1 Introduction

In 2020 the Federal Reserve adopted a new strategy for monetary policy that officials have called “average inflation targeting,” or AIT. The strategy is focused on a 2% target for the inflation rate, but it does not always seek to hit that target in the short run. Instead, according to the official statement of the strategy, “following periods when inflation has been running persistently below 2 percent, appropriate monetary policy will likely aim to achieve inflation moderately above 2 percent for some time”(Federal Reserve 2020). This intentional overshooting is a departure from conventional inflation targeting, in which policymakers always aim for 2% inflation, regardless of history.

When the new strategy was announced, Fed officials outlined their rationale in several speeches, which emphasized the interplay of two factors: the effective lower bound [ELB] on interest rates, and the need to anchor inflation expectations at 2%. The ELB leads to periods when inflation is stuck below 2%, and therefore, as Vice Chair Clarida (2020) reasoned,

> [I]f policy seeks only to return inflation to 2 percent following a downturn in which the ELB has constrained policy, an inflation-targeting monetary policy will tend to generate inflation that averages less than 2 percent, which, in turn, will tend to put persistent downward pressure on inflation expectations.

According to Chair Powell (2020), the economy could fall into an “adverse cycle of ever-lower inflation and inflation expectations” that would destabilize the economy. Periods with inflation above 2% are needed to achieve an average of 2% and keep expectations anchored.

A sizable body of research has studied the properties of average inflation targeting. This work includes analyses of New Keynesian macroeconomic models such as Mertens and Williams (2019, 2020) and Budianto, Nakata, and Schmidt (2023), and simulations of the Federal Reserve’s FRB/US model such as Arias et al. (2020) and Hebden et al. (2020). This work has helped to clarify the conditions under which AIT is effective at stabilizing the economy, and it has shed light on the optimal size and duration of inflation-target overshoots. This paper seeks to contribute to this literature.[1]

Specifically, we seek the simplest possible model that captures the factors that policymakers cite when they advocate AIT. The model is a variation on the "backward-looking" macro model of the Romer (2019) textbook in which we introduce the lower bound on interest rates and assume that expected inflation is anchored at 2% if inflation averages 2%. Our model complements other macro models that have the virtues of firmer microeconomic foundations and/or greater quantitative realism, but at the cost of greater complexity. We aim for an analysis that is more formal and rigorous than a speech by a Fed official, yet more transparent than a typical study of optimal monetary policy.

We consider two versions of our model. In the first, expectations are fully anchored: if the central bank produces 2% inflation on average, then expected inflation is constant at 2%. For that case, we prove the optimality of a simple policy rule: whenever the central bank is not constrained by the ELB, it aims for a short-run inflation target that is constant and exceeds 2%. The reason for targeting inflation above 2% is the one given in practice by Fed officials: to offset periods when the ELB pushes inflation below 2%. (While our result is simple, the proof of optimality is non-trivial because of the non-linearity in the model arising from the ELB.)

In the second version of the model, anchoring is imperfect: expected inflation responds somewhat to fluctuations in actual inflation. That case is more complex and must be analyzed numerically. Once again, optimal policy often aims at a short-run inflation target above 2%, but this target varies over time. For some parameter values, the short-run target is highest after an episode of low inflation and falls after that. This feature of optimal policy bears some resemblance to the Fed's actual strategy, which aims for temporarily high inflation after a period with inflation below 2%.

[1] Other work on AIT includes Nessén and Vestin (2005), Amano et al. (2020), Eo and Lie (2020), and Honkapohja and McClung (2024).

2 The Model

This section presents our model and the problem facing policymakers.

2.1 The IS and Phillips Curves

Our assumptions about output and inflation behavior follow Romer (2019, Chapter 12), which draws on Svensson (1997) and Ball (1999):

$$y_t = \lambda y_{t-1} - \beta(r_{t-1} - r^*) + \epsilon_t, \quad 0 < \lambda < 1, \beta > 0 \tag{1}$$

$$r_t = i_t - \pi_t^e \tag{2}$$

$$\pi_t = \pi_{t-1}^e + \alpha y_{t-1} + \eta_t, \quad \alpha > 0 \tag{3}$$

where y is the output gap (the percentage deviation of output from potential), i and r are the nominal and real interest rates, r^* is the neutral real rate, π^e is expected inflation, and ϵ and η are shocks with mean zero and bounded distributions: $\epsilon \in [-\bar{\epsilon}, \bar{\epsilon}]$ and $\eta \in [-\bar{\eta}, \bar{\eta}]$ for some $\bar{\epsilon}, \bar{\eta} > 0$.

In this backward-looking model, equation (1) is a dynamic IS equation: output depends on lagged output, the lagged real interest rate, and a demand shock. Equation (2) is the definition of the real interest rate. Equation (3) is a Phillips curve: inflation depends on lagged expected inflation (capturing the idea that current prices were set in advance), lagged output, and an inflation or supply shock.

Under these assumptions, the real interest rate affects output with a one-period lag, and output affects inflation with a one-period lag. Combining these lags, it takes two periods for an interest-rate adjustment by the central bank to affect inflation; thus policymakers in period t can set an inflation target for $t+2$ but not for t or $t+1$. If we interpret a period as a year, the model's time lags are roughly consistent with those found in empirical work (e.g., Romer and Romer 2004).

2.2 The Behavior of Expectations

Analyses of equations (1)–(3), such as Romer's, often complete the model by assuming that inflation expectations are static: $\pi_t^e = \pi_t$. In that case, equation (3) becomes a relation between the output gap and the change in inflation—an accelerationist Phillips curve—and the real interest rate is the nominal rate minus current inflation. However, while this treatment of expectations was arguably realistic in the past, there is a growing consensus that the Fed's commitment to a 2% inflation target has anchored expectations at that level (e.g., Yellen 2019). We therefore assume that expectations are anchored.

We consider two versions of anchored expectations. In the first, full anchoring, we assume that the central bank has announced a long run inflation goal of 2% and committed itself to producing 2% on average. As a result, expected inflation is fixed at that level:

FULLY ANCHORED EXPECTATIONS

$$\pi_t^e = 2 \quad \text{for all } t. \tag{4}$$

In the second version of our assumption, the central bank has again announced a 2% goal, but anchoring is not perfect: expected inflation deviates somewhat from 2% in response to movements in actual inflation. Specifically:

PARTIALLY ANCHORED EXPECTATIONS

$$\pi_t^e = \gamma(2) + (1-\gamma)\pi_t \quad 0 \le \gamma \le 1. \tag{5}$$

Here, expected inflation is a weighted average of 2% and the current inflation rate. The parameter γ indicates how strongly expectations are anchored. When $\gamma = 0$, the specification reduces to static expectations.

Some research suggests that fully anchored expectations are a good approximation to reality since the 1990s (e.g., Blanchard 2016). Other work, however, finds that movements in actual inflation have had some effect on expected inflation (e.g., Ball, Leigh, and Mishra 2022).

The variable π_t^e is a subjective expectation of price setters, and is not necessarily rational. In what follows, we also derive mathematical expectations of inflation, using standard notation. For example, $\mathrm{E}[\pi]$ is the unconditional mean of the inflation rate, and $\mathrm{E}_t[\pi_{t+1}]$ is the expected value of inflation at $t+1$ conditional on information at t.

2.3 The Policy Problem

We take it as given that the central bank has chosen 2% as its long run inflation goal and committed to achieving the goal on average. Therefore, in considering alternative policy rules, we restrict attention to rules under which $\mathrm{E}[\pi]$, the unconditional expectation of the inflation rate, equals 2%. This restriction on policy underlies our assumption of anchored expectations; as stressed by Fed officials, it would be implausible to assume that expectations are anchored at 2% if average inflation differed from 2%.

Presumably the 2% goal was chosen based on an analysis of the costs and benefits of different levels of inflation, but we do not model this decision. Instead, we ask what rule among the class that produces 2% average inflation is best at stabilizing the economy. In assuming rather than deriving a 2% goal, we follow the formal strategy review that led the Fed to adopt AIT. In introducing the review, Vice Chair Clarida (2019) noted that Congress has assigned goals including price stability to the Fed, and said: "Our review this year will take this statutory mandate as given and will also take as given that inflation at a rate of 2 percent is most consistent over the longer run with the Congressional mandate."

Policymakers choose a rule for setting the nominal interest rate i_t as a function of the state of the economy Ω_t. There is an effective lower bound on i_t, which we set to zero for simplicity. For technical reasons, we also impose an upper bound on i_t, i^-, which is arbitrarily large. This bound is never binding under the optimal policy that we derive. (The upper bound on i_t makes it easier to apply dynamic programming methods to our policy problem.)[2]

The central bank chooses the rule $i_t(\Omega_t)$ to minimize a weighted sum of the variances of output and inflation. Policymakers face two kinds of constraints: the restriction that

[2]In proving our results about optimal policy, we show that they hold as long as i^- is set above a certain level that depends on the model's parameters.

$\mathrm{E}[\pi] = 2$, and the bounds on i_t. We can write the policy problem as:

$$\min_{i_t(\Omega_t)} (1-\mu)\,\mathrm{Var}(y) + \mu\,\mathrm{Var}(\pi), \quad 0 \leq \mu \leq 1$$

$$\text{s.t. } \mathrm{E}[\pi] = 2$$

$$0 \leq i_t \leq i^- \text{ for all } t$$

The parameter μ determines the weights on output and inflation variances.

