policy when agents must learn about the general-equilibrium implications of policy. The former takes an adaptive learning approach, while the latter combines adaptive learning with level-k thinking. Dupraz, Le Bihan, and Matheron (2024) analyze make-up policy when agents have finite planning horizons. Budianto, Nakata, and Schmidt (2023) allows for both FIRE and a Gabaix (2020) style of myopic expectations, and compares delegated loss functions for a central bank optimizing under discretion. In contrast to these papers, we derive optimal target criteria under commitment analytically, which shows that a particular kind of FAIT framework can capture the features of optimal policy at the ZLB that this prior work explores.

Finally, there is a related adaptive learning literature that examines whether optimal target criteria derived under FIRE generate equilibria that are learnable. This literature finds that it depends on how the optimal target criteria is implemented. Optimal target criteria do not imply unique interest rate rules. Evans and Honkapohja (2003), Preston (2006), Orphanides and Williams (2007), and Evans and McGough (2010) show that different interest rate rules that react to different endogenous quantities, which implement the same equilibrium under FIRE, may have different stability properties under least squares learning. Honkapohja and Mitra (2020) study price level targeting under learning and varying credibility. The optimal targeting criteria we study here, however, are derived taking learning into account as a constraint, making expectational stability an endogenous concern of the policymaker. Our robust policy analysis finds versions of the FAIT framework that ensure expectational stability even when the policy rule is misspecified.

2 THE MODEL

We study the standard New Keynesian environment of Woodford (2003) or Galí (2015) with households that maximize utility choosing sequences of consumption and labor supply and firms that face sticky prices and engage in monopolistic competition. Following Preston (2005) and Eusepi and Preston (2018), the aggregate consumption and pricing decisions of agents under arbitrary expectations in this environment may be log-linearly approximated by

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

$$\pi_t = \hat{\mathbb{E}}_t \sum_{T=t}^{\infty} (\theta\beta)^{T-t} \left[\kappa x_T + (1-\theta)\, \beta \pi_{T+1} + \mu_T \right] \tag{2}$$

where the output gap is defined as

$$x_t = y_t - y_t^n.$$

The natural level of output, y_t^n, and natural real rate of interest, r_t^n, are defined as the values that would occur in the equivalent flexible-price economy and are exogenous.[5] Inflation is denoted by π_t with θ the complementary probability of price reset, i_t is the nominal interest rate controlled by the central bank, r_t^n is an exogenous real rate of interest, u_t is a cost-push shock.

Appendix C provides more details and the derivations for all the main results in this paper using the equations (1) and (2). However, to increase the interpretability of our main results, we further simply the aggregate Phillips curve to

$$\pi_t \;\; = \;\; \beta\hat{\mathbb{E}}_t\pi_{t+1} + \kappa x_t + u_t, \tag{3}$$

where the expectations operator remains undefined. This aggregate relationship (3) is implied directly by (2) under several different expectation theories that we approximate and our conclusions are not sensitive to this choice (in contrast to the choice of IS curve representation).[6]

Beliefs of the private sector. We put forward a general model of belief formation. We do not intend for our formulation to be a theory of expectation formation on its own; however, many papers have done so using similar approaches such as in Brock and Hommes (1997), Branch and McGough (2009), and Cole and Martínez-García (2023). Our goal is to write beliefs in such a way that for specific calibrations we recover the aggregate implications of distinct expectation theories put forward in the literature.

We assume that aggregate expectations are homogeneous and are a combined forecast of FIRE and adaptive learning expectations, where the weight placed on FIRE belief is a fraction $\lambda \in [0,1]$ and the complementary weight of $1-\lambda$ is placed on adaptive learning beliefs. This follows Gibbs (2017) and Gibbs and Kulish (2017) and assumes that the heterogeneity in forecasts exists within the representative household and firm so that the expectations that matter for consumption, labor supply, and pricing decisions are homogeneous. In other words, we maintain a representative agent assumption where the representative decision-maker takes a weighted average of forecasts from two different

---

[5] Output in the flexible-price economy is independent of expectation formation and a function of exogenous shocks (e.g. Preston, 2005). The real interest rate in the flexible-price economy is a function of current and expected exogenous shocks.

[6] We do not use a two period approximation of the IS curve because we make different assumptions about the formation of output-gap and interest-rate expectations. In the two-period 'Euler equation' representation, nominal interest-rate expectations do not even appear. See Preston (2005) for a detailed discussion of the two representations.

models: the correct structural model, and a reduced-form model which is re-estimated each period.

Aggregate expectations of future inflation and output gaps are

$$\hat{\mathbb{E}}_t \pi_T = \lambda \mathbb{E}_t \pi_T + (1-\lambda)\mathbb{E}_t^l \pi_T \tag{4}$$

$$\hat{\mathbb{E}}_t x_T = \lambda \mathbb{E}_t x_T + (1-\lambda)\mathbb{E}_t^l x_T \tag{5}$$

where $\mathbb{E}_t$ is the FIRE forecast and $\mathbb{E}_t^l$ denotes the learning forecast. We assume that adaptive learning forecasts inflation and the output gap using an unobserved components state space model, which is estimated using a steady-state Kalman Filter. Inflation and the output are assumed to be driven by an unobserved persistent component, with an autoregressive coefficient of ρ, and an i.i.d. component. Learning expectations use observations up to period $t-1$ to estimate the persistent component, where ω_{t-1}^z, $z \in \{\pi, x\}$, denotes period t's state estimate. Learning expectations use this estimate to forecast future variables z_T, $T \geq t+1$. The learning expectations for π_T and x_T, $T \geq t+1$, are therefore given by

$$\mathbb{E}_t^l \pi_T = \rho^{T-t} \omega_{t-1}^{\pi} \tag{6}$$

$$\mathbb{E}_t^l x_T = \rho^{T-t} \omega_{t-1}^{x}. \tag{7}$$

where the steady-state gain coefficient g reflects the beliefs of the relative variance of the perceived unobserved persistent and transitory shocks in the state space model. This procedure implies the updating rules

$$\omega_t^{\pi} = \rho \omega_{t-1}^{\pi} + \rho g \left(\pi_t - \omega_{t-1}^{\pi}\right) \tag{8}$$

$$\omega_t^{x} = \rho \omega_{t-1}^{x} + \rho g \left(x_t - \omega_{t-1}^{x}\right). \tag{9}$$

We assume for simplicity that the parameters relevant for beliefs - λ, ρ and g - are the same for both households and firms, and for both the inflation and output. Importantly, the learning expectations are always misspecified in the presence of shocks.

For expectations of the nominal interest rate, we assume that everyone has rational expectations over its path.[7] This is equivalent to assuming that the policymaker communicates its expected policy rate path each period, and that path is believed by all agents. The distinction between the formation of nominal interest-rate expectations and the formation of output gap and inflation expectations is common in the existing literature on make-up policy.[8] If the policymaker communicates its expected path for the

---

[7] Note that we still allow for drift in longer-term *real* interest rates due to drifting beliefs about future inflation.

[8] Farhi and Werning (2019) assume that households perfectly observe future policy rates, but deduce the consequences for inflation and output gap expectations using level-k reasoning. Similarly, Dupraz,

nominal interest rate, then this path is immediately available to households and firms, whereas determining the future output gap and inflation implications of that path involves much more sophisticated general-equilibrium reasoning. In Appendix B, we show how results are affected when this assumption is relaxed.

Empirical estimates Although our general form of expectations is not meant to be a novel theory of how people form expectations, it is shown to fit aggregate survey data well. When fitted to survey data of unions' one- two-year ahead inflation expectations, Beckers and Brassil (2022) find that around one-quarter of union officials' expectations are suitably proxied by rational expectations ($\lambda = 0.26$) and that those who behave as *adaptive learners* learn quickly ($g = 0.7$). This simple expectation formation process captured most of the variation in surveyed expectations ($R^2 = 0.76$).

Using the Melbourne Institute's *Consumer Attitudes, Sentiments and Expectations in Australia* Survey, Brassil, Haidari, Hambur, Nolan, and Ryan (2024) find a similar share of rationally formed expectations ($\lambda = 0.22$) but find the household *adaptive learners* learn much slower than unions ($g = 0.11$). Brassil, Haidari, Hambur, Nolan, and Ryan (2024) also explore the household-level data to assess the robustness of these results. While there is variation across cohorts, around one-quarter rationally formed expectations and g between 0.1 and 0.2 was a consistent finding.

Not only does this general form of expectations do a good job of capturing Australian inflation expectation formation in-sample, the model also provided an accurate prediction of how inflation expectations would respond to the recent high inflation episode (see Graph 7 in Hunter, 2024).

Beliefs of the policymaker. We assume throughout that the central bank possesses rational expectations and seeks to maximize the standard quadratic loss function

$$\mathcal{W} \equiv -\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\pi_t^2 + \alpha x_t^2\right) \tag{10}$$

which sees the central bank trade-off inflation and output volatility. The loss function (10) is a quadratic approximation to household welfare when $\alpha = \kappa/\theta$ if we assume a representative household and non-distorted steady state. More generally, our results can be seen as characterising optimal policy for a central bank with a 'narrow mandate' to stabilise aggregate outcomes (in the language of McKay and Wolf, 2023), without

---

Le Bihan, and Matheron (2024) model households with finite-planning horizons but who save and borrow based on financial prices determined by intermediaries with fully rational expectations over nominal interest rates. Quantitative evaluations of make-up policy by Federal Reserve Board staff have have used versions of their semi-structural model (FRB-US) in which financial-market participants have model-consistent expectations, even while households and firms may not be rational (Bernanke, Kiley, and Roberts, 2019; Hebden, Herbst, Tang, Topa, and Winkler, 2020).

any mandate to target distributional outcomes or a different average level of output.[9] When choosing policy, the policymaker takes household and firm decisions, and beliefs as constraints:

$$\pi_t = \lambda\beta\mathbb{E}\pi_{t+1} + (1-\lambda)\beta\rho\omega^{\pi}_{t-1} + \kappa x_t + u_t \tag{11}$$

$$\begin{aligned} x_t = {} & (1-\beta)\sum_{T=t}^{\infty}\beta^{T-t}\left(\lambda\mathbb{E}_t x_{T+1} + (1-\lambda)\rho^{T+1-t}\omega^{x}_{t-1}\right) \\ & - \frac{1}{\sigma}\sum_{T=t}^{\infty}\beta^{T-t}\left(\mathbb{E}_t i_T - \lambda\mathbb{E}_t\pi_{T+1} - (1-\lambda)\rho^{T+1-t}\omega^{\pi}_{t-1} - \mathbb{E}_t r^n_T\right) \end{aligned} \tag{12}$$

$$\omega^{\pi}_t = \rho\omega^{\pi}_{t-1} + \rho g\left(\pi_t - \omega^{\pi}_{t-1}\right) \tag{13}$$

$$\omega^{x}_t = \rho\omega^{x}_{t-1} + \rho g\left(x_t - \omega^{x}_{t-1}\right). \tag{14}$$

The policymaker's problem. The optimal policy problem is to choose the sequence

$$\{\pi_t, x_t, \omega^{\pi}_t, \omega^{x}_t, i_t\}_{t=0}^{\infty}$$

that maximises the objective function (10) subject to the constraints (11) – (14) imposed by private-sector equilibrium behaviour. We assume that the policymaker is credible and pursues optimal policy *from a timeless perspective*, in that they do not seek to exploit the exogeneity of the initial conditions. Instead, the policymaker in the initial period acts according to the rule that would have been optimal had they committed to it in advance.[10]

Approximating aggregates dynamics of other bounded rationality models. The payoff for the particular way we model beliefs is in capturing the aggregate implications of many different bounded rationality models in the literature through λ and g. This is because a broad range of deviations from FIRE are in the aggregate equivalent to some form of over-discounting of the rational expectations (which we capture via λ) and/or an amplified dependence on lagged outcomes (which we capture via g). For example, consider first the case where $\lambda \in (0,1]$ and $g = 0$. Equations (11) and (12) reduce

---

[9] If households and firms are actually heterogeneous in their beliefs, then there is additional dispersion in consumption, labor supply, and pricing decisions that enter the approximation to household utility. One way to view this heterogeneity is that it is similar to whether monetary policy should consider deviations from the distorted steady state or the efficient steady state, where in the latter there is an implicit fiscal policy at work that takes care of these other inefficiencies in the economy. To the extent that disagreement is not fundamental but simply a question of information and education, it begs the question of whether monetary policy should consider it as an additional objective or whether new policy tools should be used to target this source of inefficiency. This question is beyond the scope of this paper and an interesting question for future work.

[10] In Appendix B, we use a simple two-period model to explore what optimal policy under discretion would look like with our hybrid specification for expectation formation.

to

$$\begin{aligned} \pi_t &= \lambda\beta\mathbb{E}_t\pi_{t+1} + \kappa x_t + u_t \\ x_t &= (\beta + \lambda(1-\beta))\mathbb{E}_t x_{t+1} - \frac{1}{\sigma}(i_t - \lambda\mathbb{E}_t\pi_{t+1} - r_t^n), \end{aligned}$$

where λ functions as an attenuation parameter on expectations. Set $\lambda = 1$ and we recover the FIRE equations. Set $\lambda \in (0,1)$ and we approximately recover either the aggregate model of Gabaix (2020) described in his Proposition 2 or the aggregate environment implied by Angeletos and Lian (2018) described in Proposition 10 of their appendix with a common information friction. Angeletos, Guerreiro, and Zhang (2025) show that this type of attenuation is implied in aggregate, and is observationally equivalent, across several additional bounded rationality theories including level-k reasoning models of the type considered by Farhi and Werning (2019).

When $\lambda = 0$ and $g \in (0,1)$, the system reduces to a model of adaptive learning:

$$\begin{aligned} \pi_t &= \beta\rho\omega_{t-1}^{\pi} + \kappa x_t + u_t \\ x_t &= \frac{1-\beta}{1-\beta\rho}\rho\omega_{t-1}^{x} + \frac{1}{\sigma(1-\beta\rho)}\rho\omega_{t-1}^{\pi} - \frac{1}{\sigma}\mathbb{E}_t\sum_{T=t}^{\infty}\beta^{T-t}(i_T - r_t^n), \end{aligned}$$

which captures many of the key features of Evans and Honkapohja (2003), Molnár and Santoro (2014) and Eusepi, Giannoni, and Preston (2018).

Lastly, when $\lambda \in (0,1)$ and $g \in (0,1]$, the model captures many of the aggregate implications of the myopia and anchoring framework of Angeletos and Huo (2021), the sticky information model of Mankiw and Reis (2002), or joint level-k reasoning and adaptive learning model of Evans, Gibbs, and McGough (2025). For example, when $g = 1$, the model captures the Hybrid New Keynesian Phillips curve property implied by Angeletos and Huo (2021):

$$\pi_t = \lambda\beta\mathbb{E}_t\pi_{t+1} + (1-\lambda)\beta\rho\pi_{t-1} + \kappa x_t + u_t$$

By nesting the aggregate implications of these competing theories, we can compare optimal policy within a shared framework.[11]

---

[11] Another prominent form of bounded rationality studied in the literature is Diagnostic Expectations. See, for example, Bordalo, Gennaioli, and Shleifer (2018), Bianchi, Ilut, and Saijo (2024a), and Bianchi, Ilut, and Saijo (2024b). We do not study these expectations directly; however, we note that our framework for expectations can generate a similar kind of general equilibrium overreaction when λ is close to one and $g > 0$, i.e., expectations initially respond to a larger degree than if all expectations were rational. This occurs because the rational component takes into account the learning expectations, which will drift when $g > 0$.

3 Optimal Policy: The Unconstrained Case

We start with the unconstrained optimal policy problem given by equations (10) - (14). The policymaker is free to choose any policy today and into the future that they would like. Appendix A contains the full derivations for all results; here we summarize the key equations to build intuition.

Monetary policy in this case is a powerful tool even with bounded rationality because the path of interest rates is common knowledge. If the central bank only faces the demand shock, then it is able to perfectly stabilize the economy as in the standard FIRE case.

Result 1 (Demand Irrelevance)**.** *Demand is not a constraint when the path of interest rates is known and credible. Demand shocks may be perfectly offset with current and promised changes in the interest rate incurring no welfare loss.*

The trade-off for the central bank, therefore, lies with cost push shocks.

To characterize optimal policy for cost push shocks, we take the first-order conditions with respect to π_t and x_t for the central banks' policy problem and eliminate the Lagrange multiplier on the Phillips curve:

$$\frac{\alpha}{\kappa}x_t = -\frac{1}{1-\lambda L}\pi_t + \rho g\frac{1}{1-\lambda L}\mu_t^{\omega\pi} \tag{15}$$

where L is the lag operator and $\mu_t^{\omega\pi}$ is the Lagrange multiplier on the updating rule for learners' inflation expectations. The term $\mu_t^{\omega\pi}$ represents the benefit of increasing learners' inflation expectations and is determined by the first-order condition with respect to ω_t^{π}:

$$\mu_t^{\omega\pi} = \beta\rho(1-g)\mathbb{E}_t\mu_{t+1}^{\omega\pi} - \beta^2(1-\lambda)\rho\mathbb{E}_t\left(\frac{1}{1-\lambda L}\pi_{t+1} - \rho g\frac{1}{1-\lambda L}\mu_{t+1}^{\omega\pi}\right). \tag{16}$$

Solving (16) for $\mu_t^{\omega\pi}$ gives

$$\mu_t^{\omega\pi} = -(1-\lambda)\rho\Omega\frac{1-\lambda L}{1-\zeta_1 L}\mathbb{E}_t\frac{1}{1-\zeta_2^{-1}L^{-1}}\frac{1}{1-\lambda L}\pi_{t+1} \tag{17}$$

where $\Omega \equiv \frac{\beta^2\zeta_2^{-1}}{\rho\beta(1-g)+(1-\lambda)g\rho^2\beta^2} < 1$ and ζ_1 and ζ_2 are the roots of the characteristic polynomial associated with (16) (see Appendix A for derivation of results for ζ_1 and ζ_2).

The smaller root ζ_1 reflects the effectiveness of make-up commitments on current inflation. In other words, it measure how useful such commitments are for improving the trade-off between output and inflation the central bank faces in the current period. We have $\zeta_1 \in [\lambda, 1]$, and ζ_1 is increasing in λ, g, and ρ. If $\lambda \in \{0, 1\}$ or $g = 0$, then ζ_1 is just equal to the rational share of expectations: $\zeta_1 = \lambda$. But when there is both forward-looking expectations and learning, they interact to strengthen the effect of make-up commitments, so $\zeta_1 > \lambda$. For example, any increase in rational inflation expectations raises current inflation, which implies that the learned component of expectations will be higher next period. This implies higher future inflation, and further raises rational inflation expectations.

The reciprocal of the explosive root ζ_2^{-1} reflects the persistence of the effect of policy on future inflation (via beliefs). We have $\zeta_2^{-1} \in (0, \rho\beta(1-g) + (1-\lambda)g\rho^2\beta^2]$, and ζ_2^{-1} is decreasing in λ and g, and increasing in ρ. That is, policy has a more persistent effect on inflation via beliefs when (i) the backward-looking share is higher, (ii) beliefs are stickier and less updated in response to future outcomes, and (iii) current outcomes are extrapolated further into the future. Optimal policy then requires looking further into the future and responding more pre-emptively.

Result 2 (Unconstrained optimal target criterion)**.** *The optimal target criterion is*

$$\frac{\alpha}{\kappa}x_t = -\underbrace{\frac{1}{1-\lambda L}\pi_t}_{WAIT} - \underbrace{(1-\lambda)g\rho^2\Omega\frac{1}{1-\zeta_1 L}\mathbb{E}_t\frac{1}{1-\zeta_2^{-1}L^{-1}}\frac{1}{1-\lambda L}\pi_{t+1}}_{Preemption}, \tag{18}$$

which is composed of two distinct parts:

1. *a weighted average inflation target (WAIT)*
2. *preemptive policy actions based on expected future trade-offs between inflation and output gaps due to drifts in learning beliefs.*

The WAIT term represents a commitment to the forward-looking rational agents. The policymaker promises to adjust the output gap in response to a weighted average of past inflation. Shocks today that generate misses in the inflation target imply future policy adjustments, lessening the equilibrium impact of any shock.

The weight that policy should place on past inflation outcomes depends on the effect of changes in the rational inflation expectation on current inflation. If expectations are fully rational, $\lambda = 1$, then this pass-through is equal to the central bank's discount factor, β. Optimal policy then places an equal weight on current and all past inflation outcomes, i.e. price-level targeting. If expectations are less-than-fully rational, $\lambda < 1$, then the passthrough of the rational inflation expectation is lower and the effect of make-up commitments on current inflation is weaker. Optimal policy then places increasingly less weight on past inflation outcomes.[12]

The preemption term represent the ability of policy to change future trade-offs via backward-looking adaptive learning expectations if $g > 0$. By affecting current inflation,

---

[12] In our model, this passthrough is lower because of deviations of FIRE. But the passthrough would also be lower even with FIRE expectations if the private sector had a lower discount factor than the central bank. Optimal policy would then take a similar WAIT form (Blake, 2001; Jensen and McCallum, 2002; Eskelinen, Gibbs, and McClung, 2024).

