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

Anthony Brassil† Christopher G. Gibbs‡ Callum Ryan*

17 May 2025

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 behavioural expectations that nests rational and adaptive learning beliefs as special cases, and approximates several alternative expectations theories - naive, myopic, and level-k. We show that optimal policy is well-approximated by a single policy framework across all cases: Flexible Average Inflation Targeting. Policy targets a weighted average of inflation plus adjustments for belief persistence and constraints faced by the central bank (information frictions 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, Matthew Read, and Jonathan Hambur for helpful comments and suggestions. We thank seminar participants at the 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, and the 2023 Reserve Bank of Australia's Quantitative 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 alternative expectation theories. We use this general environment to derive optimal target criteria under commitment for a welfare maximizing central bank 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 expectation formation process and provide recommendations that are independent of how expectations are formed.

We find that the optimal target criterion is always a form of Flexible Average Inflation Targeting (FAIT). We show that the *average* in FAIT should be a weighted average of current and past inflation with declining weights over time. Monetary policymakers should respond most forcefully to current inflation, while also committing to partially 'make up' for past inflation misses. But the weighted average inflation target is also *flexible* because when a policymaker's ability to achieve this target is in question, it should seek to influence expectations more aggressively than usual. A policymaker should commit to stronger-than-usual make-up policy when they are unable to control the output gap to implement the current target, such as when it has imperfect information about the state of the economy or the ZLB binds. In this case, policy should seek to make up for past output gaps as well as past inflation. In addition, policymakers should preemptively deviate from the target to steer expectations favourably when they anticipate future constraints, such as drifting inflation expectations or the ZLB. Especially with respect to the ZLB constraint, policy is a "use it or lose it" proposition for all the expectations theories that we capture.

Central to understanding the advantage of a FAIT policy is recognizing a symmetry between (i) the history dependence that a central bank should optimally commit to when agents' expectations are near the full information rational expectations (FIRE) benchmark, and (ii) the history dependence in agents' own expectations that emerges under adaptive learning. In the former case, it is optimal to manage agents' expectations by committing to make up for past misses to targets by setting current policy based on a weighted average of past outcomes. In the latter case, it is optimal to respond preemptively to expected future outcomes by influencing agents' expectations, which are themselves a weighted average of past outcomes.In both cases, optimal policy requires flexible deviations from the weighted average target to more aggressively influence expectations when policy is constrained. Therefore, despite the different rationales for policy in each

case, the optimal policy shares several key features.

Our agnostic expectations framework reconciles two seemingly divergent views on optimal policy under bounded rationality. On the one hand, Eusepi and Preston (2018) and Eusepi, Giannoni, and Preston (2018) show that optimal monetary policy under adaptive learning shares many features with the full information optimal commitment policy – although it is far less effective than in the former – despite little ability for policy commitments to shape beliefs. On the other hand, Angeletos and Lian (2018) and Gabaix (2020) argue that when agents are myopic or engage in low levels of reasoning, optimal policy must depart from the full-information commitment benchmark, since commitments have only a weak effect on expectations. We show that these conflicting conclusion actually point to the same policy framework: FAIT. Moreover, many other varieties of novel bounded rationality assumptions, which we capture in our model, point to the same optimal policy.

We model expectations as a weighted average of adaptive learning and FIRE, where the rational expectations take into account the learning expectations, and where the desired path of the policy rate is credible and known to all. We show that by varying the weight placed on rational expectations relative to learning and by varying the speed of the learning, that we capture relevant dimensions of several types of expectation formation models within a standard New Keynesian economy. The expectation models include adaptive learning as in Evans and Honkapohja (2001) and Preston (2005), behavioral expectations as in Gabaix (2020), and key elements of level-k reasoning as in Angeletos and Lian (2018), Farhi and Werning (2019), Angeletos and Huo (2021) and Evans, Gibbs, and McGough (2025). Each of these different expectations models lies as a special case that we can recover or closely approximate within our more general setting, and allows for characterization of optimal policy in novel settings. The general target criteria that we derive constitute a formal theory for a FAIT policy framework.

We demonstrate the robustness of FAIT policy by running two horse races comparing a range of different inflation and output gap policy targets. We first show that we can approximate the optimal target criteria using simple weighted average criteria over past inflation and output gap outcomes, which are superior to price-level targeting, inflation targeting, or arithmetic average inflation targeting using a fixed window when agents are boundedly rational. 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 to maximize average welfare across different economies with expectations ranging from FIRE to full adaptive learning. We find that simple implementations of FAIT that target geometrically declining averages of inflation (WAIT) and the output gap (which captures some aspects of flexibility required in optimal policy) are the most robust, generating similar losses to the fully optimal policy regardless of how expectations are formed.

Literature review. We take a novel 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 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.[1] Interest rate rules are model-specific and the mapping between policy framework and an interest-rate rule is not clearcut. In some cases, different rules can generate the same equilibrium outcomes. For example, Eskelinen, Gibbs, and McClung (2024) show that unconditional optimal policy in standard New Keynesian model can be implemented with an interest rate rule that responds to a lagged interest rate term, similar to interest rate smoothing, or to a weighted average of past inflation, similar to some average inflation targeting specifications in the literature. In other cases, slightly different rules generate drastically different equilibria. For example, Honkapohja and McClung (2024) point out several examples of the fragility of interest rate rule specification that try to capture average inflation targeting (AIT) under RE and adaptive learning. They show that small changes to the number of periods included in the inflation average can lead to drastically different policy outcomes. 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 hopes 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 policy rate necessary to achieve their goals. Communication of policy in such a framework is, therefore, forward-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.

[1]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.

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.

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.[2] We can approximate the key elements of these models for expectations within our general framework, and extend the optimal policy analysis 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 a 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) analyse make-up policy when agents have finite

[2]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 price and consumption dispersion, altering the micro-founded loss function. We are instead interested in capturing how a broad range of different expectations models affect optimal aggregate outcomes, so we assume the standard loss function.

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 optimising 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.

Our general framework also lets us approximate other models for expectations that feature some kind of myopic and a backward-looking learning process, for which optimal target criteria have not been derived. For example, Angeletos and Lian (2018), Farhi and Werning (2019) and Angeletos, Huo, and Sastry (2021) all propose models with some combination of level-k reasoning, heterogeneous private information, and/or learning. Importantly, we characterize the unconstrained, imperfect central bank information, and ZLB cases under all of these different expectations assumptions. Because our framework nests a range of different specifications, we are also able to find the robustly optimal policy framework, in circumstances where policymakers do not know the true model for expectations.

