WORKING PAPERS

N° 1389

December 2022

"A Theory of Dynamic Inflation Targets"

Christopher Clayton and Andreas Schaab

Toulouse School of Economics

A Theory of Dynamic Inflation Targets

Christopher Clayton[*] Andreas Schaab[†]

November 2022

Abstract

Should central banks' inflation targets remain set in stone? We study a dynamic mechanism design problem between a government (principal) and a central bank (agent). The central bank has persistent private information about structural shocks. Firms learn the state from the central bank's reports and form inflation expectations accordingly. A *dynamic inflation target* implements the full-information commitment allocation: the central bank is delegated the authority to adjust its own target as long as it does so one period in advance. Both the level and flexibility of the dynamic inflation target respond to persistent shocks. Target flexibility is set to correct the time consistency problem, while the target level provides the correct incentives for target adjustments. An informational divine coincidence arises: the central bank's incentives to misreport its persistent private information to manipulate firm and government beliefs exactly offset each other under the mechanism. We apply our theory to study lower bound spells, a declining natural interest rate, and a flattening Phillips curve. We leverage our framework to study longer-horizon time consistency problems and speak to practical policy questions of inflation target design.

JEL codes: E52, D82

Keywords: inflation targeting, persistent private information, dynamic mechanism design, monetary policy, time consistency, dynamic inflation targets, informational divine coincidence

[*]Yale School of Management. Email: christopher.clayton@yale.edu. Website: http://christopherdclayton.com

[†]Toulouse School of Economics. Email: andreas.schaab@tse-fr.eu. Website: https://andreasschaab.com.

We are grateful to Manuel Amador, Marios Angeletos, Fabrice Collard, Emmanuel Farhi, Xavier Gabaix, Marina Halac, Sam Hanson, Christian Hellwig, Matteo Maggiori, Alp Simsek, Jeremy Stein, Ludwig Straub, Guillaume Sublet, Pierre Yared, Alex Zentefis, and seminar participants at CIREQ, Harvard, TSE, and Yale SOM for helpful comments and suggestions.

1 Introduction

Since their inception in the early 1990s, many central banks' inflation targets have evolved substantially. For example, the Bank of New Zealand has announced at least four major updates to its target definition since 1990.[1] The Bank of Canada undergoes regular reviews of its inflation target at 5-year intervals. In 2020 and 2021, the U.S. Federal Reserve and the European Central Bank both updated their inflation target frameworks.[2] Overall, central banks have exercised substantial discretion over target adjustments during this period.

In academic discourse, an important motivation for inflation targets is the interaction between a time consistency problem and central bank private information: commitment to a rule corrects inflationary bias while flexibility to set inflation allows the central bank to respond to private information about economic shocks.[3] Prior work has shown that a static inflation target solves this commitment-versus-flexibility trade-off in static environments or when shocks are uncorrelated (Walsh, 1995; Athey et al., 2005). These results motivate inflation targets as desirable mechanisms but do not speak to the empirical regularity that central banks regularly update their targets. When deliberating target adjustments, central banks in practice often invoke persistent economic change, which presupposes that shocks are correlated over time.[4] Recent debate on persistent changes in r^* and the slope of the Phillips curve—both difficult to measure in practice—highlights the importance of central bank persistent private information.

In this paper, we study a dynamic monetary policy game in the presence of persistent shocks and private information. As in previous work, the central bank faces a time consistency problem; unlike in previous work, persistent shocks make the central bank's private information persistent. This gives rise to additional information frictions because firms learn about the persistent state from the central bank, which they use to form inflation expectations. Our main result is that a time-varying, *dynamic inflation target* mechanism implements the efficient, full-information commitment allocation. The dynamic inflation target is a two-parameter mechanism, featuring both a *target level* and a *target flexibility*. Together, they serve the dual role of correcting the time consistency problem and the information frictions that emerge with persistent private information. A key property of our mechanism is an *informational divine coincidence*: under the dynamic inflation target, the central bank's incentive to misreport its information in order to bias firm inflation expectations

[1] The Bank of New Zealand's initial target postulated an inflation band of 0-2%. The band was revised in 1996 to 0-3% and again in 2002 to 1-3%. Another revision in 2012 added an explicit focus on the 2% target midpoint (McDermott and Williams, 2018).

[2] In August 2020, the Fed concluded a long-term strategic review by adopting a target that aims to "achieve inflation that averages 2% over time" (Powell, 2020). The ECB concluded a similar strategic review in July 2021, moving from a one-sided "below but close to 2%" inflation target to a symmetric one. At the same time, commentators have also suggested an upward revision in the inflation target level to 3 or 4% (Blanchard et al., 2010; Ball, 2014; Krugman, 2014).

[3] There is much empirical support for the existence of central bank private information. For example, see Romer and Romer (2000), Kuttner (2001), Gürkaynak et al. (2005), Campbell et al. (2012), Krishnamurthy and Vissing-Jorgensen (2012), and Lucca and Moench (2015) among many others.

[4] The strategic review that preceded the Fed's target adjustment in 2020 was partly motivated by the persistent decline in r^* and the accompanying concern about future lower bound spells (Clarida, 2019).

downwards to stimulate output is exactly offset by its incentive to bias the government's expectations upwards to reduce future penalties. Our paper generalizes the canonical work on inflation targets to environments with persistent private information.

Our infinite-horizon model features persistent economic shocks and general social preferences over inflation and output. Firms determine the current inflation-output relationship based on their expectations about next-period inflation, giving rise to a forward-looking Phillips curve. The standard time consistency problem emerges (Kydland and Prescott, 1977; Barro and Gordon, 1983). Neither firms nor the government observe the underlying economic state, which is persistent private information of a central bank that sets monetary policy under discretion. A Ramsey government (principal) designs a transfer/punishment mechanism to incentivize the central bank's (agent) policy decisions. Transfer/punishment mechanisms are important components of the inflation targeting framework—practical analogs include Congressional scrutiny, reputational risk, or firing (not reappointing) the central banker (Walsh, 1995; Halac and Yared, 2022).[5] The central bank's behavior under the mechanism reveals its persistent private information to both the government and firms. Firms in turn use this information to form inflation expectations, updating their beliefs about the distribution of future shocks and the conduct of future policy. An incentive compatible mechanism must account for both the time consistency problem of the central bank and its strategic incentive to use information revelation to influence firms' inflation expectations.

We develop our main result in Section 3: a *dynamic inflation target* mechanism implements the full-information Ramsey commitment allocation. This mechanism is incentive compatible—it overcomes both the central bank's time consistency problem and the strategic misreporting problem that arises under persistent private information. Formally, the dynamic inflation target is a two-parameter slope-intercept transfer rule,

$$T_t = -b_{t-1}(\pi_t - \tau_{t-1}).$$

The central bank faces a linear penalty for inflation, π_t, in excess of a *target level*, τ_{t-1}, with the slope of the penalty representing the *target flexibility*, b_{t-1}. The linear penalty for inflation is set so that the central bank internalizes the marginal cost of inflation in the prior period, which resolves the time consistency problem. Crucially, our mechanism implicitly delegates to the central bank the authority to update its own target—both level and flexibility—as long as it does so *one period in advance*. That is, the target parameters for date t are set at date $t-1$. The central bank takes as given its target (b_{t-1}, τ_{t-1}) at date t and can only make adjustments for the next period. Intuitively, the central bank internalizes its future time consistency problem when updating the target one

[5] In the U.S., for example, this process is multifaceted. The central bank Chair is directly held accountable by Congress in the form of bi-annual, as well as extraordinary, Congressional testimonies. Public hearings and independent scrutiny are also used more widely (Svensson, 2010). New Zealand allows for firing the central banker (Felix Hüfner, 2004; Halac and Yared, 2022). Delegation frameworks that call for bounds on inflation (Athey et al., 2005; Waki et al., 2018) can be thought of as a transfer/punishment mechanism where sufficiently large penalties are imposed for exceeding the bounds, and no penalties are imposed within the bounds.

period in advance.

At the heart of our paper is an *informational divine coincidence*: the dynamic inflation target overcomes the central bank's incentives to strategically misreport its private information. Intuitively, the central bank would benefit from biasing firm beliefs *downwards* in order to improve the contemporaneous inflation-output tradeoff. Our mechanism sets the target level equal to government inflation expectations, which provides a counteracting force: misreporting downwards becomes costly because it lowers the target level and raises expected future penalties from the mechanism. The central bank's incentive to bias firm expectations downwards is therefore exactly offset by its incentive to bias government expectations upwards.

We develop three applications of our theory in Section 4. Each is motivated by recent empirical evidence and monetary policy debates on structural change in the U.S., emphasizing in each case the relevance of persistent private information.