Table 1: Special Cases

Cases	Target criteria	Policy
FIRE, $\lambda = 1$	$\frac{\alpha}{\kappa} x_t = -p_t$	PLT
Myopia/Level-k, $g = 0$	$\frac{\alpha}{\kappa} x_t = -\frac{1}{1-\lambda L} \pi_t$	WAIT
Adaptive learning, $\lambda = 0$	$\frac{\alpha}{\kappa} x_t = -\pi_t - g\rho^2\beta^2 \mathbb{E}_t \frac{1}{1-\zeta_2^{-1} L^{-1}} \pi_{t+1}$	IT + preemptive
Fixed beliefs, $\lambda = 0, g = 0$	$\frac{\alpha}{\kappa} x_t = -\pi_t$	IT

Notes: Optimal target criteria in special cases. PLT stands for 'price-level targeting', IT stands for 'inflation targeting'.

policy can influence backward-looking inflation beliefs and shift the location of the future Phillips curve. Additional policy today can preempt future undesirable movements in the trade-off between inflation and output, or engineer more favourable ones. For example, if policymakers expect above-target inflation in the future, even under the optimal management of that future trade-off (e.g. due to persistent cost-push shocks or drifts in inflation beliefs), then optimal policy requires pre-emptively aiming for lower inflation now, in order to lower the backward-looking component of inflation expectations.

Importantly, a key reason why policymakers would expect a future trade-off is because backward-looking beliefs might drift away from target, and these beliefs are themselves weighted averages of past inflation. In practice, this preemption term therefore also resembles WAIT. This means that it remains desirable to target a weighted average of inflation across a broad range of λ. We illustrate this explicitly in Section 3.2.

Lastly, a key feature of the optimal target criterion (18) is that it encapsulates the target criteria from a range of different primitive assumptions about how agents form beliefs as special cases. Table 1 summarizes the target criterion for the special cases discussed in the previous section. The table includes a fixed belief case of $\lambda = 0$ and $g = 0$, which is of interest because it recovers the same target criterion that is implied under FIRE for discretionary policy. This case reveals that so long as expectations can be influenced, i.e., $g > 0$ and/or $\lambda > 0$, it makes sense to deviate from the standard FIRE optimal discretion policy, i.e., from inflation targeting.

3.1 MECHANICS OF OPTIMAL POLICY

To see how policy works in practice in response to a shock, let the cost-push shock follow a first-order autoregressive process such that $E_t u_T$ depends only on u_t for $T \geq t$. Recall from eq. (17) that the Lagrange multiplier on learners' inflation expectations $\mu_t^{\omega\pi}$ is a function of past and current expectations of inflation at different horizons. In Appendix A, we show that, in the optimal equilibrium solution, this multiplier $\mu_t^{\omega\pi}$ can be written

$$\mu_t^{\omega\pi} = (1-\lambda)\left(a_x x_t - a_\omega \rho g \pi_t - a_\omega \rho (1-g) \omega_{t-1}^{\pi} - a_u u_t\right) \tag{19}$$

for some positive constants a_x, a_ω and a_u.[13] Sensibly, optimal policy aims for lower inflation expectation drifts, ω_t^π, when the inflation is high, ω_t^π is currently high, or the cost push shock is positive, and vice versa. This desire for lower inflation though is tempered by a desire to close the output gap, which give rise to the positive relationship with x_t. Using (19), we can write the optimal target criterion (18) (in equilibrium) as

$$\frac{\alpha}{\kappa} x_t = -\phi \frac{1}{1-\gamma L} \left(\pi_t + \nu_\omega \omega_{t-1}^\pi + \nu_u u_t \right) \tag{20}$$

where $\gamma \equiv \frac{\lambda}{1-(1-\lambda)\rho g a_x \frac{\kappa}{\alpha}}$, $\phi \equiv \frac{1+(1-\lambda)\rho^2 g^2 a_\omega}{1-(1-\lambda)\rho g a_x \frac{\kappa}{\alpha}}$, $\nu_\omega \equiv \frac{(1-\lambda)\rho^2 g(1-g) a_\omega}{1+(1-\lambda)\rho^2 g^2 a_\omega}$, and $\nu_u \equiv \frac{(1-\lambda)\rho g a_u}{1+(1-\lambda)\rho^2 g^2 a_\omega}$.

Equation (20) reveal that optimal policy adjusts the output gap in response to an exponentially-weighted average of current and past inflation, ω_t^π, and cost-push shocks, where the weight on lag j is γ^j. In addition, relative to FIRE, the central bank increases its response to current inflation while also responding to learners' inflation expectations ω_{t-1}^π – preemption – because both affect future inflation due to drifting inflation expectations. The central bank also must responds separately to the cost-push shock u_t (beyond its effect on current inflation). How aggressive or accommodative policy is with respect to past inflation, expectations, or the shock depends on λ and g, i.e., how quickly these interventions show up in expectations.

Figure 1 shows the optimal response to a cost push shock as λ is varied compared to policy under an inflation targeting (IT) regime ($\gamma = \nu_\omega = \nu_u = 0$ and $\phi = 1$). IT here is equivalent to the optimal targeting criteria implied under RE discretion.

Optimal policy for wide range of λ's implies initially a muted response followed by a more aggressive policy in the medium run compared to IT. In the high λ cases, policy reflects the make-up quality of WAIT. Higher inflation now means a higher interest rate in the future even as inflation falls. The expectation of this policy is capitalized into current expectations blunting the impact of the shock and lessening the need to initially raise interest rates aggressively.

In the low λ cases, policy reflects the different way a shock propagates when there is learning. Backward-looking agents' expectation are less sensitive to the shock on impact because their expectations are anchored to the past. This means the shock has less of an impact initially, which allows for a muted interest rate response by policymakers. However, the learning process endogenously propagates the shock as expectations are

---

[13] Giannoni and Woodford (2017) show that if an optimal equilibrium exists, then this target criterion will implement it as the unique bounded equilibrium. We do not seek to prove that an optimal equilibrium necessarily exists for all parameter values. But given results in Eusepi, Giannoni, and Preston (2018) and Eusepi, Giannoni, and Preston (forthcoming), we conjecture that the parameter space for which an optimal equilibrium in this model does not exist would be very small (if it exists at all). Non-existence of the optimal equilibrium tends to occur when the lagged effects of policy are sufficiently stronger than their immediate effects, and the central bank's discount factor is sufficiently far below 1 (see Appendix B). This is unlikely to occur in this model for plausible values of the discount factor. Non-existence has also not occurred in any numerical exercise, which have covered a broad range of parameterisations.

revised. This leads to a growing inflation impulse over time because of rising ω_t^π. The optimal policy response to this growing impulse is a higher interest rates in the medium-run to constrain inflation and output.

The optimal policy, therefore, for a wide-range in λ follows the same prescription: muted on impact and more aggressive over the medium run. This is true even though the rationale for why the central bank should enact such a policy changes as λ is varied.

Optimal policy only depart from this pattern when λ is close to zero. Here, the WAIT term is nearly absent and preemption is the dominant feature of policy. Stronger policy is required in these cases initially to blunt the propagation of the shock in beliefs. This represents the flexibility required by optimal policy in FAIT. If a policymaker expects expectations to drift, i.e., become unanchored, in response to a shock, then more aggressive policy is required now.

For moderate values of λ, however, the preemption motive is secondary to WAIT. Figure 2 illustrates how preemption affects policy when $\lambda = 0.5$ as we vary g making expectations more or less responsive to the shock. The figure shows WAIT ($\gamma = \lambda$, $\phi = 1$, and $\nu_\omega = \nu_u = 0$) compared to optimal policy. The difference between WAIT and optimal policy is preemption. The WAIT policy delivers lower interest rates and real rates in each case. Preemption on the other hand increases the current response as g increases to counter the greater influence of the shock in causing drifts in learned expectations.

There are of course other scenarios where even with moderate values of λ that preemption becomes more central. For example, when shocks are very persistent or the central bank expects to be constrained in the future such as when the ZLB binds, which we explore in Section 4.

3.2 The optimal average inflation target

The continued importance of WAIT relative to preemption shown in Figure 2 when λ is small is surprising. It occurs because the separation between WAIT and the preemption term masks a more fundamental congruence between the two objectives. Because adaptive learning beliefs are weighted averages of past inflation, the preemption term, when expectations are solved out, also correspond to weighted average of inflation but with different weights. To see this note that the expectations of the adaptive learning beliefs evolve as a weighted average of current and past inflation,

$$\omega_{t-1}^\pi = \frac{\rho g}{1-\rho(1-g)L}\pi_{t-1}.$$

Preemption, therefore, is an optimal response to a weighted average of past inflation. Therefore, it is possible to rewrite optimal policy as a WAIT policy with different weights but that shares a similar geometric decay over time.

Figure 1: IT vs Optimal Policy

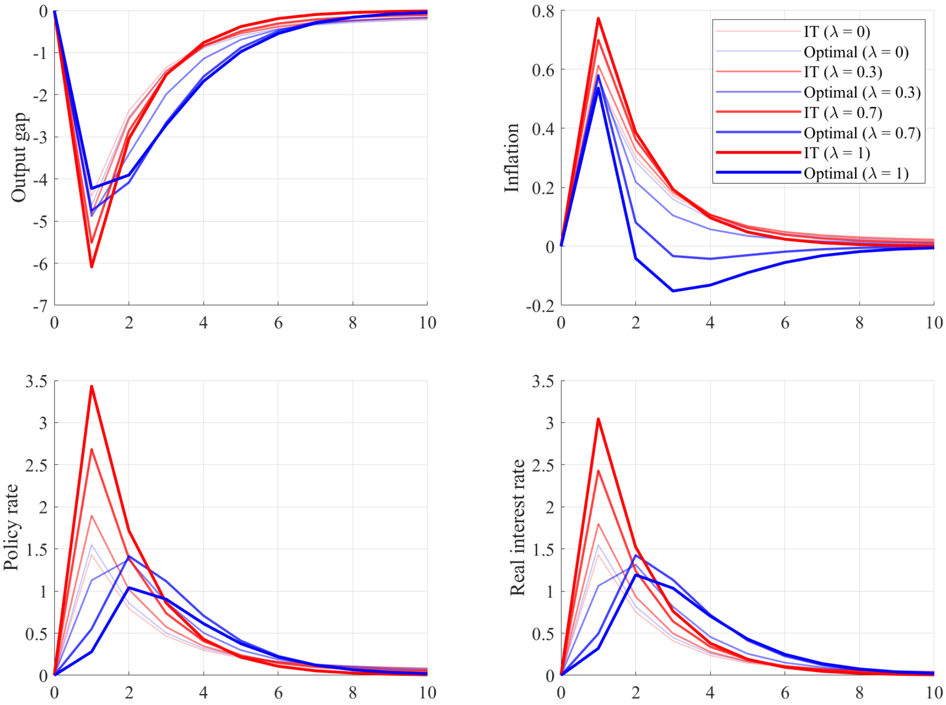

Notes: Response to a 1% cost-push shock, which occurs in period 1. Parameter values are $\rho = 0.95$, $g = 0.08$, $\beta = 0.99$, $\sigma = 1$, $\kappa = 0.1$, $\alpha = \kappa/7.87$

Figure 2: WAIT vs Optimal Policy

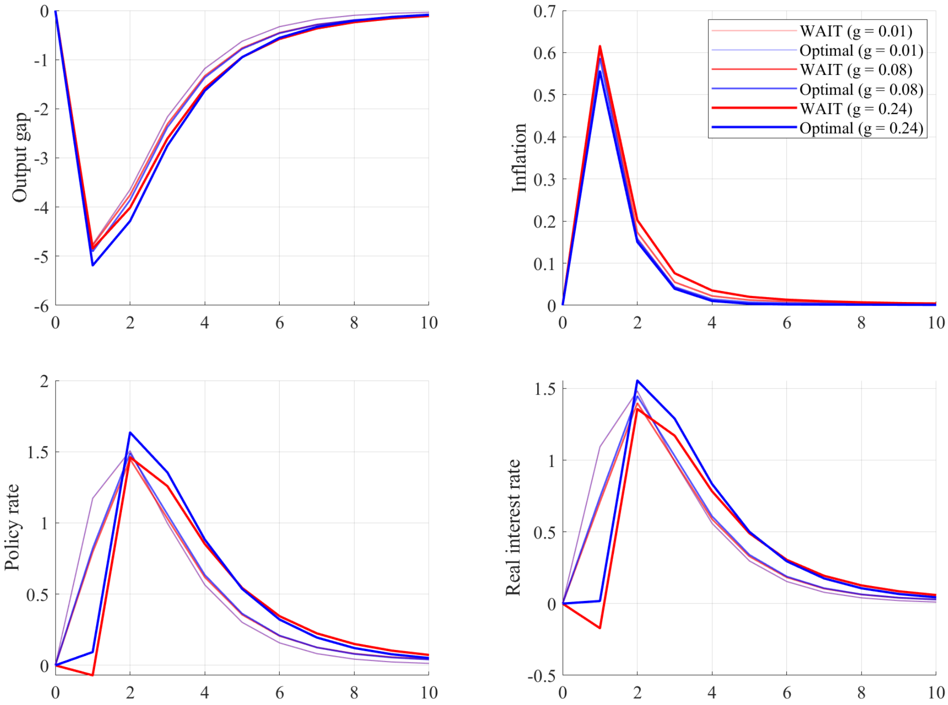

Notes: Response to a 1% cost-push shock, which occurs in period 1. Parameter values are $\lambda = 0.5$, $\rho = 0.95 =$, $\sigma = 1$, $\kappa = 0.1$, $\alpha = \kappa/7.87$

To illustrate, substitute the weighted average expression above for adaptive learning beliefs into the rewritten optimal target criterion (20) and rearrange:

$$x_t = -\frac{\kappa\phi}{\alpha}\left(\frac{\psi}{1-\gamma L} + \frac{1-\psi}{1-\rho(1-g)L}\right)\pi_t - \frac{\kappa\phi}{\alpha}\nu_u\frac{1}{1-\gamma L}u_t \tag{21}$$

where $\psi \equiv 1 - \frac{\rho g}{\rho(1-g)-\gamma}\nu_\omega$. The average inflation measure that a policymaker should target is a weighted average of two exponentially-weighted averages of current and past inflation.

Figure 3 illustrates the similarities in the weights by plotting the term in front of π_t from (21) for different values of λ and g. On the left side, we plot the optimal weights for different values of λ when $g = 0$ and $\omega_t^\pi = 0$. This calibration reflects models with myopia like Gabaix (2020) and it has no preemption component to policy. The high weight placed on current inflation reflects α, the welfare weight on the output gap, which here is set to the welfare theoretic value of κ/θ. Higher values of α scale these weights profiles down while lower values raise them. When $\lambda = 1$, we have the FIRE optimal policy, price-level targeting. Equal weight is given to all past inflation outcomes. When $\lambda = 0$, there is no value to commitment and we get the RE discretion outcome.

On the right side of Figure 3, we add a positive gain ($g > 0$) activating the preemption term. The optimal target criterion again responds to a weighted average of inflation. Preemption causes the weights to be higher depending on the specific level of g and λ. However, the general profile the weights remains the same with a geometric decline over time. Optimal policy tracks a qualitatively similar weighted average of inflation when most agents have FIRE expectations and when most agents do not. In Section 5, we consider whether there is a single calibration for this weighted average that performs well regardless of how expectations are formed and show that indeed one exists.

In addition, it is important to note that the profile of optimal weights are very different to the arithmetic averages that many researcher have studies when modeling average inflation targeting such as in Eo and Lie (2020), Jia and Wu, 2023, or Honkapohja and McClung (2024). Arithmetic averages imply uniform weights over a fixed number of lags and zero thereafter. We show in Section 5. that arithmetic average do a poor job approximating the optimal weights, while simple geometric averages approximate optimal weights well, even when the geometric weights are significantly misspecified relative to the optimal weights.

4 OPTIMAL POLICY: THE CONSTRAINED CASE

We now turn to the case where the policymaker faces a constraint that prevents it from implementing a desired level of the output gap period-by-period. We consider two ways in which this might occur: 1) imperfect information, and 2) the zero lower bound on

Figure 3: Optimal Weights

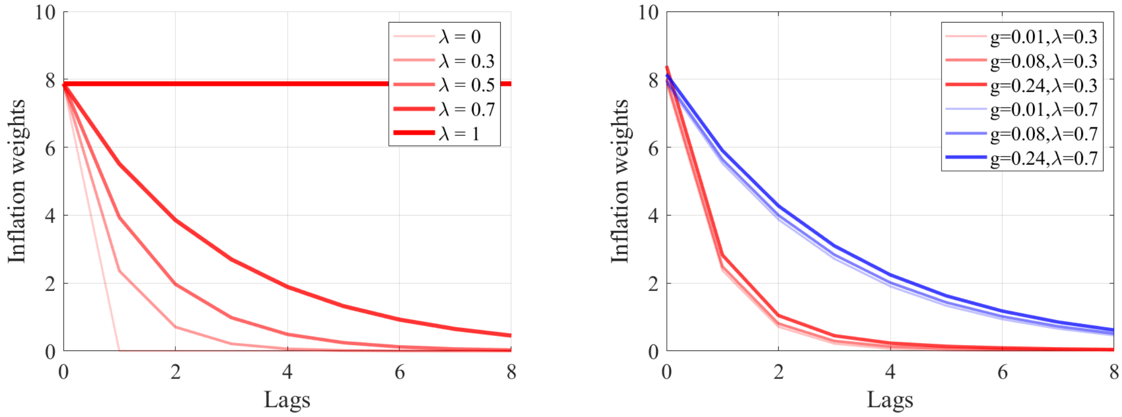

Notes: Plots coefficients in the power series in the first term of (20), which represent the weight placed on current and lags of inflation when implementing optimal policy. Parameter values are $\rho = 0.95$, $\beta = 0.99$, $\sigma = 0.5$, $\kappa = 0.1$, $\alpha = \kappa/7.87$

nominal interest rates. Again, Appendix A provides detailed derivations, while we focus here on the key equations and intuition.

4.1 IMPERFECT CENTRAL BANK INFORMATION

Let I_t^{CB} denote the central bank's information set in period t and let I_t denote the full information set. Let $z_{T|t}$ denote the best estimate of any variable z_T conditional on I_t^{CB}, while $\mathbb{E}_t z_T$ continues to denote the expectation conditional on the full information set I_t. We now assume

$$I_{t-1} \subset I_t^{CB} \subset I_t.$$

In other words, the central bank does not possess all information in the current period necessary to implement optimal policy. The nominal interest rate i_t must be set each period based on incomplete information I_t^{CB}, i.e.

$$i_t = i_{t|t}. \tag{22}$$

For our purposes, it does not matter what information is missing. The key is that any missing information relevant to policy represents a constraint on current policy choices.

The optimal policy problem is to choose the sequence $\{\pi_t, x_t, \omega_t^\pi, \omega_t^x, i_t\}_{t=0}^\infty$ that maximizes the welfare function (10) subject to the constraints (11) to (14) imposed by private-sector equilibrium behaviour and the informational constraint (22).[14]

---

[14] The private sector continues to make decisions based on the full information set, I_t. This means that when an unexpected shocks occurs, their expectations of future policy will adjust before the central bank moves the actual policy rate. This is what underlies the 'automatic stabilizer' mechanism discussed below.

The first-order condition of the policymaker's problem with imperfect information with respect to i_t is

$$\begin{aligned} \mu_t^i &= \frac{1}{\sigma}\sum_{j=0}^{t}\mu_{t-j}^x && (23) \\ \mu_{t|t}^i &= 0 && (24) \end{aligned}$$

where μ_t^i is the Lagrange multiplier on the information constraint (22). Condition (23) shows that the benefit of lowering the nominal interest rate is given by the benefit of raising the output gap in the current period, and in previous periods through the fact that agents anticipate future policy responses. Condition (24) shows that the central bank should try to set the nominal interest rate so that there is no benefit to increasing or decreasing it, given the information on hand. The main difference from the unconstrained case is that demand shocks are no longer irrelevant.

Result 3 (Demand Relevance)**.** *The central bank cannot perfectly offset demand shocks when*

$$I_{t-1} \subset I_t^{CB} \subset I_t.$$

In the unconstrained case, when $I_t^{CB} = I_t$, then $\mu_t^i = \mu_{t|t}^i = 0$ and so $\mu_t^x = 0$. A shock to the IS curve is irrelevant. When $I_t^{CB} \subset I_t$, the best the central bank can do is $\mu_{t|t}^i = 0$, and so $\mu_t^i \neq 0$ and $\mu_t^x \neq 0$ in general, which makes the IS curve a constraint on policy.

To see how this affects optimal policy, we take the expectation of (23) conditional on I_t^{CB} and using $\mu_{t-j|t}^x = \mu_{t-j}^x$ for all $j \geq 1$ gives

$$\mu_{t|t}^x = -\sum_{j=1}^{t}\mu_{t-j}^x, \tag{25}$$

which implies that the central bank commits to an i_t that balances the estimated costs of an off-target output gap in period t, $\mu_{t|t}^x$, against the reduced losses in the periods before t due to expectations of i_t, $-\sum_{j=1}^{t}\mu_{t-j}^x$. Equation (25), therefore, describes a new intertemporal trade-off, where the central bank cannot set the desired interest rate it would like, ex post, while considering ex ante how its corrections for any errors it makes interact with private sector expectations.