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) with sticky prices and monopolistically competitive firms. However, we depart from the standard environment by assuming that expectations need not be rational and may encompass a wide range of beliefs formation strategies. To accommodate this generality, we log-linearly approximate aggregate output and inflation dynamics as

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

$$\pi_t = \beta \hat{\mathbb{E}}_t \pi_{t+1} + \kappa x_t + u_t, \quad (2)$$

where x_t is the output gap, π_t is inflation, i_t is the nominal interest rate controlled by the central bank, r_t^n is an exogenous real rate of interest (which could reflect movements in demand and productivity), u_t is a cost-push shock, and $\hat{\mathbb{E}}_t$ represents potentially non-rational expectations. Equation (1) is the infinite-horizon IS curve, which captures household demand under arbitrary beliefs. Under FIRE, this equations collapses to the familiar two-period representation. However, the form here distinguishes the effect of longer-term income and real interest-rate expectations on consumption decisions, which are critical when the central bank is constrained in the present.[3]

For Equation (2), the Phillips curve, we use the common two-period approximation, which approximates the effect on pricing decisions of expected inflation and output gaps in all future periods using just the one-period-ahead inflation expectation. Appendix C provides the full infinite-horizon form of the Phillips curve (also called the anticipated-utility solution in the adaptive learning literature; under FIRE, it is equivalent to the two-period form), and replicates all of our main results in that setting. We choose to use the two-period approximation for the Phillips curve in the main analysis because the key results are not sensitive to the choice (in contrast to the choice of IS curve representation) and are easier to interpret. It also allows us to more closely nest/approximate the reduced form Phillips curve formulations of the myopia and anchoring model of Angeletos and Huo (2021), the hybrid NK Phillips curve, and older works in the adaptive learning literature such Evans and Honkapohja (2003) and Molnár and Santoro (2014).

Beliefs of the private sector. We put forward a general notion 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 distinct expectation theories put forward in the literature.

We assume that some fraction $\lambda \in [0,1]$ of agents possess FIRE, and a fraction $1-\lambda$ are *adaptive learners*, in that they use a statistical model to forecast future outcomes by extrapolating from past observations.[4] Aggregate expectations of future inflation and

[3] This is especially important because we are going to 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.

[4] We focus on how forward- or backward-looking monetary policy should be to achieve optimal aggreg-

output gaps are

$$\hat{\mathbb{E}}_t \pi_T = \lambda \mathbb{E}_t \pi_T + (1-\lambda)\mathbb{E}_t^l \pi_T \tag{3}$$

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

where $\mathbb{E}_t$ is the FIRE forecast and $\mathbb{E}_t^l$ denotes the learners' forecast.

We assume that adaptive learners forecast inflation and the output gap using an unobserved components state space model, which they estimate 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. Learners 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. Learners use this estimate to forecast future variables z_T, $T \geq t+1$. Their expectations for π_T and x_T, $T \geq t+1$, are therefore given by

$$\mathbb{E}_t^l \pi_T = \rho^{T-t} \omega_{t-1}^{\pi} \tag{5}$$

$$\mathbb{E}_t^l x_T = \rho^{T-t} \omega_{t-1}^{x}. \tag{6}$$

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

$$\omega_t^{\pi} = \rho \omega_{t-1}^{\pi} + \rho g \left(\pi_t - \omega_{t-1}^{\pi}\right) \tag{7}$$

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

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.

For expectations of the nominal interest rate, we assume that everyone has rational expectations over its path.[5] 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.[6] If the policymaker communicates its expected path for the nominal

ate outcomes and do not consider distributional consequences that may occur due to heterogeneity in expectations. We maintain a representative agent assumption where the representative decision-maker takes a weighted average of forecasts from two different models: the correct structural model, and a reduced-form model which is re-estimated each period. Therefore, the representative agent assumption is maintained. This is the approach taken in Gibbs (2017) and Gibbs and Kulish (2017).

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

[6]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,

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. Even if the central bank's communication of its interest-rate expectations is not perfect, the effect of interest-rate expectations on household decisions will often flow through fixed rates and asset prices, which are determined by sophisticated financial-market participants who are highly forward-looking and likely to form similar expectations to the central bank.[7]

Beliefs of the policymaker. We assume throughout that the central bank possesses FIRE and seeks to maximize the standard welfare function

$$\mathcal{W} \equiv -\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\pi_t^2 + \alpha x_t^2\right) \tag{9}$$

which is a quadratic approximation to household welfare when $\alpha = \kappa/\theta$. When choosing policy, the policymaker takes household decisions and beliefs as constraints:

$$\pi_t = \lambda\beta\mathbb{E}\pi_{t+1} + (1-\lambda)\beta\rho\omega_{t-1}^{\pi} + \kappa x_t + u_t \tag{10}$$

$$\begin{aligned} x_t = (1-\beta)\sum_{T=t}^{\infty} \beta^{T-t}\left(\lambda\mathbb{E}_t x_{T+1} + (1-\lambda)\rho^{T+1-t}\omega_{t-1}^{x}\right) \\ -\frac{1}{\sigma}\sum_{T=t}^{\infty}\beta^{T-t}\left(\mathbb{E}_t i_T - \lambda\mathbb{E}_t\pi_{T+1} - (1-\lambda)\rho^{T+1-t}\omega_{t-1}^{\pi} - \mathbb{E}_t r_T^n\right) \end{aligned} \tag{11}$$

$$\omega_t^{\pi} = \rho\omega_{t-1}^{\pi} + \rho g\left(\pi_t - \omega_{t-1}^{\pi}\right) \tag{12}$$

$$\omega_t^{x} = \rho\omega_{t-1}^{x} + \rho g\left(x_t - \omega_{t-1}^{x}\right). \tag{13}$$

The policymaker's problem. 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 maximises the welfare function (9) subject to the constraints (10) – (13) 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

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).

[7] There are clearly interesting questions about what occurs when the central bank lacks credibility over the path of interest rates that it plans to implement. See, for example, Eusepi, Giannoni, and Preston (forthcoming) and Eusepi, Giannoni, and Preston (2018). Our results represent a best case scenario with respect to the central bank credibility regarding the policy rate path. In Appendix B, we discuss how our results might be affected if policy-rate expectations are not entirely rational.

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.[8]

Nesting/Approximating other models. The payoff for the particular way we model beliefs is in nesting the structural equations implied by most of the prominent expectations theories in the literature. For example, when

- FIRE ($\lambda = 1$): the system reduces to structural equations implied by full information rational expectations:

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

- Myopia and Level-k ($0 < \lambda < 1$ and $g = 0$): the system reduces to a model that is "over-discounted" such as in Gabaix (2020) or Angeletos and Lian (2018):

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

- Hybrid NK Phillips curve ($0 < \lambda < 1$ and $g = 1$): the system reduces to a model with a hybrid New Keynesian Phillips curve, or one that features elements of myopia and anchoring as in Angeletos and Huo (2021):

$$\begin{aligned} \pi_t &= \lambda\beta \mathbb{E}_t \pi_{t+1} + (1-\lambda)\beta\rho\pi_{t-1} + \kappa x_t + u_t \\ x_t &= (1-\beta)\sum_{T=t}^{\infty} \beta^{T-t} \left(\lambda \mathbb{E}_t x_{T+1} + (1-\lambda)\rho^{T+1-t} x_{t-1}\right) \\ &\quad -\frac{1}{\sigma}\sum_{T=t}^{\infty} \beta^{T-t} \left(\mathbb{E}_t i_T - \lambda \mathbb{E}_t \pi_{T+1} - (1-\lambda)\rho^{T+1-t}\pi_{t-1} - \mathbb{E}_t r_T^n\right) \end{aligned}$$