When confronting the effective lower bound, central banks have recently resorted to unconventional policy instruments, focusing largely on forward guidance and asset purchases. Some commentators have raised the question whether target adjustments can serve as an additional unconventional policy instrument. Our theory provides a natural framework to ask this question. It is well known that optimal monetary policy under commitment features history dependence at the lower bound, keeping interest rates low even after the economy exits the liquidity trap. This policy is implemented with an infinite sequence of promises, or forward guidance. We show that a dynamic inflation target can implement the commitment solution even in the presence of persistent private information. The optimal target adjustment raises both level and flexibility. Importantly, implementing the commitment solution relies only on one-period iterated commitments to a dynamic inflation target, which replace the long-horizon forward guidance commitment.[6]

In our second application, we study how the dynamic inflation target responds to a decline in the natural rate of interest r^* in the presence of an occasionally-binding effective lower bound on interest rates. With mounting empirical evidence for a historically low natural rate after the Great Recession (Laubach and Williams, 2016), this question has received ample attention. Many observers in the U.S. have explicitly advocated for an increase in the Federal Reserve's inflation target level (Blanchard et al., 2010). We show that a decline in r^* leads to an increase in the dynamic inflation target's level as well as, more surprisingly, to an increase in its flexibility. While academic and policy discourse has largely focused on implications for the optimal target level, our theory suggests an equally important role for adjustments of target flexibility. A second key insight of our analysis is that the presence of an occasionally-binding lower bound constraint may even lead to a sign switch in optimal target flexibility in steady state: When the probability of lower bound spells is sufficiently large, the benefits from higher inflation expectations can dominate the standard time

[6] Dynamic inflation target adjustments thus present an alternative implementation of forward guidance in the context of discretionary monetary policy. They serve much the same "commitment" role as asset purchases in Bhattarai et al. (2019). To the extent that long-horizon central bank promises lack perfect credibility in practice, dynamic target adjustments could therefore support forward guidance.

consistency problem, leading to *too little* inflation under discretion.

Our third application studies the implications of a flattening Phillips curve. The changing slope of the U.S. Phillips curve has garnered much attention since the Great Recession (Blanchard, 2016; Galí and Gambetti, 2019; Rubbo, 2020; Del Negro et al., 2020), and the ensuing debate has engulfed monetary policy discourse in recent years (Brainard, 2015). We show that a persistent flattening of the Phillips curve leads to a *decrease* in both the level and flexibility of the dynamic inflation target. Intuitively, a flattening Phillips curve increases the sensitivity of output to expected inflation and thus exacerbates the time consistency problem. In the context of the decline in r^* and the flattening of the Phillips curve, much recent policy discourse seems to have stressed the benefits of raising the target level and allowing for more flexibility.[7] While these target adjustments are indeed optimal in response to a decline in r^*, a flattening of the Phillips curve pushes in the opposite direction in both dimensions. These results have important policy implications if the flattening of the Phillips curve proves persistent.

A dynamic inflation target allows the central bank to adjust its own target *one period in advance*. To consider the implications of our result for policy design in practice, a natural question emerges: How long is a period and what is the appropriate horizon for target adjustments? We generalize our theory in Section 5 in the necessary dimensions to tackle this question. We consider forward-looking models where output depends on forecasts of inflation for the following K periods. A longer-horizon time consistency problem emerges. We show that a *K-horizon dynamic inflation target* implements the Ramsey allocation. It takes the form of a two-parameter transfer rule and parallels our baseline dynamic inflation target: its target flexibility equals the total time consistency problem over the last K periods, and its target level equals a weighted average of inflation forecasts for date t made over the last K periods. The informational divine coincidence continues to hold.

We introduce the *commitment curve*, which characterizes the duration and persistence of the promises the central bank makes to improve the contemporaneous inflation-output tradeoff. The commitment curve formally represents the size of the commitment the central bank makes at date t for all future periods $t + k$. The flatter the commitment curve, the more important long-horizon commitments are relative to short-horizon commitments.

Our main application in this environment characterizes the determinants of the appropriate horizon for target adjustments in practice. We consider a generalized New Keynesian Phillips Curve that emerges when linearizing the standard Calvo model around a steady state with positive trend inflation (Ascari, 2004; Ascari and Sbordone, 2014). We show that the commitment curve's shape is that of quasi-hyperbolic discounting (Laibson, 1997): The central bank makes a disproportionately large commitment for the next period, as well as an exponentially decaying sequence of commitments over longer horizons. We show that almost all long-horizon promises occur over a five-year horizon, suggesting that a five-year adjustment window like that of the Bank of Canada

[7] In fact, the Federal Reserve has adopted an average inflation target in August 2020, which arguably reflects an increase in target flexibility.

can capture all desirable long-horizon promises.

Finally, we study extensions of our model to incorporate different information structures (Section 6.1), costly mechanism enforcement (Section 6.2), and preference differences between the government and central bank (Appendix C.2). We show that a penalized adjustment process for the dynamic inflation target implements the Ramsey allocation when some firms are informed about the economic state. We also show that costly enforcement and preference disagreement imply optimal policies that parallel the insights obtained in our baseline model.

While monetary policy is the primary focus of this paper, our results could be applied more broadly to principal-agent settings where "moving goal posts" are desirable due to a combination of persistent private information and time consistency problems arising through expectations.[8]

Related literature. The paper most closely related to ours is Halac and Yared (2014). They study optimal delegation mechanisms in a fiscal policy framework with persistent private information and time inconsistency due to quasi-hyperbolic discounting. By contrast, we study persistent private information in the monetary policy context with transfers/punishments, where time inconsistency results from a forward-looking Phillips curve. Our environment features novel informational frictions because firms learn the state from the central bank's report, which has important implications for the design of the mechanism and gives rise to the informational divine coincidence. We build on the dynamic mechanism design literature with persistent private information. In particular, Pavan et al. (2014) provide conditions for implementability in a general principal-agent framework with transfers and persistent shocks.[9] We deploy these techniques to study central bank inflation targets.

We also build on the literature that studies transfer/punishment mechanisms in the monetary policy context.[10] Walsh (1995) shows that an inflation target is an optimal mechanism in a static context with transferable utility. The linear form of this static target follows the same intuition as the within-period linear form of our dynamic target. In the dynamic context, it is well understood that the full-information Ramsey allocation can be implemented with a linear inflation penalty whose slope is the recursive multiplier on the Phillips curve implementability condition (Marcet and Marimon, 2019; Svensson, 1997b; Svensson and Woodford, 2004). Dávila and Schaab (2022) extend the recursive multiplier approach to and study central bank targets in a heterogeneous-agent New Keynesian model. Our contribution is to study the impact of persistent private information in a principal-agent environment. Our framework provides a novel role for the target level in over-

[8] For example, the sovereign debt literature commonly features a time consistency problem that arises because long-term debt prices depend on the government's future fiscal policy decisions.

[9] A similar first-order approach is found, for example, in Farhi and Werning (2013).

[10] A large literature considers time inconsistency. For example, see Kydland and Prescott (1977), Barro and Gordon (1983), Canzoneri (1985), Rogoff (1985), Cukierman and Meltzer (1986), and Persson and Tabellini (1993) among many others. More broadly, there has been a long tradition considering the implications of private information for the design of policy. For example, see Backus and Driffill (1985), Sleet (2001), and Angeletos et al. (2006) among many others.

coming the incentives of the central bank to strategically reveal its persistent private information.[11] Halac and Yared (2022) study the trade-off between instrument-based and target-based rules in a framework with socially costly penalties. Finally, a related literature studies transfer mechanisms in the context of quasi-hyperbolic agents (DellaVigna and Malmendier, 2004; Galperti, 2015; Beshears et al., 2020).

A closely related literature studies the delegation approach rather than the transfer approach. Athey et al. (2005) studies a dynamic monetary policy framework with independent shocks and shows the optimal mechanism features static bounds on inflation. Waki et al. (2018) extends this framework to incorporate a New Keynesian Phillips curve and independent shocks, showing that the optimal mechanism consists of history dependent bounds on inflation. Amador et al. (2006), Halac and Yared (2018), and Sublet (2022) study delegation mechanisms to control quasi-hyperbolic agents. Our contribution to this literature is to characterize the optimal transfer/penalty mechanism in a setting with persistent shocks and private information, which gives rise to novel information frictions because firms learn the state from the central bank.

Our policy applications connect to several literatures. We build on prior work studying optimal monetary policy during lower bound spells (Eggertsson and Woodford, 2003; Werning, 2011). A long literature studies the optimal rate of inflation (Schmitt-Grohé and Uribe, 2010). Recent work has investigated whether a decline in r^* could quantitatively justify a higher inflation target level (Coibion et al., 2012; Kiley and Roberts, 2017; Andrade et al., 2018; Eggertsson et al., 2019). In our paper, we take as given that persistent structural shocks can alter the welfare implications of inflation and, consequently, the socially desired rate of inflation. We ask if and how a central bank should respond to such shocks—in the presence of persistent private information and time consistency problems—by adjusting its inflation target.

