Average Inflation Targeting: Time Inconsistency and Intentional Ambiguity*

Chengcheng Jia
Federal Reserve Bank of Cleveland

Jing Cynthia Wu
University of Notre Dame and NBER

Current draft: November 8, 2021

Abstract

We study the implications of the Fed's new policy framework of average inflation targeting (AIT) and its ambiguous communication. The central bank has the incentive to deviate from its announced AIT, which improves the trade-off between inflation and real activity, and implement inflation targeting ex post to maximize social welfare. We show two motives for ambiguous communication about the horizon over which the central bank averages inflation as a result of time inconsistency. First, it is optimal for the central bank to announce different horizons depending on the state of the economy. Second, ambiguous communication helps the central bank gain credibility.

Keywords: average inflation targeting, time inconsistency, ambiguous communication.

1 Introduction

At the 2020 Jackson Hole Economic Policy Symposium, Federal Reserve Chair Jerome Powell announced a revision to the Fed's long-run monetary policy framework, replacing inflation targeting (IT) with average inflation targeting (AIT) to achieve its dual mandate; see Powell (2020). However, a critical component that is missing from the Fed's communication is the horizon over which it targets average inflation at 2%; see Kozicki's (2019) discussion at the Fed Listens event in 2019 and Brunnermeier's (2021) panel discussion at the Jackson Hole Economic Policy Symposium in 2021. Our paper investigates the implications of such an ambiguous communication.

We focus on two key issues of AIT: time inconsistency and ambiguous communication. We show AIT is not time consistent, which is the case for any path-dependent policy; see, for example, Eggertsson and Woodford (2003). By convincing the private sector about its intention to implement AIT, the central bank can improve the trade-off between inflation and real activity, which is captured by the Phillips curve. Ex post, the central bank has the incentive to deviate from its communication and implement IT instead to improve social welfare. This time-inconsistent strategy leads to two motives for ambiguous communication. First, the optimal horizon of AIT announcement is time dependent, and ambiguous communication allows such flexibility. Second, the central bank's time-inconsistent strategy is welfare-improving provided that the central bank is able to convince private agents with its communication. But can it? We show although the central bank implements IT de facto, AIT believers can have a smaller forecast error on the output gap after a cost-push shock. Moreover, announcing AIT without clearly specifying its horizon helps the central bank further gain credibility and improve welfare compared with the case with clear communication.

We are the first paper that studies AIT and deviates from the full-information rational expectations (FIRE) framework. We introduce AIT into an otherwise textbook three-equation New Keynesian model by modifying the central bank's objective function. For a standard model, the central bank minimizes a quadratic loss function in inflation and the output

gap, both of which are valued at the current period. AIT replaces current inflation in the objective function with average inflation over L periods between $t - L + 1$ and t.[1] In the FIRE benchmark, the central bank minimizes its loss subject to a forward-looking Phillips curve. Although the central bank is discretionary and solves its problem period by period, equilibrium is path dependent: inflation oscillates around its steady state after a cost-push shock.

The Phillips curve captures the trade-off between inflation and real activity the central bank faces. Although AIT does not change the forward-looking Phillips curve, it tilts its reduced form in a favorable way: compared with IT, AIT is associated with a Phillips curve with a smaller intercept and slope. This works through the expectations channel of AIT: when inflation is above its target today, the expected inflation next period will be below the target. A lower expected inflation lowers inflation today through the expectation term in the forward-looking Phillips curve.

In the zero lower bound (ZLB) environment, AIT not only flattens the Phillips curve, but also alters the IS curve. The IS curve is vertical for IT but is downward sloping for AIT, and this difference arises from the expectation terms in the IS curve. Whether the central bank implements IT or AIT, the ZLB features a below-target inflation and a negative output gap. However, AIT delivers a better equilibrium than IT with less negative inflation and output gap. This again works through the expectations channel of AIT: when the current inflation is below the target, the private sector expects inflation next period to be above the target, which generates higher demand and inflation at present. Consequently, AIT could be an effective tool to combat a liquidity trap.

Although AIT presents the central bank with a better inflation-output trade-off by tilting the Phillips curve in a favorable way, it does not necessarily yield higher welfare, because its objective function is different from social welfare. After the central bank convinces the private sector about its intention to implement AIT, it has an incentive to deviate from its

---

[1]We define AIT as flexible average inflation targeting, which also puts weight on the output gap. This definition is similar to flexible inflation targeting in the literature.

communicated objective and optimize social welfare instead. This strategy improves welfare but is time inconsistent.

The time-inconsistent nature of AIT motivates the central bank to make ambiguous announcements regarding the horizon over which it targets average inflation for two reasons. First, it is optimal for the central bank to announce different horizons for AIT depending on economic conditions. Specifically, the optimal strategy is to announce the largest feasible horizon when a cost-push shock hits and announce two-period AIT thereafter. Instead of making self-conflicting announcements, the central bank can make communication ambiguous, which does not commit itself to a fixed horizon.

Second, the central bank's strategy to announce AIT but implements IT ex post could be welfare-enhancing, but is it sustainable? Or can agents learn the truth in the long run? We answer these questions through social learning, where agents meet and update their beliefs based on their performances. We allow two layers of heterogeneity. First, agents have different beliefs about the horizon L over which the central bank averages inflation. Second, agents with the same belief about L observe different private signals about economic fundamentals, and we use this as a device to capture uncertainty. By contrast, the central bank has perfect information.

We show ambiguous communication helps the central bank gain credibility and improve social welfare in the long run, which constitutes the second motive for ambiguous communication. This result does not depend on whether the central bank is credible initially, how agents weigh forecast errors between inflation and the output gap, or if agents observe the underlying economic fundamentals perfectly. On the other hand, the optimal L is unclear. Two competing forces are at play: the largest L delivers the best Phillips curve but the least followers. However, the central bank need not necessarily face this trade off by communicating ambiguously, which delivers a larger following and better welfare than clear communication with any L.

The rest of the paper proceeds as follows: Section 1.1 draws a connection to the existing

4

literature. Section 2 sets up the FIRE benchmark. Section 3 discusses time inconsistency and the flexibility motive for ambiguous communication, and the discussion assumes the central bank can convince the private sector with its communication. Via social learning, Section 4 assesses this assumption and focuses on the credibility motive of ambiguous communication. Finally, Section 5 concludes.

1.1   Literature

Only a few papers in the literature focus on the Federal Reserve's new policy framework of AIT. Coibion et al. (2020) empirically assess how AIT changes expectations using survey data, whereas we do so using a theoretical framework. Also using a model, Amano et al. (2020) examine the optimal degree of history dependence under AIT and find it is relatively short. Their paper focuses on a simple instrument rule, whereas we introduce AIT into the central bank's optimization problem. In that sense, Nessén and Vestin (2005) is the closest to our paper. However, their paper studies the case with FIRE, whereas our paper focuses on time inconsistency and ambiguous communication, which are more relevant to current policy discussions.

Several other papers discuss AIT as one of many alternative strategies when many advanced economies face a low interest rate environment. Mertens and Williams (2019a,b) show AIT can mitigate the effects of the ZLB by raising inflation expectations when inflation is low. Svensson (2020) compares average-inflation targeting with annual inflation targeting, price-level targeting, temporary price-level targeting, and nominal-GDP targeting, and discusses its advantages. Hebden et al. (2020) discuss the robustness of AIT to alternative assumptions about the slope of the Phillips curve and the uncertainty of economic slack. Andrade et al. (2021) show AIT can work as an alternative to a higher inflation target for the euro area, which faces a low r^*. Papell and Prodan (2020) show how the new policy framework is consistent with an alternative policy rule forward guidance.

Our paper is also related to the literature on agent-based modeling. We focus on its

applications in economics, which refers to it as social learning or social dynamics. Burnside, Eichenbaum, and Rebelo (2016) use social learning to study boom and bust in housing markets. Our paper is related to Hachem and Wu (2017) in the sense that both papers study inflation. It is also related to Arifovic, Bullard, and Kostyshyna (2013) in the sense that both papers use a New Keynesian model. The difference is that we focus on the Fed's new framework of AIT.

Most of the discussion in the literature on the ZLB centers around conventional monetary policy (e.g., Eggertsson and Woodford (2003)), quantitative easing (e.g., Gertler and Karadi (2011), Gertler and Karadi (2013), Carlstrom, Fuerst, and Paustian (2017), and Sims and Wu (2020a), and Sims, Wu, and Zhang (2021)), forward guidance (e.g., Levin et al. (2010)), or the negative interest rate policy (e.g., Wu and Xia (2020) and Ulate (2021)). A few papers study a mix of policy tools at the ZLB (e.g., Swanson (2018), Wu and Zhang (2019), Sims and Wu (2020b), Levin and Sinha (2020), and Swanson (2021)). Different from these papers, we focus on AIT, which is the newest addition to central banks' toolkit.

Our paper is also broadly related to the literature on alternative expectations formation that deviates from the standard assumption of FIRE; for an extensive literature review, see Coibion, Gorodnichenko, and Kamdar (2018). One approach assumes agents are rational but have information frictions in the form of information rigidity (Mankiw and Reis (2002)), partial information (Angeletos and La'o (2020) and Jia (2019)), or rational inattention (Sims (2003)). Coibion and Gorodnichenko (2012) test these forms of information friction with a variety of survey forecasts. Another approach assumes agents learn in an adaptive way (Evans and Honkapohja (1999) and Orphanides and Williams (2005)). The learning part of our paper broadly fits into this part of the literature, but we use social learning instead.