- Adaptive learning ($\lambda = 0$ and $0 < g < 1$): the system reduces to a model of adaptive learning:

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

By nesting these competing assumptions, we can compare optimal policy within a shared

[8] 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.

framework.[9]

3 Optimal Policy: The Unconstrained Case

We start with the unconstrained optimal policy problem given by equations (9) - (13). 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:

$$\frac{\alpha}{\kappa}x_t = -\frac{1}{1-\lambda L}\pi_t + \rho g \frac{1}{1-\lambda L}\mu_t^{\omega\pi} \tag{14}$$

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}:

$$\mu_t^{\omega\pi} = \beta\rho(1-g)\mathbb{E}_t\mu_{t+1}^{\omega\pi} - \beta^2(1-\lambda)\rho\mathbb{E}_t\left(\frac{1}{1-\lambda L}\pi_{t+1} - g\frac{1}{1-\lambda L}\mu_{t+1}^{\omega\pi}\right). \tag{15}$$

Solving (15) for $\mu_t^{\omega\pi}$ gives

$$\mu_t^{\omega\pi} = -(1-\lambda)\rho\Omega\frac{1-\lambda L}{1-\zeta_1 L}\mathbb{E}_t\frac{1}{1-\zeta_2^{-1}L^{-1}}\frac{1}{1-\lambda L}\pi_{t+1} \tag{16}$$

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

[9] 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 $g > 0$.

associated with (15).[10]

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

$$\frac{\alpha}{\kappa}x_t = -\underbrace{\frac{1}{1-\lambda L}\pi_t}_{WAIT} - \underbrace{(1-\lambda)g\rho^2\Omega\frac{1}{1-\zeta_1 L}\mathbb{E}_t\frac{1}{1-\zeta_2^{-1}L^{-1}}\frac{1}{1-\lambda L}\pi_{t+1}}_{Preemption}, \tag{17}$$

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 preemption term represent the management of backward-looking adaptive learning expectations and is present if $g > 0$. The policymaker forms a beliefs about how inflation expectations are likely to evolve in the future given current and past misses in inflation, which are propagated through the adaptive learning component of beliefs. Drifts in inflation change the trade-off between future inflation and output gaps. Additional policy today can preempt future undesirable movements in this trade-off, or engineer more favourable ones.

A key feature of the optimal target criterion (17) 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. Optimal policy in this case is WAIT + preemption. The table, however, 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.

[10]That is, ζ_1 and ζ_2 are the solutions to the quadratic equation

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

For reasonable calibrations, ζ_1 is generally very close to λ and reflects the effectiveness of make-up commitments on current-period inflation. The reciprocal of the explosive root ζ_2^{-1} reflects the persistence of the effect of current policy on future inflation (via beliefs). It is decreasing in λ and g, and increasing in ρ.

Table 1: Special Cases

Cases	Target criteria	Policy
FIRE, $\lambda = 1$	$\frac{\alpha}{\kappa} x_t = -p_t$	PLT
Hybrid NK PC, $0 < \lambda < 1$, $g = 1$	Equation (17)	WAIT + premptive
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. No terms from optimal target criteria drop out for Hybrid NK PC parameters. PLT stands for 'price-level targeting', IT stands for '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. (16) 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

$$\mu_t^{\omega\pi} = (1-\lambda)\left(a_x x_t - a_\omega \rho g \pi_t - a_\omega \rho (1-g)\omega_{t-1}^{\pi} - a_u u_t\right) \tag{18}$$

for some positive constants a_x, a_ω and a_u. 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 (18), we can write the optimal target criterion (17) (in equilibrium) as

$$\frac{\alpha}{\kappa} x_t = -\phi \frac{1}{1-\gamma L}\left(\pi_t + \nu_\omega \omega_{t-1}^{\pi} + \nu_u u_t\right) \tag{19}$$

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 (19) 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.

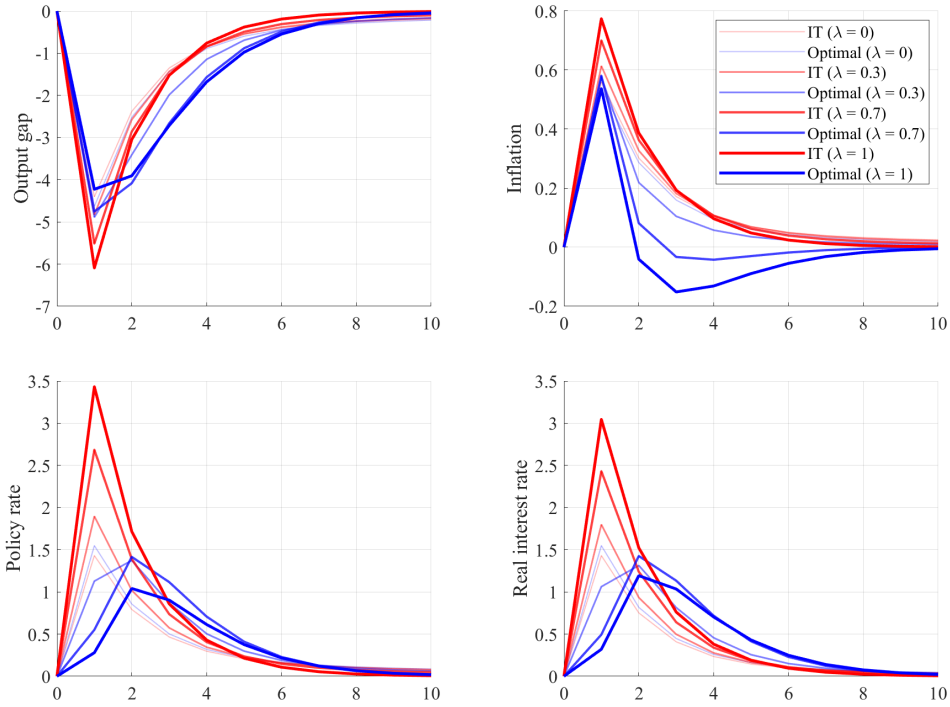
In the low λ cases, policy reflects the different way a shock propagates when there is learning. Backward-looking agents' expectation are less sensitive to the shock on impact because their expectations are anchored to the past. This means the shock has less of an impact initially, which allows for a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