Taking the first-order conditions with respect to π_t and x_t, eliminating the Lagrange multiplier on the Phillips curve, and substituting in $\mu_t^x = \sigma\left(\mu_t^i - \mu_{t-1}^i\right)$ from (23) gives

$$\frac{\alpha}{\kappa}x_t + \frac{\sigma}{\kappa}\mu_t^i = -\frac{1}{1-\lambda L}\pi_t + \rho g\frac{1}{1-\lambda L}\mu_t^{\omega\pi} + \left(\frac{\lambda}{\beta}\frac{\kappa}{\sigma}\frac{1}{1-\lambda L} + \frac{\lambda}{\beta}(1-\beta) + 1\right)\frac{\sigma}{\kappa}\mu_{t-1}^i + \frac{\rho g}{\kappa}\mu_t^{\omega x}, \tag{26}$$

which generalizes equation (15).

The information constraint implies three additional terms. The new term in μ_t^i on the left-hand side is zero in expectation and represents the extent to which a shock causes the central bank to 'miss' its intended target for the output gap. The first new term on the right-hand side - in μ_{t-1}^i - is an additional commitment requirement of optimal policy. It indicates how the central bank should respond to past misses caused by unexpected shocks, i.e., additional make-up policy. The commitment to respond to misses in the future shifts current expectations, thereby dampening the costs of the shock. In this way, commitment under imperfect information acts as a kind of 'automatic stabilizer'. Finally, the Lagrange multipliers on learners' inflation and output gap expectations are now

$$\begin{aligned}\mu_t^{\omega\pi} &= -(1-\lambda)\rho\Omega\frac{1-\lambda L}{1-\zeta_1 L}\mathbb{E}_t\frac{1}{1-\zeta_2^{-1}L^{-1}}\frac{1}{1-\lambda L}\pi_{t+1} \\ &\quad + (1-\lambda)\rho\Omega\left(\frac{\lambda}{\beta}\frac{1}{1-\zeta_2^{-1}\lambda}\frac{1}{1-\zeta_1 L} - \frac{1}{\beta(1-\beta\rho)}\frac{1-\lambda L}{1-\zeta_1 L}\right)\mu_t^i \qquad (27) \\ \mu_t^{\omega x} &= -(1-\lambda)\frac{(1-\beta)\rho\beta}{1-\beta\rho}\sigma\mu_t^i. \qquad (28)\end{aligned}$$

The new components in these expressions, i.e., the terms in μ_t^i, dampen the extent of optimal make-up policy, by accounting for its effect on learners' future expectations.

Substituting the Lagrange multipliers (27) and (28) into (26), we can derive the optimal target criterion under information constraints.

Result 4 (Information constrained optimal target criterion)**.** *Recall from eq. (18) that the optimal target criterion under full information is* $\tilde{\pi}_t = 0$*, where* $\tilde{\pi}_t$ *is defined as*

$$\tilde{\pi}_t \equiv \frac{\alpha}{\kappa}x_t + \underbrace{\frac{1}{1-\lambda L}\pi_t}_{WAIT} + \underbrace{(1-\lambda)g\rho^2\Omega\frac{1}{1-\zeta_1 L}\mathbb{E}_t\frac{1}{1-\zeta_2^{-1}L^{-1}}\frac{1}{1-\lambda L}\pi_{t+1}}_{Preemption}.$$

Then the optimal target criterion under imperfect central bank information is

$$\tilde{\pi}_{t|t} = \pi_t^*, \quad \text{where} \quad \pi_t^* \equiv \frac{1}{1+(1-\lambda)g\rho^2\Theta}P(L)\Delta_{t-1} \quad \text{and} \quad \Delta_t \equiv \pi_t^* - \tilde{\pi}_t \qquad (29)$$

with $\Theta \equiv \frac{(1-\beta)\beta}{1-\beta\rho} + \Omega\frac{1}{\beta(1-\beta\rho)}\frac{\kappa}{\sigma} - \Omega\frac{\lambda}{\beta}\frac{1}{1-\zeta_2^{-1}\lambda}\frac{\kappa}{\sigma}$ *and*

$$\begin{aligned}P(L) \;&\equiv\; \underbrace{\frac{\lambda}{\beta}\frac{\kappa}{\sigma}\frac{1}{1-\lambda L} + \frac{\lambda}{\beta}(1-\beta) + 1}_{Make-up\ policy} + \\ &\underbrace{(1-\lambda)g\rho^2\Omega\frac{\kappa}{\sigma}\left(\frac{\lambda}{\beta}\frac{1}{1-\zeta_2^{-1}\lambda}\left(\frac{\lambda}{1-\lambda L}+\zeta_1\right) - \frac{\zeta_1}{\beta(1-\beta\rho)}\right)\frac{1}{1-\zeta_1 L}}_{Interaction\ of\ preemption\ and\ make-up}. \qquad (30)\end{aligned}$$

The optimal target criterion under imperfect information (29) generalizes the unconstrained target criterion (18). It differs from (18) in two ways: (i) the central bank does its best to implement it given its information set I_t^{CB} (i.e. it can only control $\tilde{\pi}_{t|t}$, not $\tilde{\pi}_t$), and (ii) its target for $\tilde{\pi}_{t|t}$ is not always zero, but adjusts based on past errors in implementing optimal policy - Δ_t - due to the information constraint.

The optimal target criterion though remains a flexible average inflation target. The difference here is in the definition of flexibility. The central bank still retains the same WAIT and preemption incentives but now must commit to make up for any errors in implementing policy that arise from the information constraint.

Make-up policy works through three channels represented by the three added terms in the first line of (30). The first term reflects the effect on inflation expectations of a commitment to make-up for an error. Rational forward-looking agents respond today to the expected make-up policy they anticipate will occur, which dampens the impact of the shock. The second term represents the effect of this commitment on output gap expectations, which affect the current output gap through the permanent income hypothesis with $(1-\beta)$ representing the marginal propensity to consume from a change in current period income. The final term represents the effect of commitment on nominal interest-rate expectations. Committing to offset misses shifts the expected nominal yield curve and further ameliorates the effect of any unexpected shock. The first two term are both multiplied by λ, which signifies their dependence on forward-looking expectations. By contrast, the third term is multiplied by one since all agents incorporate the optimal path of interest rates into their beliefs.

The second line of (30) reflects the change in the cost and benefit to make-up policy when there are learning expectations. Policy actions, like shocks, are instantiated into learning agents' beliefs when output and inflation respond and must be accounted for when implementing policy. To aid intuition, consider the case when prices are fixed, $\kappa = 0$. The optimal target criterion reduces to

$$x_{t|t} = x_t^*$$

where x_t^* is

$$x_t^* \equiv \frac{1}{1+(1-\lambda)g\rho^2\frac{(1-\beta)\beta}{1-\beta\rho}}\left(\frac{\lambda}{\beta}(1-\beta)+1\right)\Delta_{t-1}$$

with $\Delta_t \equiv x_t^* - x_t$. The central bank should set its best estimate of the output gap $x_{t|t}$ equal to the moving target x_t^*, which is determined by the previous period's miss, Δ_{t-1}. The extent to which the central bank should commit to offset past misses depends on the strength of (i) the forward-looking income expectations channel, given by $\frac{\lambda}{\beta}(1-\beta)$, (ii) the nominal yield curve channel, given by the 1 in the parentheses, and (iii) the effect of learners' drifting income expectations, given by $(1-\lambda)g\rho^2\frac{(1-\beta)\beta}{1-\beta\rho}$. This last term lowers the response to a given miss as g rises or λ falls. Make-up policy itself is a source of variation in this economy when agents are learning, which generates further drifts in beliefs that must be preempted.[15]

4.1.1 Mechanics of optimal policy: imperfect information

Figure 4 shows how optimal policy compares with IT in response to an unexpected cost push shock under the same calibration explored in Figure 1. In contrast to the unconstrained case, the optimal policy is now more aggressive than IT policy. IT policy 'lets bygones be bygones' and does not respond to the failure to recognise the shock on impact, whereas optimal policy seeks to 'make up' for this miss. The anticipation of the aggressive policy acts as an automatic stabilizer when the shock occurs. Forward-looking agents understand the aggressive response is coming and that inflation will be lower in the future, which lowers inflation today. When λ is high, optimal policy is only a little different from the unconstrained case. This is because even under perfect information, optimal policy required only a small contemporaneous response, with a larger and persistent subsequent response. The information friction is therefore not too costly and only a small increase in make-up policy is required. When λ is low, the information friction makes more of a difference. With perfect information, the required contemporaneous response was large and the role of make-up policy was smaller. Imperfect information prevents the contemporaneous response, but introduces a substantial role for make-up policy because policymakers now want to influence interest-rate expectations, which are forward-looking even when $\lambda = 0$. The optimal policy prescription therefore ends up being similar across λ in the imperfect information case.

Figure 5 shows the dynamics of optimal policy versus IT for an unexpected demand shock (i.e. an increase in the natural rate r_t^n). In the full information case, this shock is perfectly offset and has no effect on the output gap or inflation. Information frictions, however, changes the prediction. Demand shocks generate inflation. As in the cost push shock case, optimal policy is now more aggressive than required under IT, which causes inflation to undershoot following the shock. Once again, under imperfect information, optimal policy is similar for both high and the low λ cases.

4.1.2 Optimal average inflation target: imperfect information

Like the unconstrained case, a FAIT policy is robust because we can capture most of WAIT, preemption, and the additional make-up policy using only a weighted average of past outcomes. In contrast to the unconstrained case, though, optimal policy requires tracking a weighted average over both inflation and the output gap.

---

[15] Eusepi, Giannoni, and Preston (forthcoming) explore this feature of learning models in depth, focusing on the transmission of short-run interest rate to long-run interest rate expectations, which is absent in our model. In Appendix B, we discuss how our results might be affected if some agents form policy-rate expectations by learning, instead of rationally.

Figure 4: Imperfect information IT vs Optimal Policy: cost push shock

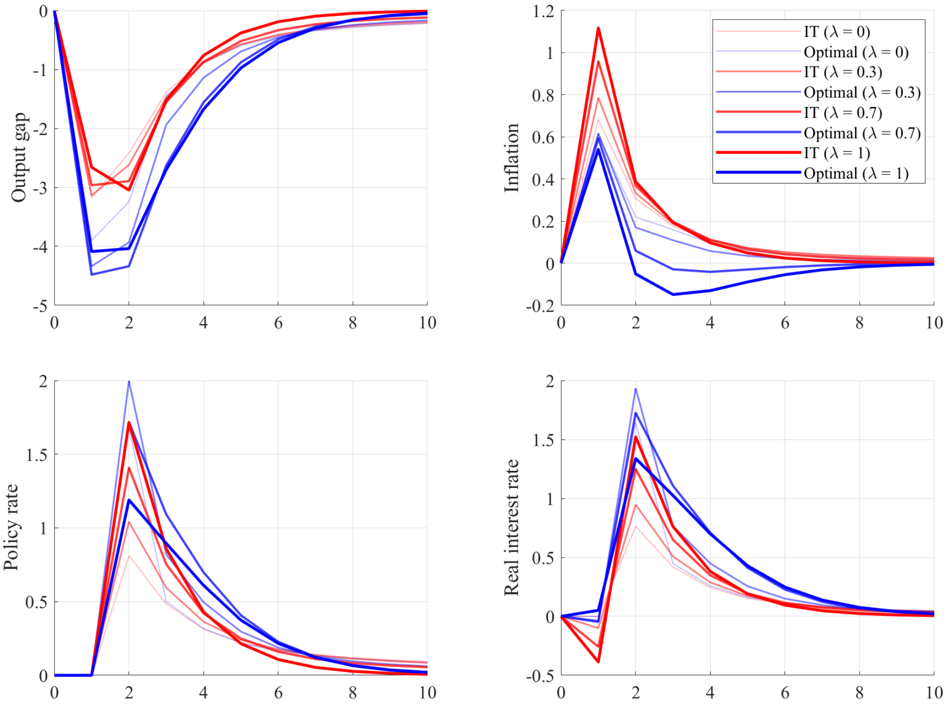

Notes: Response to a 1% cost-push shock, which occurs in period 1. Parameter values are $\rho = 0.95$, $g = 0.08$, $\beta = 0.99$, $\sigma = 1$, $\kappa = 0.1$, $\alpha = \kappa/7.87$.

Figure 5: Imperfect information IT vs Optimal Policy: Natural rate shock

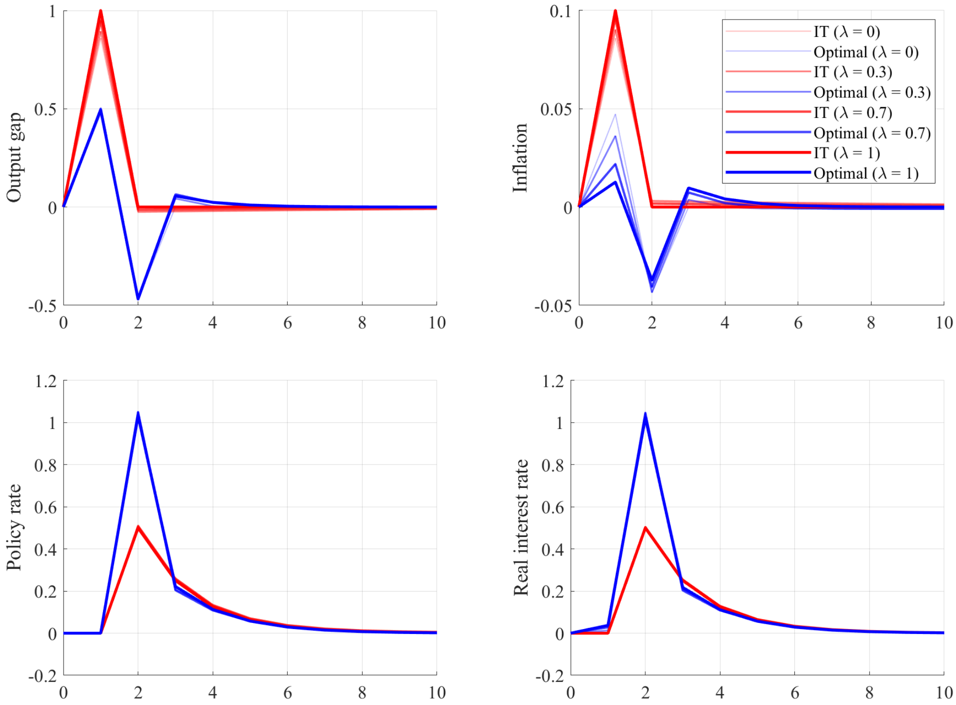

Notes: Response to a 1% natural interest rate shock, which occurs in period 1. Parameter values are $\rho = 0.95$, $g = 0.08$, $\beta = 0.99$, $\sigma = 1$, $\kappa = 0.1$, $\alpha = \kappa/7.87$.

Figure 6: Optimal Weights: Imperfect Information

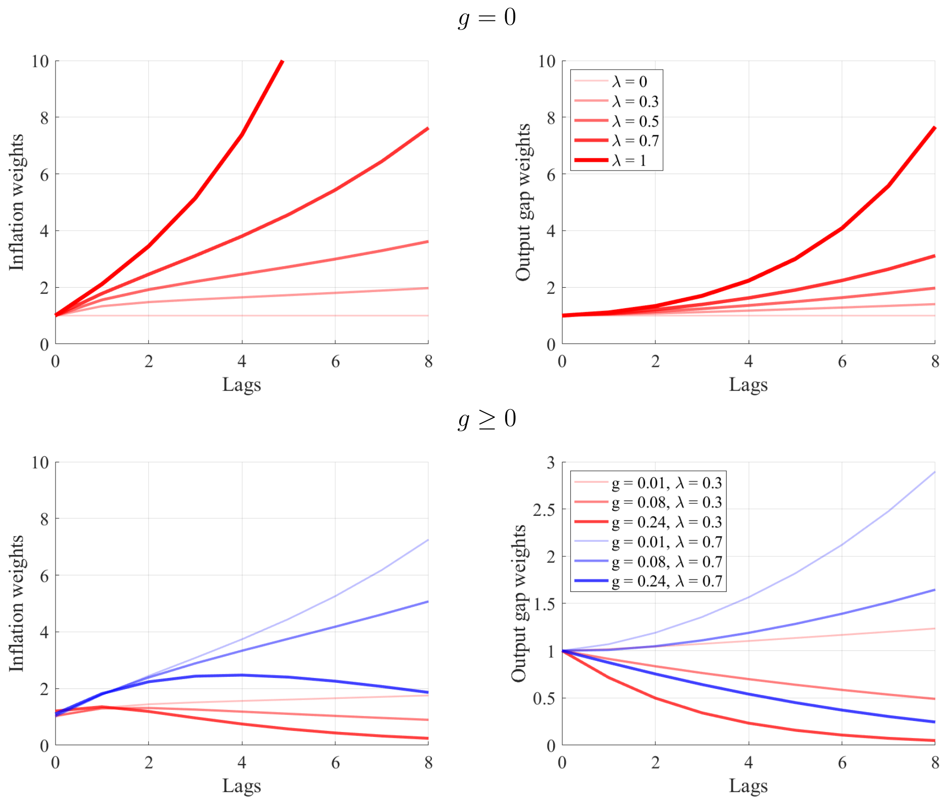

Notes: Optimal weights on past inflation and output gap under imperfect information found in $Q_\pi(L)$ and $Q_x(L)$ in (32). Parameter values are $\beta = 0.99$, $\sigma = 1$, $\kappa = 0.1$, $\alpha = \kappa/7.87$.

To illustrate, we substitute $\Delta_{t-1} \equiv \pi^*_{t-1} - \tilde{\pi}_{t-1}$ into (30), rearrange, and then substitute the result into (29) to arrive at

$$\mathbb{E}_{t|t} \frac{1}{1 - \frac{1}{1+(1-\lambda)g\rho^2\Theta} P(L)L} \tilde{\pi}_t = 0 \tag{31}$$

where $\mathbb{E}_{t|t}$ denotes the expectation conditional on I_t^{CB}. This formulation makes clear that the 'moving target' described in (29) and (30) is equivalent to trying to set some weighted average of current and past $\tilde{\pi}_t$ equal to zero. Since the central bank cannot implement the full information optimal target criterion $\tilde{\pi}_t = 0$ in every period, it instead tries to implement it on average over time.

As in the full information case, we can combine the optimal policy conditions with the equilibrium conditions to evaluate the expectation terms in (31) (which are inside $\tilde{\pi}_t$). This gives us a description of optimal policy just in terms of past inflation and output gaps, akin to Equation (21). Assuming again that the cost-push shock is a first-order autoregressive process,[16] Appendix A shows that, following an analogous approach to the derivation of (21), we can write optimal policy under imperfect information in equilibrium as

$$\mathbb{E}_{t|t} \left(Q_x(L)x_t + Q_\pi(L)\pi_t + Q_u u_t \right) = 0 \tag{32}$$

where $Q_x(L)$, $Q_\pi(L)$ and $Q_u(L)$ are power series in L and are defined in Appendix A. The coefficients in the power series Q_x and Q_π give the precise weight that the central bank places on current and past values of the output gap and inflation when implementing optimal policy.

Figure 6 plots these weights for different values of λ and g. The top row of figures mirrors the left panel of Figure 3. It shows the case where there is no updating of learning beliefs ($g = 0$), which reduces the model to one featuring myopia as in Gabaix (2020). Optimal policy requires increasing weights in this case. This reflects the requirement of additional make-up policy. The bottom row of figure shows the $g \geq 0$ case. Additional make-up policy is required but it is moderated by the need to adjust for drifts in expectations. Importantly, when g is high, the optimal weights eventually decline in a similar way to the unconstrained case. It is this declining feature of the weights that makes FAIT a robust framework across a wide range of scenarios. We show in Section 5 that tracking weighted averages of inflation and output with declining weights captures most of the benefits of optimal policy.

---

[16] Only the unexpected component of the IS curve shock is relevant to optimal policy. The central bank can fully offset any expected component, just as it does in the full information case.

4.2 Optimal policy with the zero lower bound

Optimal policy with the zero lower bound is very similar to optimal policy under imperfect information. Both constraints mean that the central bank sometimes cannot set the output gap it wants, in which case it cannot implement the full information optimal target criterion. As with imperfect information, the central bank should respond to this ZLB constraint by using commitment to influence the forward-looking expectations that enter the IS equation - i.e. household's expectations over future output, inflation, and nominal interest rates. The key difference from the imperfect information case is that losses caused by a binding ZLB are not zero in expectation in the presence of mean zero shocks drawn from a symmetric distribution. Therefore, the preemptive channel of policy has a more significant role. The central bank should seek to preemptively raise learners' output and inflation expectations if it expects the ZLB to bind in the future. This is the use it or lose it character of optimal policy when confronted with the ZLB. It is optimal to move expectations preemptively before the constraint binds.