2 Full-Information Rational Expectations Benchmark

In this section, we lay out a benchmark model that characterizes the new policy framework of average inflation targeting under the FIRE assumption. After describing the model and its equilibrium in Section 2.1, we compare it with a standard model of inflation targeting in Section 2.2. We discuss AIT's implications for the Phillips curve in Section 2.3 and for the ZLB in 2.4. We deviate from the FIRE assumption and discuss the key issues of the paper, namely, time inconsistency and ambiguous communication, in Sections 3 and 4.

2.1 A Model of Average Inflation Targeting

We model average inflation targeting with the following objective function of the central bank: a weighted sum of squared average inflation over L periods and the squared output gap:

$$\mathbb{L}_t^{cb}(L) = \frac{1}{2}\left(\left(\frac{\pi_t + \pi_{t-1} + ...\pi_{t-L+1}}{L}\right)^2 + \lambda^{cb}(L)\hat{y}_t^2\right) + \beta\, \mathbb{E}_t\, \mathbb{L}_{t+1}^{cb}(L), \qquad (2.1)$$

where $\hat{y}_t$ is the output gap, defined as the log deviation of real output from its natural level, and π_t is the rate of inflation. $\lambda^{cb}(L) > 0$ is the relative weight between the output gap and average inflation, and it could potentially depend on the horizon of AIT L.

In general, the central bank's objective function is different from the social welfare, whose period loss is defined as

$$\mathcal{L}_t = \frac{1}{2}\left(\pi_t^2 + \lambda\hat{y}_t^2\right), \qquad (2.2)$$

where λ is derived from the second-order approximation of the household's utility function (Rotemberg and Woodford, 1999). They only coincide when $L = 1$ and $\lambda^{cb}(1) = \lambda$.

The central bank minimizes its loss in (2.1) by choosing inflation and output gap subject to a standard forward-looking Phillips curves (see, e.g., Woodford (2011), Clarida, Galí, and

Gertler (1999), and Galí (2015)):

$$\pi_t = \beta \, \mathbb{E}_t \, \pi_{t+1} + \kappa \hat{y}_t + u_t, \tag{2.3}$$

where β is the discount factor, and κ depends on nominal rigidity. We assume the cost-push shock, u_t, is the only shock in the economy, which introduces a trade-off between inflation and the output gap.

Our central bank is discretionary because it solves the optimization problem period by period. However, unlike a standard model, the equilibrium is path dependent:

$$\pi_t = a_{\pi,1}^{(L)} \pi_{t-1} + \ldots + a_{\pi,L-1}^{(L)} \pi_{t-L+1} + b_\pi^{(L)} u_t \tag{2.4}$$

$$\hat{y}_t = a_{y,1}^{(L)} \pi_{t-1} + \ldots + a_{y,L-1}^{(L)} \pi_{t-L+1} + b_y^{(L)} u_t. \tag{2.5}$$

For analytic tractability and to gain some intuition, we start with two-period AIT, and we extend to multi-period in Sections 3.2 and 4. For two-period AIT ($L = 2$), the coefficients are given by the following set of fixed-point equations:

$$a_\pi \equiv a_{\pi,1}^{(2)} = -\vartheta \frac{\kappa}{4(1 - \beta a_\pi)} \tag{2.6}$$

$$b_\pi \equiv b_\pi^{(2)} = \vartheta \frac{\lambda^{cb}}{\kappa} \tag{2.7}$$

$$a_y \equiv a_{y,1}^{(2)} = \frac{1}{\kappa} (1 - \beta a_\pi) \, a_\pi \tag{2.8}$$

$$b_y \equiv b_y^{(2)} = \frac{1}{\kappa} (1 - \beta a_\pi) \, b_\pi - \frac{1}{\kappa}, \tag{2.9}$$

where $\lambda^{cb} \equiv \lambda^{cb}(2)$. See Appendix A.1 for derivations and expression for ϑ.

We characterize the equilibrium with the following proposition:

Lemma 1 $a_\pi < 0$, $a_y < 0$.

Proof: See Appendix B.

Lemma 1 says higher inflation leads to lower inflation and a lower output gap in the next

Figure 1: Impulse responses to a cost-push shock

(a) AIT   (b) IT

Notes: X-axis: horizon. Units on the y-axis: percentage points. Blue solid lines: inflation; red dashed lines: output gap. Left panel: IT; right panel: AIT. $\lambda^{cb}(2) = \lambda$, $u_t = 1$.

period. This lemma, especially a negative $a_\pi = \frac{\partial \pi_t}{\partial \pi_{t-1}}$, leads to an important result – the expectations channel of AIT, which we discuss in Section 2.3. We corroborate the result of a negative a_π with the left panel of Figure 1, which plots impulse responses of inflation and the output gap to a 1% cost-push shock.[2] The cost-push shock introduces a trade-off for the central bank by increasing inflation and decreasing the output gap. When the central bank implements AIT, inflation oscillates around zero. In particular, the central bank tightens monetary policy in the second period, which leads to a negative inflation rate and a negative output gap.

2.2 Comparison to Inflation Targeting

We define inflation targeting (IT) as the textbook version of optimal discretionary policy (see, e.g., Woodford (2011) and Galí (2015)). Its objective function is the social welfare, which is the present discounted value of the period loss function in (2.2).

---

[2]Parameters are calibrated in line with the New Keynesian literature in an annual frequency. For details, see Appendix C.

The equilibrium for IT is given by

$$\pi_t = b_\pi^{(1)} u_t = \left(\kappa + \frac{\lambda}{\kappa} \right)^{-1} \frac{\lambda}{\kappa} u_t \tag{2.10}$$

$$\hat{y}_t = b_y^{(1)} u_t = \left[\frac{1}{\kappa} \left(\kappa + \frac{\lambda}{\kappa} \right)^{-1} \frac{\lambda}{\kappa} - \frac{1}{\kappa} \right] u_t. \tag{2.11}$$

See Appendix A.2 for derivations.

Unlike the AIT equilibrium in (2.4) - (2.5), IT equilibrium in (2.10) - (2.11) does not depend on the lagged inflation. On the right side of Figure 1, we plot the impulse responses under IT. After the first period, the economy is back to the steady state, which is in contrast to the oscillation in the left panel.

2.3 Phillips Curve

This section investigates how AIT affects the Phillips curve, which captures the central bank's available trade-off between inflation and the output gap. We begin with an important property of the model, which is a direct result of Lemma 1:

Proposition 1 *The Expectations Channel implies lower (higher) positive (negative) inflation under AIT than that under IT for a given output gap.*

Proof: See Appendix B.

For AIT, the expectations channel works as follows: the private sector forms its expectation about future inflation conditional on current inflation. When inflation is positive, the expected inflation next period is negative per Lemma 1, which in turn lowers inflation today via the expectation term in the forward-looking Phillips curve (2.3). By contrast, IT does not have such an expectations channel.

As a result of the expectations channel, although IT and AIT share the same structural form of the Phillips curve as in (2.3), the reduced-form Phillips curves are different. For

AIT, it is

$$\pi_t = \frac{\kappa}{1 - \beta a_\pi} \hat{y}_t + \frac{1}{1 - \beta a_\pi} u_t, \tag{2.12}$$

and the reduced-form Phillips curve for IT is

$$\pi_t = \kappa \hat{y}_t + u_t. \tag{2.13}$$

For (2.12), see derivations in Appendix A.1. The derivation of (2.13) is trivial because $\mathbb{E}_t \pi_{t+1} = 0$.

Comparing the Phillips curves under AIT in (2.12) and under IT in (2.13) leads to the following proposition:

Proposition 2 *AIT implies a smaller slope of the Phillips curve than IT.*

Proof: See Appendix B.

Why does AIT flatten the Phillips curve? We use the right panel of Figure 2 for illustration (a similar argument can be made using the left panel). We start with where the Phillips curves of IT and AIT intersect, which corresponds to the origin in this case. Points A and B are both associated with more expansionary policy and have a 0.1% higher output gap than the origin. The difference between them is that point A is on the IT Phillips curve, whereas point B is on the AIT Phillips curve. The direct effect on the forward-looking Phillips curve (2.3) is that both A and B have inflation that is κ percent higher than the origion. But they differ in terms of the indirect effect through the expectations channel described in Proposition 1. IT does not have such an indirect effect. By contrast, with an AIT policy, higher inflation today lowers expected inflation in the next period, which feeds back to lower inflation today. This expectations channel makes point B lower than point A. Therefore, AIT is associated with a smaller slope of the Phillips curve than IT.