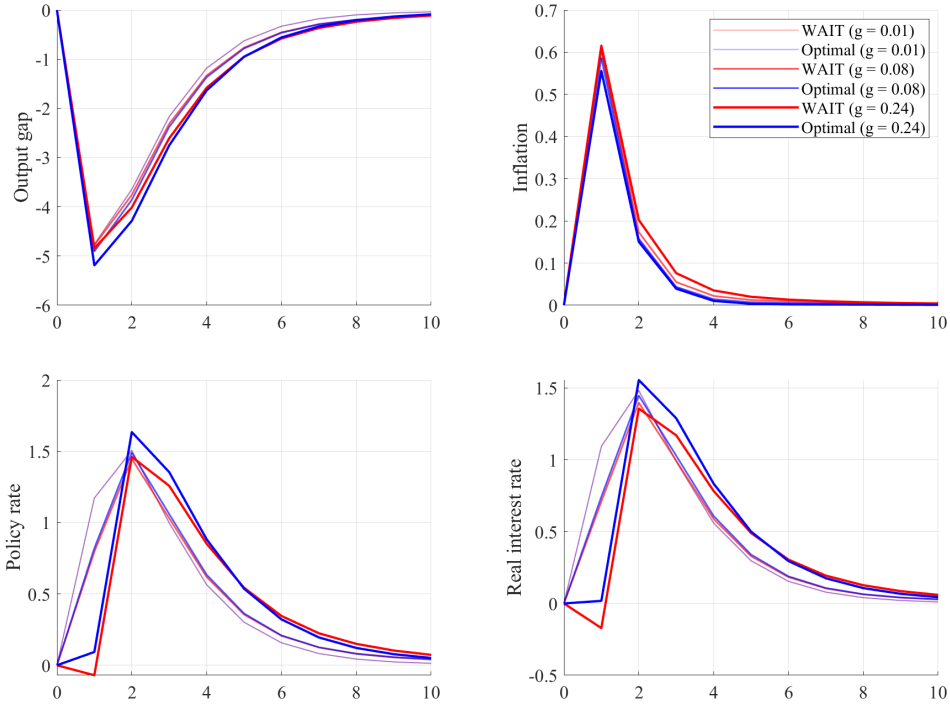
There are of course other scenarios where even with moderate values of λ that preemption becomes more central. For example, when shocks are very persistent or the central bank expects to be constrained in the future such as when the ZLB binds, which we explore in Section 4.

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 $\lambda = 0.5$, $\rho = 0.95$=, $\sigma = 1$, $\kappa = 0.1$, $\alpha = \kappa/7.87$

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. The reason it occurs is that the expectations of adaptive learners in our setup 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. It turns out that that weighted average implied by the preemption component of optimal policy in equilibrium shares many features with the original WAIT component of optimal policy. In fact, it is possible to rewrite optimal policy as a simple WAIT policy with different weights that share a similar geometric decay over time.

To illustrate, substitute the weighted average expression above for adaptive learning beliefs into the rewritten optimal target criterion (19) and rearrange:

$$x_t = -\frac{\kappa\phi}{\alpha}\left(\frac{\psi}{1-\gamma L} + \frac{1-\psi}{1-\rho(1-g)L}\right)\pi_t - \frac{\kappa\phi}{\alpha}\nu_u\frac{1}{1-\gamma L}u_t \tag{20}$$

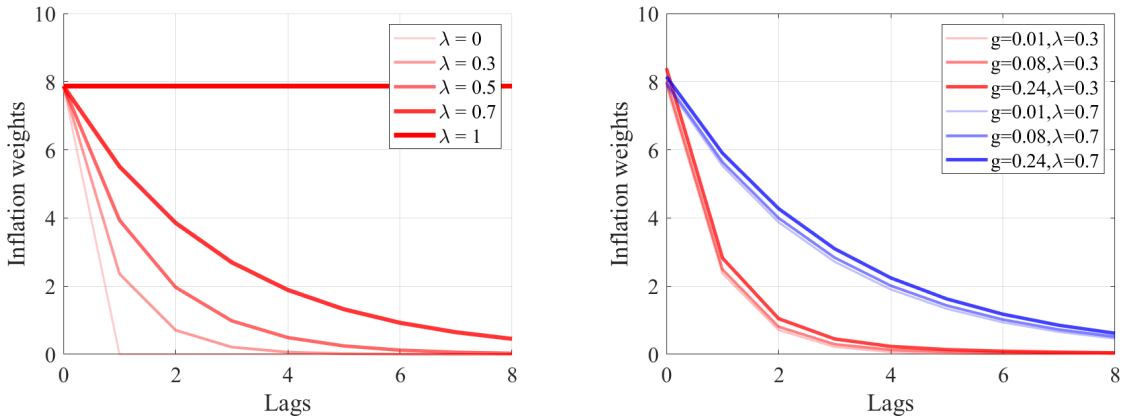
where $\psi \equiv 1 - \frac{\rho g}{\rho(1-g)-\gamma}\nu_\omega$. The average inflation measure that a policymaker should target is a weighted average of two exponentially-weighted averages of current and past inflation.

Figure 3 illustrates the similarities in the weights by plotting the term in front of π_t from (20) 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.

In addition, it is important to note that the profile of optimal weights are very different

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 $\rho = 0.95$, $\beta = 0.99$, $\sigma = 0.5$, $\kappa = 0.1$, $\alpha = \kappa/7.87$

to the arithmetic averages that many researcher have studies when modeling average inflation targeting such as in Eo and Lie (2020) 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 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.

$$i_t = i_{t|t}. \tag{21}$$

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 (9) subject to the constraints (10) to (13) imposed by private-sector equilibrium behaviour and the informational constraint (21).[11]

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

$$\mu_t^i = \frac{1}{\sigma}\sum_{j=0}^{t}\mu_{t-j}^x \tag{22}$$

$$\mu_{t|t}^i = 0 \tag{23}$$

where μ_t^i is the Lagrange multiplier on the information constraint (21). Condition (22) 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 (23) 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 (22) conditional on I_t^{CB} and using $\mu_{t-j|t}^x = \mu_{t-j}^x$ for all $j \geq 1$ gives

$$\mu_{t|t}^x = -\sum_{j=1}^{t}\mu_{t-j}^x, \tag{24}$$

which implies that the central bank commits to an i_t that balances the estimated costs

[11] 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.

of an off-target output gap in period t, $\mu^x_{t|t}$, against the reduced losses in the periods before t due to expectations of i_t, $-\sum_{j=1}^{t}\mu^x_{t-j}$. Equation (24), 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 $\mu^x_t = \sigma\left(\mu^i_t - \mu^i_{t-1}\right)$ from (22) gives

$$\frac{\alpha}{\kappa}x_t + \frac{\sigma}{\kappa}\mu^i_t = -\frac{1}{1-\lambda L}\pi_t + \rho g\frac{1}{1-\lambda L}\mu^{\omega\pi}_t + \left(\frac{\lambda}{\beta}\frac{\kappa}{\sigma}\frac{1}{1-\lambda L} + \frac{\lambda}{\beta}(1-\beta) + 1\right)\frac{\sigma}{\kappa}\mu^i_{t-1} + \frac{\rho g}{\kappa}\mu^{\omega x}_t, \tag{25}$$

which generalizes equation (14).