2.4 Optimal Inflation Targeting

In the above policy problem, the central bank's instrument is the interest rate i_t. However, given the current state of the economy, the IS equation implies a negative, one-to-one relationship between i_t and $\mathrm{E}_t[y_{t+1}]$, the expectation of output in the next period. This fact and the Phillips curve imply a negative, one-to-one relationship between i_t and $\mathrm{E}_t[\pi_{t+2}]$, the expectation of inflation in two periods. Therefore, following Romer, we can reinterpret the policy problem in period t as the choice of an inflation target for $t+2$: $\hat{\pi}_t = \mathrm{E}_t[\pi_{t+2}]$. A rule for this target implicitly defines an interest-rate rule that implements it.

The restriction of i_t to $[0, i^-]$ implies that the inflation target $\hat{\pi}_t$ is restricted to a range $[\underline{\pi}_t, \bar{\pi}_t]$. Since the relation between i_t and π_t is negative, the lower bound $\underline{\pi}_t$ is the level of $\hat{\pi}_t$ when $i_t = i^-$, and the upper bound $\bar{\pi}_t$ is the level when $i_t = 0$. The bounds on $\hat{\pi}_t$ vary over time because expected inflation depends on the state of the economy as well as on i_t. The lower bound on $\hat{\pi}_t$ never binds under optimal policy. Expressions for the upper bound are derived below for the cases of full and partial anchoring.

The policymaker chooses a rule for the inflation target, $\hat{\pi}_t(\Omega_t)$, to minimize the loss function given the bounds on $\hat{\pi}_t$ and the restriction on average inflation:

$$\min_{\hat{\pi}_t(\Omega_t)} (1-\mu)\,\mathrm{Var}(y) + \mu\,\mathrm{Var}(\pi), \quad 0 \leq \mu \leq 1$$

$$\text{s.t. } \mathrm{E}[\pi] = 2$$

$$\underline{\pi}(\Omega_t) \leq \hat{\pi}_t \leq \bar{\pi}(\Omega_t) \text{ for all } t$$

3 Optimal Policy with Full Anchoring

For the case of full anchoring, the optimal policy is simple to state:

Proposition 1. *With fully anchored expectations, the optimal rule for the short-run inflation target is* $\hat{\pi}_t = \min\{\pi^*, \bar{\pi}_t\}$ *, where* π^* *is a constant and* $\pi^* > 2\%$*.*

This result is proved in the Appendix. It says that the central bank targets a fixed inflation rate $\pi^* > 2\%$ whenever that is feasible. When π^* exceeds the upper bound $\bar{\pi}_t$, the central bank targets $\bar{\pi}_t$.

This result should make sense. The bound $\bar{\pi}_t$ is sometimes below 2%, so the target $\hat{\pi}_t$ must sometimes be below 2%. To satisfy the constraint that average inflation equals 2%, $\hat{\pi}_t$ must sometimes exceed 2%. Many specific policies would satisfy the constraint, but the one that best stabilizes the economy is to set $\hat{\pi}_t$ equal to the constant π^* or as close to π^* as possible. The value of π^* is defined by the constraint on average inflation and can be derived numerically.

While the optimal policy is simple to state, the proof of Proposition 1 is non-trivial because of the non-linearity arising from the bound on $\hat{\pi}_t$ and also because of the unconventional constraint on average inflation. In the proof, we first consider a dynamic programming problem in which policymakers minimize the discounted losses from output and inflation fluctuations, and then let the discount factor approach one to find the solution to our undiscounted problem. We impose the constraint on average inflation by adding a penalty on the level of inflation to the dynamic programming problem.

We can better understand the optimal policy by deriving the upper bound on the inflation target, $\bar{\pi}_t$. This bound is the expectation of inflation at $t + 2$ if the interest rate i_t is set at its lower bound of zero. From equations (1)–(3) and the assumption of fully anchored expectations, we derive:

$$\bar{\pi}_t = C + \alpha\lambda y_t, \tag{6}$$

where the constant $C = 2 + \alpha\beta(2 + r^*)$. Equation (6) says that $\bar{\pi}_t$ depends positively on

Figure 1: Optimal Policy with Full Anchoring

(a) Inflation Target $\hat{\pi}(y_t)$

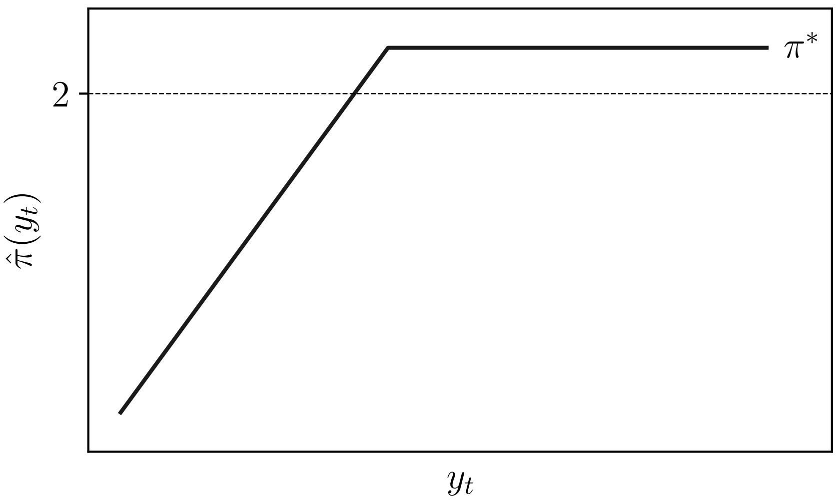

(b) Interest-Rate Rule $i(y_t)$

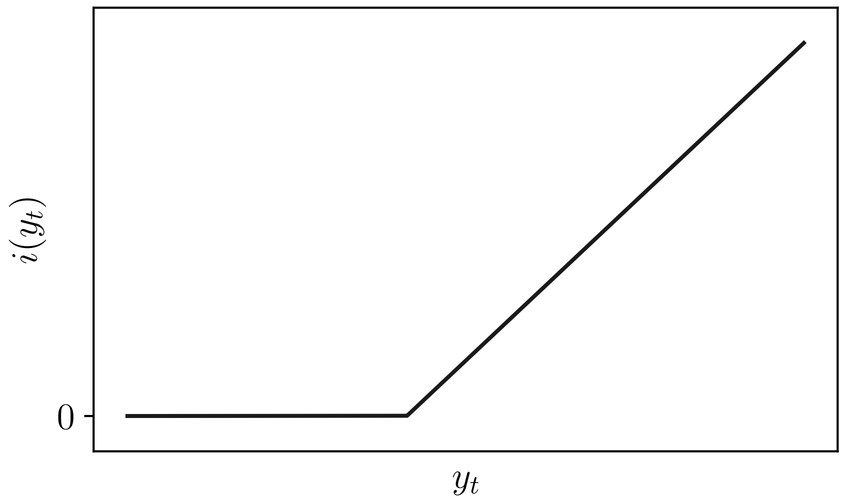

the output gap y_t. A higher y_t loosens the bound on the inflation target because it raises future output and inflation through the IS equation and Phillips curve.[3]

Our results determine the inflation target $\hat{\pi}_t$ as a function of y_t, as shown in Figure 1(a). When y_t is low (which occurs when there are adverse shocks to the IS equation), the bound $\bar{\pi}_t$ is binding. In that regime, the inflation target is increasing in y_t because the bound is increasing in y_t. When y_t is above a critical level (specifically, $(\pi^* - C)/\alpha\lambda$), the bound exceeds π^* and the target is set at π^*.

[3]With fully anchored expectations, substituting $i_t = 0$ into the IS equation yields $\mathrm{E}_t[y_{t+1}] = \lambda y_t + \beta(r^* + 2)$. Substituting that equation into the Phillips curve yields $\mathrm{E}_t[\pi_{t+2}]$ when $i_t = 0$, which is the expression in equation (6).

We can also derive the interest-rate rule that implements optimal inflation targeting:

$$i_t = \max\{i_t^*, 0\} \tag{7}$$

where i_t^* is the interest rate such that $\hat{\pi}_t = \mathrm{E}_t[\pi_{t+2}] = \pi^*$. The policymaker sets $i_t = i_t^*$ when $i_t^* \geq 0$ and $i_t = 0$ when $i_t^* < 0$, which keeps i_t as close as possible to i_t^* and $\hat{\pi}_t$ as close as possible to π^*. We can derive i_t^* from equations (1)–(3) and the assumption of full anchoring:

$$i_t^* = C' + \frac{\lambda}{\beta} y_t \tag{8}$$

where $C' = 2 + r^* + (2 - \pi^*)/(\alpha\beta)$.

Figure 1(b) shows the optimal interest rate as a function of y_t. When y_t is low, policy is constrained by the zero bound. When policy is unconstrained, i_t is increasing in y_t: a higher y_t implies a higher $\mathrm{E}_t[\pi_{t+2}]$ for a given i_t, so a higher i_t is needed to offset this effect and keep $\mathrm{E}_t[\pi_{t+2}]$ at π^*.

Notice that the optimal interest rate does not depend on the inflation rate. This result reflects the assumption of fully anchored expectations, which implies that current inflation has no effect on future inflation.

4 Optimal Policy with Partial Anchoring

This section considers our model with imperfectly anchored expectations: expected inflation is an average of the 2% target and current inflation. We generalize the proposition about optimal policy and then analyze the model numerically to learn more.