With a flatter Phillips curve, we next argue AIT presents a better trade-off between inflation and the output gap than IT:

11

Figure 2: Phillips curves and equilibria

(a) Positive cost-push shock   (b) Positive lagged inflation
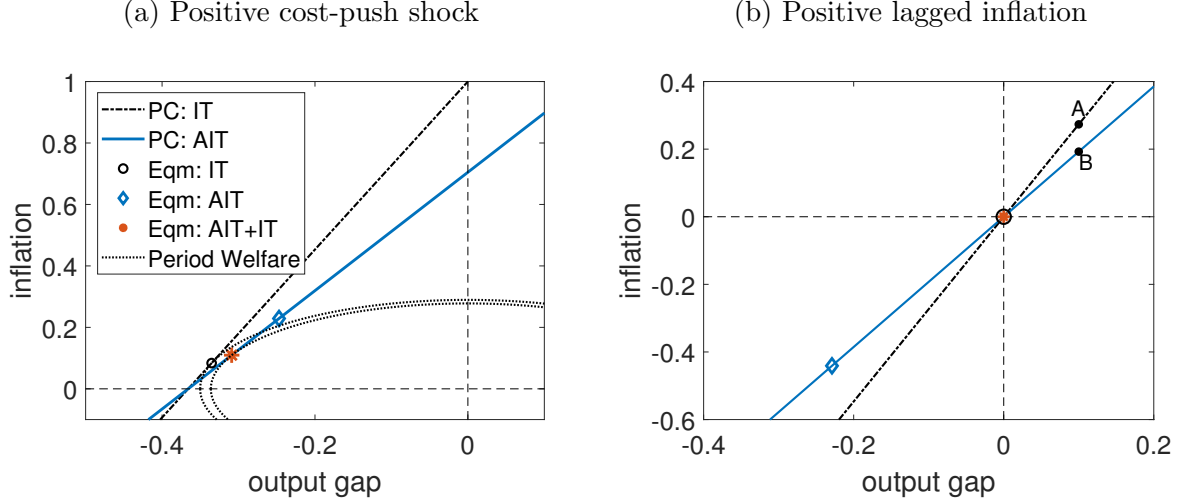
Notes: Left: $u_t = 1, \pi_{t-1} = 0$; right: $\pi_{t-1} = 1, u_t = 0$. Units: percentage points. Lines are Phillips curves. Black circle: IT equilibrium; blue diamond: AIT equilibrium with commitment; red star: the equilibrium where the private sector forms expectations based on the AIT announcement but the central bank deviates to implement IT ex post. The ellipses are period social welfare. $\lambda^{cb}(2) = \lambda$.

Proposition 3 *AIT yields a better available trade-off for the central bank between inflation and the output gap.*

Proof: See Appendix B.

This is a key result of the paper: AIT tilts the Phillips curve in a favorable way – closer to the origin in the relevant space. We show the intuition of this result in two steps, and for each step, we focus on one state variable. First, we start with a cost-push shock

Lemma 2 *The Phillips curve under AIT has*

- *the same x-intercept as*

- *and a smaller absolute value of the y-intercept than*

that under IT after a cost-push shock.

Proof: See Appendix B.

The intercept (which we refer to as the *y*-intercept in the proposition) of the Phillips curve captures the equilibrium when the output gap is stabilized after a cost-push shock.

12

The two upward-sloping lines in the left panel of Figure 2 are Phillips curves after a positive cost-push shock, which leads to a positive inflation rate and hence a positive intercept. When the central bank implements AIT, the expectations channel that we discussed in Proposition 1 works as follows: positive inflation in the current period leads to negative expected inflation in the next period, which in turn lowers current inflation. Therefore, the intercept of the Phillips curve is smaller under AIT than under IT. The x-intercept of the Phillips curve represents the equilibrium when inflation is completely stabilized after a cost-push shock. Zero inflation this period implies zero expected inflation and a zero expected output gap next period for both policies, eliminating any feedback from the expectations to the current period. Therefore, the two policies share the same x-intercept.

Next, we turn to the case with non-zero lagged inflation. IT is not path dependent, and therefore, lagged inflation does not introduce a trade-off between inflation and the output gap. Although lagged inflation is a state variable for AIT, it does not introduce an additional trade-off between inflation and the output gap for the two-period case (see the lines in the right panel of Figure 2).[3] We summarize this result in the following lemma:

Lemma 3 *The Phillips curves under two-period AIT and IT both cross the origin after non-zero lagged inflation.*

Proof: See Appendix B.

What drives this result? The private sector forms its expectations next period conditional on current inflation only. Therefore, lagged inflation does not introduce an indirect effect via the expectations channel.

Although AIT allows a better available trade-off between inflation and the output gap, it does not guarantee an improvement in welfare, because the central bank's objective function is different from the social welfare. In Section 3, we discuss a welfare-enhancing strategy in which the central bank implements a policy that is different from its announced AIT.

---

[3]We will show later that this is not the case for $L > 2$.

2.4 Implications of the Zero Lower Bound

The ZLB has plagued many advanced economies, including the US, Japan, the euro area, and many other countries, for over a decade. It is also a primary reason for the introduction of AIT; for example, see Powell (2020) and Diwan et al. (2020). This section takes a detour from the main theme of the paper, namely, time inconsistency and ambiguous communication, which we return to in Sections 3 and 4, and discusses the implications of the ZLB.

We model the ZLB with an interest rate peg and a negative shock to the natural rate of interest, for which we need to introduce the IS curve:

$$\hat{y}_t = \mathbb{E}_t \, \hat{y}_{t+1} - \frac{1}{\gamma} \left(\underline{i} - r_t^n - \mathbb{E}_t \, \pi_{t+1} \right), \tag{2.14}$$

where the nominal interest rate is $\underline{i} = 0$ and r_t^n is the natural rate of interest. γ is the inverse of the intertemporal elasticity of substitution. We model a one-period ZLB, which corresponds to one year in our calibration.

When the ZLB does not bind, the optimal interest rate policy tracks the natural rate, which offsets any effect the natural rate shock might have on the equilibrium.However, the central bank cannot achieve the dual mandate when the ZLB binds, because it cannot lower the nominal interest rate further to compensate a negative shock to the natural rate.

Under IT, $\mathbb{E}_t \, \pi_{t+1} = \mathbb{E}_t \, \hat{y}_{t+1} = 0$. Therefore, the IS curve becomes

$$\hat{y}_t = \frac{1}{\gamma} r_t^n. \tag{2.15}$$

For two-period AIT,

$$\mathbb{E}_t \, \pi_{t+1} = a_\pi \pi_t \tag{2.16}$$

$$\mathbb{E}_t \, \hat{y}_{t+1} = a_y \pi_t. \tag{2.17}$$

Figure 3: Zero lower bound

Notes: $r_t^n = -1$. Units: percentage points. The upward-sloping lines are the Phillips curves; the downward-sloping and vertical lines are the IS curves; the circle and diamond are the equilibria. Black represents IT and blue represents AIT. $\lambda^{cb}(2) = \lambda$.

Therefore, the IS curve becomes

$$\hat{y}_t = \left(a_y + \frac{1}{\gamma} a_\pi \right) \pi_t + \frac{1}{\gamma} r_t^n. \tag{2.18}$$

Figure 3 plots the implications of the ZLB. The upward-sloping lines are the Phillips curves, the vertical dotted black and downward-sloping dashed blue lines are the IS curves, and the circle and diamond are the equilibria. Black represents IT and blue is AIT. For illustration, we abstract from the cost-push shock. Our results go through as long as the cost-push shock is not too positive and the equilibria stay in the same quadrant.

Similar to our discussion in Proposition 2, the AIT Phillips curve is flatter than the IT curve. The IS curve for IT is vertical per (2.15), whereas that for AIT is downward sloping per (2.18) and Lemma 1. Moreover, the two IS curves cross the horizontal axis at the same point. Whether the central bank implements IT or AIT, the equilibrium entails negative inflation and negative output gap, which together are the hallmark of the ZLB.

However, AIT achieves a better equilibrium than IT with less negative inflation and

less negative output gap. This again works through the expectations channel discussed in Proposition 1. The current inflation is below the target. With AIT, the private sector expects inflation next period to be above the target, which in turn raises both demand and inflation this period. Therefore, AIT could be a useful tool for an economy that is stuck in a liquidity trap.

3   Time Inconsistency and Flexibility Motive for Ambiguity

This section deviates from the FIRE framework in Section 2, demonstrates the time inconsistency of AIT, and argues that a strategy whereby the central bank announces AIT but implements IT ex post dominates both IT and AIT with commitment. To optimally implement this time-inconsistent strategy, the central bank wants to announce AIT with different horizons depending on the state of the economy. Ambiguous communication allows such flexibility.

In this section, we assume the central bank can successfully convince the private sector of its intention to implement AIT, although it deviates from its announcement ex post. In Section 4, we assess this assumption and investigate whether the central bank can fool the public consistently.

3.1   Time Inconsistency

Once the central bank convinces the private sector of its intended policy rule, the Phillips curve is fixed. The central bank then can improve welfare ex post by picking the point on the Philips curve that minimizes the period loss of social welfare in (2.2), and we formalize this idea in the following corollary:

Corollary 1 *The strategy whereby the central bank announces AIT but implements IT ex*

post dominates both IT and AIT.

The equilibrium for this strategy is given by

$$\pi_t = \left(\frac{\kappa}{1-\beta a_\pi} + \lambda\frac{1-\beta a_\pi}{\kappa}\right)^{-1}\frac{\lambda}{\kappa}u_t \tag{3.1}$$

$$\hat{y}_t = \left[\left(\frac{\kappa}{1-\beta a_\pi} + \lambda\frac{1-\beta a_\pi}{\kappa}\right)^{-1}\frac{\lambda}{\kappa}\frac{1-\beta a_\pi}{\kappa} - \frac{1}{\kappa}\right]u_t. \tag{3.2}$$

See Appendix A.3 for derivations.

We explain the intuition with Figure 2, where an equilibrium is the tangent point between an objective function, which could be the AIT objective or social welfare, and the constraint imposed by a Phillips curve. We discuss one state variable at a time.

First, after non-zero lagged inflation, equations (3.1) - (3.2) suggest the central bank deviates from AIT to the equilibrium of dual stabilization $\pi_t = x_t = 0$ (the red star in the right panel of Figure 2). This new equilibrium has the same zero loss as the IT equilibrium (black circle), which improves welfare from the AIT equilibrium (blue diamond) that involves a positive loss.

Next, in response to a cost-push shock, in general, the central bank picks a different equilibrium than the equilibrium under AIT with commitment described in (2.4) - (2.5), or the red star could be different from the blue diamond in the left panel of Figure 2. The new equilibrium (the red star) is on the tangent point between the more favorable Phillips curve (the blue solid line) and the period loss of social welfare (the ellipses). Therefore, it dominates both IT, which is on the worse Phillips curve (the black dashed line), and AIT, where the central bank minimizes the AIT loss in (2.1) but not the period welfare.