2 Model

Our economy is populated by a government, a monetary authority or central bank, and a continuum of small firms. The central bank learns about persistent changes in the state of the economy. It uses this private information, which we also refer to as the central bank's *type*, to set monetary policy under discretion. The central bank is subject to a time consistency problem in the tradition of Kydland and Prescott (1977) and Barro and Gordon (1983): Firms determine the relationship between inflation and output in a forward looking manner, which gives rise to a Phillips curve. The government (principal) designs a mechanism to control the inflation policies of the central bank (agent), taking as given the price-setting behavior of firms.

Time is infinite and discrete, indexed by $t = 0, 1, \ldots$ We summarize allocations by inflation

[11] More broadly, several papers have extended the Marcet and Marimon (2019) approach for providing a recursive representation of a planner's problem to environments with moral hazard and incomplete information (Messner et al., 2012; Mele, 2014; Pavoni et al., 2018). Our contribution is to study the problem of a principal designing a mechanism for an agent, rather than giving a recursive representation to the principal's problem.

$\pi_t \in [\underline{\pi}, \overline{\pi}]$ and output $y_t \in [\underline{y}, \overline{y}]$. There is a state of the economy, $\theta_t \in \Theta = [\underline{\theta}, \overline{\theta}]$, that follows a Markov process described by the conditional transition density $f(\theta_t|\theta_{t-1})$. The central bank observes the state θ_t at the beginning of t (i.e., θ_t is central bank private information) and is tasked with setting inflation for that period. Firms do not observe the state but form posterior beliefs μ_t on its distribution based on behavior of the central bank in that period.[12] We denote by $\mathbb{E}_t[\pi_{t+1} \mid \mu_t]$ firms' expectation of next-period inflation, given their posterior beliefs μ_t about the current state θ_t. Firms' price setting determines output based on future inflation expectations, giving rise to a "Phillips curve"[13]

$$y_t = F_t(\pi_t, \mathbb{E}_t[\pi_{t+1} \mid \mu_t]). \tag{1}$$

Because shocks are persistent, inflation expectations $\mathbb{E}_t[\pi_{t+1} \mid \mu_t]$ depend on firms' beliefs about the future conduct of monetary policy and the distribution of future shocks θ_{t+1}.[14]

The per-period social welfare function for the central bank and government over inflation and output is $\mathcal{U}_t(\pi_t, y_t, \theta_t)$. To simplify exposition, we internalize the Phillips curve relationship (1) and write reduced-form preferences as $U_t(\pi_t, \mathbb{E}_t[\pi_{t+1}|\mu_t], \theta_t) = \mathcal{U}_t(\pi_t, F_t(\pi_t, \mathbb{E}_t[\pi_{t+1}|\mu_t]), \theta_t)$. The lifetime social welfare function of the central bank and government over inflation can then be written as

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t \, U_t(\pi_t, \mathbb{E}_t[\pi_{t+1} \mid \mu_t], \theta_t), \tag{2}$$

where β is the discount factor.

In Section 4, we develop applications of our theory that make use of the canonical New Keynesian Phillips curve and loss function at a distorted steady state. We choose particular shocks θ_t that are motivated by empirical evidence and recent policy debates on important structural changes.

[12] There is a long tradition in macroeconomics to motivate and study monetary policy games when the central bank has private information (Sargent and Wallace, 1975; Barro and Gordon, 1983; Canzoneri, 1985; Rogoff, 1985; Walsh, 1995; Athey et al., 2005). There is much empirical support for central bank private information. Romer and Romer (2000) show that the difference between the Federal Reserve's private inflation forecasts and commercial inflation forecasts is a significant predictor of commercial forecast errors. Lucca and Moench (2015) document sizable excess returns on U.S. equities leading up to scheduled Federal Open Market Committee (FOMC) meetings, implying substantial private information content in FOMC announcements. Krishnamurthy and Vissing-Jorgensen (2012) find strong empirical support for a signaling channel of unconventional monetary policy, whereby asset purchases between 2009 and 2012 worked to a large extent by conveying private information to financial market participants. Kuttner (2001) and Gürkaynak et al. (2005) show that FOMC announcements are associated with price effects that are not due to changes in the policy rate itself. Campbell et al. (2012) show that asset prices and commercial macroeconomic forecasts respond strongly to the information content in FOMC announcements.

[13] Although we use linear expectations $\mathbb{E}_t \pi_{t+1}$, it is straightforward to adapt our framework to nonlinear expectations. For example, suppose that we had $y_t = F_t(\pi_t, \mathbb{E}_t g_{t+1}(\pi_{t+1}))$ for a nonlinear function g_t. Then define $\pi_t^* = g_t(\pi_{t+1})$, define the Phillips curve as $y_t = F_t^*(\pi_t^*, \mathbb{E}_t \pi_{t+1}^*) = F_t(g_t^{-1}(\pi_t^*), \mathbb{E}_t \pi_{t+1}^*)$, and similarly for the preference function. More generally if we have $\mathbb{E}_t g_{t+1}(\pi_{t+1}, y_{t+1})$, then we can define a new variable $\pi_t^* = g_t(\pi_t, y_t)$, and define the problem over (π_t^*, y_t) where $y_t = Y_t(\pi_t^*, \mathbb{E}_t \pi_{t+1}^*)$, where Y_t solves $Y_t(\pi_t^*, \mathbb{E}_t \pi_{t+1}^*) = F_t(g_t^{-1}(\pi_t^* | Y_t(\pi_t^*, \mathbb{E}_t \pi_{t+1}^*)), \mathbb{E}_t \pi_{t+1}^*)$.

[14] A key concern of this Phillips curve relationship is a Lucas critique—firms' price-setting behavior may change in response to changes in the monetary policy regime, such as target changes (L'Huillier and Schoenle, 2019). Our Phillips curve relationship is robust to a Lucas critique provided that expected future (next period) inflation is sufficient for determining how changes in future policies affect firm behavior. For example, higher expected inflation may lead firms to increase the frequency with which they update prices, altering the slope of the Phillips curve.

2.1 Benchmark: Full-Information Ramsey Allocation

We begin by providing a benchmark allocation for efficiency. In particular, we characterize the efficient allocation that arises when: (i) the central bank has full commitment (Ramsey problem); and (ii) firms have full information, i.e., they observe the shock at date t. Given full information, firms' posterior beliefs are the degenerate distribution which places all mass on θ_t, which we denote by $\mu_t = \theta_t$, abusing notation slightly. We refer to this allocation as the *full-information Ramsey allocation*. It provides an efficiency benchmark that respects the Phillips curve relationship between inflation and output determined by firms.

Proposition 1 (Full-Information Ramsey Allocation)**.** *The full-information Ramsey allocation is characterized by*

$$\frac{\partial U_t}{\partial \pi_t} = \nu_{t-1}\,, \qquad \text{where } \; \nu_{t-1} = \begin{cases} -\frac{1}{\beta}\frac{\partial U_{t-1}}{\partial \mathbb{E}_{t-1}(\pi_t|\theta_{t-1})} & \textit{for } t \geq 1 \\ 0 & \textit{for } t = 0 \end{cases} \tag{3}$$

The optimality condition for inflation at date t equates the marginal utility from inflation, $\partial U_t/\partial \pi_t$, with the marginal (dis)utility from the effect of inflation on previous period's output, summarized by ν_{t-1}. The left-hand side (LHS) of equation (3) is date t adapted, whereas the right-hand side (RHS) is date $t-1$ adapted. Therefore, the RHS is constant from the perspective of time t, implying that the marginal (flow) utility from inflation is constant at date t in histories θ^t proceeding from the same history θ^{t-1}.

The wedge ν_{t-1} is a sufficient statistic for the shock history θ^{t-1} in determining the Ramsey allocation rule $\pi_t, \pi_{t+1}, \ldots$ for inflation.[15] In other words, the Ramsey allocation from dates t and onward can be calculated with the knowledge of the wedge ν_{t-1}, without knowing the exact shock history θ^{t-1} that gave rise to it. Note that since the economy starts at $t = 0$, then $\nu_{-1} = 0$.

It is helpful to contrast the full-information commitment (Ramsey) allocation of Proposition 1 with the full-information discretion (Markov) policy. Under discretion, the central bank finds it optimal to set $\partial U_t/\partial \pi_t = 0$ state by state. In particular at date t, the central bank neglects the impact of inflation on the previous period's Phillips curve, which no longer serves as a constraint of the problem. This results in inflationary bias and reflects a standard Barro and Gordon (1983) time consistency problem. ν_{t-1} is precisely the wedge between the full-information Ramsey and Markov allocations. It reflects the severity of the central bank's time consistency problem. We therefore refer to ν_{t-1} as the *inflationary bias* of the central bank at time t. In the presence of persistent shocks, this inflationary bias is potentially time-varying.