To derive optimal policy with the ZLB, we replace the information constraint (22) with the lower bound constraint

$$i_t > \bar{i}. \tag{33}$$

The optimal policy problem is then to choose the sequence $\{\pi_t, x_t, \omega_t^\pi, \omega_t^x, i_t\}_{t=0}^\infty$ that maximises welfare (10) subject to the constraints (11) to (14) imposed by private-sector equilibrium behaviour and the lower bound constraint (33).

The first-order conditions are the same as in the imperfect information case, except that (24) is replaced with the complementary slackness condition

$$\mu_t^i \geq 0, \quad \mu_t^i \left(i_t - \bar{i}\right) = 0. \tag{34}$$

Just as μ_t^i previously represented the extent to which imperfect information constrained the central bank, it now represents the losses caused by the ZLB. It is positive when the ZLB is binding and zero otherwise.

Condition (26) is the same as before, but the expressions for $\mu_t^{\omega\pi}$ and $\mu_t^{\omega x}$ in (27) and (28) take a more general form, because $\mathbb{E}_t \mu_T^i$ is no longer necessarily equal to zero for all $T \geq t+1$:

$$\begin{aligned}
\mu_t^{\omega\pi} &= -(1-\lambda)\rho\Omega \frac{1-\lambda L}{1-\zeta_1 L} \mathbb{E}_t \frac{1}{1-\zeta_2^{-1}L^{-1}} \frac{1}{1-\lambda L} \pi_{t+1} \\
&\quad + (1-\lambda)\rho\Omega \frac{1-\lambda L}{1-\zeta_1 L} \mathbb{E}_t \frac{1}{1-\zeta_2^{-1}L^{-1}} \left(\frac{\lambda}{\beta} \frac{L}{1-\lambda L} + \frac{1}{\beta(1-\beta\rho)}(1-L) \right) \mu_{t+1}^i \qquad (35) \\
\mu_t^{\omega x} &= (1-\lambda) \frac{(1-\beta)\rho\beta}{1-\beta\rho} \mathbb{E}_t \frac{1-L}{1-\beta\rho(1-g)L^{-1}} \sigma \mu_{t+1}^i. \qquad (36)
\end{aligned}$$

If $\mathbb{E}_t\mu^i_T = 0$ for all $T \geq t+1$, then (35) and (36) collapse to (27) and (28). The second line of (35) indicates how the central bank can use learners' inflation expectations to lower the real interest rate when the ZLB binds. The first term in the parentheses represents the effect of raising learners' inflation expectations on the inflation expectations of the rational forward-looking agents. This term, therefore, has both preemptive and history-dependent aspects. The second term in the parentheses represents the effect of shifting learners' inflation expectations on future real interest rates. Similarly, (36) represents the effect of shifting learners' output expectations for future output gaps.

Conditions (26), (34), (35) and (36) define optimal policy with the ZLB constraint.

Result 5 (Optimal target criterion with ZLB)**.** *Recall that the optimal target criterion under full information and without the ZLB is* $\tilde{\pi}_t = 0$*, where we defined*

$$\tilde{\pi}_t \equiv \frac{\alpha}{\kappa}x_t + \underbrace{\frac{1}{1-\lambda L}\pi_t}_{WAIT} + \underbrace{(1-\lambda)g\rho^2\Omega\frac{1}{1-\zeta_1 L}\mathbb{E}_t\frac{1}{1-\zeta_2^{-1}L^{-1}}\frac{1}{1-\lambda L}\pi_{t+1}}_{Preemption}$$

Then the optimal target criterion when the ZLB can bind is

$$\tilde{\pi}_t = \pi^*_t \tag{37}$$

$$\textit{where} \quad \pi^*_t \equiv \frac{1}{1+(1-\lambda)g\rho^2\Theta}\left(P(L)\Delta_{t-1} + \mathbb{E}_t P_1(L^{-1})\Delta_{t+1} + P_2(L)\mathbb{E}_{t-1}P_3(L^{-1})\Delta_t\right)$$

$$\textit{and} \quad \Delta \equiv \pi^*_t - \tilde{\pi}_t$$

where

$$\begin{aligned} P(L) \;\equiv\; & \underbrace{\frac{\lambda}{\beta}\frac{\kappa}{\sigma}\frac{1}{1-\lambda L} + \frac{\lambda}{\beta}(1-\beta) + 1}_{Make-up\ policy} + \\ & \underbrace{(1-\lambda)g\rho^2\Omega\frac{\kappa}{\sigma}\left(\frac{\lambda}{\beta}\frac{1}{1-\zeta_2^{-1}\lambda}\left(\frac{\lambda}{1-\lambda L}+\zeta_1\right) - \frac{\zeta_1}{\beta(1-\beta\rho)}\right)\frac{1}{1-\zeta_1 L}}_{Interaction\ of\ preemption\ and\ make-up} \end{aligned}$$

is the same as (30) and

$$\begin{aligned} P_1(L^{-1}) &= (1-\lambda)g\rho^2\Omega\frac{\kappa}{\sigma}\left(\frac{\lambda}{\beta}\frac{\zeta_2^{-1}}{1-\zeta_2^{-1}\lambda}+\frac{1-\zeta_2^{-1}}{\beta(1-\beta\rho)}\right)\mathbb{E}_t\frac{1}{1-\zeta_2^{-1}L^{-1}} \\ &\quad \underbrace{+(1-\lambda)g\rho^2\frac{(1-\beta)\beta}{1-\beta\rho}\left(1-\beta\rho(1-g)\right)\mathbb{E}_t\frac{1}{1-\beta\rho(1-g)L^{-1}}}_{Insurance\ for\ expected\ bind} \end{aligned} \tag{38}$$

$$P_2(L) = \zeta_1\frac{1}{1-\zeta_1 L} \tag{39}$$

$$P_3(L^{-1}) = \underbrace{(1-\lambda)g\rho^2\Omega\frac{\kappa}{\sigma}\left(\frac{\lambda}{\beta}\frac{\zeta_2^{-1}}{1-\zeta_2^{-1}\lambda}+\frac{1-\zeta_2^{-1}}{\beta(1-\beta\rho)}\right)\mathbb{E}_t\frac{1}{1-\zeta_2^{-1}L^{-1}}}_{Insurance\ for\ expected\ bind}. \tag{40}$$

Optimal policy commits to responding to pass misses (see $P(L)$), exactly as in the information constraint case, and it adjusts the target π_t^* in response to expected future target misses ($P_1(L^{-1})$) and past expectations of target misses ($P_2(L)$ and $P_3(L^{-1})$). The net effect is even more make-up policy is promised.

The terms $P_1(L^{-1})$ and $P_3(L^{-1})$ capture the insurance nature of optimal policy when agents are learning. Eusepi, Gibbs, and Preston (2024) coin the term the *insurance principle*, which says that when the duration of the zero lower bound is uncertain, and agents are learning, that larger front-loaded forward guidance promises are optimal. These policies are too stimulatory when the shock is short-lived but highly effective when the shock is long-lived because of the positive contribution they make to agents' drifting beliefs. The expectations operator present in these two terms captures the idea that policy must be set according to the expectation of how long the ZLB binds, which may not match the ex post realized duration of the shock. The more learning that is present (lower λ), and the faster learning happens (higher g), the more the central bank should respond now to these expectations of future binds of the ZLB.

To more clearly see how the terms in Result 5 relate to make-up and insurance policies, consider again the case where prices are fixed, $\kappa = 0$. Optimal policy implies $x_t = x_t^*$ where

$$\begin{aligned} x_t^* \;\equiv\; & \frac{1}{1+(1-\lambda)g\rho^2\Theta}\underbrace{\left(\frac{\lambda}{\beta}(1-\beta)+1\right)\Delta_{t-1}}_{Make-up\ policy} \\ & +\frac{(1-\lambda)g\rho^2\Theta}{1+(1-\lambda)g\rho^2\Theta}\left(1-\beta\rho(1-g)\right)\underbrace{\mathbb{E}_t\frac{1}{1-\beta\rho(1-g)L^{-1}}\Delta_{t+1}}_{Insurance\ principle} \end{aligned}$$

with $\Delta_t \equiv x_t^* - x_t$. The first line is identical to the target criterion in the imperfect information case with fixed prices. The second line represents the additional requirement of optimal policy to lift agents' expectations preemptively when the ZLB is expected to

Figure 7: Optimal policy at the ZLB

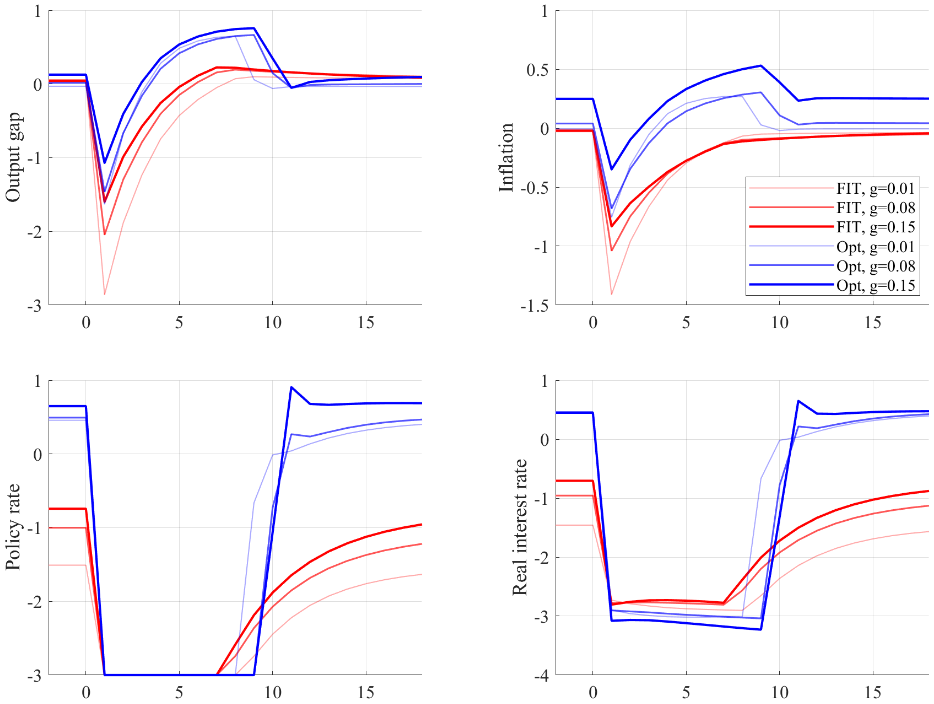

Notes Response to a 1.5% natural rate shock, which occurs in period 1 and has persistence 0.75. Parameter values are $\lambda = 0.25$, $\rho = 0.95$, $\sigma = 1$, $\kappa = 0.1$, $\alpha = \kappa/7.87$.

bind, or continue to bind. Policy should move in anticipation of a binding of the ZLB to raise expectations to preempt the losses that are expected to accrue while the ZLB binds.

Figure 7 compares optimal policy to IT in the presence of the ZLB.[17] First, note that even before the shock hits, optimal policy maintains a small positive output gap and inflation rate. This is the preemptive component of optimal policy: given that the ZLB could bind with some probability in any period, it is optimal to ensure that learners' inflation and output gap expectations are somewhat above target. For larger g, shocks have greater influence over learners' expectations, and so it is optimal to 'run the economy hotter'. Optimal policy also holds the policy rate at the ZLB for longer than IT, generating a positive output gap and inflation overshoot. This is the make-up component. As a result of these two components, the downturn when the shock hits is much smaller, even though there is some welfare loss before and after the shock. Interestingly, even though optimal policy involves a more expansionary stochastic steady state, the policy rate is higher, because inflation expectations are higher.

Figure 8 shows how the preemptively higher stochastic steady-state inflation target varies with λ, g, and the mean neutral rate, r^n.[18] As expected, the inflation target is

---

[17] To solve the model with the ZLB, we used a policy function iteration method, drawing on code in Sijmen Duineveld's PROMES toolbox (Duineveld, 2021).

[18] Here, when we say 'steady-state inflation target', we are referring to the stochastic steady state, meaning the value that inflation converges to in the absence of shocks, but where agents still expect that shocks could occur in the future.

Figure 8: Stochastic steady state for inflation with the ZLB

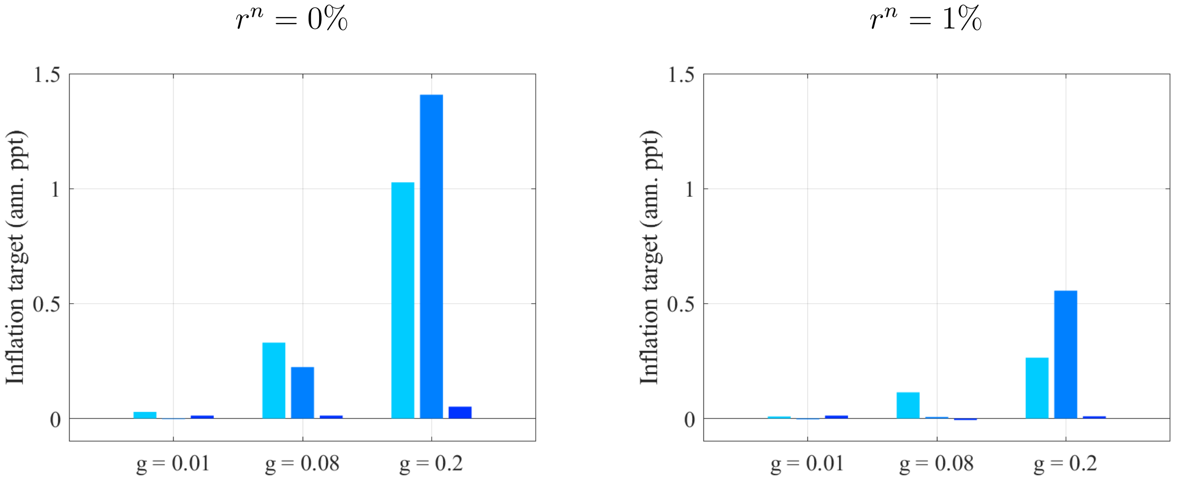

Notes We assume that the deterministic inflation steady state is 2%, so e.g. $r^n = 1\%$ implies that the ZLB is 3 percentage points below steady state. The standard deviation of r_t^n is 4 percentage points. Parameter values are $\rho = 0.95$, $\sigma = 1$, $\kappa = 0.1$, $\alpha = \kappa/7.87$.

higher when g is higher, which requires more preemptive policy, and when r^n is lower, which increases the likelihood of hitting the ZLB. When g is high, the inflation target is highest for intermediate values of λ, reflecting the amplification that can occur between forward-looking and backward-looking expectations through optimal preemptive policy.

5 COMPARING TARGET CRITERIA

The optimal target criteria we have derived require complex considerations on the part of the central bank where policy must respond to competing objectives in precise ways. The target criteria also requires a knowledge of the underlying expectation formation process in order to implement. In this section, we ask how robust are the policy lessons we draw from the optimal criteria. In other words, how flexible can you be as a flexible average inflation targeter?

To answer this questions, we first look at how closely simple weighted average inflation targeting schemes can perform relative to the optimal target criteria, and how do these perform against known target criteria like inflation or price-level targeting. Second, we study how these simplified optimal target criteria perform when a central bank does not know how expectations are formed and must calibrate their target criterion to maximize the average welfare for a range of different belief formation processes.

We allow for cost-push shocks and imperfect central bank information in this analysis so the policy cannot perfectly offset all demand and productivity shocks. We consider this to be the most general and realistic case.[19] Appendix D provides results for the

---

[19] Alternatively, we could also allow for an occasionally-binding ZLB. However, as we show in Result 5, imperfect central bank information captures key features of the ZLB constraint, in that both prevent the central bank from perfectly controlling the output gap. The main difference is that target misses due

Table 2: Simple target criteria

Name	Form	Parameters
IT	$x_t = -\psi \pi_t$	ψ
PLT	$x_t = -\psi p_t$	ψ
k-yr AIT	$x_t = -\psi \sum_{j=0}^{4k-1} \pi_{t-j}$	ψ
WAIT	$x_t = -\psi \frac{1}{1-\gamma L} \pi_t$	ψ, γ
WAIT + WAXT	$\frac{1}{1-\gamma_x L} x_t = -\psi \frac{1}{1-\gamma L} \pi_t$	ψ, γ, γ_x
IT + pre-emptive	$x_t = -\psi \pi_t - \psi_f \mathbb{E}_t \frac{1}{1-\gamma_f L^{-1}} \pi_{t+1}$	ψ, ψ_f, γ_f
WAIT + pre-emptive	$x_t = -\psi \frac{1}{1-\gamma L} \pi_t - \psi_f \mathbb{E}_t \frac{1}{1-\gamma_f L^{-1}} \pi_{t+1}$	$\psi, \gamma, \psi_f, \gamma_f$
WAIT + WAXT + pre-emptive	$\frac{1}{1-\gamma_x L} x_t = -\psi \frac{1}{1-\gamma L} \pi_t - \psi_f \mathbb{E}_t \frac{1}{1-\gamma_f L^{-1}} \pi_{t+1}$	$\psi, \gamma, \gamma_x, \psi_f, \gamma_f$

Notes The parameters are optimized to minimize welfare loss when implementing each rule.

unconstrained case, where cost-push shocks are the only source of losses.[20]

5.1 Approximating optimal policy with simple target criteria.

Table 2 presents a list of simple target criteria that capture key features of the optimal target criteria in the unconstrained and constrained cases, as well as other target criteria commonly studied in the literature. The list includes inflation targeting, price-level targeting, and simple average inflation targeting rules (‘k-yr AIT’). Informed by the analytical results in the previous sections, we also include rules that incorporate some combination of an IT or weighted AIT component (‘WAIT’); a weighted average output-gap targeting component (‘WAXT’); and a preemptive component that is a forward-looking weighted average of inflation forecasts (‘preemptive’).

For each of the target criteria listed in Table 2, we find the optimal parameters that minimize welfare loss and compare it to the loss achieved under the fully optimal target criteria. For the calibration, we set $\sigma = 1$ and $\kappa = 0.1$ (equivalent to a Calvo stickiness probability of around 0.7). For the expectation formation parameters, we use a baseline of $\lambda = 0.5$, $g = 0.08$, and $\rho = 0.95$. Finally, we set the persistence of the shocks to $\rho_u = \rho_{rn} = 0.5$ with the standard deviation of the neutral rate shock at ten times the standard deviation of the cost-push shock (similar to Giannoni, 2014). For simplicity, we let $\beta \to 1$, so that the loss function is $\mathcal{L} = \text{Var}(\pi_t) + \alpha \text{Var}(x_t)$, with $\alpha = \kappa/7.87$.

Table 3 presents the results. The gains from optimal policy relative to IT are large, at 36.5% of the loss under IT. These gains can be almost entirely achieved by following a rule that responds to a weighted average of both past inflation and past output gaps (the ‘WAIT + WAXT’ rule). Significant gains come from making up for past output gaps,

---

to the ZLB are non-zero in expectation, so allowing the ZLB when optimising the simple target criteria outlined below would likely result in a stronger preemptive response to expected future outcomes.

[20] The main difference is that the optimal degree of make-up policy is lower when the central bank has perfect control of the output gap, as the analytical results in previous sections suggest. Compared to the imperfect information case, the optimal weight on lagged inflation and output gap outcomes is smaller, PLT underperforms IT, and IT in general is closer to optimal policy. That said, Appendix C show that allowing for an infinite-horizon Phillips curve somewhat increases the optimal degree of make-up policy in the unconstrained case, making it more like the imperfect information case.

Table 3: Performance of simple target criteria - imperfect information case

	Loss (% rel. to IT)	ψ	γ	γ_x	ψ_f	γ_f
5-yr AIT	24.6	27.4				
IT	0	20.9				
2-yr AIT	-4.5	27.9				
IT + preemptive	-14.8	0.00			1.60	1.75
1-yr AIT	-15.7	15.6				
PLT	-23.8	20.3	1			
WAIT	-25.1	17.4	0.75			
WAIT + preemptive	-25.1	17.4	0.75		0.00	*undef.*
WAIT + WAXT	-35.4	12.3	0.94	0.89		
WAIT + WAXT + preemptive	-35.4	12.3	0.94	0.89	0.00	*undef.*
Optimal	-36.5					

Notes Welfare losses are normalized relative to the welfare under IT. The parameters for each rule are optimized to minimize welfare. The optimal value of a parameter is undefined (labelled *undef.*) if the parameter has no effect on loss at the optimum.

because imperfect information prevents the central bank from perfectly controlling the output gap. Recall that lagged output gap terms enter the optimal target criteria under imperfect information, eq. (29) (as with the ZLB), unlike the unconstrained case, eq. (18). The decay parameters in the WAIT and WAXT are 0.94 and 0.89, so substantial weight is placed on inflation and output outcomes in the past. The PLT rule outperforms IT. Simple AIT rules with short windows perform worse than the weighted-average rules.

The result support the idea that the average in FAIT should be a weighted average. Importantly, though, the simple target criteria all require a strong response to current inflation, which is seen by the optimized values of ψ and γ. Simple AIT rules require large ψ values as well to compensate for the $1/k$ weight applied to current inflation. However, it is not enough to offset the wrong type of average.