3.2   Flexibility Motive of Ambiguous Communication

The time-inconsistent nature of AIT, which is discussed in Section 3.1, leads to a flexibility motive of ambiguous communication; that is, the central bank would like to announce AIT of different horizons depending on the state of the economy. As discussed, the central bank

always maximizes the period welfare ex post regardless of its announcement. Therefore, comparing equilibria between announcing AIT with different horizons L is equivalent to comparing their Phillips curves, which are characterized by[4]

$$\pi_t = \frac{\kappa}{1 - \beta a_{\pi,1}^{(L)}} \hat{y}_t + \sum_{l=1}^{L-2} \frac{\beta a_{\pi,l+1}^{(L)}}{1 - \beta a_{\pi,1}^{(L)}} \pi_{t-l} + \frac{1}{1 - \beta a_{\pi,1}^{(L)}} u_t. \tag{3.3}$$

3.2.1 Optimal Multi-Period AIT

To make our comparison across different L meaningful, we need to first discipline $\lambda^{cb}(L)$, which determines how much weight the central bank puts on the output gap. Our results that compare IT with two-period AIT thus far hold for any $\lambda^{cb}(2)$. But $\lambda^{cb}(L)$ is no longer inconsequential when we compare AIT of different horizons.

We choose the optimal $\lambda^{cb,*}(L)$ for each L such that it minimizes the unconditional welfare loss

$$\mathbb{E}_0 \mathcal{L}_t = \frac{1}{2} \left(var[\pi_t] + \lambda var[\hat{y}_t] \right), \tag{3.4}$$

where unconditional variances of inflation and the output gap are calculated using the AIT equilibrium conditions in (2.4) - (2.5). We use the optimal $\lambda^{cb,*}(L)$ hereafter.

The optimal $\lambda^{cb,*}(L)$ is decreasing in L; see the blue dots in Figure 4. We can show the intuition for why $\lambda^{cb,*}(2) < \lambda^{cb,*}(L = 1) = \lambda$ with the left panel of Figure 2. The AIT equilibrium (blue diamond) is associated with higher inflation and a smaller output gap than the equilibrium where the central bank announces AIT but implements IT (red star). That is because the AIT central bank in (2.1) puts more weight on the output gap relative to current inflation than the social welfare when we assign $\lambda^{cb}(2) = \lambda$. To compensate for this effect, the optimal AIT needs to assign a weight that is smaller than λ on the output gap.

For a general L, the weight that the central bank's period loss puts on current inflation is

---

[4]We solve the linear-quadratic optimization by the solution algorithm developed in Oudiz and Sachs (1985) and further discussed in Söderlind (1999) and Svensson (2007).

Figure 4: $\lambda^{cb}(L)$ for multi-period AIT

Notes: $\lambda^{cb}(L)$ (y-axis) as a function of L (x-axis). Blue dots: $\lambda^{cb}(L) = \lambda^{cb,*}(L)$; red diamonds: $\lambda^{cb}(L) = \lambda/L^2$.

$1/L^2$, which decreases in L; see (2.1). To compensate for this, $\lambda^{cb,*}(L)$ needs to decrease in L as well. In Figure 4, we compare $\lambda^{cb,*}(L)$ with $\lambda^{cb}(L) = \lambda/L^2$. They are slightly different because π_t also enters $\mathbb{E}_t \, \mathbb{L}^{cb}_{t+1}(L)$, which makes $\lambda^{cb,*}(L)$ and λ/L^2 not identical. However, like the intuition, all our results go through if we set $\lambda^{cb}(L)$ to λ/L^2 instead of $\lambda^{cb,*}(L)$.

3.2.2 Ambiguous Communication and Optimal AIT

Suppose a cost-push shock hits the economy at t. The central bank's optimal strategy is to announce the largest feasible L at time t and announce $L = 2$ subsequently. To show why, we begin with period t, for which the intercept of the Phillips curve (3.3), $\frac{1}{1-\beta a^{(L)}_{\pi,1}} u_t$, is a decreasing function of L for a positive shock; see Figure 5. In other words, the Phillips curve with the largest L ($L = 4$ years in this example) is associated with the strongest effect of the expectations channel and the best available trade-off between inflation and the output gap. Hence, the largest L forms the optimal horizon for the AIT announcement at time t.

At time $t + 1$, it is optimal for an AIT central bank to announce $L = 2$. In this case, the Phillips curve does not have an intercept; see (3.3). Consequently, we can achieve the

Figure 5: Phillips curve for multi-period AIT

equilibrium of zero inflation and zero output gap. By contrast, if the central bank announces an $L > 2$, a non-zero past inflation causes a non-zero intercept in the Phillips curve, which implies non-zero loss in the period social welfare. At time $t + h$, the central bank can announce any $L \in \{2, ..., h + 1\}$ and implement IT to achieve dual stability.

However, switching between different L does not seem to be attainable under reasonable assumptions. We argue in Section 4 that agents might believe the central bank when its announcement is different from the actual implementation. But agents will not believe the central bank when it changes its announcements all the time. The fact that the central bank would like the flexibility to change its announcements about the length of AIT over time but cannot argues for ambiguous communication.

4 Social Learning and Credibility Motive of Ambiguous Communication

Section 3 demonstrates AIT's time-inconsistent nature and discusses the flexibility motive for ambiguous announcement about the horizon of AIT. In Section 3, we assume the central bank has full credibility. This section assesses this assumption and focuses on the following questions: Can agents eventually learn the truth about the central bank's behavior? Is the central bank's welfare-enhancing strategy sustainable and beneficial over time in light of learning? Section 4.1 sets up the environment of social learning. Section 4.2 argues ambiguous communication improves the central bank's credibility, and Section 4.3 explains when and why AIT can forecast better than IT, which is the central bank's de facto strategy. Section 4.4 shows our results are robust regarding initial credibility and uncertainty about underlying economic conditions.

4.1 Social Learning

We model social learning similar to Arifovic, Bullard, and Kostyshyna (2013) and Hachem and Wu (2017). Beliefs are rule-based. L_{max} groups of beliefs exist, $L \in \{1, 2, ...L_{max}\}$, and beliefs can also be heterogeneous within each group. Agents update their beliefs via tournament selection and mutation.

The sequence of events is drawn in Figure 6. Before entering period t, each agent has a belief about L, and we label it L_t^i. When agents enter period t, a shock u_t is realized, and agent i observes it with a private signal u_t^i:

$$u_t^i = u_t + v_t^i, \qquad v_t^i \sim N(0, \sigma_v^2). \tag{4.1}$$

Within-group heterogeneity is present when $\sigma_v \neq 0$.[5] An agent believes all other agents in

---

[5]We will discuss the role of uncertainty, captured by σ_v, in Section 4.4.2.

Figure 6: The sequence of events

the economy have the same information set that he/she does, and forms his/her expectations about variables at time t using equations (2.4) - (2.5):

$$\mathbb{E}_t^i \pi_t = a_{\pi,1}^{(L_t^i)} \pi_{t-1} + ... + a_{\pi,L_t^i-1}^{(L_t^i)} \pi_{t-L_t^i+1} + b_\pi^{(L_t^i)} u_t^i \tag{4.2}$$

$$\mathbb{E}_t^i \hat{y}_t = a_{y,1}^{(L_t^i)} \pi_{t-1} + ... + a_{y,L_t^i-1}^{(L_t^i)} \pi_{t-L_t^i+1} + b_y^{(L_t^i)} u_t^i. \tag{4.3}$$

Next, agent i forms expectations about inflation in the next period:

$$\mathbb{E}_t^i \pi_{t+1} = a_{\pi,1}^{(L_t^i)} \mathbb{E}_t^i \pi_t + ... + a_{\pi,L_t^i-1}^{(L_t^i)} \pi_{t-L_t^i+2}. \tag{4.4}$$

The average expectation is

$$\bar{\mathbb{E}}_t \pi_{t+1} = \frac{1}{N} \sum_i \mathbb{E}_t^i \pi_{t+1}, \tag{4.5}$$

where N is the total number of agents.

Therefore, the Phillips curve in equation (2.3) becomes

$$\pi_t = \beta \bar{\mathbb{E}}_t \pi_{t+1} + \kappa \hat{y}_t + u_t. \tag{4.6}$$

The central bank now picks π_t and $\hat{y}_t$ to minimize the period loss $\mathcal{L}_t$ defined in (2.2) subject

to the Phillips curve in equation (4.6). The realized laws of motion for π_t and $\hat{y}_t$ are different from equations (2.4) - (2.5).

Tournament selection works as follows: at each time period t, agents are randomly selected to meet in pairs. When two agents meet, they update their beliefs by comparing forecast errors, which are defined as

$$\varepsilon_t^i = |\, \mathbb{E}_t^i \, \pi_t - \pi_t| + w \,|\, \mathbb{E}_t^i \, \hat{y}_t - \hat{y}_t|, \tag{4.7}$$

where w is the weight agents put on the forecast error on the output gap.[6] When two agents are from the same group, that is, $L_t^i = L_t^j$, they stay in this group. When two agents come from two different groups, suppose agent i has a smaller forecast error $\varepsilon_t^i < \varepsilon_t^j$. Agent i stays in his/her current group, whereas agent j switches. At the end of period t, agents can mutate, that is, randomly switch to another rule. Mutation is a standard device in the literature to avoid the algorithm to get stuck in a corner solution prematurely. This step further updates the belief to L_{t+1}^i, which will be used in time $t+1$. Parameter values are in Appendix C.