This inflationary bias under discretion motivates studying how the government can design a mechanism to control the behavior of the central bank. Such a mechanism must respect the asymmetric information problem that stems from the central bank's persistent private information.

[15] Equivalently, we can give a recursive representation to the Ramsey problem (Marcet and Marimon 2019).

2.2 Mechanism Structure

Our framework is a principal-agent problem in which the central bank privately observes the state of the economy θ_t and then sets inflation under discretion. Because θ_t is private information and the central bank has a time consistency problem, the government (principal) designs a mechanism to control the decision making process of the central bank (agent). The mechanism the government establishes can specify transfers (or punishments) T_t based on inflation policy.

Although explicit monetary transfers are one interpretation, the practical analogs of the control mechanism T_t may be closer to policies such as Congressional scrutiny, reputational risk, or firing (not reappointing) the central banker (Walsh 1995, Svensson 2010, Halac and Yared 2022). For example, a central bank that is awarded high T_t may face a low degree of Congressional scrutiny in its policy determination.

The lifetime preferences of the central bank over social welfare and transfers are given by

$$\mathbb{E}\sum_{t=0}^{\infty}\beta^t\Big[U_t(\pi_t,\mathbb{E}_t[\pi_{t+1}|\mu_t],\theta_t)+T_t\Big]. \tag{4}$$

Our main focus will be on characterizing a mechanism that implements the full-information Ramsey allocation. Such a mechanism is optimal when there is no social cost of implementing the mechanism, as we assume here. In Section 6, we study the case where transfers are not neutral from the perspective of the government.

The mechanism requires the central bank to make a report of the observed shock at date t. We denote the reported type $\tilde{\theta}_t$ and say that reporting is *truthful* when $\tilde{\theta}_t = \theta_t$. We study direct and *full-transparency* mechanisms, under which the central bank truthfully reports its type each period.[16] Full transparency implies that there is no pooling of central bank types in reporting in a manner that shrouds the private information. Along the equilibrium path, agents' posterior will therefore be the degenerate distribution at the reported type, or $\mu_t = \tilde{\theta}_t$. Note that we abuse notation here because μ_t is a full distribution in general.[17]

We denote by Θ^t the space of shock histories up to date t. A mechanism in our model is a mapping from the history of reported types into a transfer and allocation, given by (π_t, T_t) : $\Theta^t \to \mathbb{R}^2$. Although the date t allocations depend on the entire history of reported types, we will show state space reduction results that allow us to characterize sufficient statistics for information histories.

[16] Once we restrict to full transparency, the Revelation Principle as usual allows us to focus on mechanisms where the central bank truthfully reports its type.

[17] Restricting attention to full transparency mechanisms is not without loss of generality. In principle, the government could want to pool central bank types to manipulate firms' posterior beliefs. By considering mechanisms under which the central bank truthfully reveals its type, we assume away such motivations. Given that central bank transparency has become an increasingly prominent focal point over the last two decades, we view the full transparency benchmark as important and realistic (Powell, 2019).

2.3 Incentives, Time Consistency, and Information

At every date t, the central bank makes a report $\tilde{\theta}_t$ of its true type θ_t. We define the value function of the central bank under the mechanism (π, T) by

$$\mathcal{W}_t(\theta^t) = \max_{\tilde{\theta}_t} \left\{ T_t + U_t\left(\pi_t, \mathbb{E}_t\left[\pi_{t+1}|\tilde{\theta}_t\right], \theta_t\right) + \beta \mathbb{E}_t \left[\mathcal{W}_{t+1}(\theta^{t-1}, \tilde{\theta}_t, \theta_{t+1}) \middle| \theta_t \right] \right\},$$

where θ^{t-1} is the history of reported types whereas θ_t is always the current true type. Note that π_t and T_t are functions of the current reported type and the history of reported types. The incentive constraint of the central bank at date t with history θ^{t-1} is

$$\begin{aligned} &U_t(\pi_t, \mathbb{E}_t[\pi_{t+1}|\theta_t], \theta_t) + T_t + \beta \mathbb{E}_t \left[\mathcal{W}_{t+1}(\theta^{t+1}) \middle| \theta_t \right] \\ &\geq U_t(\tilde{\pi}_t, \mathbb{E}_t[\tilde{\pi}_{t+1}|\tilde{\theta}], \theta_t) + \tilde{T}_t + \beta \mathbb{E}_t \left[\mathcal{W}_{t+1}(\theta^{t-1}, \tilde{\theta}, \theta_{t+1}) \middle| \theta_t \right] \end{aligned} \tag{5}$$

for all t, θ^t, and $\tilde{\theta}$, and where we denote $\tilde{\pi}_t = \pi_t(\theta^{t-1}, \tilde{\theta})$ and so on. Equation (5) is the truthful reporting incentive constraint: at date t, a central bank should find it preferable to truthfully report its type θ_t as opposed to reporting any alternate type $\tilde{\theta} \in \Theta$. As usual, incentive compatibility is characterized using a one-shot deviation along a path of truthful reporting, which is why the continuation value includes the true continuation type. The global incentive constraint (5) is high-dimensional, as there is an incentive constraint for each $\tilde{\theta} \in \Theta$ and every history $\theta^t \in \Theta^t$. As usual, we will employ a first order approach to incentive compatibility in deriving results (e.g., Pavan et al. 2014, Farhi and Werning 2013). The required envelope condition associated with global incentive compatibility—derived in the proof of our main result in Appendix A—is given by

$$\frac{\partial \mathcal{W}_t(\theta^t)}{\partial \theta_t} = \frac{\partial U_t\left(\pi_t, \mathbb{E}_t\left[\pi_{t+1}|\theta_t\right], \theta_t\right)}{\partial \theta_t} + \beta \mathbb{E}_t \left[\mathcal{W}_{t+1}(\theta^{t+1}) \frac{\partial f(\theta_{t+1}|\theta_t)/\partial \theta_t}{f(\theta_{t+1}|\theta_t)} \middle| \theta_t \right] \tag{6}$$

where, for clarity, $\frac{\partial U_t(\pi_t, \mathbb{E}_t[\pi_{t+1}|\theta_t], \theta_t)}{\partial \theta_t}$ is the derivative of U_t in the direct type θ_t, but *not* including the Phillips curve expectation (which is based on the reported type). The familiar integral incentive constraint is obtained by integrating and iterating forward (see the proof of Proposition 13 for this representation).

The global incentive constraint (5) and its envelope formulation (6) reveal three principal driving forces of the model. The first two are conventional forces. First, there is a standard Barro and Gordon (1983) time consistency problem, marked by the absence of any terms that capture the impact of inflation at date t on the Phillips curve at date $t-1$.[18]

Second, there are *information rents* the central bank earns from its persistent private information

[18] This follows the standard Barro and Gordon (1983) logic: When the central bank considers which type $\tilde{\theta}_t$ to report in period t under discretion, it does not consider the implications of its actions on past price-setting decisions of firms.

(Pavan et al., 2014). There are two components to this information rent (equation 6). The first is the static information rent, $\partial U_t / \partial \theta_t$. It captures the gain in welfare that the central bank achieves from an increase in its type θ_t, while holding fixed its report. An incentive compatible allocation must maintain this information rent to ensure that a central bank with a higher type does not report a lower type. The second is the dynamic information rent from shock persistence: revealing θ_t gives up private information about the distribution of future shocks, captured by the term $\frac{\partial f(\theta_{t+1}|\theta_t)}{\partial \theta_t}$. If high θ_t on average leads to high θ_{t+1} and high continuation values $\mathcal{W}_{t+1}$, then the information rent earned by θ_t is higher because the central bank knows it will receive high continuation values even without changing its report. This means that incentive compatibility requires awarding the central bank more for reporting higher values $\tilde{\theta}_t$ today. If shocks are not persistent, then the dynamic information rent is zero.

The third and novel force in our model is that the central bank has an incentive to manipulate firm beliefs. Firms form inflation expectations, which appear in the Phillips curve, based on their beliefs about next period's shock distribution. The central bank's report today affects these beliefs, i.e., $\mathbb{E}_t[\pi_{t+1} \mid \tilde{\theta}_t]$. Global incentive compatibility (5) reflects that a change in reported type alters the central bank's current flow utility indirectly by changing firms' inflation expectations. In the conventional New Keynesian framework, an increase in expected inflation generally lowers current flow utility. The central bank therefore has an incentive to bias firm expectations *downward* in order to improve the contemporaneous inflation-output trade-off. The fact that expectations are formed based on the *reported* type also means that the central bank does not earn an information rent from this channel. Formally this is seen in the fact that the first information rent term on the RHS of equation (6) is only the direct derivative in the type, and does not include a term for inflation expectations.

3 Dynamic Inflation Target

In this section, we show the main result of our paper: A "*dynamic inflation target*" mechanism can implement the full-information Ramsey allocation when the target is set by the central bank *one period in advance*. This mechanism overcomes the time consistency and informational problems we identified in Section 2.3.