5.2 ROBUST OPTIMAL POLICY UNDER PARAMETER UNCERTAINTY.

The previous results show that simpler targeting rules can approximate the optimal target criteria. However, the simple targeting rules require the policymaker to know λ, g, and ρ. Here we study what happens when the policymaker does not know these values and instead must choose a target criterion that is optimized to maximize welfare on average when λ, g, and ρ can take on any value in economically meaningful ranges, which we approximate with a grid. In particular, we consider the values: $\lambda \in \{0, 0.25, 0.5, 0.75, 1\}$, $g \in \{0, 0.01, 0.08, 0.25, 1\}$ and $\rho \in \{0.95, 0.99\}$, which covers the range of emperical estimates discussed in Section 2.[21]

Table 4 shows the average welfare and the optimized parameter values for each target criterion. In addition, we show the average loss that would be achieved if the optimal criterion was implemented at each point of the grid, which represent the minimum loss achievable. WAIT rules perform best, with decay parameters of around 0.75 to 0.87.

---

[21] We set the other parameters at the baseline calibration described in the previous section.

Table 4: Policy under parameter uncertainty - imperfect information case

	Average Loss (% rel. to IT)	ψ	γ	γ_x	ψ_f	γ_f
5-yr AIT	118.0	1.4				
2-yr AIT	19.8	9.2				
IT	0	20.9				
IT + preemptive	-1.3	14.9			13.0	1.52
1-yr AIT	-16.6	20.7				
PLT	-22.4	16.8	1			
WAIT	-24.0	15.1	0.75			
WAIT + preemptive	-24.0	15.1	0.75		0.00	*undef.*
WAIT + WAXT	-31.0	14.5	0.87	0.69		
WAIT + WAXT + preemptive	-31.1	13.3	0.87	0.66	6.8	0.00
Opt	-35.7					

Notes The parameters reports are set to maximize the average loss across all points in the grid. Average welfare losses are normalized relative to the welfare under IT. The optimal value of a parameter is undefined (labelled *undef.*) if the parameter has no effect on loss at the optimum.

Making up for past output gaps as well via a WAXT component substantially lowers the loss. The PLT rule significantly outperforms IT. The 5-yr and 2-yr AIT rules always perform poorly.

To show what lies behind the average losses in the table, Figure 9 plots the loss under a selection of target criteria from Table 4 for different values of the expectation formation parameters λ and g. The parameters in the policy rules are held fixed at the robustly optimal values set to minimize average loss and not optimized for each λ and g shown in the figure. The blacked dash line shows the optimal target criteria optimized at each point to show how close the simple rules with the average coefficients can come to the fully optimal policy.

The figures reflect the nesting of different expectations assumptions of our approach. The IT rule and PLT rules perform similarly when λ is low (when make-up commitments have little effect on inflation or output gap expectations) or g is high (when the sensitivity of learners' beliefs increases the cost of make-up commitments relative to their benefit).[22] When λ is high, make-up commitments are more powerful and the IT rule performs poorly.

In contrast, the WAIT rules generally perform well for all λ and g, and including a WAXT component consistently lowers the loss for all parameters. It is important to reiterate that this is the case even though we have fixed the coefficients of the simple WAIT and WAXT rules to their robustly optimal values in Table 4, so it is a fair comparison with the IT and PLT rules. Only in the optimal policy benchmark do we allow the policy rule parameters to change as we vary λ and g. The performance of the fixed-coefficient simple WAIT and WAXT rules is still remarkably close to the optimal benchmark, which illustrates the robustness of this form of FAIT.

---

[22] PLT still slightly outperforms IT even when $\lambda = 0$ because nominal interest-rate expectations still respond to make-up commitments. If the central bank has perfect control of the output gap, then IT outperforms PLT significantly when $\lambda = 0$ (see Appendix D).

Figure 9: Loss under different parameter values - imperfect information case

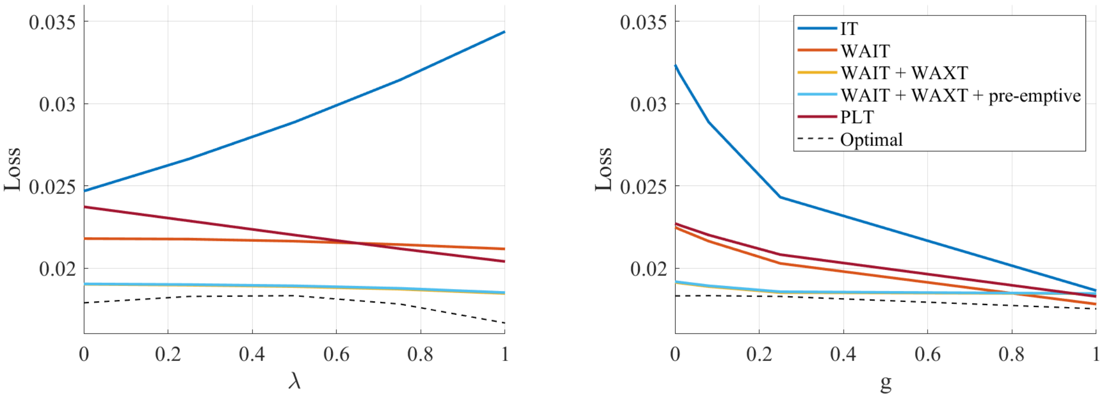

Notes: The parameters in the policy rules are fixed at the robustly optimal values presented in Table 4. For the other parameters (other than the one changing), we set $\lambda = 0.5$, $g = 0.08$, $\rho = 0.95$, $\kappa = 0.1$, $\sigma = 1$, $\rho_{rn} = \rho_u = 0.5$, $\sigma_{rn} = 1$ and $\sigma_u = 0.1$.

5.3 Discussion

The reason that simple FAIT target criteria are robust is that once λ and g are in the interior of their theoretical ranges, and expectations are both backward-looking and forward-looking, the features of optimal policy are the same. Optimal policy requires a weighted average inflation target and some flexible component that captures preemption and make-up. Preemption appears to be satisfied in most cases simply by placing high weight on current inflation and output realizations as found in the optimised ψ and γ coefficients in the WAIT + WAXT rules, while additional make-up policy is captured by including the WAXT term.

6 Conclusion

It is widely recognized that rational expectations is a strong assumption. Many researchers and policymakers have worried that policy recommendations that rely on rational expectations may not be robust. A large literature has shown that this concern is warranted in a variety of settings. We show that, in a framework encompassing a broad set of theories for expectation formation, Flexible Average Inflation Targeting (FAIT) is a robust policy strategy. This is the case whether the central bank is unconstrained, faces information constraints, or if the ZLB may bind.

Distinct from existing characterizations of FAIT, we make precise what *flexible* and *average* mean, and how they depend on expectation formation and constraints faced by the central bank. Policy should target a *weighted* average of current and past inflation, with weights that decline over time.

The optimal flexibility consists of preemptive actions and additional make-up policy that move away from the policy implied by the weighted average of inflation alone. When constraints are anticipated but not yet binding, or shocks are anticipated, optimal policy calls for preemptive deviations from the weighted-average target. If targets are missed today because the central bank is constrained (by information constraints or the ZLB), then policy should do more to make-up for these misses in the future.

A Appendix - Derivations

A.1 Unconstrained optimal policy

The unconstrained optimal policy problem is to maximise (10) subject to (11) - (14). The first-order conditions for this problem are

$$\pi_t + \mu_t^\pi - \lambda \mu_{t-1}^\pi - \frac{\lambda}{\beta\sigma}\sum_{j=0}^{t} \mu_{t-1-j}^x - \rho g \mu_t^{\omega\pi} = 0 \tag{A1}$$

$$\alpha x_t - \kappa \mu_t^\pi + \mu_t^x - \lambda\left(\frac{1}{\beta} - 1\right)\sum_{j=0}^{t} \mu_{t-1-j}^x - \rho g \mu_t^{\omega x} = 0 \tag{A2}$$

$$-(1-\lambda)\beta^2\rho \mathbb{E}_t \mu_{t+1}^\pi - (1-\lambda)\frac{\beta\rho}{1-\beta\rho}\frac{1}{\sigma}\mathbb{E}_t \mu_{t+1}^x + \mu_t^{\omega\pi} - \beta\rho(1-g)\mathbb{E}_t \mu_{t+1}^{\omega\pi} = 0 \tag{A3}$$

$$-(1-\lambda)\frac{(1-\beta)\beta\rho}{1-\beta\rho}\mathbb{E}_t \mu_{t+1}^x + \mu_t^{\omega x} - \beta\rho(1-g)\mathbb{E}_t \mu_{t+1}^{\omega x} = 0 \tag{A4}$$

$$\frac{1}{\sigma}\sum_{j=0}^{t} \mu_{t-j}^x = 0. \tag{A5}$$

From (A5), we have $\mu_t^x = 0$. This is Result 1 (Demand Irrelevance). When policy is unconstrained, the IS curve is not a binding constraint. The central bank can generate whatever output gap it wants.

Substituting $\mu_t^x = 0$ into (A4) gives $\mu_t^{\omega^x} = 0$. Learners' output gap expectations enter only the IS curve. Therefore, when policy is unconstrained, there is no reason for the central bank to try to influence learners' output gap expectations.[23]

Combining (A1) and (A2) to eliminate μ_t^π gives equation (15) in the main text. Similarly, using (A1) to eliminate μ_t^π in (A3) gives equation (16). The remainder of the derivation for Result 18 (unconstrained optimal target criterion) is in the main text.

A.1.1 Results for ζ_1 and ζ_2

The characteristic equation for (16) is

$$A\zeta^2 - B\zeta + C = 0$$

---

[23] This conclusion would not hold if we had an infinite-horizon Phillips curve, as in Eusepi, Giannoni, and Preston (2018), because output gap expectations would then enter the Phillips curve.

where $A = \rho\beta(1-g)+(1-\lambda)g\rho^2\beta^2$, $B = 1+\lambda\rho\beta(1-g)$ and $C = \lambda$. This equation has two real non-negative solutions. Substituting $\zeta = 1$ into the LHS gives $-(1-\lambda)(1-\beta\rho)(1+\rho g)$, so $\zeta_1 \leq 1$ (with equality if $\lambda = 1$) and $\zeta_2 > 1$. Substituting $\zeta = \lambda$ into the LHS gives $\lambda^2(1-\lambda)g\rho^2\beta^2$, so $\zeta_1 \geq \lambda$, with equality if $\lambda \in \{0,1\}$ or $g = 0$. Therefore, $\zeta_2^{-1} \leq \rho\beta(1-g)+(1-\lambda)g\rho^2\beta^2$.

Differentiating the characteristic equation with respect to any parameter and rearranging gives

$$\zeta' = \frac{A'\zeta^2 - B'\zeta + C'}{B - 2A\zeta}$$

where ζ is a solution to the quadratic equation. Since $\zeta_1 < \frac{B}{2A} < \zeta_2$, the sign of ζ' is determined by the numerator.

First consider g:

$$\frac{\partial A}{\partial g}\zeta^2 - \frac{\partial B}{\partial g}\zeta + \frac{\partial C}{\partial g} = \left(-\rho\beta + (1-\lambda)\rho^2\beta^2\right)\zeta^2 + \lambda\rho\beta\zeta = \rho\beta\zeta\left(\frac{\lambda}{1-(1-\lambda)\rho\beta} - \zeta\right).$$

Substituting $\frac{\lambda}{1-(1-\lambda)\rho\beta}$ into the quadratic gives $-\frac{\lambda(1-\lambda)\rho\beta(1-\rho\beta)}{(1-(1-\lambda)\rho\beta)^2} \leq 0$. Therefore $\zeta_1 \leq \frac{\lambda}{1-(1-\lambda)\rho\beta} < \zeta_2$, which implies that both ζ_1 and ζ_2 are increasing in g.

Now consider λ:

$$\frac{\partial A}{\partial \lambda}\zeta^2 - \frac{\partial B}{\partial \lambda}\zeta + \frac{\partial C}{\partial \lambda} = -g\rho^2\beta^2\zeta^2 - \rho\beta(1-g)\zeta + 1 = 1 - \rho\beta\zeta\left(1 - g\left(1-\rho\beta\zeta\right)\right)$$

This expression is positive for ζ_1, because $\zeta_1 < 1$. It is less than or equal to zero for ζ_2, because $\zeta_2 \geq \frac{1}{\rho\beta}$. Therefore, both ζ_1 and ζ_2 are increasing in λ.

Finally consider ρ:

$$\begin{aligned}\frac{\partial A}{\partial \rho}\zeta^2 - \frac{\partial B}{\partial \rho}\zeta + \frac{\partial C}{\partial \rho} &= \left(\beta(1-g) + 2(1-\lambda)g\rho\beta^2\right)\zeta^2 - \lambda\beta(1-g)\zeta \\ &= \beta(1-g)\zeta(\zeta-\lambda) + 2(1-\lambda)g\rho\beta^2\zeta\end{aligned}$$

which is non-negative for ζ_1 and positive for ζ_2, since $\zeta_2 > \zeta_1 \geq \lambda$. Therefore, ζ_1 is increasing in ρ and ζ_2 is decreasing.

A.1.2 Derivation of eq. (20) First, we rewrite $\mu_t^{\omega\pi}$ in terms of expected output gaps by substituting (15) into the second right-hand side term in (16) and iterating forward to get

$$\mu_t^{\omega\pi} = (1-\lambda)\rho\beta^2\mathbb{E}_t\frac{1}{1-\rho\beta(1-g)L^{-1}}\frac{\alpha}{\kappa}x_{t+1}.$$

Then, inverting the optimal target criterion (18) gives

$$\pi_t = -\frac{\alpha}{\kappa}x_t + \lambda\frac{\alpha}{\kappa}x_{t-1} + (1-\lambda)g\rho^2\beta^2\mathbb{E}_t\frac{1}{1-\rho\beta(1-g)L^{-1}}\frac{\alpha}{\kappa}x_{t+1} \tag{A6}$$

The equilibrium system of (11), (13) and (A6) has the following minimum state variable solution

$$\mu_t^{\omega\pi} = (1-\lambda)\left(a_x x_t - a_\omega\omega_t^\pi - a_u u_t\right) = (1-\lambda)\left(a_x x_t - a_\omega\rho g\pi_t - a_\omega\rho(1-g)\omega_{t-1}^\pi - a_u u_t\right)$$

for some constants a_x, a_ω and a_u, where the second equality uses the law of motion of ω_t^π.

A.2 CONSTRAINED OPTIMAL POLICY

Suppose we add to (11) - (14) some constraint on the policy rate. This could be an information constraint (22), or a ZLB constraint (33). For this optimal policy problem, the first four first-order conditions, (A1) - (A4), remain the same. But instead of (A5), we have (23), reproduced here:

$$\frac{1}{\sigma}\sum_{j=0}^{t}\mu_{t-j}^x - \mu_t^i = 0 \tag{23}$$

where μ_t^i is the Lagrange multiplier on the policy rate constraint.[24]

Using (A1) to eliminate μ_t^π from (A2) - (A4) and then substituting in $\sigma(1-L)\mu_t^i = \mu_t^x$ from (23) gives (26), (35) and (36), which we reproduce here:

$$\frac{\alpha}{\kappa}x_t + \frac{\sigma}{\kappa}\mu_t^i = -\frac{1}{1-\lambda L}\left(\pi_t - \frac{\lambda}{\beta}\frac{\kappa}{\sigma}\frac{\sigma}{\kappa}\mu_{t-1}^i - \rho g\mu_t^{\omega\pi}\right) + \left(\lambda\left(\frac{1}{\beta}-1\right)+1\right)\frac{\sigma}{\kappa}\mu_{t-1}^i + \frac{\rho g}{\kappa}\mu_t^{\omega x} \tag{26}$$

$$\begin{aligned}
\mu_t^{\omega\pi} &= -(1-\lambda)\rho\Omega\frac{1-\lambda L}{1-\zeta_1 L}\mathbb{E}_t\frac{1}{1-\zeta_2 L^{-1}}\frac{1}{1-\lambda L}\left(\pi_{t+1} - \frac{\lambda}{\beta}\frac{\kappa}{\sigma}\frac{\sigma}{\kappa}\mu_t^i\right) \\
&\quad + (1-\lambda)\frac{\rho\Omega}{\beta(1-\beta\rho)}\frac{\kappa}{\sigma}\frac{1-\lambda L}{1-\zeta_1 L}\mathbb{E}_t\frac{1-L}{1-\zeta_2 L^{-1}}\frac{\sigma}{\kappa}\mu_{t+1}^i \qquad (35) \\
\mu_t^{\omega x} &= (1-\lambda)\frac{(1-\beta)\beta\rho}{1-\beta\rho}\mathbb{E}_t\frac{1-L}{1-\beta\rho(1-g)L^{-1}}\frac{\sigma}{\kappa}\mu_{t+1}^i. \qquad (36)
\end{aligned}$$

Now substitute (35) and (36) into (26) and collect all the non-μ_t^i terms in $\tilde{\pi}_t$, which is

---

[24] A similar first-order condition would apply if, instead of a constraint on the policy rate, the loss function contained a policy-rate stability or smoothing term. Then μ_t^i would be replaced with the derivative of the (intertemporal) loss function with respect to i_t. The structure of the optimal target criterion is therefore identical in all these cases up to the definition of μ_t^i.

defined in Result (4) in the main text. This gives

$$\begin{aligned}
\tilde{\pi}_t + \frac{\sigma}{\kappa}\mu_t^i = {} & \frac{\lambda}{\beta}\frac{\kappa}{\sigma}\frac{1}{1-\lambda L}\frac{\sigma}{\kappa}\mu_{t-1}^i \\
& + (1-\lambda)g\rho^2\Omega\frac{\lambda}{\beta}\frac{\kappa}{\sigma}\frac{1}{1-\zeta_1 L}\mathbb{E}_t\frac{1}{1-\zeta_2 L^{-1}}\frac{1}{1-\lambda L}\frac{\sigma}{\kappa}\mu_t^i \\
& + (1-\lambda)g\rho^2\Omega\frac{1}{\beta(1-\beta\rho)}\frac{\kappa}{\sigma}\frac{1}{1-\zeta_1 L}\mathbb{E}_t\frac{1-L}{1-\zeta_2 L^{-1}}\frac{\sigma}{\kappa}\mu_{t+1}^i \\
& + \left(\lambda\left(\frac{1}{\beta}-1\right)+1\right)\frac{\sigma}{\kappa}\mu_{t-1}^i \\
& + (1-\lambda)g\rho^2\frac{(1-\beta)\beta}{1-\beta\rho}\mathbb{E}_t\frac{1-L}{1-\beta\rho(1-g)L^{-1}}\frac{\sigma}{\kappa}\mu_{t+1}^i \\
= {} & \left\{\frac{\lambda}{\beta}\frac{\kappa}{\sigma}\frac{1}{1-\lambda L}+\lambda\left(\frac{1}{\beta}-1\right)+1\right\}\frac{\sigma}{\kappa}\mu_{t-1}^i \\
& + (1-\lambda)g\rho^2\left\{\Omega\frac{\lambda}{\beta}\frac{\kappa}{\sigma}\frac{1}{1-\zeta_2\lambda}\frac{1}{1-\zeta_1 L}\frac{1}{1-\lambda L}-\Omega\frac{1}{\beta(1-\beta\rho)}\frac{\kappa}{\sigma}\frac{1}{1-\zeta_1 L}-\frac{(1-\beta)\beta}{1-\beta\rho}\right\}\frac{\sigma}{\kappa}\mu_t^i \\
& + (1-\lambda)g\rho^2\frac{(1-\beta)\beta}{1-\beta\rho}\left(1-\beta\rho(1-g)\right)\mathbb{E}_t\frac{1}{1-\beta\rho(1-g)L^{-1}}\frac{\sigma}{\kappa}\mu_{t+1}^i \\
& + (1-\lambda)g\rho^2\Omega\frac{\kappa}{\sigma}\left\{\frac{\lambda}{\beta}\frac{1}{1-\zeta_2\lambda}\zeta_2+\frac{1}{\beta(1-\beta\rho)}(1-\zeta_2)\right\}\mathbb{E}_t\frac{1}{1-\zeta_2 L^{-1}}\frac{\sigma}{\kappa}\mu_{t+1}^i \\
& + (1-\lambda)g\rho^2\Omega\frac{\kappa}{\sigma}\left\{\frac{\lambda}{\beta}\frac{1}{1-\zeta_2\lambda}\zeta_2+\frac{1}{\beta(1-\beta\rho)}(1-\zeta_2)\right\}\zeta_1\frac{1}{1-\zeta_1 L}\mathbb{E}_{t-1}\frac{1}{1-\zeta_2 L^{-1}}\frac{\sigma}{\kappa}\mu_t^i \\
= {} & -(1-\lambda)g\rho^2\Theta\frac{\sigma}{\kappa}\mu_t^i + P(L)\frac{\sigma}{\kappa}\mu_{t-1}^i + \mathbb{E}_t P_1(L^{-1})\frac{\sigma}{\kappa}\mu_{t+1}^i + P_2(L)\mathbb{E}_{t-1}P_3(L^{-1})\frac{\sigma}{\kappa}\mu_t^i
\end{aligned}$$

where Θ and $P(L)$ are defined in Result (4) and $P_1(L^{-1})$, $P_2(L)$ and $P_3(L^{-1})$ are defined in Result (5). Finally, define

$$\begin{aligned}
\pi_t^* &\equiv \frac{1}{1+(1-\lambda)g\rho^2\Theta}\frac{\sigma}{\kappa}\left(P(L)\mu_{t-1}^i + \mathbb{E}_t P_1(L^{-1})\mu_{t+1}^i + P_2(L)\mathbb{E}_{t-1}P_3(L^{-1})\mu_t^i\right) \\
\Delta_t &\equiv \pi_t^* - \tilde{\pi}_t = \left(1+(1-\lambda)g\rho^2\Theta\right)\frac{\sigma}{\kappa}\mu_t^i.
\end{aligned}$$

So we have

$$\tilde{\pi}_t + \Delta_t = \pi_t^* \equiv \frac{1}{1+(1-\lambda)g\rho^2\Theta}\left(P(L)\Delta_{t-1} + \mathbb{E}_t P_1(L^{-1})\Delta_{t+1} + P_2(L)\mathbb{E}_{t-1}P_3(L^{-1})\frac{\sigma}{\kappa}\Delta_t\right). \tag{A7}$$

To get the optimal target criterion under imperfect information, (4), take the expectation of (A7) with respect to the central bank's information set. Since Δ_t is just a multiple of μ_t^i, (24) implies $\Delta_{t|t} = 0$ and $\mathbb{E}_t\Delta_{t+k} = 0$ for any $k \geq 1$.