4.2 Credibility Motive of Ambiguous Communication

This section argues ambiguous communication helps improve the central bank's credibility. We define the fraction of agents in each group as $p^{(L)}$.

The left panel of Figure 7 plots the fraction of AIT believers when the central bank announces AIT with a clear horizon $L \in (2, 3, 4)$ for two, three, or four years. In each case, there are only two groups of beliefs: IT and L-period AIT, or $p^{(1)} + p^{(L)} = 1$. The red dashed, the yellow dash-dotted, and purple dotted lines represent a central bank that announces two-period, three-period, or four-period AIT, respectively. All our results are averaged over 1,000 simulations. At each point in time, the fraction of AIT believers is the

---

[6]We use equal weight $w = 1$ in the benchmark case and discuss the implications of different weighting schemes in Section 4.3.

Figure 7: Fraction of AIT believers

(a) Clear communication
(b) Ambiguous communication
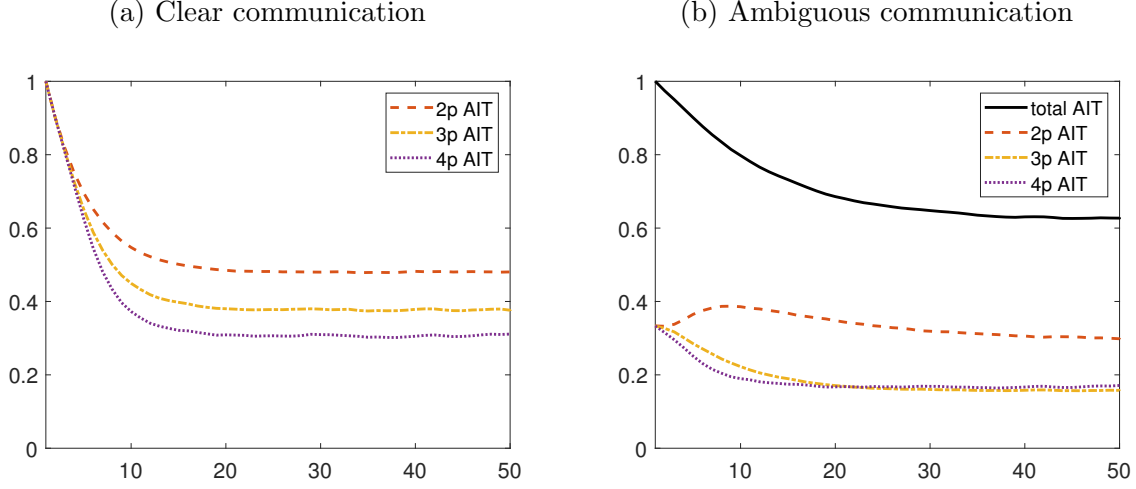
Notes: X-axis: horizon. Left panel: L is clearly communicated; right panel: AIT is announced without explicitly specifying L. Red dashed lines: fraction of two-period AIT believers, yellow dash-dotted lines: fraction of three-period AIT believers; purple dotted lines: fraction of four-period AIT believers. Black solid line in the right panel: fraction of total AIT believers. All the results are averaged over 1,000 simulations. $\lambda^{cb}(L) = \lambda^{cb,*}(L)$.

highest when the central bank announces two-period AIT, and lowest when the central bank announces four-period AIT. In this section, we endow the central bank with full credibility, and we show the robustness of our results in Section 4.4.1 when the central bank starts with no or partial credibility initially.

The right panel of Figure 7 plots the evolution of beliefs under ambiguous communication, where the central bank only announces AIT as its tool to manage inflation but does not specify its length. At the beginning, an equal fraction of agents believe in AIT with different lengths: $p^{(2)} = p^{(3)} = p^{(4)}$. Over time, these fractions evolve through social learning with $p^{(1)} + p^{(2)} + p^{(3)} + p^{(4)} = 1$. The red dashed, yellow dash-dotted, and purple dotted lines capture the fraction of agents in each AIT group, and the black solid line is the sum of the three, which shows the fraction of total AIT believers.

Comparing the black solid line in the right panel with each of the lines in the left panel, we find ambiguous communication allows a larger fraction of agents to believe in the AIT strategy that the central bank announces, even though it implements IT ex post. Therefore,

Figure 8: Welfare loss
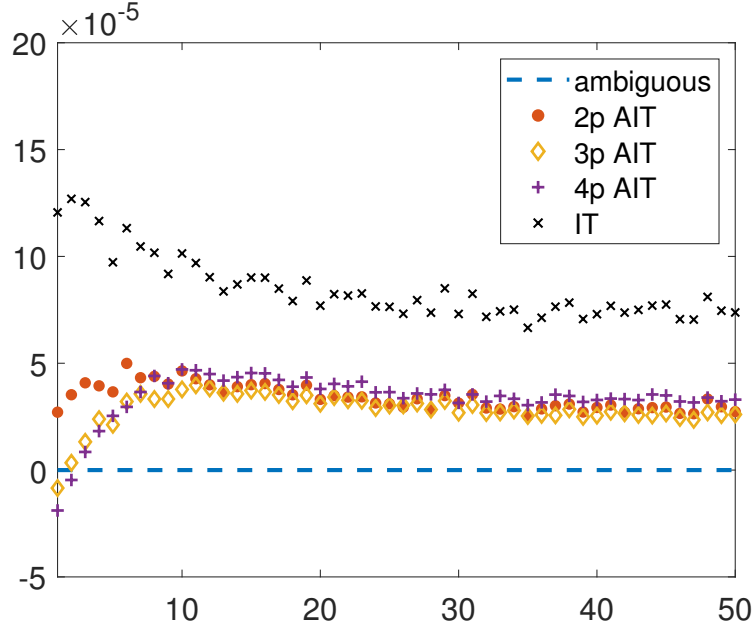
Notes: X-axis: horizon. y-axis: welfare loss. Blue dashed line: ambiguous communication, which is normalized to zero. Red dots: the central bank announces two-period AIT; yellow diamonds: the central bank announces three-period AIT; purple plus signs: the central bank announces four-period AIT; black crosses: IT. All the results are averaged over 1,000 simulations. $\lambda^{cb}(L) = \lambda^{cb,*}(L)$

ambiguous communication helps the central bank gain more credibility in the long run.

Comparing the three color lines in the right panel, similar to the left panel, more agents believe in two-period AIT than in three- or four-period AIT. Notice that in the left panel, the three colors represent three different types of announcements. For each announcement, only the corresponding color is relevant, whereas, in the right panel, there is only one type of announcement - ambiguous announcement - but agents can have different beliefs about L in this scenario.

We have established that ambiguous communication can help the central bank gain a larger following of its communicated AIT than clear communication. But does it improve welfare in the long run? In Figure 8, we plot the period loss of social welfare for five different announcements. We take the ambiguous communication in the blue dashed line as the benchmark and plot loss in other cases relative to it. Therefore, the blue dashed line

coincides with the horizontal axis by construction.

Initially, four-period AIT, which is represented by the purple "+"s, has the smallest loss. This result is consistent with what we find in Section 3.2: it is optimal for the central bank to announce the largest L on impact of the shock. But the advantage of announcing a large L disappears quickly, because four-period AIT has a smaller following over time; see Figure 7. At some point, we see a smaller welfare loss if the central bank announces $L = 3$ instead of $L = 4$.[7] This finding is consistent with the flexibility motive we discussed in Section 3.2. Two competing forces are at play: a larger L is associated with a more favorable Phillips curve but fewer followers. However, the central bank does not have to face this trade off and can take advantage of ambiguous communication, which is associated with the smallest loss for all but the initial few periods; see the blue dashed line.

Comparing IT in the black "x"s with the rest of the colors, announcing IT and sticking to it, which is the operating framework before introducing AIT, is universally the worst in terms of welfare. This result is consistent with our argument in Corollary 1.

4.3 Why Can AIT Forecast Better When the Truth Is IT

In Section 4.2, we showed AIT believers' forecast could potentially outperform IT believers'. The question is why, especially when the central bank always implements IT ex post.

We use Figure 9 to illustrate the intuition.[8] The red star marks the realized equilibrium after a positive cost-push shock in which the agents believe the central bank's announcement about two-period AIT but the central bank implements IT subsequently. The black circle and blue diamond mark the forecasted equilibria of the IT believers and AIT believers. Neither AIT believers nor IT believers forecast perfectly; or the red star does not overlap with either the blue diamond or black circle. The IT believers have a smaller forecast error on inflation, whereas AIT believers have a smaller forecast error on the output gap.