Equation (3), along with its sufficient statistic implications, suggests a mechanism that uses the transfer rule T_t to penalize inflation deviations from a target. An inflation target of this form seeks to correct the time consistency problem by incentivizing the central bank to set inflation close to the target. In the presence of persistent structural shocks, however, the target itself might need to be adjusted over time to accommodate a changing efficient level of inflation. That is, the optimal inflation rate may drift far from the central bank's target in a persistent manner, implying large potential gains from letting the target adjust. The commitment-flexibility trade-off that motivated the inflation target in the first place may itself be subject to structural change. Indeed, this is precisely reflected in the time variation of the full-information Ramsey allocation in the presence of persistent θ_t shocks.

3.1 Inflation Targets as Dynamic Mechanisms

We look over a class of mechanisms defined by the affine transfer rule

$$T_t = -b_{t-1}(\pi_t - \tau_{t-1}). \tag{7}$$

We say that τ_{t-1} is the *level* of the transfer, and b_{t-1} is the *slope* of the punishment for increasing inflation.[19] This class of mechanisms specifies affine transfers at date t based on inflation at date t, where the affine function parameters (b_{t-1}, τ_{t-1}) are determined at date $t-1$.

We define in particular a class of mechanisms with affine transfer rules under which the level is expected next-period inflation.

Definition 2 (Dynamic Inflation Target)**.** A *dynamic inflation target* is an affine transfer rule mechanism whose *target level* equals expected inflation, $\tau_{t-1} = \mathbb{E}_{t-1}[\pi_t \mid \tilde{\theta}_{t-1}]$, and whose *target flexiblity* is the slope b_{t-1}.

Under our proposed dynamic inflation target, two things happen when the central bank reports its type θ_t at date t. First, its type report maps into a contemporaneous inflation policy π_t, which in turn generates a transfer T_t based on the target parameters (b_{t-1}, τ_{t-1}) specified in the previous period. The mechanism establishes a target in the sense that $\tau_{t-1} = \mathbb{E}_{t-1}[\pi_t | \tilde{\theta}_{t-1}]$; that is, the level of the mechanism is always equal to expected inflation. Second, the report also maps into target parameters (b_t, τ_t) for the transfer rule in the next period, i.e., the new target. In sum, the mechanism is a mapping $(\pi_t, b_t, \tau_t) : \Theta^t \to \mathbb{R}^3$ from the history of reported types into inflation for the current period and the target for the next period. In this sense, we can also think of the central bank as directly choosing inflation and its own future target, represented by (π_t, b_t, τ_t), from among the set of triples that follow from the same history θ^{t-1} of reported types prior to date t.

The *target level* τ_{t-1} and *target flexibility* b_{t-1} capture two distinct facets of the inflation target mechanism. The target level is the level of inflation the central bank is expected to hit on average. The target flexibility characterizes how severe the punishment is when the central bank exceeds its inflation target. An increase in target level means that the central bank incurs lower penalties for higher average inflation. An increase in target flexibility means the central bank incurs lower penalties for higher marginal inflation.

Our main result is that this dynamic inflation target implements the full-information Ramsey allocation in a locally incentive compatible mechanism. Moreover, it admits a key state space reduction property.

[19] The structure of our target mechanism and, in particular, its linearity are similar to Walsh (1995) in a static setting.

Proposition 3 (Dynamic Inflation Target Implements Efficient Allocation)**.** *A dynamic inflation target implements the full-information Ramsey allocation in a locally incentive compatible mechanism, with target flexibility $b_{t-1} = \nu_{t-1}$. The target (τ_{t-1}, b_{t-1}) is a sufficient statistic at date t for the history θ^{t-1} of past types.*

Proposition 3 shows that the full-information Ramsey allocation can be implemented by a simple dynamic inflation target. Inflation always meets the target level in expectation, that is $\tau_{t-1} = \mathbb{E}_{t-1}\pi_t$, while the target flexibility is set to the inflationary bias, $b_{t-1} = \nu_{t-1}$. The inflation target prescribed by our mechanism is dynamic in the sense that both its level and flexibility are time-varying.

Intuitively, the mechanism serves two roles: It uses the inherited target from the prior period to correct the time consistency problem in the central bank's contemporaneous inflation choice, and it provides incentives for correctly updating the target for the next period. The form of the inflation target follows the well-known logic from the static setting (Walsh, 1995). Since ν_{t-1} is the central bank's inflationary bias, the mechanism provides the correct incentives for the inflation choice by assigning a penalty $b_{t-1} = \nu_{t-1}$ for raising inflation. This means the target's flexibility is used as the means of correcting inflationary bias that arises from the fact that firm inflation expectations affect contemporaneous output.

In the presence of persistent shocks, the inflation target must be updated to accommodate persistent changes in the full-information Ramsey allocation. Proposition 3 yields two key insights. First, the central bank optimally resets its target one period in advance. That is, when the central bank observes a persistent shift in the efficient inflation level, it adjusts its inflation target for the *next* period in response to this shift. The current target, on the other hand, remains in effect for the current period and governs contemporaneous inflation policy. Second, both the target level τ_{t-1} and the target flexibility b_{t-1} are subject to change when the target is updated.

Dynamic target adjustments under our mechanism are best understood in relation to the underlying frictions discussed in Section 2.3. Consider first a change in the target flexibility. When the central bank updates b_t in period t—to go into effect and govern inflation policy in period $t+1$—it internalizes that expectations about future inflation affect output today via the Phillips curve. In other words, even though the central bank takes the behavior of its future self as given under discretion, it understands that the target it sets in period t will constrain the inflation policy of its future self in period $t+1$. The central bank consequently internalizes its future time consistency problem and corrects it by setting the appropriate penalty, $b_t = \nu_t$, for its future self—one period in advance.[20]

Our mechanism uses changes in the target level, τ_t, on the other hand, to overcome the core informational frictions of our model, in particular the central bank's incentives to manipulate firm and government beliefs in the presence of persistent private information. While it is surprising that

[20] This is also similar to the static setting, where the central bank is willing "ex ante" to set up a targeting mechanism for itself. It is also closely related to the literature on optimal mechanisms to control present bias (e.g. Amador et al. 2006), where agents are willing to set up mechanisms to control their own time consistency problems.

a simple dynamic inflation target is able to account for these complex effects, the affine transfer rule of our mechanism is designed so that the two information forces exactly offset each other. We call this important property of our mechanism *informational divine coincidence*.

To illustrate, consider the effect of a perturbation in inflation expectations on the central bank's lifetime value. The two relevant terms in the central bank's Bellman equation are $U_t(\pi_t, \mathbb{E}_t[\pi_{t+1} \mid \tilde{\theta}_t], \theta_t) + \beta\nu_t\mathbb{E}_t[\pi_{t+1} \mid \tilde{\theta}_t]$, where the latter comes from next-period's inflation target T_{t+1}. The indirect effect of a marginal perturbation in the central bank's report, $d\tilde{\theta}_t$, through inflation expectations is given by

$$
\begin{aligned}
&\frac{\partial}{\partial \mathbb{E}_t[\pi_{t+1}|\tilde{\theta}_t]}\left[U_t(\pi_t, \mathbb{E}_t[\pi_{t+1}|\tilde{\theta}_t], \theta_t) + \beta\nu_t\mathbb{E}_t[\pi_{t+1}|\tilde{\theta}_t]\right]\frac{\partial \mathbb{E}_t[\pi_{t+1}|\tilde{\theta}_t]}{\partial\tilde{\theta}_t}\\
=&\left(\frac{\partial U_t(\pi_t, \mathbb{E}_t[\pi_{t+1}|\tilde{\theta}_t], \theta_t)}{\partial \mathbb{E}_t[\pi_{t+1}|\tilde{\theta}_t]} + \beta\nu_t\right)\frac{\partial \mathbb{E}_t[\pi_{t+1}|\tilde{\theta}_t]}{\partial\tilde{\theta}_t}\\
=&\left(-\beta\nu_t + \beta\nu_t\right)\frac{\partial \mathbb{E}_t[\pi_{t+1}|\tilde{\theta}_t]}{\partial\tilde{\theta}_t}\\
=&0
\end{aligned}
$$

where the third line follows from Proposition 1. In economic terms, the central bank wishes to bias *downward* the inflation expectations of firms in order to economize on the Phillips curve relationship and improve the contemporaneous inflation-output trade-off. By setting next period's target level to also equal inflation expectations, i.e., $\tau_t = \mathbb{E}_t\pi_{t+1}$, the government provides the central bank with a distinct incentive to bias *upward* inflation expectations: increasing expected inflation raises the target level and so reduces average future penalties for high inflation. The marginal benefit of this upward bias is equal to the target flexibility, ν_t. But under the full-information Ramsey allocation, the target flexibility is precisely equal to the inflationary bias. Thus these two forces exactly offset each other.