4.1 Optimal Policy

In what follows, an important variable is $\mathrm{E}_t[\pi_{t+1}]$, the expectation of inflation in the next period. To simplify notation, we denote this variable by X_t. In period t, when policymakers chooses an inflation target $\hat{\pi}_t$, they take X_t as given because they cannot affect inflation until $t + 2$. Using the Phillips curve (3) and the partial anchoring assumption (5), we can write X_t in terms of current output and inflation:

$$X_t = \mathrm{E}_t[\pi_{t+1}] = 2\gamma + (1-\gamma)\pi_t + \alpha y_t. \tag{9}$$

With this notation, we can establish the following about optimal policy (this result is proved in the Appendix):

Proposition 2. *With partially anchored expectations and $\gamma > \alpha\beta/(\alpha\beta + 1 - \lambda)$, the optimal rule for the short-run inflation target is $\hat{\pi}_t = \min\{\pi^*(X_t), \bar{\pi}_t\}$ for some function $\pi^*(\cdot)$, and $\pi^*(X_t) > 2$ for some X_t.*

Notice first that the Proposition applies only if the degree of anchoring γ (the weight on the inflation target in the equation for π_t^e) exceeds a bound of $\alpha\beta/(\alpha\beta + 1 - \lambda)$. If anchoring is weak enough that this condition fails, we cannot rule out the possibility that a very bad sequence of shocks will send the economy into a disinflationary spiral in which falling output and expected inflation reinforce each other forever. In this case, no policy rule produces finite variances of output and inflation. When the restriction on γ holds, however, expectations are pulled toward 2% strongly enough that policy can make inflation and output stationary.[4]

When the bound on γ is satisfied, the optimal policy is a variation on the one under full anchoring. The inflation target $\hat{\pi}_t$ is the minimum of the bound $\bar{\pi}_t$ and a level π^* that policymakers choose if unconstrained. This level must sometimes or always exceed 2% to offset periods when the bound pushes $\hat{\pi}_t$ below 2%.

The major difference from the full-anchoring case is that π^*, the target when policymakers are unconstrained, is no longer a constant. Instead it is a function of the state of the economy, in particular $X_t = \mathrm{E}_t[\pi_{t+1}]$.

Current output and inflation affect the inflation target through X_t. They also affect the bound that constrains the target. Using equations (1)–(3) and (5), and setting $i_t = 0$,

[4]When economists study models with a lower bound on interest rates, they often rule out disinflationary spirals by assuming that sufficiently adverse outcomes trigger an emergency fiscal stimulus (e.g., Kiley and Roberts 2017). If emergency fiscal policy were added to our model, one could derive optimal monetary policy without any restriction on the degree of anchoring.

we can derive:

$$\bar{\pi}_t = [\alpha(\lambda + 1 - \gamma)]y_t + [(1-\gamma)(\alpha\beta + 1 - \gamma)]\pi_t + C'' \quad (10)$$

where $C'' = \alpha\beta r^* + \gamma(2\alpha\beta + 4) - 2\gamma^2$. In this expression, the coefficients on y_t and π_t are positive. Higher output or higher inflation loosens the bound on the inflation target because, with imperfect anchoring, they both raise future output and inflation.[5]

The nature of optimal policy depends on the function $\pi^*(X_t)$. We have not been able to derive general results about this function beyond the result that π^*_t must sometimes exceed 2. In what follows, we analyze the function numerically.

4.2 The Behavior of the Short-Run Inflation Target

For given parameter values, we use numerical methods to derive a close approximation to the $\pi^*(\cdot)$ function in the optimal policy rule. Specifically, we consider a version of the policy problem with discounting but a discount factor very close to one, and solve the Bellman equation using optimistic policy iteration (Bertsekas 2013). See the Appendix for details.

We find that the $\pi^*(\cdot)$ function varies greatly depending on parameter values. However, in all the cases we have examined, the function has one of three basic shapes: it is monotonically increasing, monotonically decreasing, or U-shaped with a minimum at some level of X_t.

We illustrate these possibilities with the parameterization of the model shown in Table 1. We assume a fixed set of coefficients for the model's equations, and fixed distributions of shocks that roughly approximate normal distributions.[6] We generate different shapes of $\pi^*(\cdot)$ by varying the parameter μ, which determines the weights on

[5]With partial anchoring, substituting $i_t = 0$ into the IS equation yields $\mathrm{E}_t[y_{t+1}] = \lambda y_t + \beta(1-\gamma)\pi_t + \beta(2\gamma + r^*)$. The Phillips curve implies $\mathrm{E}_t[\pi_{t+2}] = \mathrm{E}_t[\pi^e_{t+1}] + \alpha\,\mathrm{E}_t[y_{t+1}]$. From the partial anchoring assumption and the Phillips curve, $\mathrm{E}_t[\pi^e_{t+1}] = (1-\gamma)\alpha y_t + (1-\gamma)^2\pi_t + 4\gamma - 2\gamma^2$. Combining these results yields the expression for $\bar{\pi}_t$.

[6]For both the shock ϵ in the IS equation and the shock η in the Phillips curve, we discretize a normal distribution $\mathcal{N}(0, 4^2)$ on a grid of 9 points ranging from -4 to 4 standard deviations using the method in Tauchen (1986).

Table 1: Parameterization of the Partially Anchored Model

α	β	λ	γ	μ
0.4	2.0	0.1	0.5	0.5, 0.9, 1.0

Figure 2: Optimal $\pi^*(X_t)$

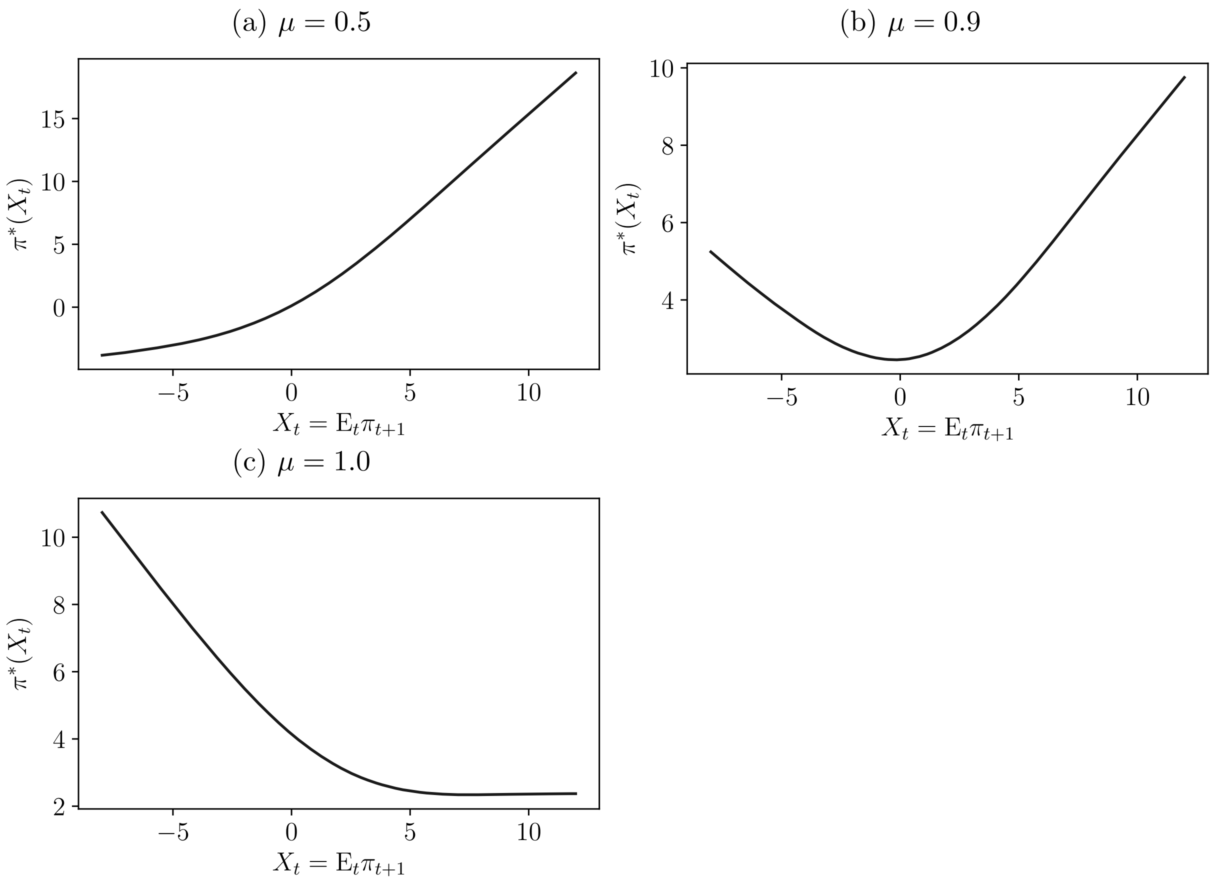

output and inflation variance in policymakers' objective function. Figure 2 shows the $\pi^*(\cdot)$ function for three values of μ (0.5, 0.9 and 1.0), which produce the three possible shapes.

The $\pi^*(\cdot)$ function can have various shapes because X_t, the expected value of next period's inflation rate, influences the policy problem in complex ways. The Appendix discusses the relevant factors and the roles of various parameters in detail. Here we seek to provide some intuition about why π^* can be either increasing or decreasing in X_t.

One factor is the output variance term in policymakers' objective function. When anchoring is imperfect—expected inflation π^e is tied to current inflation—large changes in inflation require movements in output. Policymakers can stabilize output by keeping their target for inflation in two periods, $\hat{\pi}_t = \mathrm{E}_t[\pi_{t+2}]$, close to the expectation of inflation in one period, $X_t = \mathrm{E}_t[\pi_{t+1}]$. This factor tends to make π^* an increasing function of X_t. It is the dominant factor when output has a large weight in the objective function, i.e., when μ is low.

Another factor can make π^* a decreasing function of X_t: the future risk of hitting the zero bound on interest rates. For some parameter values, either a higher X_t or a higher $\hat{\pi}_t$—that is, a higher expectation of inflation at $t+1$ or a higher target for $t+2$—raises inflation at $t+3$ and later. As a result, X_t and π^* are substitutes for reducing zero-bound risk, and that fact tends to reduce the optimal π^* when X_t is high.