Although the central bank implements IT ex post, AIT believers could have smaller

---

[7]Notice we interpret time in a relative sense instead of an absolute sense and focus our discussion on the long run. Although we calibrate structural parameters to an annual frequency, the time to converge depends

26

Figure 9: Forecasted and realized equilibria

Notes: $u_t = 1, \pi_{t-1} = 0$. Units: percentage points. Black circle: the equilibrium forecasted by the IT believers; blue diamond: the equilibrium forecasted by the AIT believers; red star: the realized equilibrium where the private sector forms expectations based on the two-period AIT announcement but the central bank deviates to implement IT ex post. Lines are Phillips curves. The ellipses are period social welfare. $\lambda^{cb}(2) = \lambda^{cb,*}(2)$.

forecast error if agents put more weight on the output gap. We further elaborate this point in Figure 10, where we extend the results in Figure 7 with two extreme cases: in the top panels, agents put all the weight on inflation, or $w = 0$; in the bottom panels, agents put all the weight on the output gap, or $w \to \infty$. Similar to Figure 7, the left panels are associated with a central bank that announces the horizon of AIT explicitly, whereas in the right panel, the central bank announces AIT without specifying the horizon.
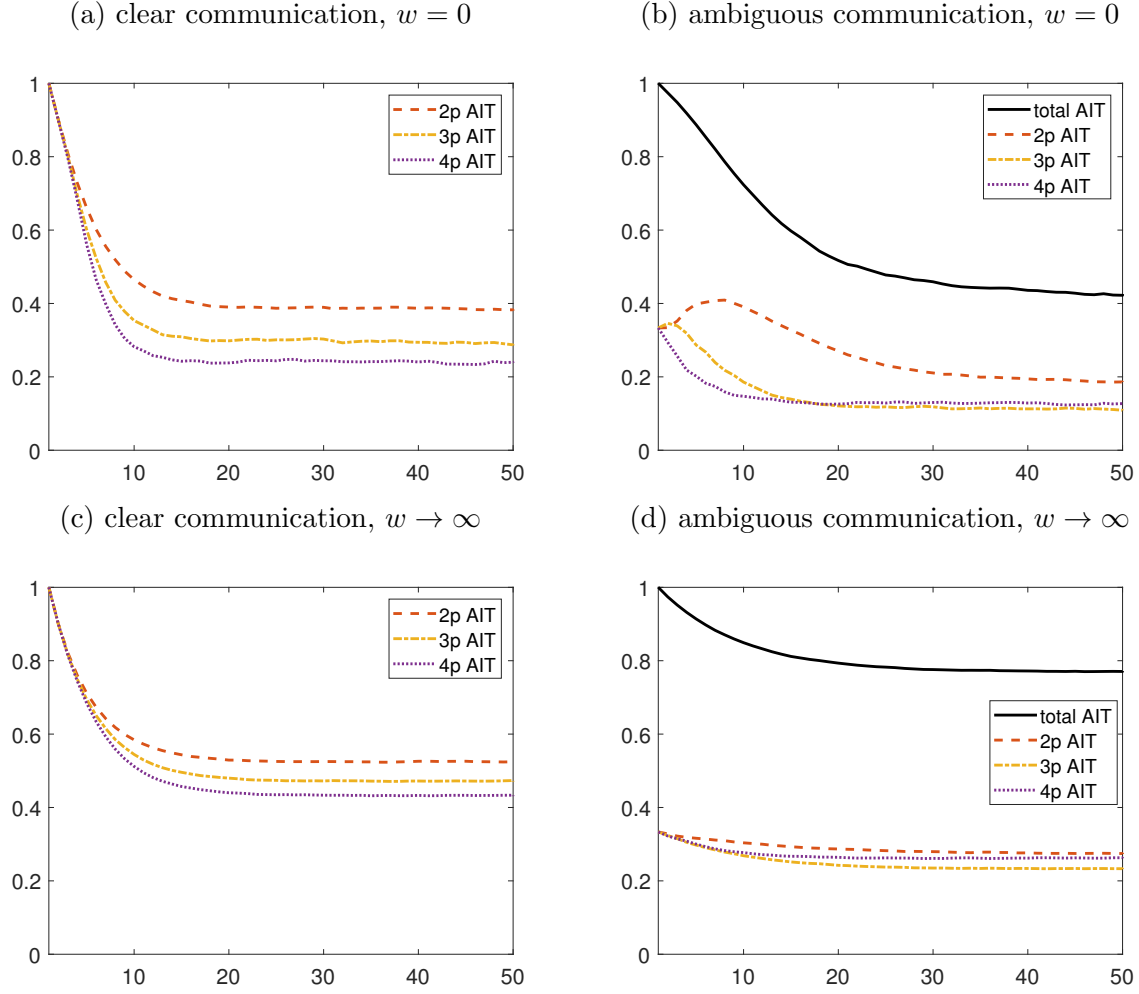
Comparing the top panels with the bottom panels, we find AIT has a larger following when agents put more weight on the output gap; that is, larger w. This result holds regardless of whether the central bank communicates the horizon of AIT clearly or ambiguously, which confirms the intuition provided by Figure 9.

Figure 10 also serves as a robustness check to the results in Figure 7. We compare the

---

on the specifics of social learning, which are not the focus of the paper.

[8]Figure 9 and the left panel of Figure 2 only differ by $\lambda^{cb}(2)$ and scale.

Figure 10: Different weights in the forecast error

(a) clear communication, $w = 0$

(b) ambiguous communication, $w = 0$

(c) clear communication, $w \to \infty$

(d) ambiguous communication, $w \to \infty$

Notes: X-axis: horizon. Left panels: L is clearly communicated; right panels: AIT is announced without explicitly specifying L. Top panels: $w = 0$; bottom panels: $w \to \infty$. Red dashed lines: fraction of two-period believers, yellow dash-dotted lines: fraction of three-period believers; purple dotted lines: fraction of four-period AIT believers. Black solid line in the right panels: fraction of total AIT believers. All the results are averaged over 1,000 simulations. $\lambda^{cb}(L) = \lambda^{cb,*}(L)$

left panels with the right panels. In the long run, the central bank always arrives with more credibility when using ambiguous communication than when using clear communication; compare the black lines in the right panel with the colored lines in the left panel. This result further confirms the credibility motive of ambiguous communication and shows the qualitative results in Figure 7 do not depend on how agents weigh forecast errors between inflation and the output gap.

Figure 11: Fraction of AIT believers with different initial credibility

Notes: X-axis: horizon. Blue solid line: the central bank starts with full credibility; red dashed line: the central bank starts with partial credibility; yellow dotted line: the central bank starts with zero credibility. All the results are averaged over 1,000 simulations. $\lambda^{cb}(L) = \lambda^{cb,*}(L)$

4.4 Robustness

4.4.1 Initial Credibility

Our analysis so far in this section starts with a central bank that is endowed with full credibility, which is a stark assumption. For example, Coibion et al. (2020) use survey data to show the private sector does not understand the impact of AIT soon after its introduction. In this section, we assess what happens when the central bank is partially credible or not credible initially.

In Figure 11, we plot the evolution of the fraction of AIT believers with different levels of initial credibility. The blue solid line starts with 100% of agents believing AIT, which replicates the black line in the right panel of Figure 7. The red dashed line starts with partial credibility where 50% of agents believe AIT, and the yellow dotted line starts with no credibility. In the long run, all three lines converge to the same point where over 60% of agents believe AIT. In summary, the initial credibility does not matter for the long run.

Figure 12: Fraction of IT/AIT believers

(a) No uncertainty   (b) Low uncertainty   (c) High uncertainty

Notes: X-axis: horizon. Black dashed lines: fraction of IT believers $p^{(1)}$; blue solid lines: fraction of AIT believers $p^{(2)}$. Left: no uncertainty with $\sigma_v = 0$; middle: low uncertainty with $\sigma_v = 0.01$; right: high uncertainty with $\sigma_v = 0.1$. All the results are averaged over 1,000 simulations. $\lambda^{cb}(L) = \lambda^{cb,*}(L)$.

4.4.2 Uncertainty

In this section, we show uncertainty about underlying economic conditions does not change results qualitatively although it does quantitatively. Figure 12 plots three scenarios: no uncertainty is present in the left panel, or $\sigma_v = 0$. The middle panel is associated with low uncertainty, and the right panel with high uncertainty. We compare IT with two-period AIT, and the fractions of their believers are in black dashed lines for IT and blue solid lines for AIT. When uncertainty increases, more agents believe AIT in the long run, and convergence takes longer.

However, the following results are robust regardless of whether agents observe the underlying shocks with uncertainty. First, AIT has followers even with no uncertainty. In the left panel of Figure 12, a non-trivial fraction of agents believe the AIT announcement; that is, the fraction is higher than the mechanical mutation probability. Second, ambiguous communication allows the central bank to have more credibility than clear communication, regardless of the value of σ_v.[9]

---

[9]To save space, we do not show plots.

5   Conclusion

Our paper studies the implications of AIT. We focus on two key issues: time inconsistency and ambiguous communication. AIT can improve the available trade-off between inflation and real activity that the central bank faces, but a more favorable Phillips curve does not automatically improve social welfare. To improve social welfare, the central bank has the incentive to deviate from its communicated objective and implement IT ex post. The time-inconsistent nature calls for two motives for ambiguous communication. First, the optimal strategy for the central bank is to announce different horizons for AIT at different points in time, and ambiguous communication offers an implementable alternative to such self-conflicting announcements. Second, the time-inconsistent strategy is welfare improving, assuming the central bank can convince the private sector of its intention to implement AIT. We assess this assumption using social learning. We show ambiguous communication helps the central bank gain credibility and improve welfare in the long run despite AIT being time inconsistent. These results do not rely on whether the central bank has credibility initially, how agents weigh forecast errors between inflation and the output gap, or if agents observe economic shocks with uncertainty.

References

Amano, Robert, Stefano Gnocchi, Sylvain Leduc, and Joel Wagner. 2020. "Average Is Good Enough: Average-Inflation Targeting and the ELB." Federal Reserve Bank of San Francisco Working Paper 2020-21.

Andrade, Philippe, Jordi Galí, Hervé Le Bihan, and Julien Matheron. 2021. "Should the ECB Adjust its Strategy in the Face of a Lower r^*?" CREI working paper.

Angeletos, George-Marios and Jennifer La'o. 2020. "Optimal monetary policy with informational frictions." *Journal of Political Economy* 128 (3):1027–1064.

Arifovic, Jasmina, James Bullard, and Olena Kostyshyna. 2013. "Social Learning and Monetary Policy Rules." *The Economic Journal* 123 (567):38–76.