This informational divine coincidence is central to our mechanism. It arises because firms and the government have the same information sets, i.e., firms learn from the mechanism in exactly the same way as the government does. This information structure is critical for a simple dynamic inflation target to be able to implement the full-information Ramsey allocation. In Section 6.1, we study the implications of alternative information structures for the design of dynamic inflation targets.

Finally, a key source of tractability for our mechanism is that the current target (ν_{t-1}, τ_{t-1}) is a sufficient statistic for the entire history θ^{t-1} of shock realizations. This means that our mechanism admits a recursive formulation where the date t state variables are the inherited target, (ν_{t-1}, τ_{t-1}), and the current state, θ_t. This sufficient statistic property follows precisely because the target flexibility ν_{t-1} summarizes the inflationary bias from the previous period, while the target level τ_{t-1} summarizes a form of promised utility to the central bank for truthfully revealing its persistent type. This property greatly reduces the knowledge required for the central bank to adjust its target: the central bank only needs to know its current target and not the history under which that target arose.

We characterize the first-order welfare gains of switching from a static to a dynamic inflation target in Appendix C.1.

3.2 Evolution of the Target

A key feature of our dynamic inflation target is that it can be updated over time by the central bank. This subsection characterizes the evolution of the target's flexibility and level in response to structural shocks.

Target flexibility. Combining the Ramsey first-order condition (3) with the definition of ν_t, we obtain the law of motion for target flexibility

$$\nu_t = \delta_t \left(\nu_{t-1} - \frac{\partial \mathcal{U}_t}{\partial \pi_t} \right), \tag{8}$$

where the derivative $\frac{\partial \mathcal{U}}{\partial \pi_t}$ holds output fixed, and where $\delta_t = \frac{-\partial y_t / \partial \mathbb{E}_t \pi_{t+1}}{\beta \partial y_t / \partial \pi_t}$ measures the relative effects of inflation expectations and current inflation on current output. For the standard New Keynesian Phillips Curve, we have $\delta_t = 1$. $\delta_t < 1$ implies contemporaneous inflation has a larger effect on output than inflation expectations, while $\delta_t > 1$ implies contemporaneous inflation has a smaller effect.[21]

When $\delta_t = 1$, the evolution of target flexibility is $\nu_t = \nu_{t-1} - \frac{\partial \mathcal{U}_t}{\partial \pi_t}$. The mechanism starts from the flexibility afforded in the current period, ν_{t-1}, and then adjusts ν_t *upward* as the central bank incurs greater *disutility* from inflation today, $\frac{\partial \mathcal{U}_t}{\partial \pi_t}$. Intuitively, if the central bank is willing to incur greater disutility from inflation, then the value from stimulating output must be higher. But when $\delta_t = 1$ and the effects of contemporaneous and future inflation on output are the same, then the cost of future inflation must also be high. The high cost of future inflation means the time consistency problem is large, leading the central bank to adopt a less flexible target for the next period.

When $\delta_t \neq 1$, target adjustment is scaled by the relative pass-through of current and future inflation to output. If $\delta_t < 1$, current inflation has a larger impact on output than future inflation. The time consistency problem is then less severe than under the standard Phillips curve, and the mechanism imposes an increasingly more flexible target over time. By contrast if $\delta_t > 1$, future inflation has a larger impact on output, the time consistency problem is more severe, and the target becomes less flexible over time.

[21] See Werning (2022) for a recent treatment of the pass-through of inflation expectations.

Target level. The response of the target level to a marginal increase in the structural shock is

$$\frac{d\tau_t}{d\theta_t} = \underbrace{\mathbb{E}_t\left[\pi_{t+1}\frac{\partial f(\theta_{t+1}|\theta_t)/\partial\theta_t}{f(\theta_{t+1}|\theta_t)}\right]}_{\text{Expectations}} + \underbrace{\frac{\partial\nu_t}{\partial\theta_t}\mathbb{E}_t\left[\frac{\partial\pi_{t+1}}{\partial\nu_t}\middle|\theta_t\right]}_{\text{Target Flexibility Adjustment}} . \tag{9}$$

The first effect, "expectations," reflects that the probability measure over future states changes in response to the shock. If a higher θ_t raises the probability of high-inflation states, then the target level τ_{t-1} increases as well. The expectations effect therefore implies that, in response to persistent shocks, the target intercept can change even when the target slope remains constant. If shocks were fully transitory, on the other hand, the probability measure would not be affected and no adjustment in the target level would be required.

The second effect, "target slope adjustment," reflects the extent to which a change in target flexibility passes through to optimal future inflation. In the natural case where $\frac{\partial\pi_t}{\partial\nu_{t-1}} < 0$, an increase in target flexibility ($\frac{\partial\nu_{t-1}}{\partial\theta_{t-1}} < 0$) is accompanied by an increase in target level, and vice versa. In economic terms, if a structural shock leads to an increase in flexibility (reduction in ν_{t-1}), then the central bank will find it optimal to generate higher average levels of inflation in the next period, since the penalty for exceeding the target has been reduced. Firms anticipate this, so that inflation expectations increase for a given state θ_{t-1} and associated conditional density f. As a result, the target level τ_{t-1}, which is set equal to firm inflation expectations, also increases.

In sum, when the economy experiences a structural shock θ_t, both components of the target may be affected. The flexibility of the target responds to the shock if it leads to a fundamental change in either the central bank's motivation to generate excess inflation or the nature of the time consistency problem. The level of the target is affected directly by changes in expectations but also indirectly if target flexibility is adjusted.

4 Applications

This section presents three applications of our theory. Each stresses in a different way the relevance of persistent private information for the design and conduct of monetary policy.[22]

In our applications, we study special cases of a linearized New Keynesian model. The model consists of a standard New Keynesian Phillips Curve (NKPC), given by

$$\pi_t = \beta\mathbb{E}_t\pi_{t+1} + \kappa_t y_t, \tag{10}$$

and a dynamic IS equation, given by

$$y_t = \mathbb{E}_t y_{t+1} - \frac{1}{\sigma}\Big(i_t - \mathbb{E}_t\pi_{t+1} - r_t^* + \epsilon_t\Big), \tag{11}$$

[22] Appendix B provides additional applications and also revisits these applications with costly enforcement (Section 6.2).

where π_t and y_t denote inflation and the output gap, respectively. This formulation of the model potentially allows for time variation in several parameters. We think of $\{\kappa_t, r_t^*, \epsilon_t\}$ as exogenously specified stochastic processes, given some initial conditions $\{\kappa_0, r_0^*, \epsilon_0\}$. In particular, κ_t is the slope of the Phillips curve, r_t^* denotes the natural rate of interest, and ρ_t captures a demand shock.[23]

4.1 Lower Bound Spells: Target Adjustments as Unconventional Policy

When the economy is at the effective (zero) lower bound, which we refer to as a "lower bound spell", the central bank loses its conventional policy instrument (short-term interest rates). Historically, central banks have then resorted to unconventional policy, focusing largely on forward guidance and asset purchases. Some commentators have explicitly raised the question whether changes in the targeting framework could and should be seen as a potential additional unconventional monetary policy instrument. Our theory provides a natural framework to ask this question.[24]

Zero lower bound spells are commonly represented by a constraint $i_t \geq 0$ (Eggertsson and Woodford, 2003; Werning, 2011). Consider a canonical loss function at a distorted steady state, $\mathcal{U}(\pi_t, y_t) = -\frac{1}{2}\pi_t^2 - \frac{1}{2}\alpha y_t^2 + \lambda y_t$. When explicitly accounting for the zero lower bound constraint, $i_t \geq 0$, social welfare can be associated with the Lagrangian $\mathbb{E}\sum_{t=0}^{\infty}\beta^t\left[\mathcal{U}(\pi_t, y_t) + \theta_t i_t\right]$. The Lagrange multiplier θ_t can be interpreted as the shadow value of being able to set negative nominal rates. In other words, when the economy falls into a liquidity trap, the shadow value on policies that push the economy away from the constraint rises—for example by raising inflation expectations, lowering current output, or raising future expected output.

In this application, we represent the mechanism design problem directly over the reduced-form loss function $\mathcal{U}_t(\pi_t, y_t) + \theta_t i_t$, which encodes the shadow value of being able to set negative rates directly into utility. A positive innovation to θ_t qualitatively captures the same economics as an explicit lower bound spell: the shadow value of higher nominal interest rates or, as we show below, higher inflation expectations rises. We associate a persistent lower bound spell with a persistently high shadow value θ_t.