The information constraint implies three additional terms. The new term in μ^i_t 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^i_{t-1} - 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

$$\begin{aligned}\mu^{\omega\pi}_t = &-(1-\lambda)\rho\Omega\frac{1-\lambda L}{1-\zeta_1 L}\mathbb{E}_t\frac{1}{1-\zeta_2^{-1}L^{-1}}\frac{1}{1-\lambda L}\pi_{t+1}\\ &+(1-\lambda)\rho\Omega\left(\frac{\lambda}{\beta}\frac{1}{1-\zeta_2^{-1}\lambda}\frac{1}{1-\zeta_1 L} - \frac{1}{\beta(1-\beta\rho)}\frac{1-\lambda L}{1-\zeta_1 L}\right)\mu^i_t \end{aligned}\tag{26}$$

$$\mu^{\omega x}_t = -(1-\lambda)\frac{(1-\beta)\rho\beta}{1-\beta\rho}\sigma\mu^i_t. \tag{27}$$

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

Substituting the Lagrange multipliers (26) and (27) into (25), we can derive the optimal target criterion under information constraints.

Result 4 (Information constrained optimal target criterion)**.** *Recall from eq. (17) 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

$$\tilde{\pi}_{t|t} = \pi_t^*, \quad \textit{where} \quad \pi_t^* \equiv \frac{1}{1+(1-\lambda)g\rho^2\Theta} P(L)\Delta_{t-1} \quad \textit{and} \quad \Delta_t \equiv \pi_t^* - \tilde{\pi}_t \tag{28}$$

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*

$$\begin{aligned} P(L) \equiv & \underbrace{\frac{\lambda}{\beta}\frac{\kappa}{\sigma}\frac{1}{1-\lambda L} + \frac{\lambda}{\beta}(1-\beta) + 1}_{\textit{Make-up policy}} + \\ & \underbrace{(1-\lambda)g\rho^2\Omega\frac{\kappa}{\sigma}\left(\frac{\lambda}{\beta}\frac{1}{1-\zeta_2^{-1}\lambda}\left(\frac{\lambda}{1-\lambda L}+\zeta_1\right) - \frac{\zeta_1}{\beta(1-\beta\rho)}\right)\frac{1}{1-\zeta_1 L}}_{\textit{Interaction of preemption and make-up}}. \end{aligned} \tag{29}$$

The optimal target criterion under imperfect information (28) generalizes the unconstrained target criterion (17). It differs from (17) 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 (29). 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 (29) 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

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

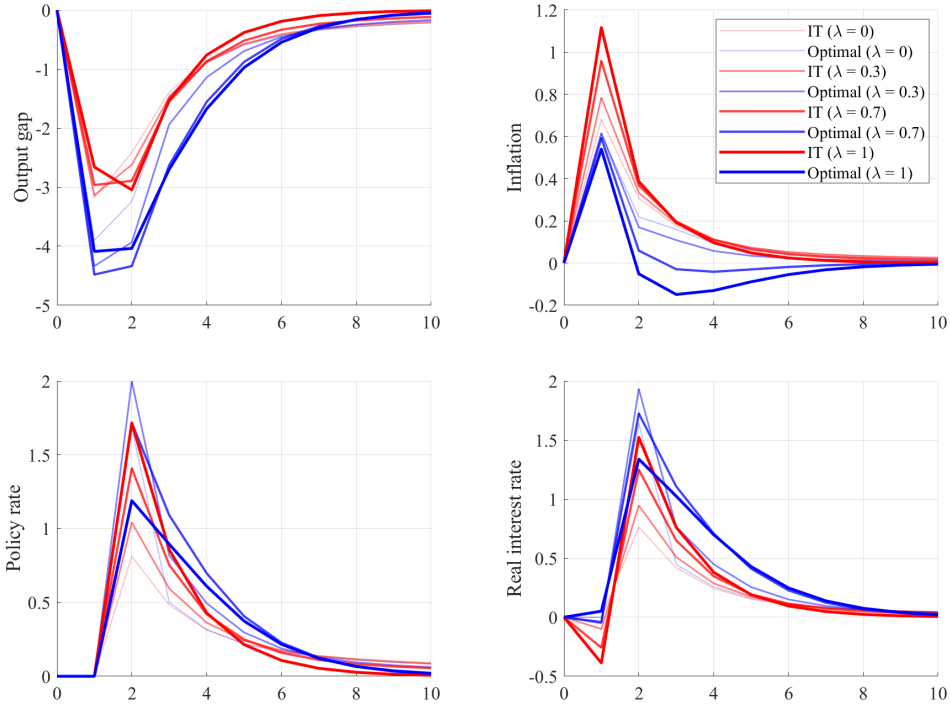
with $\Delta_t \equiv x_t^* - x_t$. The central bank should set its best estimate of the output gap $x_{t|t}$ equal to the moving target x_t^*, which is determined by the previous period's miss, Δ_{t-1}. The extent to which the central bank should commit to offset past misses depends on the strength 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.[12]

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

[12] 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 $\rho = 0.95$, $g = 0.08$, $\beta = 0.99$, $\sigma = 1$, $\kappa = 0.1$, $\alpha = \kappa/7.87$.

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

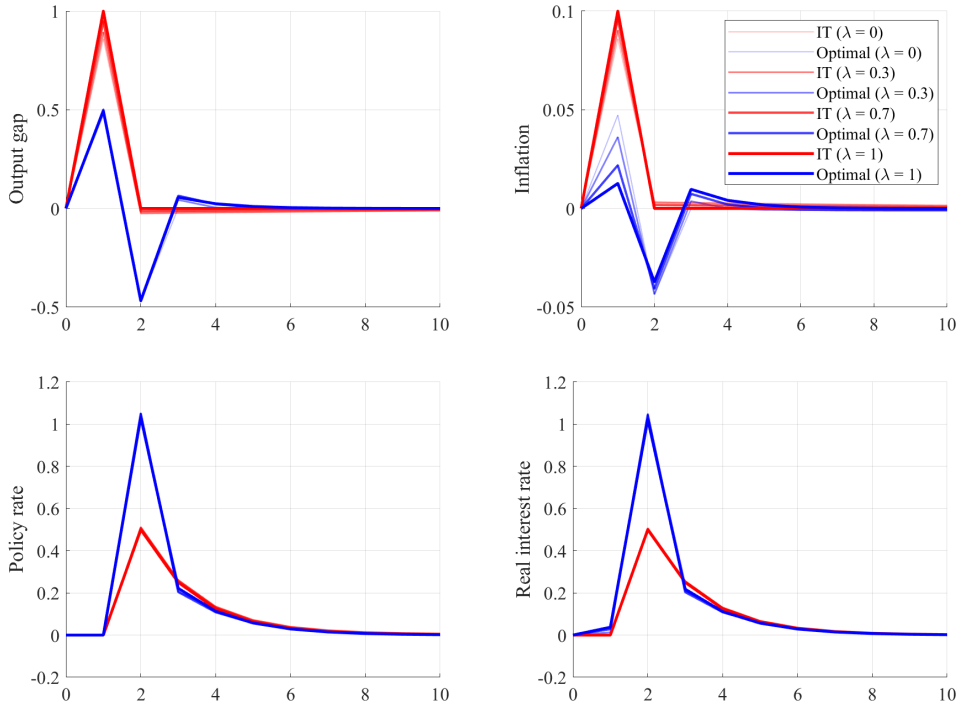
Like the unconstrained case, a FAIT policy is robust because we can capture most of WAIT, preemption, and the additional make-up policy using only a weighted average of past outcomes. In contrast to the unconstrained case, though, optimal policy requires tracking a weighted average over both inflation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