Brunnermeier, Markus. 2021. "Averaging and Anchoring: Is there an Inflation Dilemma?" *Jackson Hole Economic Policy Symposium "Macroeconomic Policy in an Uneven Economy", Jackson Hole, Wyoming.*

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo. 2016. "Understanding booms and busts in housing markets." *Journal of Political Economy* 124 (4):1088–1147.

Carlstrom, Charles T., Timothy S. Fuerst, and Matthias Paustian. 2017. "Targeting Long Rates in a Model with Segmented Markets." *American Economic Journal: Macroeconomics* 9 (1):205–42.

Clarida, Richard, Jordi Galí, and Mark Gertler. 1999. "The science of monetary policy: a new Keynesian perspective." *Journal of economic literature* 37 (4):1661–1707.

Coibion, Olivier and Yuriy Gorodnichenko. 2012. "What can survey forecasts tell us about information rigidities?" *Journal of Political Economy* 120 (1):116–159.

Coibion, Olivier, Yuriy Gorodnichenko, and Rupal Kamdar. 2018. "The formation of expectations, inflation, and the phillips curve." *Journal of Economic Literature* 56 (4):1447–91.

Coibion, Olivier, Yuriy Gorodnichenko, Edward S Knotek II, and Raphael Schoenle. 2020. "Average inflation targeting and household expectations." National Bureau of Economic Research Working Paper 27836.

Diwan, Renuka, Sylvain Leduc, Thomas M Mertens et al. 2020. "Average-Inflation Targeting and the Effective Lower Bound." *FRBSF Economic Letter* 2020 (22).

Eggertsson, Gauti B and Michael Woodford. 2003. "The Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity* :139–211.

Evans, George W and Seppo Honkapohja. 1999. "Learning dynamics." *Handbook of macroeconomics* 1:449–542.

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

Gertler, Mark and Peter Karadi. 2011. "A Model of Unconventional Monetary Policy." *Journal of Monetary Economics* :17–34.

———. 2013. "QE 1 vs. 2 vs. 3 . . .: A Framework for Analyzing Large-Scale Asset Purchases as a Monetary Policy Tool." *International Journal of Central Banking* 9(S1):5–53.

Hachem, Kinda and Jing Cynthia Wu. 2017. "Inflation announcements and social dynamics." *Journal of Money, Credit and Banking* 49 (8):1673–1713.

Hebden, James, Edward Herbst, Jenny Tang, Giorgio Topa, and Fabian Winkler. 2020. "How Robust Are Makeup Strategies to Key Alternative Assumptions?" .

Jia, Chengcheng. 2019. "The informational effect of monetary policy and the case for policy commitment." .

Kozicki, Sharon. 2019. "Monetary Policy Strategies for the Federal Reserve: Discussion of Practical Considerations." *A Fed Listens Event "Conference on Monetary Policy Strategy, Tools, and Communication Practices",* Chicago, IL.

Levin, Andrew, David López-Salido, Edward Nelson, and Tack Yun. 2010. "Limitations on the Effectiveness of Forward Guidance at the Zero Lower Bound." *International Journal of Central Banking* 6 (1):143–189.

Levin, Andrew T. and Arunima Sinha. 2020. "Pitfalls of Make-Up Strategies for Mitigating the Effective Lower Bound." *Cato Journal* .

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

Mertens, Thomas M and John C. Williams. 2019a. "Monetary policy frameworks and the effective lower bound on interest rates." In *AEA Papers and Proceedings*, vol. 109. 427–32.

Mertens, Thomas M. and John C. Williams. 2019b. "Tying down the anchor: monetary policy rules and the lower bound on interest rates." *FRB of New York Staff Report* (887).

Nessén, Marianne and David Vestin. 2005. "Average inflation targeting." *Journal of Money, credit and Banking* 37 (5):837–863.

Orphanides, Athanasios and John C. Williams. 2005. "Inflation scares and forecast-based monetary policy." *Review of Economic Dynamics* 8 (2):498–527.

Oudiz, Gilles and Jeffrey Sachs. 1985. "International policy coordination in dynamic macroeconomic models." In *International economic policy coordination.* Cambridge University Press, 274–330.

Papell, David H and Ruxandra Prodan. 2020. "Policy Rule Forward Guidance Following the COVID-19 Recession."

Powell, Jerome H. 2020. "New Economic Challenges and the Fed's Monetary Policy Review." *Jackson Hole Economic Policy Symposium "Navigating the Decade Ahead: Implications for Monetary Policy",* Jackson Hole, Wyoming.

Rotemberg, Julio J and Michael Woodford. 1999. "Interest rate rules in an estimated sticky price model." In *Monetary policy rules*. University of Chicago Press, 57–126.

Sims, Christopher A. 2003. "Implications of rational inattention." *Journal of monetary Economics* 50 (3):665–690.

Sims, Eric and Jing Cynthia Wu. 2020a. "Are QE and Conventional Monetary Policy Substitutable?" *International Journal of Central Banking* 16 (1):195–230.

———. 2020b. "Evaluating Central Banks' Tool Kit: Past, Present, and Future." *Journal of Monetary Economics* Forthcoming.

Sims, Eric, Jing Cynthia Wu, and Ji Zhang. 2021. "The Four Equation New Keynesian Model." *Review of Economics and Statistics* Forthcoming.

Söderlind, Paul. 1999. "Solution and estimation of RE macromodels with optimal policy." *European Economic Review* 43 (4-6):813–823.

Svensson, Lars E. O. 2007. "Optimization under commitment and discretion, the recursive saddlepoint method, and targeting rules and instrument rules: Lecture notes."

Svensson, Lars E.O. 2020. "Monetary Policy Strategies for the Federal Reserve." *International Journal of Central Banking* 16 (1):133–193.

Swanson, Eric T. 2018. "The Federal Reserve Is Not Very Constrained by the Lower Bound on Nominal Interest Rates." *Brookings Papers on Economic Activity* :555–572.

———. 2021. "Measuring the effects of federal reserve forward guidance and asset purchases on financial markets." *Journal of Monetary Economics* 118:32–53.

Ulate, Mauricio. 2021. "Going Negative at the Zero Lower Bound: The Effects of Negative Nominal Interest Rates." *American Economic Review* 111 (1):1–40.

Woodford, Michael. 2011. *Interest and prices: Foundations of a theory of monetary policy.* princeton university press.

Wu, Jing Cynthia and Fan Dora Xia. 2020. "Negative Interest Rate Policy and Yield Curve." *Journal of Applied Econometrics* 35 (6):653–672.

Wu, Jing Cynthia and Ji Zhang. 2019. "A shadow rate New Keynesian model." *Journal of Economic Dynamics and Control* 107:103728.

A Derivations of Equilibria

A.1 Equilibrium of Two-Period AIT Under FIRE

We solve the equilibrium with the method of undetermined coefficients. We guess the following linear functions for π_t and $\hat{y}_t$:

$$\pi_t = a_\pi \pi_{t-1} + b_\pi u_t \tag{A.1}$$

$$\hat{y}_t = a_y \pi_{t-1} + b_y u_t, \tag{A.2}$$

which yield

$$\mathbb{E}_t\,\pi_{t+j} = a_\pi^j \pi_t \tag{A.3}$$

$$\mathbb{E}_t\,\hat{y}_{t+j} = a_y^j \pi_t. \tag{A.4}$$

Substitute the expressions for $\mathbb{E}_t\,\pi_{t+1}$ and $\mathbb{E}_t\,\hat{y}_{t+1}$ from (A.3) - (A.4) into the Phillips curve in (2.3) and rearrange, and we get (2.12). Consequently, the slope of the Phillips curve is

$$\frac{\partial \pi_t}{\partial \hat{y}_t} = \frac{\kappa}{1 - \beta a_\pi}. \tag{A.5}$$

Rearranging equation (2.12), we get

$$\hat{y}_t = \frac{1}{\kappa}\left(1 - \beta a_\pi\right)\pi_t - \frac{1}{\kappa}u_t. \tag{A.6}$$

Rewrite the central bank's objective function in (2.1) by substituting out the expectation terms using (A.3) - (A.4):

$$\mathbb{L}_t^{cb}(2) = \frac{1}{2}\left(\left(\frac{\pi_t + \pi_{t-1}}{2}\right)^2 + \lambda^{cb}\hat{y}_t^2\right) + \frac{1}{2}\sum_j \mathbb{E}_t\,\beta^j\left(\left(\frac{\pi_{t+j} + \pi_{t+j-1}}{2}\right)^2 + \lambda^{cb}\hat{y}_{t+j}^2\right) \tag{A.7}$$

$$= \frac{1}{2}\left(\left(\frac{\pi_t + \pi_{t-1}}{2}\right)^2 + \lambda^{cb}\hat{y}_t^2\right) + \frac{1}{2}\pi_t^2\sum_j \beta^j\left(\left(\frac{a_\pi + 1}{2}\right)^2 a_\pi^{2(j-1)} + \lambda^{cb}a_y^{2j}\right).$$

Minimize the central bank's objective function with respect to $\hat{y}_t$, which yields the following first-order condition (using (A.5) - (A.6)):

$$0 = \frac{1}{4}\left(\pi_t + \pi_{t-1}\right)\frac{\kappa}{1 - \beta a_\pi} + \lambda^{cb}\left(\frac{1}{\kappa}\left(1 - \beta a_\pi\right)\pi_t - \frac{1}{\kappa}u_t\right) \tag{A.8}$$

$$+ \pi_t\frac{\kappa}{1 - \beta a_\pi}\sum_j \beta^j\left(\left(\frac{a_\pi + 1}{2}\right)^2 a_\pi^{2(j-1)} + \lambda^{cb}a_y^{2j}\right).$$

Rearrange

$$0 = \left[\frac{\kappa}{4(1-\beta a_\pi)} + \lambda^{cb}\frac{1}{\kappa}(1-\beta a_\pi) + \frac{\kappa}{1-\beta a_\pi}\sum_j \beta^j \left(\left(\frac{a_\pi+1}{2} \right)^2 a_\pi^{2(j-1)} + \lambda^{cb} a_y^{2j} \right) \right] \pi_t \quad \text{(A.9)}$$
$$+ \frac{\kappa}{4(1-\beta a_\pi)}\pi_{t-1} - \lambda^{cb}\frac{1}{\kappa}u_t.$$

Define

$$\vartheta \equiv \left[\frac{\kappa}{4(1-\beta a_\pi)} + \lambda^{cb}\frac{1}{\kappa}(1-\beta a_\pi) + \frac{\kappa}{1-\beta a_\pi}\sum_j \beta^j \left(\left(\frac{a_\pi+1}{2} \right)^2 a_\pi^{2(j-1)} + \lambda^{cb} a_y^{2j} \right) \right]^{-1} \quad \text{(A.10)}$$

then

$$\pi_t = -\frac{\kappa}{4(1-\beta a_\pi)}\vartheta\pi_{t-1} + \lambda^{cb}\frac{1}{\kappa}\vartheta u_t. \quad \text{(A.11)}$$

Comparing it with (A.1), we get (2.6) - (2.7).