To get the optimal target criterion with the ZLB, (5), just recognise that Δ_t is nonzero if and only if the ZLB is binding.

A.2.1 DERIVATION OF EQ. (32)

Start the optimal policy conditions (26), (27) and (28). Together with the Phillips curve (11) and the updating rule for learners' inflation expectations (13), these five equations determine the optimal equilibrium paths of $\{\pi_t, x_t, \omega_t^{\pi}, \mu_t^{\omega\pi}, \mu_t^{\omega x}\}$ in terms of the paths of μ_t^i and the cost-push shock u_t.

As in the derivation of eq. (20), we can rewrite $\mu_t^{\omega\pi}$ in (27) in terms of expected output gaps instead of expected inflation:

$$\mu_t^{\omega\pi} = (1-\lambda)\rho\beta^2 \mathbb{E}_t \frac{1}{1-\rho\beta(1-g)L^{-1}} \frac{\alpha}{\kappa} x_{t+1} - (1-\lambda)\rho\beta^2 \left(\frac{\lambda}{\beta}(1-\beta) + 1 + \frac{1}{\beta(1-\beta\rho)} \frac{\kappa}{\sigma} \right) \frac{\sigma}{\kappa} \mu_t^i. \tag{A8}$$

Then combining (28) and(A8) and using the minimum state variable solution to evaluate expectations terms, we can write

$$\begin{aligned}
\mu_t^{\omega\pi} + (1-\lambda L)\frac{1}{\kappa}\mu_t^{\omega x} &= (1-\lambda)\rho\beta^2 \mathbb{E}_t \frac{1}{1-\rho\beta(1-g)L^{-1}} \frac{\alpha}{\kappa} x_{t+1} \\
&\quad - (1-\lambda)\rho\beta^2 \left(\frac{\lambda}{\beta}(1-\beta) + 1 + \frac{1}{\beta(1-\beta\rho)}\frac{\kappa}{\sigma} + \frac{1-\beta}{\beta(1-\beta\rho)} \right) \frac{\sigma}{\kappa}\mu_t^i \\
&\quad + \lambda(1-\lambda)\rho\beta^2 \frac{1-\beta}{\beta(1-\beta\rho)} \frac{\sigma}{\kappa} \mu_{t-1}^i \\
&= (1-\lambda)\left(a_x x_t - a_\omega \rho g \pi_t - a_\omega \rho(1-g)\omega_{t-1}^{\pi} - a_u u_t - a_{\mu,1}\mu_t^i - a_{\mu,2}\mu_{t-1}^i \right)
\end{aligned}$$

where a_x, a_ω and a_u are the same constants as in (19), and $a_{\mu,1}$ and $a_{\mu,2}$ are new constants. Combining with the updating rule for learners' inflation expectations and using (22) gives the following description of optimal policy

$$\mathbb{E}_{t|t} \frac{1}{1-\frac{\rho_0-\rho_1 L}{1-\gamma L}L} \left(\frac{\alpha}{\kappa} x_t + \phi \left(\frac{\psi}{1-\gamma L} + \frac{1-\psi}{1-\rho(1-g)L} \right) \pi_t + \phi \nu_u \frac{1}{1-\gamma L} u_t \right) = 0.$$

Or, simply $\mathbb{E}_{t|t}\left(Q_x(L)x_t + Q_\pi(L)\pi_t + Q_u u_t\right) = 0$, where $Q_x(L)$, $Q_\pi(L)$ and $Q_u(L)$ are power series in L. The new parameters are

$$\rho_0 \equiv \frac{\lambda + \frac{\lambda}{\beta}\frac{\kappa}{\sigma} + \frac{\lambda}{\beta}(1-\beta) + 1 - (1-\lambda)\rho g a_{\mu,2}\frac{\kappa}{\sigma}}{1 + (1-\lambda)\rho g a_{\mu,1}\frac{\kappa}{\sigma}} - \gamma$$

and $\rho_1 \equiv \frac{\lambda\left(\frac{\lambda}{\beta}(1-\beta)+1\right)}{1+(1-\lambda)\rho g a_{\mu,1}\frac{\kappa}{\sigma}}$. The first and third power series, $Q_x(L)$ and $Q_u(L)$, are two different weighted averages of two exponentially-weighted averages of lags. The second power series, $Q_\pi(L)$, is a weighted average of these same two exponentially-weighted averages of lags, plus a third exponentially-weighted average of lags, i.e.

$$\begin{aligned}
Q_x(L) &\equiv \frac{1}{1-\frac{\rho_0-\rho_1 L}{1-\gamma L}L} \equiv \frac{\phi_x}{1-\xi_1 L} + \frac{1-\phi_x}{1-\xi_2 L} \\
Q_\pi(L) &\equiv \frac{1}{1-\frac{\rho_0-\rho_1 L}{1-\gamma L}L} \phi \left(\frac{\psi}{1-\gamma L} + \frac{1-\psi}{1-\rho(1-g)L} \right) \equiv \phi \left(\frac{\phi_{\pi 1}}{1-\xi_1 L} + \frac{\phi_{\pi 2}}{1-\xi_2 L} + \frac{1-\phi_{\pi 1}-\phi_{\pi 2}}{1-\rho(1-g)L} \right) \\
Q_u(L) &\equiv \frac{1}{1-\frac{\rho_0-\rho_1 L}{1-\gamma L}L} \phi\nu_u \frac{1}{1-\gamma L} \equiv \phi\nu_u \left(\frac{\phi_u}{1-\xi_1 L} + \frac{1-\phi_u}{1-\xi_2 L} \right)
\end{aligned}$$

where ξ_1 and ξ_2 are the roots of $\xi^2-(\gamma+\rho_0)\xi+\rho_1$ and $\phi_{\pi 1}$, $\phi_{\pi 2}$, ϕ_x and ϕ_u are constants.

B APPENDIX - FURTHER RESULTS

B.1 OPTIMAL POLICY WITH NON-RATIONAL INTEREST-RATE EXPECTATIONS

When does learning about interest rates constrain policy? To get a sense of how learning about i_t might alter our results, here we repeat the simple example in section 2 of Eusepi, Giannoni, and Preston (forthcoming) with our hybrid model for expectations. Consider the model

$$\pi_t = \beta \hat{E}_t \pi_{t+1} + \kappa x_t \ , \ x_t = \hat{E}_t \frac{1}{1-\beta L^{-1}} \left[(1-\beta) x_{t+1} - \frac{1}{\sigma}(i_t - \pi_{t+1} - r_t^n) \right]$$

with r_t^n i.i.d.. Suppose the central bank wants to implement $\pi_t = 0$. Then the interest-rate path follows

$$i_t = r_t^n + \frac{\sigma\beta}{\kappa} \hat{E}_t \pi_{t+1} + \hat{E}_t \frac{1}{1-\beta L^{-1}} \left[\sigma(1-\beta) x_{t+1} - (\beta i_{t+1} - \pi_{t+1} - \beta r_{t+1}^n) \right]$$

Now suppose that nominal interest-rate expectations are partially rational and partially learned:

$$\begin{aligned} \hat{E}_t i_T &= \lambda \mathbb{E}_t i_T + (1-\lambda) \rho^{T-t} \omega_{t-1}^i \\ \omega_t^i &= \rho \omega_{t-1}^i + \rho g (i_t - \omega_{t-1}^i). \end{aligned}$$

For simplicity, assume all output gap and inflation expectations are fixed at steady state (the mechanism we are interested in here is the feedback between interest rates and interest-rate expectations). Substituting in expectations, the interest-rate path needed to implement π_t is

$$i_t = r_t^n - \lambda\beta \mathbb{E}_t \frac{1}{1-\beta L^{-1}} i_{t+1} - (1-\lambda) \frac{\beta\rho}{1-\beta\rho} \omega_{t-1}^i.$$

Substituting in $\omega_t = \frac{\rho g}{1-\rho(1-g)L} i_t$ and rearranging gives

$$\mathbb{E}_t \left[1 + \lambda\beta \frac{L^{-1}}{1-\beta L^{-1}} + (1-\lambda) \frac{\beta\rho^2 g}{1-\beta\rho} \frac{L}{1-\rho(1-g)L} \right] i_t = r_t^n. \tag{B1}$$

A bounded path for i_t and ω_t^i exists only when[25]

$$g < \frac{1+\rho}{\rho}\Lambda \quad \text{where } \Lambda = \frac{1}{1+\frac{(1-\lambda)\frac{\beta\rho}{1-\beta\rho}}{1-\lambda\frac{\beta}{1+\beta}}}. \tag{B2}$$

If we substitute in $\rho = 1$ and $\lambda = 0$, then this expression nests the equivalent condition (8) in Eusepi, Giannoni, and Preston (forthcoming): $g < 2(1-\beta)$.

Effect of λ: Relative to EGP, introducing rational expectations ($\lambda > 0$) affects this constraint in two opposite ways. The reduction in the share of learners means their interest-rate expectations have less of an effect (this is the $(1-\lambda)$ term in Λ). But the rational agents recognise that the central bank in the future have to offset the interest-rate expectations of learners, which changes their rate expectations, and the central bank now has to offset these rational expectations (this is the λ term in Λ). But we can rewrite Λ as

$$\Lambda = (1-\beta\rho)\left(1+\frac{\beta\rho\lambda}{1+\beta(1-(1+\rho)\lambda)}\right) \geq 1-\beta\rho.$$

Therefore, the first effect outweighs the second, and so increasing the share of rational agents eases the constraint.

Effect of ρ: If $\rho < 1$, then the constraint eases because (i) a given interest-rate surprise does not affect short-term beliefs as much or for as long, and (ii) a given short-term belief does not affect long-term interest rate expectations as much.

Optimal policy with learning about interest rates For a simple example, assume prices are fixed, and continue with the assumption that output gap expectations are fixed. The optimal policy problem is

$$\begin{aligned}
\min_{\{x_t, i_t, \omega_t^i\}} \quad & \frac{1}{2}\mathbb{E}_t \frac{1}{1-\beta L^{-1}} x_t^2 \\
\text{s.t.} \quad & \sigma x_t = -i_t - \lambda\beta\mathbb{E}_t \frac{1}{1-\beta L^{-1}} i_{t+1} - (1-\lambda)\frac{\beta\rho}{1-\beta\rho}\omega_{t-1}^i + r_t^n \\
& \omega_t^i = \rho g i_t + \rho(1-g)\omega_{t-1}^i.
\end{aligned}$$

Let $\beta \to 1$. Take the first-order conditions and eliminate multipliers to get the targeting rule

$$x_t + \lambda\frac{1}{1-L}x_{t-1} + (1-\lambda)\frac{\rho^2 g}{1-\rho}\mathbb{E}_t \frac{1}{1-\rho(1-g)L^{-1}} x_{t+1} = 0. \tag{B3}$$

---

[25] For all parameterisations, the roots of the characteristic polynomial for (B1) are real, and only one is greater than 1. A bounded path for i_t and ω_t^i exists if and only if the other root is greater than -1, which occurs when condition (B2) is satisfied.

The intuition in the main text around history dependence and pre-emption depending on the proportion of rational agents and learners extends to the case of interest-rate expectations.

The characteristic polynomial for (B3) is very closely related to the characteristic polynomial for (B1). Condition (B2) turns out to be important for implementing this optimal target criterion:[26]

- If (B2) is satisfied, then the target criterion (B3) pins down a unique path for x_t (which converges to $x_t = 0$). Substituting it into the IS curve gives a non-explosive path for i_t.
- If (B2) is not satisfied, then there are infinitely many paths for x_t that satisfy the target criterion. But, substituting into the IS curve, only one of those paths implies a bounded path for i_t and ω_t^i. Therefore, there still exists a unique bounded solution that is consistent with the optimal target criterion.

If $\beta < 1$, then there is a set of parameterisations for which these is no bounded solution consistent with the optimal target criterion. This occurs when β is sufficiently low that the optimal intertemporal trade-off requires $x_t = 0$, but g is sufficiently high than this cannot be achieved without an explosive interest-rate path. It occurs due to a tension between an impatient policymaker, who discounts the future a lot, and significant lags in the transmission of policy, which mean that current policy has significant future effects. If β is close to one, then the set of parameterisations for which this occurs is small.

B.2 Optimal policy under discretion

What would our optimal targeting rules look like under discretion? This is a difficult question to address analytically in our full model, but we can consider a two-period version of our model to get some intuition:

$$\pi_0 = \beta\lambda\mathbb{E}_0\pi_1 + \kappa x_0 + u_0 \tag{B4}$$

$$\pi_1 = \beta(1-\lambda)g\pi_0 + \kappa x_1 + u_1. \tag{B5}$$

Optimal policy under commitment We start by solving for optimal commitment policy in this two-period model so that we have a benchmark to which we can compare

---

[26] The roots of the characteristic polynomial for (B3) are the reciprocals of the roots of the characteristic polynomial for (B1) (setting $\beta = 1$ in (B1)). They are real, and at least one lies inside the unit circle. The other root lies outside the unit circle if and only if condition (B2) is satisfied.

optimal discretionary policy. The first-order conditions under commitment are

$$\begin{aligned}\pi_0 - \mu_0 + \beta^2(1-\lambda)g\mathbb{E}_0\mu_1 &= 0\\ \pi_1 + \lambda\mu_0 - \mu_1 &= 0\\ \alpha x_0 + \kappa\mu_0 &= 0\\ \alpha x_1 + \kappa\mu_1 &= 0.\end{aligned}$$

Eliminating the multipliers gives the optimal target criteria:

$$\frac{\alpha}{\kappa}x_0 = -\frac{\pi_0 + \beta^2(1-\lambda)g\mathbb{E}_0\pi_1}{1-\beta^2(1-\lambda)g\lambda} \tag{B6}$$

$$\frac{\alpha}{\kappa}x_1 = -\pi_1 - \frac{\lambda\beta^2(1-\lambda)g}{1-\lambda\beta^2(1-\lambda)g}\mathbb{E}_0\pi_1 - \frac{1}{1-\lambda\beta^2(1-\lambda)g}\lambda\pi_0. \tag{B7}$$

The heterogeneity of expectations and the interaction between rational and learners' expectations is important (i.e. the denominators increase to 1 if $\lambda = 0$ or $\lambda = 1$). Eq. (B6) shows that when $\lambda \in (0,1)$, greater weight should be placed on current and expected future inflation, because the costs of following through on make-up commitments next period are eased by influencing learners' expectations this period. Eq. (B7) shows that for $\lambda \in (0,1)$, greater weight should also be placed on past inflation and past expectations of inflation, because these make-up commitments ease the costs of past pre-emptive policy.

Optimal policy under discretion Working backwards, optimal policy under discretion in the second period is

$$\frac{\alpha}{\kappa}x_1 = -\pi_1. \tag{B8}$$

Then in the first period, the policymaker faces the constraints (B4), (B5) and (B8). The first-order conditions are

$$\begin{aligned}\pi_0 - \mu_0 + \beta^2(1-\lambda)g\mathbb{E}_0\mu_1 &= 0\\ \pi_1 + \lambda\mu_0 - \mu_1 - \mu_2 &= 0\\ \alpha x_0 + \kappa\mu_0 &= 0\\ \alpha x_1 + \kappa\mu_1 - \frac{\alpha}{\kappa}\mu_2 &= 0.\end{aligned}$$

Eliminating the multipliers gives the optimal target criterion

$$\frac{\alpha}{\kappa}x_0 = -\frac{\pi_0 + \beta^2(1-\lambda)g\mathbb{E}_0\pi_1}{1-\beta^2(1-\lambda)g\frac{\lambda}{1+\frac{\kappa^2}{\alpha}}}. \tag{B9}$$

Compare (B9) to the optimal commitment policy (B6). Optimal policy under discretion places less weight on both inflation terms. Again heterogeneity is important: this difference is only because of the interaction between rational agents and learners. If $\lambda = 0$ or $\lambda = 1$, then period 0 optimal policy is identical under commitment and discretion.

The reason why discretion involves *less* response to inflation is because the future policymaker will not be responding to past inflation. So there is no need to use learners' expectations to ease the cost of doing so. With both learning and rational expectations, there is still some reason to lift learners' inflation expectations, because that will affect future inflation and thereby influence current rational inflation expectations. But the future policymaker will attempt to offset this inflation rather than accomodate it, so the effect is less powerful.

C APPENDIX C - OPTIMAL POLICY WITH INFINITE-HORIZON PHILLIPS CURVE

In the main body, we model firms' pricing decisions by taking the standard New Keynesian Phillips curve, and substituting our alternative specification for aggregate inflation expectations into the expectation term. This is consistent with the 'Euler equation' approach to adaptive learning, and also approximates a range of other behavioural theories. But it is not consistent with the 'anticipated utility' approach to adaptive learning. In this appendix, we provide optimal policy results using an 'anticipated utility' specification for the Phillips curve:

$$\pi_t = \hat{\mathbb{E}}_t \frac{1}{1-\theta\beta L^{-1}} \left(\kappa x_t + (1-\theta)\beta\pi_{t+1}\right). \tag{C1}$$

Here the parameter θ is the Calvo probability, i.e. the probability that a firm cannot change its price in a given period. This Phillips curve follows from the firms' optimal price setting under the standard microfoundations of the New Keynesian model for any arbitrary expectations. Firms' prices depend on their expectations over the entire future sequence of inflation and output gaps, $\{\pi_T, x_T\}_{T=t}^{\infty}$ when set prices. If we assume rational expectations, then (C1) collapses to the standard one-step-ahead Phillips curve (3).

Using our specification for agents' beliefs, (4) - (7), we have

$$\begin{aligned}
\pi_t &= \lambda\mathbb{E}_t \frac{1}{1-\theta\beta L^{-1}} \left(\kappa x_t + (1-\theta)\beta\pi_{t+1}\right) + (1-\lambda)\mathbb{E}_t^l \frac{1}{1-\theta\beta L^{-1}} \left(\kappa x_t + (1-\theta)\beta\pi_{t+1}\right) \\
&= \lambda\mathbb{E}_t \frac{1}{1-\theta\beta L^{-1}} \left(\kappa x_t + (1-\theta)\beta\pi_{t+1}\right) + (1-\lambda)\left(\kappa x_t + \frac{\theta\beta\rho}{1-\theta\beta\rho}\kappa\omega_{t-1}^x + \frac{(1-\theta)\beta\rho}{1-\theta\beta\rho}\omega_{t-1}^{\pi}\right).
\end{aligned} \tag{C2}$$

The optimal policy problem is the same as in the main body, but with a different

Phillips curve, i.e. choose the sequence

$$\{\pi_t, x_t, \omega_t^\pi, \omega_t^x, i_t\}_{t=0}^\infty$$

that maximises the welfare function (10) subject to the Phillips curve (C2), the IS curve (12), the evolution of beliefs (13) and (14), and potentially some information constraint or the ZLB.

The first-order conditions are

$$\begin{aligned}
\pi_t + \mu_t^\pi - \lambda(1-\theta)\frac{1}{1-\theta L}\mu_{t-1}^\pi - \frac{\lambda}{\beta\sigma}\frac{1}{1-L}\mu_{t-1}^x - \rho g \mu_t^{\omega\pi} &= 0\\
\alpha x_t - \kappa\mu_t^\pi - \lambda\theta\kappa\frac{1}{1-\theta L}\mu_{t-1}^\pi + \mu_t^x - \frac{\lambda(1-\beta)}{\beta}\frac{1}{1-L}\mu_{t-1}^x - \rho g\mu_t^{\omega x} &= 0\\
-(1-\lambda)\frac{(1-\theta)\beta^2\rho}{1-\theta\beta\rho}\mathbb{E}_t\mu_{t+1}^\pi - (1-\lambda)\frac{\beta\rho}{\sigma(1-\beta\rho)}\mathbb{E}_t\mu_{t+1}^x + \mu_t^{\omega\pi} - \beta\rho(1-g)\mathbb{E}_t\mu_{t+1}^{\omega\pi} &= 0\\
-(1-\lambda)\frac{\theta\beta^2\rho}{1-\theta\beta\rho}\kappa\mathbb{E}_t\mu_{t+1}^\pi - (1-\lambda)\frac{(1-\beta)\beta\rho}{1-\beta\rho}\mathbb{E}_t\mu_{t+1}^x + \mu_t^{\omega x} - \beta\rho(1-g)\mathbb{E}_t\mu_{t+1}^{\omega x} &= 0
\end{aligned}$$

plus an additional condition for μ_t^x that depends on whether there is an imperfect information or ZLB constraint.