We assume that $\mathbb{E}_t\theta_{t+1} = \rho\theta_t$ for $0 \leq \rho \leq 1$. We associate $\rho = 0$ with a transitory liquidity trap, where the lower bound constraint is expected not to bind in the following period. In this application, we abstract from shocks to the slope of the Phillips curve, $\kappa_t = \kappa$, innovations in the natural rate, $r_t^* = r^*$, and demand shocks, $\epsilon_t = 0$. Substituting the NKPC (10) into the dynamic IS equation (11) then implies

$$i_t = \mathbb{E}_t\pi_{t+1} + r^* + \frac{\sigma}{\kappa}\Big[-\pi_t + (1+\beta)\mathbb{E}_t\pi_{t+1} - \beta\mathbb{E}_t\pi_{t+2}\Big]. \tag{12}$$

[23] Equations (10) and (11) also take as given a process for the nominal interest rate $\{i_t\}$. In our applications, the nominal interest rate—the conventional instrument of monetary policy—will be determined as part of the dynamic inflation target mechanism.

[24] Crucially, we implicitly abstract from asset purchases: That is, we do not allow the central bank to use any other unconventional tool that would allow it to make the lower bound constraint slack again. We assume that instruments are incomplete to such an extent that the economy experiences a lower bound spell.

This means that, after substituting out for i_t and y_t in preferences $\mathcal{U}_t(\pi_t, y_t) + \theta_t i_t$, we can represent reduced-form preferences by $U_t(\pi_t, \mathbb{E}_t\pi_{t+1}, \mathbb{E}_t\pi_{t+2}, \theta_t)$. Since $\mathbb{E}_t\pi_{t+2}$ appears in this implementability condition, the resulting time consistency problem has a horizon of more than one period. We study longer-horizon time consistency problems in Section 5, where we revisit this application for general $\sigma \neq 0$. In this application, we set $\sigma = 0$ so that the time consistency problem reverts to a single period. We can then rewrite the reduced-form utility function as

$$U_t(\pi_t, \mathbb{E}_t\pi_{t+1}, \theta_t) = -\frac{1}{2}\pi_t^2 - \frac{1}{2}\hat{\alpha}\left(\pi_t - \beta\mathbb{E}_t\pi_{t+1}\right)^2 + \hat{\lambda}\left(\pi_t - \beta\mathbb{E}_t\pi_{t+1}\right) + \theta_t\left(\mathbb{E}_t\pi_{t+1} + r^*\right)$$

where $\hat{\alpha} = \frac{\alpha}{\kappa^2}$ and $\hat{\lambda} = \frac{\lambda}{\kappa}$.[25] We now characterize the dynamic inflation target of Proposition 3 when the economy experiences a lower bound spell.

Proposition 4. *The dynamic inflation target that implements the full-information Ramsey allocation is*

$$\nu_t = \gamma_0 + \gamma_1\theta_t + \gamma_2\nu_{t-1}$$

$$\mathbb{E}_t\pi_{t+1} = \gamma_0 + (\gamma_2 - 1)\nu_t + \left(\gamma_1 + \frac{1}{\beta}\right)\rho\theta_t$$

where $\gamma_0 = \frac{\frac{\hat{\lambda}}{\hat{\alpha}}\gamma_2}{1-\beta\gamma_2} > 0$*, where* $\gamma_1 = \frac{\gamma_2}{1-\gamma_2\beta\rho}\left[\rho - \frac{1+\hat{\alpha}}{\hat{\alpha}}\frac{1}{\beta}\right] < 0$*, and where* $\gamma_2 = \frac{1+\hat{\alpha}(1+\beta)-\sqrt{(1+\hat{\alpha}(1+\beta))^2-4\hat{\alpha}^2\beta}}{2\hat{\alpha}\beta}$ *with* $0 < \gamma_2 < 1$*. Optimal inflation sets* $\pi_t = \nu_t - \nu_{t-1} + \frac{1}{\beta}\theta_t$*.*

To illustrate the economic forces that govern the dynamic inflation target mechanism, consider the following exercise: We initialize the economy at its risky steady state.[26] Formally, we consider a particular realization of the stochastic process where $\theta_t = 0$ for sufficiently many periods such that the economy and the mechanism asymptotically converge. It is straightforward to see that the target flexibility converges to $\nu_t \to \nu = \frac{\gamma_0}{1-\gamma_2} = \frac{1}{1-\gamma_2}\frac{\gamma_2}{1-\beta\gamma_2}\kappa\lambda > 0$ in this limit. In the language of Svensson (1997b), the distorted steady state $\lambda > 0$ implies that there is an *average inflationary bias*, which $\nu > 0$ corrects. Similarly, the target level converges to $\tau_t = \mathbb{E}_t\pi_{t+1} \to \tau = \gamma_0 + (\gamma_2 - 1)\nu = 0$ in the risky steady state limit. This reflects a common Ramsey intuition: with a distorted steady state, the central bank achieves a better inflation-output trade-off today by promising lower inflation tomorrow, and subsequently achieves a better inflation-output trade-off tomorrow by promising future lower

[25] In both this application and the ones that follow, the proof shows that there are two linear solutions that satisfy the first order conditions of the optimum, and we take the non-explosive solution to remain consistent with the transversality condition.

[26] We define the risky steady state of the economy under a dynamic inflation target as comprising the allocation, prices, and target parameters (τ, ν) that the model converges to if a shock sequence of $\theta_t = 0$ for all t is realized. This is distinct from the standard deterministic steady state because agents understand that the environment is stochastic. It is also distinct from the stochastic steady state, which describes the random variables that allocation, prices, and target parameters converge to in distribution as the model is simulated for a sufficiently long period of time under the ergodic stochastic process $\{\theta_t\}$.

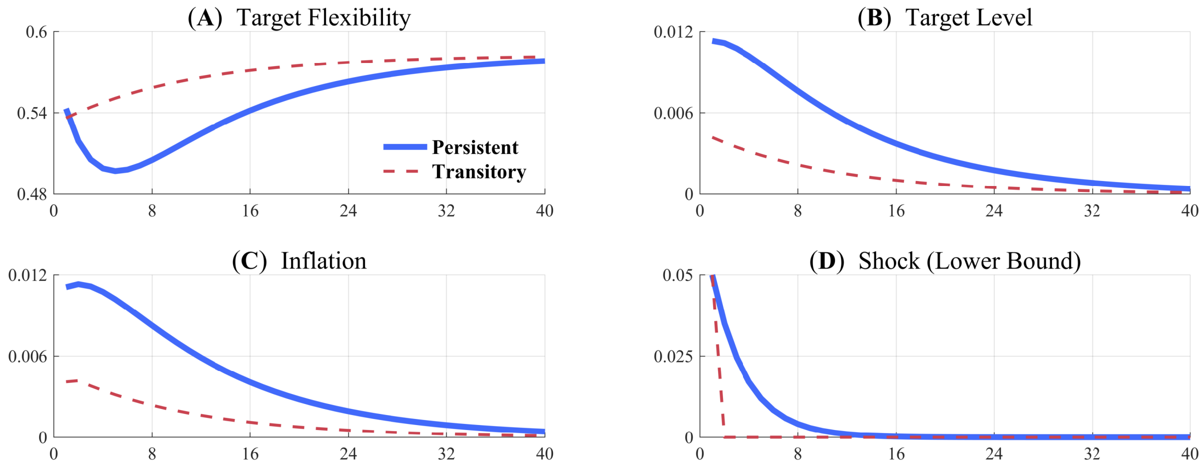

Figure 1. Impulse Responses: Lower Bound Spell

Note. Figure 1 plots the impulse responses of inflation and the dynamic inflation target after a lower bound shock $\theta_0 > 0$. Panels (A) through (D) show target flexibility, target level, inflation, and the shock, respectively. We target a quarterly calibration, staying as close as possible to Galí (2015), setting $\beta = 0.99$, $\alpha = 0.75$, and $\kappa = \frac{(1-\alpha)(1-\alpha\beta)}{\alpha}$. The blue solid line corresponds to a persistent shock ($\rho = 0.6$) and the red dashed line to a transitory shock ($\rho = 0$). In each case, we initialize the economy at the risky steady state and consider a shock at time 0.

inflation, and so on. This pushes optimal inflation under commitment towards zero in the long run, absent shock innovations. Formally, the allocation rule implies $\pi_t = \nu_t - \nu_{t-1} + \frac{1}{\beta}\theta_t \to \nu - \nu = 0$. Our dynamic inflation target implements the long-run Ramsey allocation in the risky steady state of this economy with a target level of $\tau = 0$ and a positive target flexibility $\nu > 0$ that exactly offsets the central bank's time inconsistent incentive to respond to the steady state distortion.[27]

We now initialize the economy at this risky steady state and consider a positive realization of the shock, $\theta_0 > 0$. Intuitively, we consider the economy as having entered a lower bound spell of uncertain duration at date 0. We plot the resulting impulse response functions (IRFs) under the dynamic inflation target mechanism in Figure 1.

Suppose first that the ZLB spell is purely transitory, and hence $\mathbb{E}_0\theta_1 = 0$. We consider a realization of the shock path such that $\theta_t = 0$ for all $t \geq 1$. The red-dashed line in Panel (a) of Figure 1 plots the dynamics of the target flexibility under this path.