Next, substitute (A.1) in (A.6) and rearrange,

$$\hat{y}_t = \frac{1}{\kappa}(1-\beta a_\pi)a_\pi\pi_{t-1} + \left(\frac{1}{\kappa}(1-\beta a_\pi)b_\pi - \frac{1}{\kappa} \right)u_t. \quad \text{(A.12)}$$

Comparing it with (A.2), we obtain (2.8) - (2.9).

A.2 Equilibrium of IT

The central bank minimizes the period loss of welfare in (2.2) with respect to $\hat{y}_t$, which results in the following first-order condition:

$$\pi_t\frac{\partial\pi_t}{\partial\hat{y}_t} + \lambda\hat{y}_t = 0. \quad \text{(A.13)}$$

The constraint is the IT Phillips curve in (2.13). Using the Phillips curve to substitute out $\frac{\partial\pi_t}{\partial\hat{y}_t}$ and $\hat{y}_t$ in (A.13), we obtain

$$\kappa\pi_t + \frac{\lambda}{\kappa}(\pi_t - u_t) = 0, \quad \text{(A.14)}$$

which is equivalent to (2.10). Combining it with (2.13) to cancel out π_t, we get (2.11).

A.3 Equilibrium of Two-Period AIT with Time Inconsistency

The central bank minimizes the period loss of welfare in (2.2) with respect to $\hat{y}_t$, which results in the following first-order condition:

$$\pi_t\frac{\partial\pi_t}{\partial\hat{y}_t} + \lambda\hat{y}_t = 0. \quad \text{(A.15)}$$

The private sector believes AIT announcement. Therefore, the Phillips curve is given by equation (2.12). Substituting (A.5) - (A.6) in, we obtain

$$\frac{\kappa}{1 - \beta a_\pi}\pi_t + \frac{\lambda}{\kappa}\left(1 - \beta a_\pi\right)\pi_t - \frac{\lambda}{\kappa}u_t = 0, \tag{A.16}$$

which is equivalent to (3.1). Substitute (3.1) into (A.6),

$$\hat{y}_t = \frac{1}{\kappa}\left(1 - \beta a_\pi\right)\left(\frac{\kappa}{1 - \beta a_\pi} + \lambda\frac{1 - \beta a_\pi}{\kappa}\right)^{-1}\frac{\lambda}{\kappa}u_t - \frac{1}{\kappa}u_t. \tag{A.17}$$

which is equivalent to (3.2).

B   Proofs

B.1   Proof of Lemma 1

We first prove $a_\pi < 0$ by contradiction. Suppose $a_\pi > 0$. For the system to have a stationary solution, $0 < a_\pi < 1$. Together with $0 < \beta < 1$, this implies $1 - \beta a_\pi > 0$. Together with $\kappa > 0$, this implies $a_y = \frac{1}{\kappa}\left(1 - \beta a_\pi\right)a_\pi > 0$. These results plus $\lambda^{cb} > 0$ mean $\vartheta > 0$ in (A.10). Consequently, $a_\pi = -\vartheta\frac{\kappa}{4(1 - \beta a_\pi)} < 0$, hence a contradiction. We have proved $a_\pi < 0$, which implies $1 - \beta a_\pi > 0$, and it follows that $a_y = \frac{1}{\kappa}\left(1 - \beta a_\pi\right)a_\pi < 0$.■

B.2   Proof of Proposition 1

The expectation term in the forward-looking Phillips curve (2.3) is $\mathbb{E}_t\,\pi_{t+1} = a_\pi\pi_t$ for AIT per (2.4), and $\mathbb{E}_t\,\pi_{t+1} = 0$ for IT per (2.10). Lemma 1 shows $a_\pi < 0$. Therefore $a_\pi\pi_t < 0$ when $\pi_t > 0$ and $a_\pi\pi_t > 0$ when $\pi_t < 0$.■

B.3   Proof of Proposition 2

The slope of the Phillips curve of AIT in (2.12) is $\frac{\kappa}{1 - \beta a_\pi}$, and that of the IT Phillips curve in (2.13) is κ. We have shown $a_\pi < 0$ in Lemma 1. Therefore, $1 - \beta a_\pi > 1$, and hence, $\frac{\kappa}{1 - \beta a_\pi} < \kappa$.■

B.4   Proof of Proposition 3

The AIT Phillips curve in (2.12) can be written as

$$(1 - \beta a_\pi)\pi_t^{AIT} = \kappa\hat{y}_t + u_t,$$

where we label inflation under AIT π_t^{AIT}. Compare it with the IT Phillips curve in (2.13):

$$\pi_t^{IT} = \kappa\hat{y}_t + u_t.$$

Given $\hat{y}_t$ and u_t,

$$(1 - \beta a_\pi)\pi_t^{AIT} = \pi_t^{IT}.$$

We have shown $a_\pi < 0$ in Lemma 1. Therefore, $1 - \beta a_\pi > 1$, and

$$|\pi_t^{IT}| > |\pi_t^{AIT}|.$$

∎

B.5 Proof of Lemma 2

The y-intercept of the Phillips curve of AIT in (2.12) is $\frac{1}{1-\beta a_\pi} u_t$, and that of the IT Phillips curve in (2.13) is u_t. We have shown $a_\pi < 0$ in Lemma 1. Therefore, $1 - \beta a_\pi > 1$, and hence, $\frac{1}{1-\beta a_\pi}|u_t| < |u_t|$ for $u_t \neq 0$.

Next, the x-intercept of the Phillips curve of AIT in (2.12) is characterized by $0 = \frac{\kappa}{1-\beta a_\pi}\hat{y}_t + \frac{1}{1-\beta a_\pi}u_t$, or $0 = \kappa\hat{y}_t + u_t$. For the IT Phillips curve in (2.13), it is also $0 = \kappa\hat{y}_t + u_t$. Hence, they have the same x-intercept.∎

B.6 Proof of Lemma 3

The AIT Phillips curve in (2.12) is

$$\pi_t = \frac{\kappa}{1 - \beta a_\pi}\hat{y}_t,$$

when $u_t = 0$. When $\hat{y}_t = 0, \pi_t = 0$.

Similarly, the IT Phillips curve in (2.12) is

$$\pi_t = \kappa\hat{y}_t,$$

when $u_t = 0$. When $\hat{y}_t = 0, \pi_t = 0$.∎

C Calibration

C.1 Structural parameters

We calibrate the model at an annual frequency, and our parameter values are in line with the standard New Keynesian literature; see parameters in Table C.1. The top panel displays the structural parameters we calibrate, some of which do not show up in the linear equations in our model, and the bottom panel displays implied parameter values. We set $\beta = 0.95$ to match the annual return on safe assets in the U.S. We assume a log utility for consumption, $\gamma = 1$, and a unitary Frisch elasticity of labor supply, $\varphi = 1$. We choose η, the price elasticity of demand, to be 4. We set the Calvo parameter, which governs price stickiness, to be $\frac{1}{3}$, implying the average duration between price changes is 1.5 years. We assume the cost-push shock is *i.i.d.* with a standard deviation of $\sigma_\epsilon = 0.1$.

With these underlying parameters, the slope of the Phillips curve, κ, is 2.73, and the weight on output gap stabilization in the welfare function, λ, is 0.6833.

Table C.1: Structural parameters

parameter	description	value
β	discount factor	0.95
γ	elasticity of intertemporal substitution	1
φ	elasticity of labor supply	1
η	the price elasticity of demand	10
θ	price rigidity	$\frac{1}{3}$
σ_ϵ	std. dev. of cost-push shock	0.1
κ	slope of the Phillips curve	2.73
λ	weight in welfare	0.6833

C.2 Learning parameters

We summarize the learning parameters in Table C.2. We set the number of agents in the private sector to be 1,000. We further assume the number of random meetings per period is 1,000. Our results are reported as the average from 1,000 simulations. We follow Arifovic, Bullard, and Kostyshyna (2013) by setting the mutation probability to 10%. For the parameter that governs the level of uncertainty, we use $\sigma_v = 0.1$ for high uncertainty, which is our benchmark case, 0.01 for low uncertainty, and 0 for no uncertainty.

Table C.2: Learning parameters

parameter	description	value
N	number of agents	1,000
	number of simulations	1,000
	number of meetings every period	1,000
	probability of mutation	0.1
σ_v	std. dev. in private signals	0.1 (benchmark) / 0.01 (low)/ 0 (no)