4.3 Policy After a Low-Inflation Episode

In describing AIT, Fed officials have stressed the implications for policy following a period of low inflation. After a low-inflation episode, the Fed "will likely aim to achieve inflation moderately above 2 percent for some time." Does optimal policy in our model include this kind of overshooting of the long run inflation goal?

We address this question with an exercise in which inflation is initially low. Starting from that point, we derive the path of inflation if the central bank follows the optimal policy and the realizations of shocks are zero. We interpret this "intended inflation path" as policymakers' plan for returning inflation to a higher level. Of course actual inflation will deviate from this path when shocks occur.

Specifically, we assume that the inflation rate starts at 1% in period $t = 0$. For simplicity, the output gap is zero in period 0. We derive the intended inflation path for the set of parameter values in Table 1, which imply the $\pi^*(\cdot)$ functions in Figure 2. We compare results when the parameter μ in the objective function is 1.0 and 0.5.

In this example, the variable $X_0 = \mathrm{E}_0[\pi_1]$ is determined by current output and inflation. From the Phillips curve and expectations equation,

$$\mathrm{E}_0[\pi_1] = 2\gamma + (1-\gamma)\pi_0 + \alpha y_0 = 1 + \gamma. \tag{11}$$

Figure 3: Intended Inflation Paths

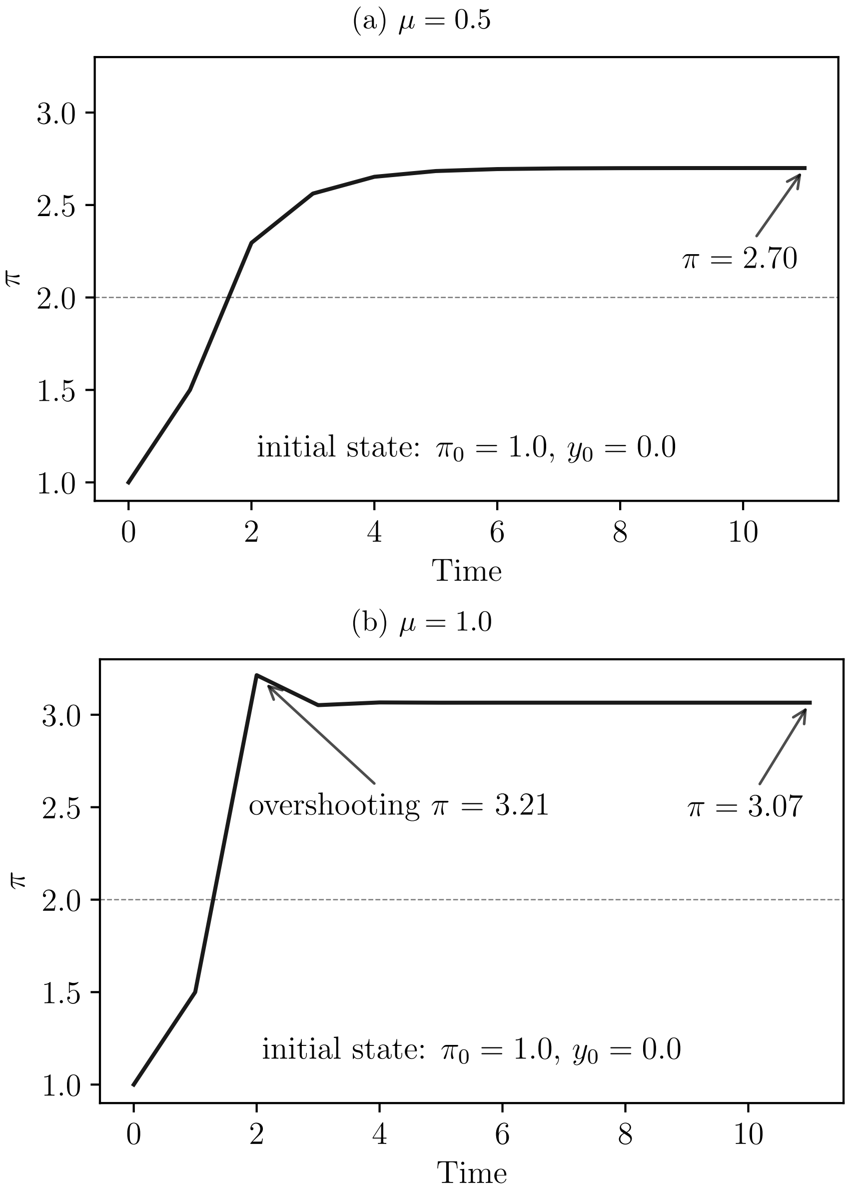

The anchoring parameter γ is 0.5, so $E_0[\pi_1]$ is 1.5. With shocks equal to zero, the actual inflation rate in period 1, π_1, is also 1.5.

Starting in period 2, the inflation rate is determined by the optimal policy rule. With zero shocks, the upper bound on the target is never binding, and actual inflation equals the target set two periods earlier: $\pi_t = \hat{\pi}_{t-2} = \pi^*(X_{t-2})$. Further, with zero shocks, $X_{t-2} = E_{t-2}[\pi_{t-1}] = \pi_{t-1}$. Combining these facts, we obtain $\pi_t = \pi^*(\pi_{t-1})$, which recursively defines the inflation path for $t \geq 2$.

Figure 3 shows the intended inflation paths for $\mu = 0.5$ and $\mu = 1.0$. In both cases, we see that inflation converges to a constant level, which is defined by the fixed point of $\pi = \pi^*(\pi)$. This level exceeds 2%: as in the case of full credibility, policy targets inflation above 2% when it can to offset periods when shocks force the target below 2%.

While inflation converges to similar levels for $\mu = 0.5$ and $\mu = 1.0$, it takes different paths in the two cases. For $\mu = 0.5$, the $\pi^*(\cdot)$ function is upward-sloping, so π_t is increasing in π_{t-1}. This implies monotonic convergence: inflation rises steadily until it levels off. This process is gradual, with inflation still rising noticeably in period 5. Faster adjustment would be sub-optimal because it would require larger output movements.

For $\mu = 1$, the $\pi^*(\cdot)$ function is downward sloping, so π_t is decreasing in π_{t-1}. This implies non-monotonic convergence: the inflation rate jumps to 3.21 in period 2 before converging to 3.07. The reason for this overshooting is that low inflation in period 0 implies a high risk of hitting the zero bound if adverse shocks occur. This risk creates an incentive to target high inflation, which reduces the risk by pulling up expected inflation π^e.[7]

The overshooting in Figure 3 bears some resemblance to the Fed's interpretation of AIT. In both the model and the Fed's rationale for its policy, temporarily high inflation can be optimal because it pulls up expected inflation and thereby reduces zero-bound risk. In the model, however, the rise in the intended inflation rate is only partly reversed: the inflation target remains above 2% indefinitely, falling farther only when it is forced down by the zero bound. The Fed aims for above-2% inflation only "for some time," which suggests that its target will return to 2% even if policy is unconstrained.

5 Conclusion

In 2025 the Federal Reserve will formally review the average-inflation-targeting strategy that it announced in 2020, so now is a good time to analyze the strategy. To that end, this paper derives optimal policy rules in a model with the key features that motivate AIT:

[7]The inflation rate continues to oscillate after period 2, but the oscillations are too small to be noticeable in the Figure.

the effective lower bound on interest rates, and the anchoring of inflation expectations. We seek to capture these features in the simplest possible model.

With fully anchored expectations, the optimal policy is to target a fixed level of inflation above 2% whenever the target is not forced lower by the bound on interest rates. When expectations are only partially anchored and respond to movements in actual inflation, the optimal short-run target varies with the state of the economy. For some parameter values, the target is highest after a period of low inflation, when a high target pulls up expected inflation and reduces the risk of hitting the interest-rate bound.

A Appendix

A.1 Proof of Proposition 1

We prove Proposition 1, which states the optimal policy under full anchoring, in several steps, each of which is summarized by its own proposition. We state these intermediate results, then show how they lead to Proposition 1, and then prove the intermediate results.

The first step is to rewrite the policy problem. The text explains that a rule for setting the policy interest rate i_t can be reinterpreted as a rule for setting an inflation target $\hat{\pi}_t = \mathrm{E}_t[\pi_{t+2}]$. A policy rule can also be interpreted as a rule for choosing a target for next period's output, $\hat{y}_t = \mathrm{E}_t[y_{t+1}]$, because the IS equation (1) implies a one-for-one relation between i_t and $\mathrm{E}_t[y_t]$ given the current state of the economy. The first step in our proof of Proposition 1 is a result that states the equivalence of the inflation targeting problem to a simpler problem of choosing an output target to minimize the variance of output:

Proposition A.1. *Consider the problem:*

$$\begin{aligned} &\min_{\hat{y}(\Omega_t)} \ \mathrm{Var}(y) \\ \text{s.t.} \ \ &\hat{y}_-(\Omega_t) \le \hat{y}_t = \hat{y}(\Omega_t) \le \lambda y_t + b \\ &\mathrm{E}[y] = 0, \end{aligned} \tag{A}$$

where $\hat{y}_-(\Omega_t) = \lambda y_t + b - \beta i^-$ and $b = \beta(2.0 + r^*)$. If $\hat{y}(\Omega_t)$ solves this problem, then the inflation targeting policy $\hat{\pi}(\Omega_t) = 2 + \alpha \hat{y}(\Omega_t)$ solves the policy problem stated in Section 2 of the text.