To illustrate, we substitute $\Delta_{t-1} \equiv \pi^*_{t-1} - \tilde{\pi}_{t-1}$ into (29), rearrange, and then substitute the result into (28) to arrive at

$$\mathbb{E}_{t|t} \frac{1}{1 - \frac{1}{1+(1-\lambda)g\rho^2\Theta} P(L)L} \tilde{\pi}_t = 0 \tag{30}$$

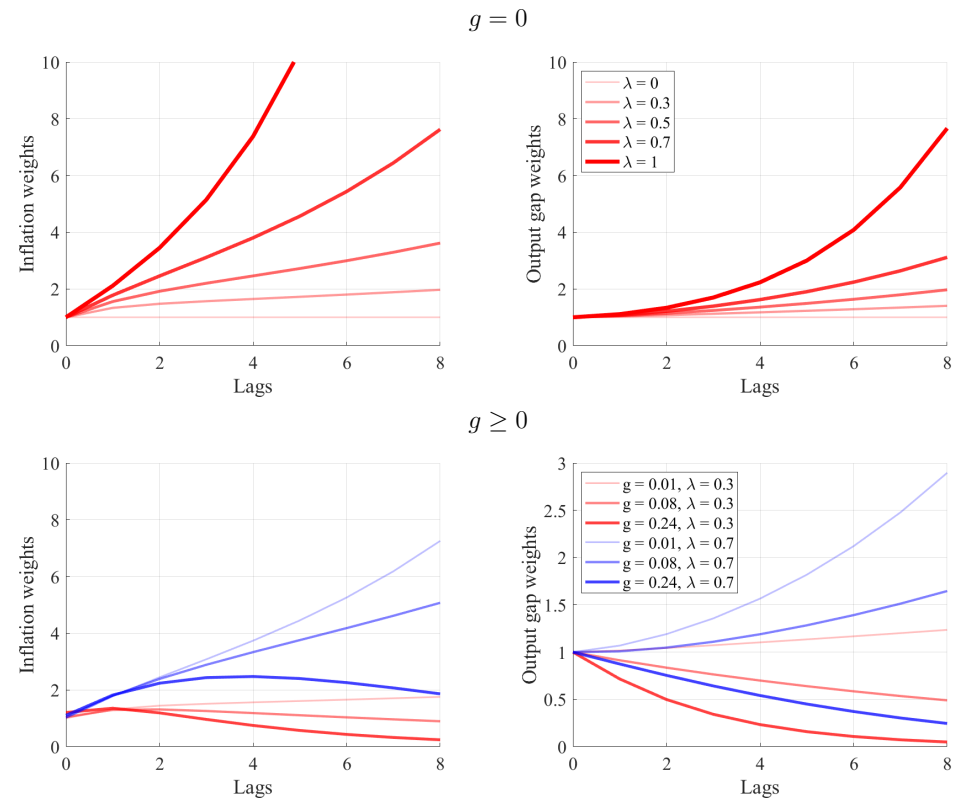
where $\mathbb{E}_{t|t}$ denotes the expectation conditional on I_t^{CB}. This formulation makes clear

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 (31). Parameter values are $\beta = 0.99$, $\sigma = 1$, $\kappa = 0.1$, $\alpha = \kappa/7.87$.

that the 'moving target' described in (28) and (29) 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 (30) (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 (20). Assuming again that the cost-push is a first-order autoregressive process,[13] Appendix A shows that, following an analogous approach to the derivation of (20), we can write optimal policy under imperfect information in equilibrium as

$$\mathbb{E}_{t|t}\left(Q_x(L)x_t + Q_\pi(L)\pi_t + Q_u u_t\right) = 0 \tag{31}$$

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.

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

[13] 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.

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 (21) with the lower bound constraint

$$i_t > \bar{i}. \tag{32}$$

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 (9) subject to the constraints (10) to (13) imposed by private-sector equilibrium behaviour and the lower bound constraint (32).

The first-order conditions are the same as in the imperfect information case, except that (23) is replaced with the complementary slackness condition

$$\mu_t^i \geq 0, \quad \mu_t^i \left(i_t - \bar{i}\right) = 0. \tag{33}$$

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

$$\begin{aligned}\mu_t^{\omega\pi} = &-(1-\lambda)\rho\Omega\frac{1-\lambda L}{1-\zeta_1 L}\mathbb{E}_t\frac{1}{1-\zeta_2^{-1}L^{-1}}\frac{1}{1-\lambda L}\pi_{t+1} \\ &+(1-\lambda)\rho\Omega\frac{1-\lambda L}{1-\zeta_1 L}\mathbb{E}_t\frac{1}{1-\zeta_2^{-1}L^{-1}}\left(\frac{\lambda}{\beta}\frac{L}{1-\lambda L}+\frac{1}{\beta(1-\beta\rho)}(1-L)\right)\mu_{t+1}^i \qquad (34)\\ \mu_t^{\omega x} = &\,(1-\lambda)\frac{(1-\beta)\rho\beta}{1-\beta\rho}\mathbb{E}_t\frac{1-L}{1-\beta\rho(1-g)L^{-1}}\sigma\mu_{t+1}^i. \qquad (35)\end{aligned}$$

If $\mathbb{E}_t\mu_T^i = 0$ for all $T \geq t+1$, then (34) and (35) collapse to (26) and (27). The second line of (34) 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, (35) represents the

effect of shifting learners' output expectations for future output gaps.