C.1 UNCONSTRAINED CASE

In the absence of imperfect central bank information or the ZLB, the additional first-order condition is $\mu_t^x = 0$. Eliminating μ_t^x and μ_t^π gives

$$\frac{\alpha}{\kappa}x_t = -\frac{1-\theta(1-\lambda)L}{1-(\lambda+\theta(1-\lambda))L}\pi_t + \frac{1-\theta(1-\lambda)L}{1-(\lambda+\theta(1-\lambda))L}\rho g\mu_t^{\omega^\pi} + \frac{\rho g}{\kappa}\mu_t^{\omega^x}. \tag{C3}$$

The multipliers $\mu_t^{\omega^\pi}$ and $\mu_t^{\omega^x}$ - i.e. the welfare effect of shifts in learners' inflation and output-gap expectations - are given by

$$\begin{aligned}
\mu_t^{\omega^\pi} &= \beta\rho(1-g)\mathbb{E}_t\mu_{t+1}^{\omega^\pi} - \beta^2(1-\lambda)\frac{(1-\theta)\rho}{1-\beta\theta\rho}\mathbb{E}_t\frac{1-\theta L}{1-(\lambda+\theta(1-\lambda))L}\left(\pi_{t+1} - \rho g\mu_{t+1}^{\omega^\pi}\right)\\
\mu_t^{\omega^x} &= \kappa\frac{\theta}{1-\theta}\mu_t^{\omega^\pi}.
\end{aligned}$$

Solving forward for $\mu_t^{\omega^\pi}$ then gives

$$\frac{1}{1-(\lambda+\theta(1-\lambda))L}\mu_t^{\omega^\pi} = -(1-\lambda)\frac{(1-\theta)\rho}{1-\beta\theta\rho}\Omega'\frac{1}{1-\zeta_1'L}\mathbb{E}_t\frac{1}{1-\zeta_2'^{-1}L^{-1}}\frac{1-\theta L}{1-(\lambda+\theta(1-\lambda))L}\pi_{t+1}$$

where $\Omega' \equiv \frac{\beta^2 \zeta_2'^{-1}}{\beta\rho(1-g)+(1-\lambda)g\beta^2\frac{(1-\theta)\rho^2}{1-\beta\theta\rho}} \equiv \frac{\beta^2\zeta_1'}{\lambda+\theta(1-\lambda)}$ and ζ_1' and ζ_2' are the roots of the characteristic polynomial

$$\left(\beta\rho(1-g)+(1-\lambda)g\beta^2\frac{(1-\theta)\rho^2}{1-\beta\theta\rho}\right)\zeta^2-\left(1+\beta\rho(1-g)\left(\lambda+\theta(1-\lambda)\right)+\theta(1-\lambda)g\beta^2\frac{(1-\theta)\rho^2}{1-\beta\theta\rho}\right)\zeta$$
$$+\left(\lambda+\theta(1-\lambda)\right)=0.$$

Using this solution to substitute for $\mu_t^{\omega^\pi}$ and $\mu_t^{\omega^x}$ in (C3) gives the optimal target criterion:[27]

$$\frac{\alpha}{\kappa}x_t=-\frac{1-\theta(1-\lambda)L}{1-(\lambda+\theta(1-\lambda))L}\pi_t-(1-\lambda)g\frac{\rho^2}{1-\beta\theta\rho}\Omega'\frac{1-\theta L}{1-\zeta_1' L}\mathbb{E}_t\frac{1}{1-\zeta_2'^{-1}L^{-1}}\frac{1-\theta L}{1-(\lambda+\theta(1-\lambda))L}\pi_{t+1}. \tag{C4}$$

This optimal target criterion has the same structure as the analogous one in the main body, (18). It includes a make-up component and a pre-emptive component, and the intuition behind it is the same. The differences are a matter of degree. First, the optimal weights on past outcomes decay more slowly in (C4). This is because the infinite-horizon Phillips curve implies that make-up commitments further into the future have a stronger effect on current-period pricing decisions than is the case with the one-period-ahead Phillips curve.[28] Second, the optimal pre-emptive response to expected future inflation is stronger. This is because policy-induced changes in learners' beliefs now have a stronger effect on future outcomes, because both their inflation and output-gap expectations influence their pricing decisions (with the one-step-ahead Phillips curve, only their inflation expectations were relevant).

Table C1 compares the performance of the simple target criteria from Table 2 under the infinite-horizon Phillips curve. The equivalent values for the one-step-ahead Phillips curve (i.e. Table D1) are in parentheses. The results are fairly similar. As the previous paragraph would suggest, the main differences are that with the infinite-horizon Phillips curve (i) the optimal weight γ on lagged inflation outcomes in the WAIT is higher, and PLT now outperforms IT, and (ii) a preemptive response to expected inflation now offers a small welfare improvement.

---

[27] An alternative form of this optimal target criterion is

$$\pi_t=-\frac{\alpha}{\kappa}x_t+\left(\lambda\frac{L}{1-\theta L}+(1-\lambda)g\rho^2\frac{\beta^2}{1-\beta\theta\rho}\mathbb{E}_t\frac{L^{-1}}{1-\beta\rho(1-g)L^{-1}}\right)\left((1-\theta)\frac{\alpha}{\kappa}x_t-\theta\pi_t\right).$$

[28] With the one-step-ahead Phillips curve, commitments about policy more than one-period-ahead still influence current-period outcomes via the inflation expectations of rational agents. But this effect occurs via a general equilibrium channel and is therefore dampened by the presence of nonrational agents. With the infinite-horizon Phillips curve, these commitments affect the pricing decisions of rational agents directly, via both their inflation and output-gap expectations, without this dampening.

Table C1: Performance of simple target criteria with infinite-horizon Phillips curve - unconstrained case

	Loss (% rel. to IT)	ψ	γ	γ_x	ψ_f	γ_f
5 yr	61.1 (96.5)	16.9 (31.3)				
2 yr	16.8 (48.8)	18.5 (13.0)				
1 yr	2.8 (7.7)	15.0 (21.5)				
IT	0 (0)	10.6 (11.6)				
IT + preemptive	-0.0 (-4.4)	10.6 (0.03)			0.00 (1.09)	*undef.* (1.50)
PLT	-5.9 (6.6)	6.6 (9.1)	1 (1)			
WAIT	-11.0 (-6.6)	6.9 (8.4)	0.70 (0.52)			
WAIT + preemptive	-12.5 (-6.6)	5.5 (8.2)	0.89 (0.52)		10.92 (0.33)	0.00 (1.02)
WAIT + WAXT	-12.5 (-6.6)	8.0 (8.4)	0.88 (0.53)	0.43 (0.01)		
WAIT + WAXT + preemptive	-12.5 (-6.6)	7.9 (8.2)	0.88 (0.52)	0.43 (0.00)	0.12 (0.33)	0.03 (1.02)
Optimal	-12.5 (-6.6)					

Notes Values in parentheses are the equivalent values with the one-step-ahead Phillips curve (from Table D1). Welfare losses are normalized relative to the welfare under IT. The parameters for each rule are optimized to minimize welfare. The optimal value of a parameter is undefined (labelled *undef.*) if the parameter has no effect on loss at the optimum.

C.2 Comparison to Eusepi, Giannoni and Preston (2018)

Eusepi, Giannoni, and Preston (2018) study optimal policy in an adaptive learning model (without imperfect information or the ZLB). They show that in this setting, the infinite-horizon Phillips curve under some calibrations implies that optimal policy is history-dependent, whereas it is not history-dependent under the one-step-ahead Phillips curve.

If we set $\lambda = 0$, our model nests theirs. Optimal policy is given by setting $\lambda = 0$ in (C4):

$$\frac{\alpha}{\kappa} x_t = -\pi_t - g \frac{\beta^2 \rho^2}{1 - \beta\theta\rho} \mathbb{E}_t \frac{1}{1 - \beta\rho \left(1 - g \frac{1-\beta\rho}{1-\beta\theta\rho}\right) L^{-1}} \pi_{t+1}.$$

Setting $\theta = 0$ in this targeting rule gives optimal policy with $\lambda = 0$ and the one-step-ahead Phillips curve, as studied by Molnár and Santoro (2014). Therefore, using the infinite-horizon Phillips curve just increases the coefficient on the average of expected future inflation (while also slightly increasing the weight on nearer-term expectations within that average compared to further-out expectations). With $\lambda = 0$, optimal policy always involves no (direct) response to any lagged variables.

The difference in the optimal equilibrium stems not from the optimal policy rules, but from the direct effect of beliefs on the Phillips curve. Only inflation beliefs enter the one-step-ahead Phillips curve, so the Phillips curve will always shift up following an inflationary cost-push shock, then return to its original position. In contrast, both inflation and output-gap beliefs enter the infinite-horizon Phillips curve. If the loss function places sufficiently low weight on the output gap term, then the central bank's aggressive response will cause the Phillips curve to shift downwards (once the initial shock has dissipated), which means that inflation will overshoot. It is the private-sector pricing decisions that are history-dependent (via beliefs) and overshoot, not the policy rule.

C.3 General case

If there is a constraint on the policy rate (e.g. due to imperfect central bank information or the ZLB), then $\sigma(1-L)\mu_t^i = \mu_t^x$, where μ_t^i is the multiplier on this constraint. Substituting this condition in for μ_t^x and eliminating μ_t^π from the first-order conditions gives

$$\begin{aligned}\frac{\alpha}{\kappa}x_t &= -\frac{1-\theta(1-\lambda)L}{1-(\lambda+\theta(1-\lambda))L}\pi_t + \frac{1-\theta(1-\lambda)L}{1-(\lambda+\theta(1-\lambda))L}\rho g \mu_t^{\omega^\pi} + \frac{\rho g}{\kappa}\mu_t^{\omega^x} &\text{(C5)}\\ &\quad + \left(\frac{\lambda}{\beta}\frac{\kappa}{\sigma}\frac{1-\theta(1-\lambda)L}{1-(\lambda+\theta(1-\lambda))L} + \lambda\left(\frac{1}{\beta}-1\right)+1\right)\frac{\sigma}{\kappa}\mu_{t-1}^i. &\text{(C6)}\end{aligned}$$

The multipliers on learners' beliefs are

$$\begin{aligned}\mu_t^{\omega\pi} &= \beta\rho(1-g)\mathbb{E}_t\mu_{t+1}^{\omega\pi} - (1-\lambda)\frac{(1-\theta)\beta^2\rho}{1-\beta\theta\rho}\mathbb{E}_t\frac{1-\theta L}{1-(\lambda+\theta(1-\lambda))L}\left(\pi_{t+1} - \frac{\lambda}{\beta}\frac{\kappa}{\sigma}\frac{\sigma}{\kappa}\mu_t^i - \rho g \mu_{t+1}^{\omega^\pi}\right)\\ &\quad + (1-\lambda)\frac{\beta\rho}{1-\beta\rho}\mathbb{E}_t(1-L)\mu_{t+1}^i\\ \mu_t^{\omega^x} &= \frac{\theta}{1-\theta}\kappa\mu_t^{\omega^\pi} + (1-\lambda)\frac{\beta\rho}{1-\beta\rho}\left((1-\beta) - \frac{\theta}{1-\theta}\frac{\kappa}{\sigma}\right)\mathbb{E}_t\frac{1}{1-\beta\rho(1-g)L^{-1}}\sigma(1-L)\mu_{t+1}^i.\end{aligned}$$

Solving forward for $\mu_t^{\omega^\pi}$ gives

$$\begin{aligned}\mu_t^{\omega^\pi} &= -(1-\lambda)\frac{(1-\theta)\rho}{1-\beta\theta\rho}\Omega'\frac{1-(\lambda+\theta(1-\lambda))L}{1-\zeta_1'L}\mathbb{E}_t\frac{1}{1-\zeta_2'^{-1}L^{-1}}\frac{1-\theta L}{1-(\lambda+\theta(1-\lambda))L}\left(\pi_{t+1} - \frac{\lambda}{\beta}\frac{\kappa}{\sigma}\frac{\sigma}{\kappa}\mu_t^i\right)\\ &\quad + (1-\lambda)\frac{\rho\Omega'}{\beta(1-\beta\rho)}\frac{\kappa}{\sigma}\frac{1-(\lambda+\theta(1-\lambda))L}{1-\zeta_1'L}\mathbb{E}_t\frac{1-L}{1-\zeta_2'^{-1}L^{-1}}\frac{\sigma}{\kappa}\mu_{t+1}^i.\end{aligned}$$

Using this solution to substitute for $\mu_t^{\omega^\pi}$ and $\mu_t^{\omega^x}$ in (C6) and collecting all the non-μ_t^i terms in $\tilde{\pi}_t$ gives

$$\begin{aligned}\tilde{\pi}_t + \frac{\sigma}{\kappa}\mu_t^i &= \frac{\lambda}{\beta}\frac{\kappa}{\sigma}\frac{1-\theta(1-\lambda)L}{1-(\lambda+\theta(1-\lambda))L}\frac{\sigma}{\kappa}\mu_{t-1}^i\\ &\quad + (1-\lambda)g\frac{\rho^2}{1-\beta\theta\rho}\Omega'\frac{\lambda}{\beta}\frac{\kappa}{\sigma}\frac{1-\theta L}{1-\zeta_1'L}\mathbb{E}_t\frac{1}{1-\zeta_2'^{-1}L^{-1}}\frac{1-\theta L}{1-(\lambda+\theta(1-\lambda))L}\frac{\sigma}{\kappa}\mu_t^i\\ &\quad + (1-\lambda)g\frac{\rho^2}{\beta(1-\beta\rho)(1-\theta)}\Omega'\frac{\kappa}{\sigma}\frac{1-\theta L}{1-\zeta_1'L}\mathbb{E}_t\frac{1-L}{1-\zeta_2'^{-1}L^{-1}}\frac{\sigma}{\kappa}\mu_{t+1}^i\\ &\quad + \left(\lambda\left(\frac{1}{\beta}-1\right)+1\right)\frac{\sigma}{\kappa}\mu_{t-1}^i\\ &\quad + (1-\lambda)g\frac{\beta\rho^2}{1-\beta\rho}\left((1-\beta) - \frac{\theta}{1-\theta}\frac{\kappa}{\sigma}\right)\mathbb{E}_t\frac{1-L}{1-\beta\rho(1-g)L^{-1}}\frac{\sigma}{\kappa}\mu_{t+1}^i\\ &= -(1-\lambda)g\rho^2\Theta'\frac{\sigma}{\kappa}\mu_t^i + P'(L)\frac{\sigma}{\kappa}\mu_{t-1}^i + \mathbb{E}_tP_2'(L^{-1})\frac{\sigma}{\kappa}\mu_{t+1}^i + P_3'(L)\mathbb{E}_{t-1}P_4'(L^{-1})\frac{\sigma}{\kappa}\mu_t^i\end{aligned}$$

where $\tilde{\pi}_t$ is the residual from the unconstrained optimal target criterion (C4) and Θ', $P(L)$, $P_1'(L^{-1})$, $P_2'(L)$ and $P_3'(L^{-1})$ are defined as

$$\begin{aligned}
\Theta' &= -\frac{\Omega'}{1-\beta\theta\rho}\frac{\lambda}{\beta}\frac{\kappa}{\sigma}\frac{1-\zeta_2'^{-1}\theta}{1-\zeta_2'^{-1}(\lambda+\theta(1-\lambda))} + \frac{\Omega'}{\beta(1-\beta\rho)(1-\theta)}\frac{\kappa}{\sigma} + \frac{\beta}{1-\beta\rho}\left((1-\beta) - \frac{\theta}{1-\theta}\frac{\kappa}{\sigma}\right) \\
P'(L) &= \frac{\lambda}{\beta}\frac{\kappa}{\sigma}\frac{1-\theta(1-\lambda)L}{1-(\lambda+\theta(1-\lambda))L} + \lambda\left(\frac{1}{\beta}-1\right) + 1 \\
&\quad + (1-\lambda)g\rho^2\Omega'\frac{\kappa}{\sigma}\left(\frac{\lambda}{\beta}\frac{1-\zeta_2'^{-1}\theta}{1-\zeta_2'^{-1}(\lambda+\theta(1-\lambda))}\left(\frac{\lambda(1-\theta)}{1-\zeta_2'^{-1}\theta}\frac{1-\theta L}{1-(\lambda+\theta(1-\lambda))L} + \zeta_1 - \theta\right)\right. \\
&\qquad\qquad \left. - \frac{\zeta_1-\theta}{\beta(1-\beta\rho)(1-\theta)}\right)\frac{1}{1-\zeta_1 L} \\
P_1'(L^{-1}) &= (1-\lambda)g\rho^2\Omega'\frac{\kappa}{\sigma}\left(\frac{\lambda}{\beta}\frac{\zeta_2'^{-1}(1-\zeta_2'^{-1}\theta)}{1-\zeta_2'^{-1}(\lambda+\theta(1-\lambda))} + \frac{1-\zeta_2'^{-1}}{\beta(1-\beta\rho)(1-\theta)}\right)\mathbb{E}_t\frac{1}{1-\zeta_2'^{-1}L^{-1}} \\
&\quad + (1-\lambda)g\rho^2\frac{\beta}{1-\beta\rho}\left((1-\beta) - \frac{\theta}{1-\theta}\frac{\kappa}{\sigma}\right)(1-\beta\rho(1-g))\mathbb{E}_t\frac{1}{1-\beta\rho(1-g)L^{-1}} \\
P_2'(L) &= (\zeta_1-\theta)\frac{1}{1-\zeta_1 L} \\
P_3'(L^{-1}) &= (1-\lambda)g\rho^2\Omega'\frac{\kappa}{\sigma}\left(\frac{\lambda}{\beta}\frac{\zeta_2'^{-1}(1-\zeta_2'^{-1}\theta)}{1-\zeta_2'^{-1}(\lambda+\theta(1-\lambda))} + \frac{1-\zeta_2'^{-1}}{\beta(1-\beta\rho)(1-\theta)}\right)\mathbb{E}_t\frac{1}{1-\zeta_2'^{-1}L^{-1}}.
\end{aligned}$$

Now define π_t^* and Δ_t exactly as in Appendix A:

$$\begin{aligned}
\pi_t^* &\equiv \frac{1}{1+(1-\lambda)g\rho^2\Theta'}\frac{\sigma}{\kappa}\left(P'(L)\mu_{t-1}^i + \mathbb{E}_t P_1'(L^{-1})\mu_{t+1}^i + P_2'(L)\mathbb{E}_{t-1}P_3'(L^{-1})\mu_t^i\right) \\
\Delta_t &\equiv \pi_t^* - \tilde{\pi}_t = \left(1+(1-\lambda)g\rho^2\Theta'\right)\frac{\sigma}{\kappa}\mu_t^i
\end{aligned}$$

so we have

$$\tilde{\pi}_t + \Delta_t = \pi_t^* \equiv \frac{1}{1+(1-\lambda)g\rho^2\Theta'}\left(P'(L)\Delta_{t-1} + \mathbb{E}_t P_1'(L^{-1})\Delta_{t+1} + P_2'(L)\mathbb{E}_{t-1}P_3'(L^{-1})\frac{\sigma}{\kappa}\Delta_t\right).$$

The optimal target criterion is then the same as with the one-step-ahead Phillips curve in Result (5), with the newly defined Θ', $P(L)$, $P_1'(L^{-1})$, $P_2'(L)$ and $P_3'(L^{-1})$.

As in the unconstrained case, the main differences between this optimal target criterion and the one with the one-step-ahead Phillips curve are that (i) the weight on lagged misses is larger, because commitments further in the future have a larger effect on current-period inflation via the Phillips curve, and (ii) the optimal preemptive response to expected future misses is larger, because the central bank has greater effect on future outcomes now that learners' beliefs about both inflation and the output gap affect inflation.

Since both the make-up and preemptive channels of policy are stronger, the interaction between them is also strengthened. But this interaction has offsetting effects, so the change in optimal policy is small.[29]

Table C2 compares the performance of the simple target criteria from Table 2 under the infinite-horizon Phillips curve with imperfect central bank information. The equivalent values for the one-step-ahead Phillips curve (i.e. Table D1) are in parentheses. The results are very similar.