The next step is to tackle the problem in Proposition A.1 by showing it is related to yet another problem. This problem modifies the loss function by removing the constraint on $\mathrm{E}[y]$ and instead introducing a penalty on that term.

Proposition A.2. *Consider the problem:*

$$\begin{aligned} &\min_{\hat{y}(\Omega_t)} \left\{ \mathrm{E}[y^2] + \theta\, \mathrm{E}[y] \right\} \\ \text{s.t.} \ \ &\hat{y}_-(\Omega_t) \le \hat{y}_t = \hat{y}(\Omega_t) \le \lambda y_t + b \end{aligned} \tag{B}$$

If there exists some θ such that a policy $\hat{y}(\Omega_t)$ solves Problem (B) *and* $\mathrm{E}[y] = 0$ *under that policy, then $\hat{y}(\Omega_t)$ also solves Problem* (A).

By Proposition A.2, we need only find an optimal policy for the above problem such that $\mathrm{E}[y] = 0$. However, Problem (B) is still difficult to solve directly. Instead, we consider a discounted version of the problem first and use dynamic programming to find a solution. We then take the limit as the discount factor approaches one and show that the limit policy solves Problem (B).

Proposition A.3. *Consider the discounted problem:*

$$\begin{aligned} \min_{\hat{y}(\Omega_t)} & \left\{ \mathrm{E} \sum_{t=0}^{\infty} \delta^t (y_t^2 + \theta y_t) \right\} \\ \text{s.t.} \quad & \hat{y}_- \le \hat{y}_t = \hat{y}(\Omega_t) \le \lambda y_t + b \end{aligned} \tag{C}$$

where $\hat{y}_- = (b - \lambda\bar{\epsilon} - \beta i^-)/(1-\lambda)$. *For any* $\theta \in \mathbb{R}$ *and* $\delta \in (0,1)$, *there exists a* $y^* \in \mathbb{R}$ *such that* $\hat{y}(\Omega_t) = \min\{y^*, \lambda y_t + b\}$ *is a unique optimal policy for Problem* (C)*.*

Notice that the lower bound on $\hat{y}_t$ in Problem (C) is now a constant, which makes it easier to apply dynamic programming. This bound is derived by first calculating y_-, the lowest value y_t can attain under $i = i^-$, and then setting $\hat{y}_- = y_- + \bar{\epsilon}$. It can be seen that this bound is smaller than $\hat{y}_-(\Omega_t)$ in any state. Given Proposition A.3 for Problem (C), we are able to establish the following result about Problem (B):

Proposition A.4. *There exist a* θ *and a* y^* *in* $\mathbb{R}$ *such that* $\hat{y}(\Omega_t) = \min\{y^*, \lambda y_t + b\}$ *solves Problem* (B) *and* $\mathrm{E}[y] = 0$*.*

Given the propositions above, we are now able to prove Proposition 1.

Proof of Proposition 1. It follows from Proposition A.4 and Proposition A.2 that there exists a y^* such that the policy rule $\hat{y}(\Omega_t) = \min\{y^*, \lambda y_t + b\}$ solves Problem (A). Proposition A.1 then implies that the solution to the original policy problem in the text is $\hat{\pi}(\Omega_t) = \min\{\pi^*, \bar{\pi}_t\}$, where $\pi^* = 2 + \alpha y^*$ and $\bar{\pi}_t = 2 + \alpha(\lambda y_t + b)$. This expression for $\bar{\pi}_t$ is equal to the expression in our main theorem about optimal policy. To establish the claim in that theorem, it remains to show that the target away from the zero bound, π^*, exceeds 2, which follows if $y^* > 0$. That condition holds because $y \le y^*$ in all states and $y < y^*$ in some states, which imply $\mathrm{E}[y] < y^*$, and $\mathrm{E}[y] = 0$. □

A.2 Proofs of Propositions A.1–A.4

Proof of Proposition A.1. From the Phillips curve, equation (3), $\mathrm{E}_t[\pi_{t+2}] = 2 + \alpha\,\mathrm{E}_t[y_{t+1}]$, or $\hat{\pi}_t = 2 + \alpha\hat{y}_t$. Therefore an output-targeting rule $\hat{y}(\Omega_t)$ is equivalent to an inflation targeting rule $\hat{\pi}(\Omega_t) = 2 + \alpha\hat{y}(\Omega_t)$. In addition, the constraint $\hat{y}_t \le \lambda y_t + b$ is equivalent to $\hat{\pi}_t \le \bar{\pi}_t$, where $\bar{\pi}_t = 2 + \alpha(\lambda y_t + b)$, which is the constraint in the problem for choosing the optimal $\hat{\pi}(\Omega_t)$. The lower bound on $\hat{y}_t$ is also equivalent to the lower bound on $\hat{\pi}_t$. Finally, the constraint $\mathrm{E}[\pi] = 2$ is equivalent to $\mathrm{E}[y] = 0$.

The objective function in the policy problem is $(1-\mu)\,\mathrm{Var}(y) + \mu\,\mathrm{Var}(\pi)$. From the Phillips curve,

$$(1-\mu)\,\mathrm{Var}(y) + \mu\,\mathrm{Var}(\pi) = (1-\mu)\,\mathrm{Var}(y) + \mu[\mathrm{Var}(\alpha y) + \mathrm{Var}(\eta)]$$

$$= [1 - \mu + \mu\alpha^2]\,\mathrm{Var}(y) + \mu\,\mathrm{Var}(\eta)$$

because η_t is uncorrelated with y_{t-1}. Minimizing this expression is equivalent to minimizing $\mathrm{Var}(y)$ because the other terms are constants. Therefore, if a policy rule $\hat{y}(\Omega_t)$, or equivalently $\hat{\pi}(\Omega_t) = 2 + \alpha\hat{y}(\Omega_t)$, solves Problem (A), then it solves the original policy problem in Proposition 1, because the objective function and constraints are equivalent. □

To simplify notation, for any output process y generated by $\hat{y}(\Omega_t)$, we denote its mean by $\mathrm{E}[\hat{y}]$ and its variance by $\mathrm{Var}(\hat{y})$.

Proof of Proposition A.2. Assume there exists a θ such that $\hat{y}(\Omega_t)$ solves Problem (B) and $\mathrm{E}[\hat{y}] = 0$. Then, $\mathrm{E}[\hat{y}^2] = \mathrm{E}\left(\hat{y} - \mathrm{E}[\hat{y}]\right)^2 = \mathrm{Var}(\hat{y})$. Suppose some other policy rule $\hat{y}'(\Omega_t)$ solves Problem (A), which implies $\mathrm{E}[\hat{y}'] = 0$ and $\mathrm{Var}(\hat{y}') = \mathrm{E}[\hat{y}'^2] < \mathrm{Var}(\hat{y}) = \mathrm{E}[\hat{y}^2]$. Then $\mathrm{E}[\hat{y}'^2] + \theta\,\mathrm{E}[\hat{y}'] < \mathrm{E}[\hat{y}^2] + \theta\,\mathrm{E}[\hat{y}]$, so $\hat{y}$ no longer solves Problem (B). This is a contradiction. □

Proof of Proposition A.3. Given the bounds on $\hat{y}$, Problem (C) is well-defined on a compact state space $[y_-, y^-]$ where $y_- = \hat{y}_- - \bar{\epsilon}$ and $y^- = (b + \bar{\epsilon})/(1 - \lambda)$. The lower bound y_- is the lowest value y_t can attain given $\hat{y}_-$. The upper bound y^- is the steady state y_t if $i_t = 0$ and $\epsilon_t = \bar{\epsilon}$ in all periods. These bounds ensure that $y_{t+1} \in [y_-, y^-]$ whenever $y_t \in [y_-, y^-]$. Since i^- can be arbitrarily large, we can make $\hat{y}_-$ arbitrarily small and assume that $\hat{y}_- < -\theta/2$. The Bellman equation is

$$v(y) = \min_{\hat{y}_- \le \hat{y} \le \lambda y + b} \{y^2 + \theta y + \delta\,\mathrm{E}\,v(\hat{y} + \epsilon)\}. \tag{12}$$

Dynamic programming arguments (see, e.g., Stokey and Lucas 1989, Chapter 9.2) imply that $v(y)$ is the value function and the optimal policy rule is given by

$$\hat{y}(y) = \operatorname*{arg\,min}_{\hat{y}_- \le \hat{y} \le \lambda y + b} \{y^2 + \theta y + \delta\,\mathrm{E}\,v(\hat{y} + \epsilon)\},$$

which is equivalent to

$$\hat{y}(y) = \operatorname*{arg\,min}_{\hat{y}_- \le \hat{y} \le \lambda y + b} \mathrm{E}\,v(\hat{y} + \epsilon).$$

Since $y^2 + \theta y$ is strictly convex and the constraint set $[\hat{y}_-, \lambda y + b]$ is convex in y, Theorem 9.8 of Stokey and Lucas (1989) implies that v is strictly convex. It then follows from Stokey and Lucas (1989, Lemma 9.5) that $\mathrm{E}\,v(\hat{y} + \epsilon)$ is strictly convex in $\hat{y}$. Therefore, there exists a unique minimizer $y^* \in [\hat{y}_-, y^- - \bar{\epsilon}]$. In other words, y^* is the optimal policy in the absence of the zero lower bound in the current period. We claim that $y^* > \hat{y}_-$ and that the optimal policy for Problem (C) is given by $\hat{y}(y) = \min\{y^*, \lambda y + b\}$.