Conditions (25), (33), (34) and (35) 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

$$\tilde{\pi}_t = \pi_t^* \tag{36}$$

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

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

where

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

is the same as (29) and

$$P_1(L^{-1}) = \underbrace{(1-\lambda)g\rho^2\Omega\frac{\kappa}{\sigma}\left(\frac{\lambda}{\beta}\frac{\zeta_2^{-1}}{1-\zeta_2^{-1}\lambda} + \frac{1-\zeta_2^{-1}}{\beta(1-\beta\rho)}\right)\mathbb{E}_t\frac{1}{1-\zeta_2^{-1}L^{-1}} + (1-\lambda)g\rho^2\frac{(1-\beta)\beta}{1-\beta\rho}\left(1-\beta\rho(1-g)\right)\mathbb{E}_t\frac{1}{1-\beta\rho(1-g)L^{-1}}}_{Insurance\ for\ expected\ bind} \tag{37}$$

$$P_2(L) = \zeta_1\frac{1}{1-\zeta_1 L} \tag{38}$$

$$P_3(L^{-1}) = \underbrace{(1-\lambda)g\rho^2\Omega\frac{\kappa}{\sigma}\left(\frac{\lambda}{\beta}\frac{\zeta_2^{-1}}{1-\zeta_2^{-1}\lambda} + \frac{1-\zeta_2^{-1}}{\beta(1-\beta\rho)}\right)\mathbb{E}_t\frac{1}{1-\zeta_2^{-1}L^{-1}}}_{Insurance\ for\ expected\ bind}. \tag{39}$$

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

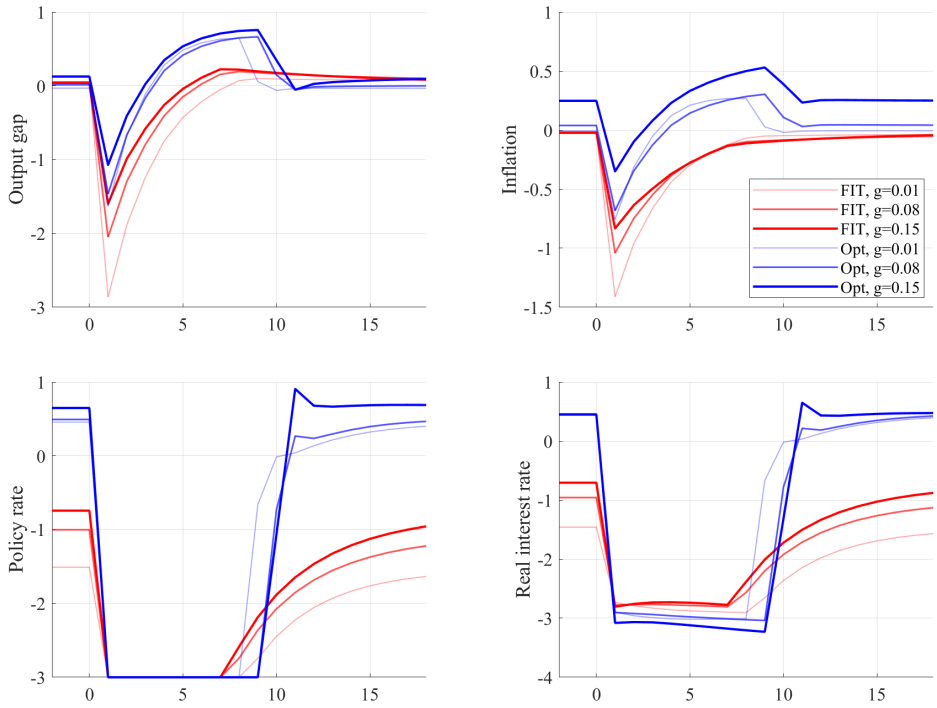
$$\begin{aligned} x_t^* \equiv & \frac{1}{1+(1-\lambda)g\rho^2\Theta} \underbrace{\left(\frac{\lambda}{\beta}(1-\beta)+1\right)\Delta_{t-1}}_{Make-up\ policy} \\ & + \frac{(1-\lambda)g\rho^2\Theta}{1+(1-\lambda)g\rho^2\Theta}\left(1-\beta\rho(1-g)\right) \underbrace{\mathbb{E}_t \frac{1}{1-\beta\rho(1-g)L^{-1}}\Delta_{t+1}}_{Insurance\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

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 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.[14] 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.

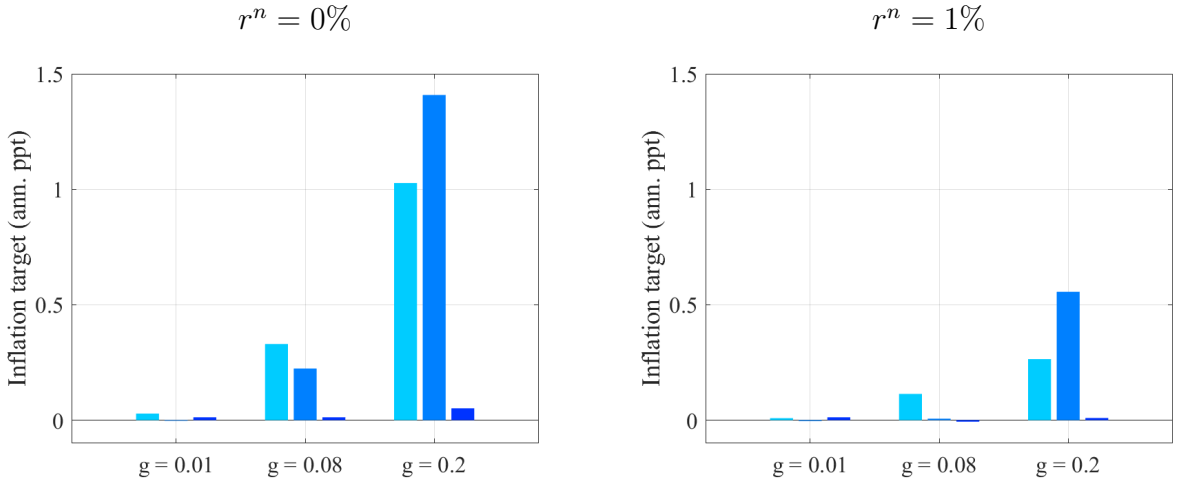
[14]To solve the model with the ZLB, we used a policy function iteration method, drawing on code in Sijmen Duineveld's PROMES toolbox (Duineveld, 2021).

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 $\lambda = 0.25$, $\rho = 0.95$, $\sigma = 1$, $\kappa = 0.1$, $\alpha = \kappa/7.87$.

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 $\rho = 0.95$, $\sigma = 1$, $\kappa = 0.1$, $\alpha = \kappa/7.87$.

Figure 8 shows how the preemptively higher stochastic steady-state inflation target varies with λ, g, and the mean neutral rate, r^n.[15] As expected, the inflation target is 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. Implementing the target criteria also requires knowledge of the underlying expectation formation process. 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 these questions, we first look at how closely simple weighted average inflation targeting schemes can perform relative to the optimal target criteria, and how they 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.[16] Appendix D provides results for the unconstrained case, where cost-push shocks are the only source of losses.[17]

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 cri-

[15] 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.

[16] 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 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.

[17] 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 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.