Table C2: Performance of simple target criteria with infinite-horizon Phillips curve - imperfect information case

	Loss (% rel. to IT)	ψ	γ	γ_x	ψ_f	γ_f
5 yr	18.2 (24.6)	23.8 (27.4)				
IT	0 (0)	20.9 (20.9)				
IT + preemptive	-2.3 (-14.8)	27.1 (0.0)			10.04 (1.60)	1.28 (1.75)
2 yr	-5.9 (-4.5)	34.3 (27.9)				
1 yr	-12.3 (-15.7)	24.4 (15.6)				
PLT	-18.8 (-23.8)	9.2 (20.3)	1.00 (1.00)			
WAIT	-20.0 (-25.1)	9.6 (17.4)	0.80 (0.75)			
WAIT + preemptive	-20.0 (-25.1)	9.6 (17.4)	0.80 (0.75)		0.00 (0.00)	*undef.* (*undef.*)
WAIT + WAXT	-29.5 (-35.4)	10.0 (12.3)	0.97 (0.94)	0.69 (0.89)		
WAIT + WAXT + preemptive	-29.5 (-35.4)	10.0 (12.3)	0.97 (0.94)	0.69 (0.89)	0.00 (0.00)	*undef.* (*undef.*)
Optimal	-30.5 (-36.5)					

Notes Values in parentheses are the equivalent values with the one-step-ahead Phillips curve (from Table 3). Welfare losses are normalized relative to the welfare under IT. The parameters for each rule are optimized to minimize welfare. The optimal value of a parameter is undefined (labelled *undef.*) if the parameter has no effect on loss at the optimum.

C.4 ROBUSTLY OPTIMAL POLICY UNDER PARAMETER UNCERTAINTY WITH THE INFINITE-HORIZON PHILLIPS CURVE

Tables C3 and C4 replicate the robustly optimal policy exercise in Tables D2 and 4, but with the infinite-horizon Phillips curve. The results are very similar.

---

[29] The dynamic interaction between rational and learners' expectations in the Phillips curve tends to strengthen the effect of both make-up and preemptive policy relative to its cost (this is captured in the first term in Θ'). But the dynamic interaction between rational interest-rate expectations and learners' inflation and output-gap expectations lowers the benefits of make-up and preemptive policy relative to cost (this is the second and third terms in Θ').

Table C3: Policy under parameter uncertainty with infinite-horizon PC - unconstrained case

	Average Loss (% rel. to IT)	ψ	γ	γ_x	ψ_f	γ_f
5 yr	5.6×10^{11} (535.6)	1.0 (1.2)				
2 yr	187.8 (136.1)	1.1 (7.4)				
1 yr	86.3 (9.7)	2.6 (19.9)				
PLT	0.5 (4.2)	8.0 (10.9)	1 (1)			
IT	0 (0)	6.2 (13.7)				
IT + preemptive	-9.9 (-0.4)	6.4 (3.0)			10.37 (14.9)	0.00 (0.42)
WAIT	-14.6 (-8.1)	7.0 (9.2)	0.53 (0.56)			
WAIT + WAXT	-15.4 (-8.1)	6.9 (9.4)	0.73 (0.60)	0.29 (0.08)		
WAIT + preemptive	-16.5 (-8.9)	5.5 (7.2)	0.74 (0.67)		11.60 (7.4)	0.00 (0.51)
WAIT + WAXT + preemptive	-16.6 (-8.9)	6.6 (7.2)	0.73 (0.67)	0.17 (0.00)	6.91 (7.4)	0.00 (0.51)
Optimal	-21.2 (-13.7)					

Notes Values in parentheses are the equivalent values with the one-step-ahead Phillips curve (from Table D2). The parameters reports are set to maximize the average loss across all points in the grid. Average welfare losses are normalized relative to the welfare under IT.

Table C4: Policy under parameter uncertainty with infinite-horizon PC - imperfect information case

	Average Loss (% rel. to IT)	ψ	γ	γ_x	ψ_f	γ_f
5 yr	2.3×10^{11} (118.0)	1.0 (1.4)				
2 yr	58.9 (19.8)	0.8 (9.2)				
1 yr	27.0 (-16.6)	2.0 (20.7)				
IT	0 (0)	4.8 (20.9)				
PLT	-16.8 (-22.4)	7.4 (16.8)	1 (1)			
IT + preemptive	-17.3 (-1.3)	6.7 (14.9)			16.92 (13.0)	1.67 (1.52)
WAIT	-22.7 (-24.0)	6.1 (15.1)	0.66 (0.75)			
WAIT + WAXT	-26.9 (-31.0)	4.7 (14.5)	0.83 (0.87)	0.28 (0.69)		
WAIT + preemptive	-27.0 (-24.0)	12.1 (15.1)	0.65 (0.75)		14.53 (0.00)	0.00 (*undef.*)
WAIT + WAXT + preemptive	-33.5 (-31.1)	8.9 (13.3)	0.86 (0.87)	0.51 (0.66)	14.54 (6.8)	0.28 (0.00)
Optimal	-38.2 (-35.7)					

Notes Values in parentheses are the equivalent values with the one-step-ahead Phillips curve (from Table 4). The parameters reports are set to maximize the average loss across all points in the grid. Average welfare losses are normalized relative to the welfare under IT. The optimal value of a parameter is undefined (labelled *undef.*) if the parameter has no effect on loss at the optimum.

D APPENDIX D - ROBUSTLY OPTIMAL POLICY IN THE UNCONSTRAINED CASE

This section replicates the numerical results in Section 5 for the unconstrained case, where the central bank has perfect control of the output gap (i.e. perfect information and no ZLB).

Table D1 presents the optimal simple target criteria (defined in Table 2) under the baseline calibration. Optimal policy reduces the loss by 6.6% compared to the loss under the IT rules. These gains can be virtually entirely achieved by following a WAIT rule

Table D1: Performance of simple target criteria - unconstrained case

	Loss (% rel. to IT)	ψ	γ	γ_x	ψ_f	γ_f
5-yr AIT	96.5	31.3				
2-yr AIT	48.8	13.0				
1-yr AIT	7.7	21.5				
PLT	6.6	9.1	1			
IT	0	11.6				
IT + preemptive	-4.4	0.03			1.09	1.50
WAIT	-6.6	8.4	0.52			
WAIT + WAXT	-6.6	8.4	0.53	0.01		
WAIT + preemptive	-6.6	8.2	0.52		0.33	1.02
WAIT + WAXT + preemptive	-6.6	8.2	0.52	0.00	0.33	1.02
Optimal	-6.6					

Notes Welfare losses are normalized relative to the welfare under IT. The parameters for each rule are optimized to minimize welfare.

with a decay parameter of 0.52 (close to the share of rational agents, $\lambda = 0.5$). Placing additional weight on past output gaps or expected future inflation offers very small, if any, benefit. The importance of preemption is small when there is a reasonable share of rational forward-looking agents. The optimal preemptive respones to expected inflation can largely be achieved by a strong response to current and past inflation (especially because the cost-push shock is AR(1)). The price level targeting rule and especially the simple AIT rules perform poorly. Compared to the general case in Table 3, the optimal weight on lagged outcomes is smaller and IT performs better, which aligns with the analytical results in earlier sections.

Table D2 shows the average welfare and the optimized parameter values for each target criterion when the policymaker is uncertain about expectations formation (i.e. across a set of different λ and g values). It also shows the average loss that would be achieved if the optimal criterion was implemented at each point of the grid, which represent the minimum loss achievable. As in the imperfect information case (Table 4), WAIT rules perform well, although the decay parameters are a bit lower at 0.56 to 0.67. Unlike the imperfect information case, there is little gain from making up for past output gaps via a WAXT component. This is consistent with the analytical optimal target criterion in Result 18, in which no lagged output gap terms appear. The PLT rule performs poorly in the unconstrained case, as do the 5-yr and 2-yr AIT rules always perform poorly.

Note that under the infinite-horizon Phillips curve, the optimal make-up weights in the unconstrained case are a bit larger (see Appendix C), so optimal policy in the unconstrained case is not as different from the constrained case.

To show what lies behind the average losses in the table, Figure D1 plots the loss under a selection of target criteria from Tables D2 for different values of the expectation formation parameters λ and g. The parameters in the policy rules are held fixed at the robustly optimal values set to minimize average loss and not optimized for each λ and g shown in the figure. The blacked dash line shows the optimal target criteria optimized

Table D2: Policy under parameter uncertainty - unconstrained case

	Average Loss (% rel. to IT)	ψ	γ	γ_x	ψ_f	γ_f
5-yr AIT	535.6	1.2				
2-yr AIT	136.1	7.4				
1-yr AIT	9.7	19.9				
PLT	4.2	10.9	1			
IT	0	13.7				
IT + preemptive	-0.4	3.0			14.9	0.42
WAIT	-8.1	9.2	0.56			
WAIT + WAXT	-8.1	9.4	0.60	0.08		
WAIT + preemptive	-8.9	7.2	0.67		7.4	0.51
WAIT + WAXT + preemptive	-8.9	7.2	0.67	0.00	7.4	0.51
Opt	-13.7					

Notes The parameters reports are set to maximize the average loss across all points in the grid. Average welfare losses are normalized relative to the welfare under IT.

at each point to show how close the simple rules with the average coefficients can come to the fully optimal policy.

The IT rule performs well when λ is low, which is the expected result under adaptive learning. The PLT rule well when λ is high, close to the RE benchmarket, but unlike the constrained case, it performs quite poorly when λ is low. This is because nominal interest-rate expectations are irrelevant when the central bank can perfectly control the output gap, so make-up commitments have zero effect when $\lambda = 0$.

As in the constrained case, the WAIT rules perform consistently well across all parameter values. This is the case even though we have fixed the weight in the WAIT rules to the robustly optimal values in D2, so it is a fair comparison with the IT and PLT rules. Only in the optimal policy benchmark do we allow the policy rule parameters to change as we vary λ and g. The performance of the fixed-coefficient simple WAIT rules is remarkably close to the optimal benchmark, illustrating the robustness of this form of FAIT.

Figure D1: Loss under different parameter values - unconstrained case

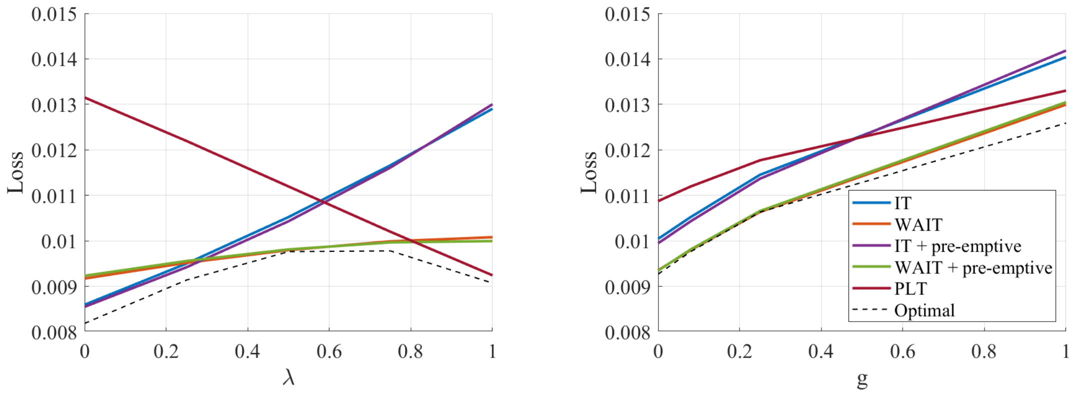

Notes: The parameters in the policy rules are fixed at the robustly optimal values presented in Table D2. For the other parameters (other than the one changing), we set $\lambda = 0.5$, $g = 0.08$, $\rho = 0.95$, $\kappa = 0.1$, $\sigma = 1$, $\rho_{rn} = \rho_u = 0.5$, $\sigma_{rn} = 1$ and $\sigma_u = 0.1$.

REFERENCES

ANGELETOS, G.-M., J. GUERREIRO, AND D. R. ZHANG (2025): "From RANK to HANK, without FIRE," Discussion paper, National Bureau of Economic Research.

ANGELETOS, G.-M., AND Z. HUO (2021): "Myopia and anchoring," *American Economic Review*, 111(4), 1166–1200.

ANGELETOS, G.-M., AND C. LIAN (2018): "Forward guidance without common knowledge," *American Economic Review*, 108(9), 2477–2512.

ARIAS, J., M. BODENSTEIN, H. T. CHUNG, T. DRAUTZBURG, AND A. RAFFO (2020): "Alternative strategies: How do they work? How might they help?," Discussion paper, Finance and Economics Discussion Series 2020-066. Washington: Board of Governors of the Federal Reserve System.

BECKERS, B., AND A. BRASSIL (2022): "Inflation Expectations in Australia," *Australian Economic Review*, 55(1), 125–135.

BENCHIMOL, J., AND L. BOUNADER (2023): "Optimal monetary policy under bounded rationality," *Journal of Financial Stability*, 67, 101151.

BERNANKE, B. S., M. T. KILEY, AND J. M. ROBERTS (2019): "Monetary Policy Strategies for a Low Rate Environment," *AEA Papers and Proceedings*, 109, 421–426.

BIANCHI, F., C. ILUT, AND H. SAIJO (2024a): "Diagnostic business cycles," *Review of Economic Studies*, 91(1), 129–162.

BIANCHI, F., C. L. ILUT, AND H. SAIJO (2024b): "Smooth diagnostic expectations," Discussion paper, National Bureau of Economic Research.

BLAKE, A. P. (2001): "A timeless perspective on optimality in forward-looking rational expectations models," Discussion paper, Citeseer.

BORDALO, P., N. GENNAIOLI, AND A. SHLEIFER (2018): "Diagnostic expectations and credit cycles," *The Journal of Finance*, 73(1), 199–227.

BRANCH, W. A., AND B. MCGOUGH (2009): "A New Keynesian model with heterogeneous expectations," *Journal of Economic Dynamics and Control*, 33(5), 1036–1051.

BRASSIL, A., Y. HAIDARI, J. HAMBUR, G. NOLAN, AND C. RYAN (2024): "How Do Households Form Inflation and Wage Expectations?," Discussion Paper 2024-07, Reserve Bank of Australia.

BROCK, W. A., AND C. H. HOMMES (1997): "A rational route to randomness," *Econometrica*, pp. 1059–1095.

BUDIANTO, F., T. NAKATA, AND S. SCHMIDT (2023): "Average inflation targeting and the interest rate lower bound," *European Economic Review*, 152, 104384.

CECIONI, M., G. COENEN, R. MOTTO, H. LE BIHAN, V. AJEVSKIS, U. ALBERTAZZI, N. GILBERT, A. AL-HASCHIMI, S. GOMES, F. BORNEMANN, ET AL. (2021): "The Ecb's Price Stability Framework: Past Experience, and Current and Future Challenges," Discussion paper, ECB Occassional Paper Series, 269.

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.

Cole, S. J., and E. Martínez-García (2023): "The effect of central bank credibility on forward guidance in an estimated New Keynesian model," *Macroeconomic Dynamics*, 27(2), 532–570.

Di Bartolomeo, G., M. Di Pietro, and B. Giannini (2016): "Optimal monetary policy in a New Keynesian model with heterogeneous expectations," *Journal of Economic Dynamics and Control*, 73, 373–387.

Dorich, J., R. R. Mendes, and Y. Zhang (2021): "The Bank of Canada's "horse race" of alternative monetary policy frameworks: Some interim results from model simulations," Discussion paper, Bank of Canada.

Duineveld, S. (2021): "Standardized Projection Algorithms to Solve Dynamic Economic Models," .

Dupraz, S., H. Le Bihan, and J. Matheron (2024): "Make-up strategies with finite planning horizons but infinitely forward-looking asset prices," *Journal of Monetary Economics*, 143, 103542.

Dupraz, S., and M. Marx (2023): "Anchoring Boundedly Rational Expectations," Discussion paper.

Eggertsson, G. B., and M. Woodford (2003): "Optimal monetary policy in a liquidity trap," *National Bureau of Economic Research.*

Eo, Y., and D. Lie (2020): "Average inflation targeting and interest-rate smoothing," *Economics Letters*, 189, 109005.

Eskelinen, M., C. Gibbs, and N. McClung (2024): "Resolving New Keynesian Puzzles," *Bank of Finland Research Discussion Paper*, (5).

Eusepi, S., M. Giannoni, and B. Preston (forthcoming): "The Short-run Policy Constraints of Long-run Expectations," *Journal of Political Economy.*

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

Eusepi, S., C. Gibbs, and B. Preston (2024): "Monetary policy as insurance," *working paper, University of Sydney.*

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., C. G. Gibbs, and B. McGough (2025): "A Unified Model of Learning to Forecast," *American Economic Journal: Macroeconomics*, 17(2), 101–133.

Evans, G. W., and S. Honkapohja (2001): "Learning and expectations in macroeconomics," in *Learning and Expectations in Macroeconomics.* Princeton University Press.

——— (2003): "Expectations and the stability problem for optimal monetary policies," *The Review of Economic Studies*, 70(4), 807–824.

Evans, G. W., and B. McGough (2010): "Implementing optimal monetary policy in New-Keynesian models with inertia," *The BE Journal of Macroeconomics*, 10(1).

Farhi, E., and I. Werning (2019): "Monetary policy, bounded rationality, and incomplete markets," *American Economic Review*, 109(11), 3887–3928.

Gabaix, X. (2020): "A behavioral New Keynesian model," *American Economic Review*, 110(8), 2271–2327.

Galí, J. (2015): *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press.

Gasteiger, E. (2014): "Heterogeneous expectations, optimal monetary policy, and the merit of policy inertia," *Journal of Money, Credit and Banking*, 46(7), 1535–1554.

——— (2021): "Optimal constrained interest-rate rules under heterogeneous expectations," *Journal of Economic Behavior & Organization*, 190, 287–325.

Giannoni, M., and M. Woodford (2005): "Optimal inflation-targeting rules," in *The Inflation-Targeting Debate*, ed. by B. Bernanke, and M. Woodford, pp. 93–172. University of Chicago Press.

Giannoni, M. P. (2014): "Optimal interest-rate rules and inflation stabilization versus price-level stabilization," *Journal of Economic Dynamics and Control*, 41, 110–129.

Giannoni, M. P., and M. Woodford (2017): "Optimal target criteria for stabilization policy," *Journal of Economic Theory*, 168, 55–106.

Gibbs, C. G. (2017): "Forecast combination, non-linear dynamics, and the macroeconomy," *Economic Theory*, 63(3), 653–686.

Gibbs, C. G., and M. Kulish (2017): "Disinflations in a model of imperfectly anchored expectations," *European Economic Review*, 100, 157–174.

Hagenhoff, T. (2021): "Optimal monetary policy, heterogeneous expectations and consumption dispersion," *Working Paper, University of Bamberg.*

Hebden, J., E. P. Herbst, J. Tang, G. Topa, and F. Winkler (2020): "How Robust Are Makeup Strategies to Key Alternative Assumptions?," *Financial and Economics Discussion Series 2020-069.*

Honkapohja, S., and N. McClung (2024): "Average inflation targeting in Calvo model with imperfect knowledge and learning," Discussion paper, Bank of Finland Research Discussion Papers.

Honkapohja, S., and K. Mitra (2020): "Price level targeting with evolving credibility," *Journal of Monetary Economics*, 116, 88–103.

Hunter, S. (2024): "Inflation Expectations – Why They Matter and How They Are Formed," Speech to Citi Australia and New Zealand Investment Conference.

Jensen, C., and B. T. McCallum (2002): "The non-optimality of proposed monetary policy rules under timeless perspective commitment," *Economics Letters*, 77(2), 163–168.

Jia, C., and J. C. Wu (2023): "Average inflation targeting: Time inconsistency and ambiguous communication," *Journal of Monetary Economics*, 138, 69–86.

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.

MCKAY, A., AND C. K. WOLF (2023): “Monetary policy and inequality,” *Journal of Economic Perspectives*, 37(1), 121–144.

MOLNÁR, K., AND S. SANTORO (2014): “Optimal monetary policy when agents are learning,” *European Economic Review*, 66, 39–62.

ORPHANIDES, A., AND J. C. WILLIAMS (2007): “Robust monetary policy with imperfect knowledge,” *Journal of monetary Economics*, 54(5), 1406–1435.

POWELL, J. H. (2025): “Monetary Policy and the Fed’s Framework Review,” Speech at “Labor Markets in Transition: Demographics, Productivity, and Macroeconomic Policy,” an economic symposium sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming, August 22, 2025. PDF: https://www.federalreserve.gov/newsevents/speech/files/powell20250822.pdf (accessed 2025-12-30).

PRESTON, B. (2005): “Learning about Monetary Policy Rules when Long-Horizon Expectations Matter,” *International Journal of Central Banking*, 1(2).

——— (2006): “Adaptive learning, forecast-based instrument rules and monetary policy,” *Journal of Monetary Economics*, 53(3), 507–535.

SVENSSON, L. E., AND M. WOODFORD (2005): “Implementing optimal policy through inflation-forecast targeting,” in *The Inflation-Targeting Debate*, ed. by B. Bernanke, and M. Woodford, pp. 19–92. University of Chicago Press.

SVENSSON, L. E. O. (2003): “What is wrong with Taylor rules? Using judgment in monetary policy through targeting rules,” *Journal of Economic Literature*, 41(2), 426–477.

WOODFORD, M. (2003): *Interest and prices.* Princeton University Press Princeton.

——— (2010): “Optimal monetary stabilization policy,” *Handbook of monetary economics*, 3, 723–828.