To see this, we first show that $y^* > \hat{y}_-$. Suppose $y^* = \hat{y}_-$. Since $\hat{y}_-$ is always in the constraint set, the optimal policy rule is given by $\hat{y}(y) = \hat{y}_-$ for all y. Since v satisfies the Bellman equation (12), $v(y)$ is given by $y^2 + \theta y + \delta C$, where $C = \mathrm{E}\, v(\hat{y}_- + \epsilon)$ is a constant. If this is the case, however,

$$\mathrm{E}\, v(\hat{y} + \epsilon) = \mathrm{E}\left[(\hat{y} + \epsilon)^2 + \theta(\hat{y} + \epsilon) + \delta C\right] = \hat{y}^2 + \theta \hat{y} + \mathrm{Var}(\epsilon) + \delta C,$$

which is decreasing on $[\hat{y}_-, -\theta/2]$ because $\hat{y}_- < -\theta/2$. This implies that $\hat{y}_-$ is not always optimal, which is a contradiction. Hence, $y^* > \hat{y}_-$.

By convexity of v, $\hat{y} \mapsto \mathrm{E}\, v(\hat{y}+\epsilon)$ is decreasing on $[\hat{y}_-, y^*]$ and increasing on $[y^*, y^- - \bar{\epsilon}]$. It follows that the optimal policy rule is $\hat{y}(y) = \min\{\lambda y + b, y^*\}$. □

Proof of Proposition A.4. The theorem is proved in the following steps.

First, we consider a variant of Problem (B):

$$\begin{aligned} &\min_{\hat{y}(\Omega_t)} \left\{\mathrm{E}[y^2] + \theta\, \mathrm{E}[y]\right\} \\ &\text{s.t. } \hat{y}_- \leq \hat{y}_t = \hat{y}(\Omega_t) \leq \lambda y_t + b \end{aligned} \tag{B'}$$

where the lower bound on $\hat{y}_-$ is fixed as in Problem (C). Problem (B') is the same as Problem (C) without discounting. By Lemma A.6, we can find the solution to (B') by taking the limit as $\delta \to 1$ of the solution to (C), given in Proposition A.3. This implies that the optimal policy for Problem (B') is given by $\hat{y}_\theta(y) = \min\{\lambda y + b, y^*(\theta)\}$. It remains to show that $\hat{y}_\theta$ also solves Problem (B) for some θ with $\mathrm{E}[\hat{y}_\theta] = 0$.

We first prove that there exists a θ such that $\mathrm{E}[\hat{y}_\theta] = 0$. Lemma A.6 shows that as we vary θ, $y^*(\theta)$ changes continuously. It follows from Lemma A.7 that $\mathrm{E}[\hat{y}_\theta]$ is continuous in θ. By Lemma A.8, there exist θ_1 and θ_2 such that $\mathrm{E}[\hat{y}_{\theta_1}] > 0$ and $\mathrm{E}[\hat{y}_{\theta_2}] < 0$. It then follows from the intermediate value theorem that there exists a θ with $\mathrm{E}[\hat{y}_\theta] = 0$.

Finally, fix θ such that $\mathrm{E}[\hat{y}_\theta] = 0$. Lemma A.9 states that the solution to Problem (B') is also the solution to Problem (B). Therefore, $\hat{y}(\Omega_t) = \min\{\lambda y_t + b, y^*(\theta)\}$ solves Problem (B). □

A.3 Sketch Proof of Proposition 2

We adopt a similar approach to the proof of Proposition 1, by considering an equivalent problem of choosing an output target $\hat{y}(\Omega_t)$ and then introducing a penalty on $\mathrm{E}[y]$ to remove the constraint on long-run inflation:

$$\begin{aligned} &\min_{\hat{y}(\Omega_t)} \left\{(1-\mu)\, \mathrm{E}[y^2] + \mu\, \mathrm{E}[(\pi - 2)^2] + \theta\, \mathrm{E}[y]\right\} \\ &\text{s.t. } \hat{y}_-(\Omega_t) \leq \hat{y}_t = \hat{y}(\Omega_t) \leq \lambda y_t + \beta(\pi_t^e + r^*) \end{aligned} \tag{13}$$

where expected inflation π_t^e follows (5) and $\hat{y}_-(\Omega_t) = \lambda y_t + \beta(\gamma(2) + (1-\gamma)\pi_t + r^*) - \beta i^-$. If for some θ the solution to (13) satisfies $\mathrm{E}[y] = 0$, then it also solves the original problem. ($\mathrm{E}[y] = 0$ is equivalent to $\mathrm{E}[\pi] = 2$.)

The discounted version of this problem is

$$\begin{aligned}&\min_{\hat{y}(\Omega_t)} \mathrm{E}\sum_{t=0}^{\infty} \delta^t \left[(1-\mu)y_t^2 + \mu(\pi_t - 2)^2 + \theta y_t\right] \\ &\text{s.t. } \ \hat{y}_-(\Omega_t) \le \hat{y}_t = \hat{y}(\Omega_t) \le \lambda y_t + \beta(\pi_t^e + r^*)\end{aligned} \tag{14}$$

The Bellman equation is given by

$$v(y,\pi) = \min_{\hat{y}_-(y,\pi)\le\hat{y}\le\bar{y}(y,\pi)} \left\{(1-\mu)y^2 + \mu(\pi-2)^2 + \theta y + \delta\,\mathrm{E}\,v\left(\hat{y}+\epsilon, X+\eta\right)\right\}, \tag{15}$$

where $\bar{y}(y,\pi) = \lambda y + \beta(\gamma(2) + (1-\gamma)\pi + r^*)$ and $X = \gamma(2) + (1-\gamma)\pi + \alpha y$. Note that we need two state variables due to partially anchored expectations. We have the following proposition similar to Proposition A.3:

Proposition A.5. *If* γ *satisfies*

$$\gamma > \frac{1}{1+\frac{1-\lambda}{\alpha\beta}}, \tag{16}$$

then Problem (14) *has a solution and the optimal policy rule is given by*

$$\hat{y}(y,\pi) = \min\left\{y^*(X), \bar{y}(y,\pi)\right\} \tag{17}$$

for some function y^**.*

Condition (16) ensures that there exist $\Pi = [\pi_-, \pi^-]$ and $Y = [y_-, y^-]$ such that $(y_t, \pi_t) \in Y \times \Pi$ for all t. To see why this is the case, we write the law of motion for y_t and π_t as:

$$\begin{bmatrix} y_{t+1} \\ \pi_{t+1} \end{bmatrix} = A \begin{bmatrix} y_t \\ \pi_t \end{bmatrix} + \begin{bmatrix} \beta(\gamma(2) + r^*) - \beta i_t + \epsilon_{t+1} \\ \gamma(2) + \eta_{t+1} \end{bmatrix}, \quad A = \begin{bmatrix} \lambda & \beta(1-\gamma) \\ \alpha & 1-\gamma \end{bmatrix}.$$

Then, the upper bounds for y_t and π_t are given by the steady state of the above system when $i_t = 0$, $\epsilon_t = \bar{\epsilon}$, and $\eta_t = \bar{\eta}$ for all t. Similarly, the lower bounds can be computed by setting $i_t = i^-$, $\epsilon_t = -\bar{\epsilon}$, and $\eta_t = -\bar{\eta}$ for all t.

These steady states exist if and only if the spectral radius of A, $\rho(A)$, is less than one. By deriving the eigenvalues of this matrix, we can show that $\rho(A) < 1$ reduces to condition (16). Then we have

$$\begin{bmatrix} y^- \\ \pi^- \end{bmatrix} = (I-A)^{-1} \begin{bmatrix} \beta(\gamma(2) + r^*) + \bar{\epsilon} \\ \gamma(2) + \bar{\eta} \end{bmatrix}, \begin{bmatrix} y_- \\ \pi_- \end{bmatrix} = (I-A)^{-1} \begin{bmatrix} \beta(\gamma(2) + r^*) - \beta i^- - \bar{\epsilon} \\ \gamma(2) - \bar{\eta} \end{bmatrix}.$$

The way we choose the bounds ensures that $(y_{t+1}, \pi_{t+1}) \in Y \times \Pi$ whenever $(y_t, \pi_t) \in Y \times \Pi$.

The existence of a compact state space allows us to apply dynamic programming techniques, which show that the value function v is strictly convex and the optimal policy without the ELB in the current period depends only on X, the expected inflation next period. In fact, $y^*(X)$ is given by

$$y^*(X) = \arg\min_{\hat{y}} \{\mathrm{E}\, v\, (\hat{y} + \epsilon, X + \eta)\} \tag{18}$$

since $\hat{y}$ does not affect current y or π.

Note that the policy rule given by (17) depends on the discount factor δ and the parameter θ. Similar to the proof of Proposition 1, we first take the limit as δ approaches one and then choose a θ such that $\mathrm{E}[y] = 0$. Those steps yield the solution $\hat{y}(y, \pi)$ for the original policy problem. From $\hat{y}$ and the Phillips curve, the optimal rule for the inflation target is given by $\hat{\pi}(y, \pi) = \min\{\pi^*(X), \bar{\pi}(y, \pi)\}$ where $\pi^*(X) = \alpha y^*(X) + (1-\gamma)X + \gamma(2)$.

A.4 Numerical Method

For the model with imperfect anchoring, we solve the discounted version of the policy problem (14) with a discount factor close to 1 ($\delta = 0.997$). By a result similar to Lemma A.6, the solution to this problem will be close to the optimal policy for the true undiscounted problem.

To solve the discounted problem for a given θ, we first solve the Bellman equation (15) using optimistic policy iteration (Bertsekas 2013, Chapter 2.5). After obtaining the value function v, we compute the policy rule $\hat{y}$ using (17) and (18).