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.

teria 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, 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. (28) (as with the ZLB), unlike the unconstrained case, eq. (17). 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 $\rho \in \{0.95, 0.99\}$.[18]

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. 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

[18] 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					

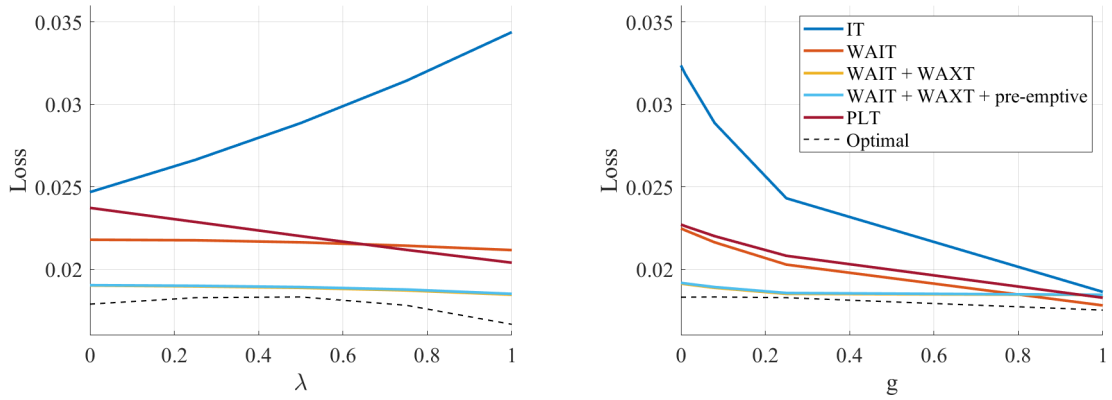
Notes The parameters reports are set to maximize the average loss across all points in the grid. Average welfare losses are normalized relative to the welfare under IT. The optimal value of a parameter is undefined (labelled *undef.*) if the parameter has no effect on loss at the optimum.

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).[19] 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.

[19]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 and many researchers and policymakers have worried that policy recommendations that rely it may not be robust. A large literature has shown that this concern in 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 is a robust policy strategy. Whether the central bank is unconstrained, faces information constraints, or the ZLB, it is optimal to target a weighted average of inflation.

In addition, the optimal target criteria we derive give clear guidance on when policy should be flexible and deviate from the policy implied by the weighted average of inflation alone. If expectations drift in ways that are inconsistent with policymakers' objectives, then preemptive actions should be taken with more policy now to re-anchor beliefs. If targets are missed today, then policy should do more to make-up for the misses in the future. If targets are expected to be missed in the future because of the constraint of

the ZLB, then policy should do more today to insure itself against the possibility of a prolonged period of constrained actions. This guidance summarizes what flexible in FAIT should mean.

A Appendix - Derivations

A.1 Unconstrained optimal policy

The unconstrained optimal policy problem is to maximise (9) subject to (10) - (13). 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.[20]

Combining (A1) and (A2) to eliminate μ_t^{π} gives equation (14) in the main text. Similarly, using (A1) to eliminate μ_t^{π} in (A3) gives equation (15). The remainder of the derivation for Result 17 (unconstrained optimal target criterion) is in the main text.

A.1.1 Derivation of eq. (19)

First, we rewrite $\mu_t^{\omega\pi}$ in terms of expected output gaps by substituting (14) into the second right-hand side term in (15) and iterating forward to get

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

[20] 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.

Then, inverting the optimal target criterion (17) 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 (10), (12) 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 (10) - (13) some constraint on the policy rate. This could be an information constraint (21), or a ZLB constraint (32). For this optimal policy problem, the first four first-order conditions, (A1) - (A4), remain the same. But instead of (A5), we have (22), reproduced here:

$$\frac{1}{\sigma}\sum_{j=0}^{t}\mu_{t-j}^x - \mu_t^i = 0 \tag{22}$$

where μ_t^i is the Lagrange multiplier on the policy rate constraint.[21]

Using (A1) to eliminate μ_t^π from (A2) - (A4) and then substituting in $\sigma(1-L)\mu_t^i = \mu_t^x$ from (22) gives (25), (34) and (35), which we reproduce here:

$$\frac{\alpha}{\kappa}x_t + \frac{\sigma}{\kappa}\mu_t^i = -\frac{1}{1-\lambda L}\left(\pi_t - \frac{\lambda}{\beta}\frac{\kappa}{\sigma}\frac{\sigma}{\kappa}\mu_{t-1}^i - \rho g\mu_t^{\omega\pi}\right) + \left(\lambda\left(\frac{1}{\beta}-1\right)+1\right)\frac{\sigma}{\kappa}\mu_{t-1}^i + \frac{\rho g}{\kappa}\mu_t^{\omega x} \tag{25}$$

$$\begin{aligned}\mu_t^{\omega\pi} = &-(1-\lambda)\rho\Omega\frac{1-\lambda L}{1-\zeta_1 L}\mathbb{E}_t\frac{1}{1-\zeta_2 L^{-1}}\frac{1}{1-\lambda L}\left(\pi_{t+1} - \frac{\lambda}{\beta}\frac{\kappa}{\sigma}\frac{\sigma}{\kappa}\mu_t^i\right)\\ &+ (1-\lambda)\frac{\rho\Omega}{\beta(1-\beta\rho)}\frac{\kappa}{\sigma}\frac{1-\lambda L}{1-\zeta_1 L}\mathbb{E}_t\frac{1-L}{1-\zeta_2 L^{-1}}\frac{\sigma}{\kappa}\mu_{t+1}^i\end{aligned} \tag{34}$$

$$\mu_t^{\omega x} = (1-\lambda)\frac{(1-\beta)\beta\rho}{1-\beta\rho}\mathbb{E}_t\frac{1-L}{1-\beta\rho(1-g)L^{-1}}\frac{\sigma}{\kappa}\mu_{t+1}^i. \tag{35}$$

Now substitute (34) and (35) into (25) and collect all the non-μ_t^i terms in $\tilde{\pi}_t$, which is

[21] 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 \\
&\quad + (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 \\
&\quad + (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 \\
&\quad + \left(\lambda\left(\frac{1}{\beta}-1\right)+1\right)\frac{\sigma}{\kappa}\mu_{t-1}^i \\
&\quad + (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 \\
&\quad + (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 \\
&\quad + (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 \\
&\quad + (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 \\
&\quad + (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, (23) 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. (31) Start the optimal policy conditions (25), (26) and (27). Together with the Phillips curve (10) and the updating rule for learners' infla-

tion expectations (12), 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. (19), we can rewrite $\mu_t^{\omega\pi}$ in (26) 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 (27) 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 (18), and $a_{\mu,1}$ and $a_{\mu,2}$ are new constants. Combining with the updating rule for learners' inflation expectations and using (21) 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[22]

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

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

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

[22]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:[23]

- 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

[23] 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.

References

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

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

Angeletos, G.-M., and 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.

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.

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.

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. G. Gibbs, and N. McClung (2024): "Resolving new keynesian puzzles," Discussion paper, Bank of Finland Research Discussion Papers.

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.

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.

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.

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.

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.