The optimal policy in our model is the solution to (14) for θ such that $\mathrm{E}[\hat{y}] = 0$. We determine the optimal policy with an iterative procedure in which we make an initial guess of θ; solve (14) to get the policy rule for that θ; calculate $\mathrm{E}[y]$ under that rule by simulating the model for 10 million periods; and choose a revised θ with Brent's root finding algorithm (Brent 2013). We repeat this process until we find a policy with $\mathrm{E}[y] = 0$.

A.5 Further Results for the Model with Imperfect Anchoring

Section 4 shows that the optimal policy rule with imperfect anchoring is $\hat{\pi}_t = \min\{\pi^*(X_t), \bar{\pi}_t\}$, where $X_t = \mathrm{E}_t[\pi_{t+1}]$. We also give examples of the $\pi^*(X_t)$ function for specific parameter values. Here we delve further into the details of what determines the shape of $\pi^*(X_t)$.

As discussed in the text, one determinant of $\pi^*(\cdot)$ is the value of μ, the weight on inflation variance in the policymaker's objective function. To build understanding, we first consider the case in which $\mu = 1$—the policymaker cares only about inflation variance—and see how the $\pi^*(\cdot)$ function is influenced by the other parameters.

The Case of $\mu = 1$ In this case, the function $\pi^*(\cdot)$ is always monotonic. Its slope can be either positive or negative and depends on a simple condition:

> Numerical Result: When $\mu = 1$, π^* is increasing (decreasing) in X if $(\lambda - \alpha\beta)$ is positive (negative).

That is, the shape of $\pi^*(\cdot)$ is determined by three of the model's parameters: the coefficients in the IS curve on lagged output (λ) and the real interest rate (β), and the slope of the Phillips curve (α).

We have not been able to prove this result analytically. We have, however, solved the model numerically for wide ranges of parameter values and found that the result always holds. We have also derived an analytical result for the knife-edge case of $\lambda = \alpha\beta$: For that case, π^* is a constant independent of X, as in the model with full credibility. (Proof omitted.)

Examining the model can help us understand this result. The choice at t of an inflation target, $\hat{\pi}_t = \mathrm{E}_t[\pi_{t+2}]$, affects the economy through two channels: It affects π_{t+2}, and it affects $\bar{\pi}_{t+1}$, the upper bound on the inflation target $\hat{\pi}_{t+1}$ that the policymaker will face at $t+1$. The behavior of $\pi^*(X_t)$ involves the somewhat complex interaction of X_t and $\bar{\pi}_{t+1}$.

To see this, note that $\bar{\pi}_{t+1}$ is a bound on the choice at $t+1$ of a target for π_{t+3}. Using the Phillips curve (3), this bound is

$$\bar{\pi}_{t+1} = \mathrm{E}_{t+1}[\pi^e_{t+2}] + \alpha\bar{y}_{t+1} \tag{19}$$

where $\bar{y}_{t+1}$ is the bound on $\hat{y}_{t+1}$, the equivalent target for y_{t+2} set at $t+1$, which is determined by the IS equation (1) with $i_{t+1} = 0$. With considerable algebra involving the IS equation, Phillips curve, and equation (5) for π^e, along with the facts that $y_{t+1} = \hat{y}_t$ plus shocks at $t+1$ and π_{t+2} equals $\hat{\pi}_t$ plus shocks at $t+1$ and $t+2$, we can write the bound $\bar{\pi}_{t+1}$ as

$$\bar{\pi}_{t+1} = (1-\gamma)\hat{\pi}_t + \alpha[\lambda\hat{y}_t + \beta(1-\gamma)X_t] + \text{constant} + \text{shocks}_{t+1} \tag{20}$$

Next, using the Phillips curve and the equation (5) for π^e, we can derive the relationship between the output target $\hat{y}_t$ and the equivalent inflation target $\hat{\pi}_t$:

$$\hat{y}_t = \frac{1}{\alpha}\hat{\pi}_t - \frac{1-\gamma}{\alpha}X_t + \text{constant} \tag{21}$$

Finally, substituting (21) into (20) leads to

$$\bar{\pi}_{t+1} = (1-\gamma+\lambda)\hat{\pi}_t + (1-\gamma)(\alpha\beta-\lambda)X_t + \text{constant} + \text{shocks}_{t+1}. \tag{22}$$

Equation (22) shows that the expectation of the future bound $\bar{\pi}_{t+1}$ is determined by $\hat{\pi}_t$ and X_t. A higher $\hat{\pi}_t$ implies both higher inflation and higher output in the future, which raises the bound. A higher expectation of inflation X_t has an ambiguous effect. It directly raises the bound for given $\hat{\pi}_t$ and $\hat{y}_t$, as indicated by its positive coefficient in (20). More subtly, it reduces the output target corresponding to a given inflation target, as shown in (21), and this lower $\hat{y}$ reduces $\bar{\pi}_{t+1}$, as shown in (20). As shown in (22), the sign of the net effect is determined by the sign of $\lambda - \alpha\beta$.

The fact that a higher $\hat{\pi}_t$ relaxes the future upper bound $\bar{\pi}_{t+1}$ creates an incentive for the policymaker to choose a high $\hat{\pi}_t$. The effect of X_t on the choice of $\hat{\pi}_t$ depends on the complementarity or substitutability of X_t and $\hat{\pi}_t$ as means of influencing $\bar{\pi}_{t+1}$. When the coefficient on X_t is negative, a higher X_t implies a tighter bound at $t+1$ for a given $\hat{\pi}_t$, and that increases the incentive to raise $\hat{\pi}_t$. Therefore, in this case, which arises for $\lambda > \alpha\beta$, π_t^* is increasing in X_t. In contrast, when the coefficient on X_t is positive, a higher X_t is a substitute for a higher $\hat{\pi}_t$ in its effect on the future bound. As a result, a higher X_t reduces the optimal target: π_t^* is decreasing in X_t.

The Case of $\mu < 1$ When output variance is added to the objective function, we find one consistent result:

> Numerical Result: Holding constant all other parameter values, a decrease in μ raises $\pi^*(X_2) - \pi^*(X_1)$ for any $X_2 > X_1$.

That is, increasing the weight on output variance makes the slope of the $\pi^*(\cdot)$ function more positive at all points. If the function is upward sloping for $\mu = 1$, it is also upward sloping and steeper for $\mu < 1$. If the function is downward sloping for $\mu = 1$, then for $\mu < 1$ it is less downward sloping and may become upward sloping at some or all points.

Once again we have not proven this result analytically, but it arises in a wide range of simulations, including the examples given in Figure 2 of the text. In the Figure, $\pi^*(\cdot)$ is downward-sloping for $\mu = 1$. As μ falls, the slope becomes more positive, resulting first in a U-shaped function and then, for even lower μ, an upward-sloping function. As discussed in the text, a stronger desire to reduce output variance (a smaller μ) gives the policymaker an incentive to target an inflation rate at $t+2$ that is close to the inflation rate at $t+1$, and that tends to make the target an upward-sloping function of $X_t = \mathrm{E}_t[\pi_{t+1}]$.

A.6 Technical Lemmas

Lemma A.6. *For any $\theta \in \mathbb{R}$ and $\delta \in (0,1)$, let $y^*(\theta, \delta)$ denote the optimal output target from Proposition A.3. Then, there exists $y^*(\theta)$ such that $\hat{y}_\theta(y) = \min\{\lambda y + b, y^*(\theta)\}$ solves Problem* (B')*. Moreover,*

(1) $y^*(\theta) = \lim_{\delta \to 1} y^*(\theta, \delta)$;

(2) $y^*(\theta)$ changes continuously with θ.

Proof. Fix θ and $\delta \in (0,1)$. Proposition A.3 guarantees the existence of an optimal policy $\hat{y}_{\theta,\delta}(y) = \min\{y^*(\theta,\delta), \lambda y + b\}$. Pick any y_0, y_0' and let $\{y_t\}$ and $\{\tilde{y}_t\}$ be generated by $\hat{y}_{\theta,\delta}$ starting from y_0 and y_0', respectively. It can be shown that there exists a $\rho \in (0,1)$ such that Corollary 11.2.22 of Stachurski (2009) applies and we have

$$\left|\mathrm{E}(y_{Nt}^2 + \theta y_{Nt}) - \mathrm{E}(\tilde{y}_{Nt}^2 + \theta \tilde{y}_{Nt})\right| \leq (1-\rho)^t M$$

for some $N \geq 1$, and $M < \infty$. Then, Lemma 11.1.28 of Stachurski (2009) implies that

$$\begin{aligned} |v(y_0) - v(y_0')| &\leq \left| \mathrm{E}\sum_{t=0}^{\infty} \delta^t (y_t^2 + \theta y_t) - \mathrm{E}\sum_{t=0}^{\infty} \delta^t (\tilde{y}_t^2 + \theta \tilde{y}_t) \right| \\ &\leq \sum_{t=0}^{\infty} \left|\mathrm{E}(y_t^2 + \theta y_t) - \mathrm{E}(\tilde{y}_t^2 + \theta \tilde{y}_t)\right| \leq \frac{\bar{M}}{\rho} \end{aligned}$$

for some $\bar{M} < \infty$. Hence, the value-boundedness assumption of Dutta (1991) is satisfied. Therefore, by Theorem 3 of Dutta (1991), $v_\theta(y) = \lim_{\delta\to1}(1-\delta)v_{\theta,\delta}(y)$ is the value function for Problem (B') and $\hat{y}_\theta(y) = \lim_{\delta\to1} \hat{y}_{\theta,\delta}(y)$, if the limit exists, is an optimal policy. Since $\hat{y}_{\theta,\delta}(y) = \min\{\lambda y + b, y^*(\theta,\delta)\}$ is increasing in y, by Corollary 1 of Dutta (1991), $\hat{y}_\theta(y) = \min\{\lambda y + b, y^*(\theta)\}$ exists for some $y^*(\theta)$.

Next we prove the continuity of $y^*(\theta)$. Pick any sequence of θ_n such that $\lim_{n\to\infty} \theta_n = \theta$. By Theorem 5.1 of Langen (1981), $y^*(\theta,\delta)$ is continuous in θ and δ for all $\theta \in \mathbb{R}$ and $\delta \in (0,1)$. By the results above, for each θ_n, there exist $\delta_{m,n} \to 1$ such that $y^*(\theta_n, \delta_{m,n}) \to y^*(\theta_n)$ and there exist $\delta_m \to 1$ such that $y^*(\theta,\delta_m) \to y^*(\theta)$. Since

$$|y^*(\theta_n) - y^*(\theta)| \leq |y^*(\theta_n) - y^*(\theta_n,\delta_{m,n})| + |y^*(\theta_n,\delta_{m,n}) - y^*(\theta,\delta_m)| + |y^*(\theta,\delta_m) - y^*(\theta)|,$$

we can choose m and n sufficiently large to make $|y^*(\theta_n) - y^*(\theta)|$ arbitrarily small. Therefore, $y^*(\theta_n) \to y^*(\theta)$. □

Lemma A.7. *Let $\hat{y}^x(y) = \min\{\lambda y + b, x\}$. Then $\mathrm{E}[\hat{y}^x]$ is continuous and strictly increasing in x.*

Proof. That $\mathrm{E}[\hat{y}^x]$ is strictly increasing in x is obvious. To show continuity, note that for any y and any $\epsilon > 0$, there exists $\tau > 0$ such that $|x - x'| < \tau$ implies that $|\hat{y}^x(y) - \hat{y}^{x'}(y)| < \epsilon$. Moreover, for any y, y', $|\hat{y}^x(y) - \hat{y}^x(y')| \leq \lambda|y - y'|$. Therefore, $|\hat{y}^x(y) - \hat{y}^{x'}(y')| \leq \lambda|y-y'| + \epsilon$. Let $\{y_t\}$ and $\{\tilde{y}_t\}$ be generated by $\hat{y}^x$ and $\hat{y}^{x'}$, respectively. Then, for any given sequence of shocks $\{\epsilon_n\}$, $|y_t - \tilde{y}_t| \leq \epsilon(1 + \lambda + \ldots + \lambda^t) = \epsilon(1-\lambda^t)/(1-\lambda)$. Therefore,

$$\begin{aligned} |\,\mathrm{E}[\hat{y}^x] - \mathrm{E}[\hat{y}^{x'}]| &\leq \lim_{n\to\infty} \frac{1}{n} \sum_{t=0}^{n-1} \epsilon \frac{1-\lambda^t}{1-\lambda} \\ &\leq \lim_{n\to\infty} \frac{1}{n} \sum_{t=0}^{n-1} \epsilon \frac{1}{1-\lambda} = \epsilon \frac{1}{1-\lambda}. \end{aligned}$$

Since ϵ is arbitrary, $\mathrm{E}[\hat{y}^x]$ is continuous in x. □

Lemma A.8. *For $\hat{y}_\theta$ given in Lemma A.6, there exist $\theta_1, \theta_2 \in \mathbb{R}$ such that $\mathrm{E}[\hat{y}_{\theta_1}] > 0$ and $\mathrm{E}[\hat{y}_{\theta_2}] < 0$.*

Proof. We can choose θ_1 sufficiently small such that $y^2 + \theta_1 y$ is strictly decreasing on the state space $[y_-, y^-]$. By Theorem 9.7 of Stokey and Lucas (1989), the value function v is also strictly decreasing. Then, the optimal policy is $\hat{y}_{\theta_1}(y) = \lambda y + b$, or equivalently, $\hat{y}_{\theta_1}(y) = \min\{\lambda y + b, y^-\}$. Similarly, we can choose θ_2 sufficiently large such that $y^*(\theta_2) = \hat{y}_-$. By our assumption on $\bar{\epsilon}$, $\hat{y}_- < 0$. It follows that $\mathrm{E}[\hat{y}_{\theta_1}] > 0$ and $\mathrm{E}(\hat{y}_{\theta_2}) < 0$. □

Lemma A.9. *If for some θ, $\hat{y}(\Omega_t)$ solves Problem (B') with $\mathrm{E}[y] = 0$, then $\hat{y}(\Omega_t)$ also solves Problem (B).*

Proof. By construction, it suffices to show that the optimal policy for Problem (B') always satisfies the constraints in Problem (B). Since $\mathrm{E}[y] = 0$ under $\hat{y}(\Omega_t) = \min\{y^*, \lambda y_t + b\}$, $y^* > 0$. Since i^- is arbitrarily large, we have

$$y^* > 0 > \lambda y_t + b - \beta i^-$$

for all t. Therefore,

$$\lambda y_t + b - \beta i^- \leq \hat{y}(\Omega_t) \leq \lambda y_t + b$$

for all t. This completes the proof. □

References

Amano, Robert A, Stefano Gnocchi, Sylvain Leduc, and Joel Wagner. 2020. *Average Is Good Enough: Average-Inflation Targeting and the ELB.* Technical report. Bank of Canada Staff Working Paper.

Arias, Jonas, Martin Bodenstein, Hess Chung, Thorsten Drautzburg, and Andrea Raffo. 2020. "Alternative Strategies: How Do They Work? How Might They Help?" *Finance and Economics Discussion Series* 2020-068 (August): 1–31.

Ball, Laurence. 1999. "Efficient Rules for Monetary Policy." *International Finance* 2 (1): 63–83.

Ball, Laurence, Daniel Leigh, and Prachi Mishra. 2022. "Understanding US Inflation During the Covid-19 Era." *Brookings Papers on Economic Activity* 2022 (2): 1–80.

Bertsekas, Dimitri P. 2013. *Abstract Dynamic Programming.* Athena Scientific Belmont, MA.

Blanchard, Olivier. 2016. "The Phillips Curve: Back to the'60s?" *American Economic Review* 106 (5): 31–34.

Brent, Richard P. 2013. *Algorithms for Minimization Without Derivatives.* Courier Corporation.

Budianto, Flora, Taisuke Nakata, and Sebastian Schmidt. 2023. "Average Inflation Targeting and the Interest Rate Lower Bound." *European Economic Review* 152:104384.

Clarida, Richard H. 2019. "The Federal Reserve's Review of Its Monetary Policy Strategy, Tools, and Communication Practices," https://www.federalreserve.gov/newsevents/speech/clarida20190222a.htm.

———. 2020. "The Federal Reserve's New Monetary Policy Framework: A Robust Evolution," https://www.federalreserve.gov/newsevents/speech/clarida20200831a.htm.

Dutta, Prajit K. 1991. "What Do Discounted Optima Converge To?: A Theory of Discount Rate Asymptotics in Economic Models." *Journal of Economic Theory* 55 (1): 64–94.

Eo, Yunjong, and Denny Lie. 2020. "Average Inflation Targeting and Interest-Rate Smoothing." *Economics Letters* 189:109005.

Federal Reserve. 2020. "2020 Statement on Longer-Run Goals and Monetary Policy Strategy," https://www.federalreserve.gov/monetarypolicy/review-of-monetary-policy-strategy-tools-and-communications-statement-on-longer-run-goals-monetary-policy-strategy.htm.

Hebden, James, Edward P. Herbst, Jenny Tang, Giorgio Topa, and Fabian Winkler. 2020. "How Robust Are Makeup Strategies to Key Alternative Assumptions?" *Finance and Economics Discussion Series* 2020-069 (August): 1–42.

Honkapohja, Seppo, and Nigel McClung. 2024. "On Robustness of Average Inflation Targeting." *Available at SSRN 4021712.*

Kiley, Michael T, and John M Roberts. 2017. "Monetary Policy in a Low Interest Rate World." *Brookings Papers on Economic Activity* 2017 (1): 317–396.

Langen, Hans-Joachim. 1981. "Convergence of Dynamic Programming Models." *Mathematics of Operations Research* 6 (4): 493–512.

Mertens, Thomas M, and John C Williams. 2019. "Monetary Policy Frameworks and the Effective Lower Bound on Interest Rates." In *AEA Papers and Proceedings,* 109:427–432.

———. 2020. "Tying Down the Anchor: Monetary Policy Rules and the Lower Bound on Interest Rates." Chap. 3 in *Strategies for Monetary Policy,* edited by John H Cochrane and John B Taylor, 103–172. Hoover Press.

Nessén, Marianne, and David Vestin. 2005. "Average Inflation Targeting." *Journal of Money, Credit and Banking,* 837–863.

Powell, Jerome H. 2020. "New Economic Challenges and the Fed's Monetary Policy Review," https://www.federalreserve.gov/newsevents/speech/powell20200827a.htm.

Romer, Christina D, and David H Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review* 94 (4): 1055–1084.

Romer, David. 2019. *Advanced Macroeconomics.* 5th ed. McGraw Hill.

Stachurski, John. 2009. *Economic Dynamics: Theory and Computation.* MIT Press.

Stokey, Nancy L, and Robert E Lucas. 1989. *Recursive Methods in Economic Dynamics.* Harvard University Press.

Svensson, Lars EO. 1997. "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets." *European Economic Review* 41 (6): 1111–1146.

Tauchen, George. 1986. "Finite state markov-chain approximations to univariate and vector autoregressions." *Economics letters* 20 (2): 177–181.

Yellen, Janet. 2019. "Former Fed Chair Janet Yellen on Why the Answer to the Inflation Puzzle Matters." *Remark, Hutchins Center on Fiscal and Monetary Policy at Brookings, October* 3.