The dynamic inflation target becomes more flexible at the lower bound, i.e., ν_0 falls since $\gamma_1 < 0$. Intuitively, the transitory lower bound spell increases the shadow value of future inflation and calls for a lower future inflation penalty. Even though the economy escapes from the lower bound at date 1, the added target flexibility is persistent and decays only at the rate $\gamma_2 < 1$. This endogenous persistence in the target response captures the standard intuition that optimal

[27] Similarly, we have $i_t \to r^*$ and $y_t \to 0$. The allocation in the risky steady state is therefore the same as in the deterministic steady state of this model. This follows from certainty equivalence under a first-order linearization.

monetary policy in a liquidity trap makes long-lived promises to keep interest rates low even after the economy moves away from the lower bound (Werning, 2011). Intuitively, promising high inflation at date 1 means that unless the central bank also promises high inflation at date 2, the economy experiences a significant output contraction at date 1. The central bank therefore smooths the output contraction by promising to maintain higher inflation for longer.

The associated increase in inflation expectations is also reflected in an upwards adjustment of the target level—see panel (b) of Figure 1. This reflects the success of the central bank in using the increased target flexibility to raise inflation expectations. It manifests in a higher inflation level in the next period. Coinciding with the gradual decay in target flexibility, the target level and realized inflation also both remain above zero even after the shock has phased out.

A persistent shock, $\rho > 0$, leads to qualitatively similar but more persistent dynamics. Target flexibility is hump-shaped under persistence, at first becoming *more* flexible as the shock phases out. The intuition comes from the evolution of ν_t in Proposition 4: A persistent lower bound spell makes it valuable to increase target flexibility beyond the initial date, which is reinforced by earlier promises of greater flexibility. These promises therefore compound in the initial phase of the lower bound spell.

Inflation target adjustments as forward guidance. The full-information Ramsey solution in a liquidity trap is well understood. In the language of Eggertsson and Woodford (2003), optimal monetary policy features history dependence and keeps interest rates low beyond the duration of the lower bound spell. This policy is implemented through an infinite sequence of promises, or "forward guidance."

Proposition 4 demonstrates how a dynamic inflation target can implement the commitment solution even in the presence of persistent private information. More importantly, however, implementing the optimal policy relies only on one-period iterated commitments to a dynamic inflation target. The central bank can implement forward guidance by adjusting its dynamic inflation target, replacing the long-horizon forward guidance commitment by a sequence of iterated one-period commitments.

Our theory provides a role for target adjustments as an unconventional monetary policy instrument: the dynamic inflation target implements forward guidance, thus serving much the same "commitment" role as asset purchases in Bhattarai et al. (2019). To the extent that long-horizon central bank promises lack perfect credibility in practice, dynamic target adjustments could therefore support forward guidance. This is not unlike the view under flexible inflation targeting—already a mainstay idea in central banking—that there may be benefits to allowing short-run flexibility around the central bank's inflation goal.

Another important insight of our theory provides a cautionary tale for such arguments, however: target adjustments must be made *one period in advance*. Indeed, determining the appropriate horizon for target adjustments becomes crucial to evaluate the implications of our theory for policy design practice; we take up this question in Section 5.

4.2 r^*

A vibrant debate has emerged on how monetary policy should respond to the decline in the natural rate of interest, or r^* (Laubach and Williams, 2016). In the presence of an effective lower bound (ELB) on interest rates, a decline in the natural rate implies that nominal rates are closer to the ELB on average and there is less room for central banks to use conventional policy during recessions. Many observers in the U.S. have advocated for an increase in the Federal Reserve's inflation target (Blanchard et al., 2010; Ball, 2014; Krugman, 2014).

This application studies a persistent fall in the natural rate of interest in the stochastic steady state of an economy that may encounter the ELB with some probability in the future. We denote movements in the natural rate by $\theta_t = r_t^* - r^*$, i.e., as deviations from a long-run value. We assume that $\mathbb{E}_t\theta_{t+1} = \rho\theta_t$ for $0 \leq \rho \leq 1$, where high ρ corresponds to persistence in natural rate movements. Finally, we set $\kappa_t = \kappa$, shutting off shocks to the slope of the Phillips curve.

We model an effective lower bound as follows. At the beginning of each date t, the central bank observes θ_t and adopts a rule for monetary policy. After the policy rule has been set for the period, an observable transitory demand shock $\varepsilon \in [\underline{\epsilon} + r^*, \overline{\epsilon} + r^*]$ is realized, with $\mathbb{E}\varepsilon = 0$. Defining $\epsilon_t = \varepsilon_t - r^*$, so that $\epsilon_t \in [\underline{\epsilon}, \overline{\epsilon}]$ and $\mathbb{E}\epsilon_t = -r^*$, we have

$$i_t = \mathbb{E}_t\pi_{t+1} + \theta_t - \epsilon_t,$$

where as in Section 4.1 we set $\sigma = 0$ in the dynamic IS equation. We represent the effective lower bound as a separable utility penalty $\lambda_0 - \lambda_1 i_t$ for negative realized nominal interest rates $i_t < 0$, with $\lambda_0, \lambda_1 \geq 0$. The expected (dis)utility from the ELB is therefore

$$v(i_t^*) = -\int_{i_t^*}^{\overline{\epsilon}} \Big(\lambda_0 - \lambda_1(i_t^* - \epsilon)\Big) f(\epsilon) d\epsilon$$

where $i_t^* \equiv \mathbb{E}_t\pi_{t+1} + \theta_t$. We assume ϵ is uniformly distributed, $f(\epsilon) = \frac{1}{\overline{\epsilon} - \underline{\epsilon}}$, so that we have

$$v(i_t^*) = -\underline{v} + \beta v_0 i_t^* - \frac{1}{2}\beta v_1 i_t^{*2},$$

for $v_0 = \frac{1}{2\beta}\frac{\lambda_0 + \lambda_1\overline{\epsilon}}{\overline{\epsilon} + r^*}$ and $v_1 = \frac{1}{2\beta}\frac{\lambda_1}{\overline{\epsilon} + r^*}$.[28] This means we have $v'(i_t^*) = \beta(v_0 - v_1 i_t^*)$. Intuitively, the effective lower bound generates welfare gains from setting $i_t^* > 0$ as this creates distance from the ELB and makes it less likely that a demand shock will push nominal rates below zero. This has implications for the design of the dynamic inflation target, as we show below.

Social preferences take the form $\mathcal{U}_t(\pi_t, y_t, i_t^*) = -\frac{1}{2}\pi_t^2 - \frac{1}{2}\alpha y_t^2 + v(i_t^*)$, reflecting both the loss from inflation and output gaps, and the loss from ELB penalties. We can then represent the

[28] We further have $\underline{v} = \lambda_0 \frac{\overline{\epsilon}}{\overline{\epsilon} - \underline{\epsilon}} + \frac{1}{2}\lambda_1 \frac{\overline{\epsilon}^2}{\overline{\epsilon} - \underline{\epsilon}}$.

reduced-form utility function as

$$U_t(\pi_t, \mathbb{E}_t\pi_{t+1}, \theta_t) = -\frac{1}{2}\pi_t^2 - \frac{1}{2}\hat{\alpha}\left(\pi_t - \beta\mathbb{E}_t\pi_{t+1}\right)^2 + v(\mathbb{E}_t\pi_{t+1} + \theta_t)$$

for $\hat{\alpha} = \frac{\alpha}{\kappa^2}$. We now characterize the dynamic inflation target in the presence of time-varying natural interest rate movements θ_t.

Proposition 5. *The dynamic inflation target that implements the full-information Ramsey allocation is*

$$\nu_t = \delta_0 + \delta_1\nu_{t-1} + \delta_2\theta_t$$
$$\frac{1}{\zeta}\tau_t = v_0 + \delta_0 + \left(\delta_1 - \frac{\hat{\alpha}\beta + v_1}{\hat{\alpha}\beta}\right)\nu_t + (\delta_2 - v_1)\rho\theta_t$$

where $\delta_0 = -\delta_1\frac{1+\hat{\alpha}(1-\beta)}{\hat{\alpha}(1-\beta\delta_1)}v_0 < 0$, *where* $\delta_1 = \frac{1+\hat{\alpha}(1+\beta)+v_1-\sqrt{(1+\hat{\alpha}(1+\beta)+v_1)^2-4\hat{\alpha}^2\beta}}{2\hat{\alpha}\beta}$ *with* $0 \leq \delta_1 \leq 1$, *where* $\delta_2 = \frac{1}{\hat{\alpha}}\frac{1+\hat{\alpha}(1-\beta\rho)}{(1-\beta\delta_1\rho)}\delta_1 v_1 < v_1$, *and where* $\zeta = \frac{\hat{\alpha}\beta}{(\hat{\alpha}\beta+v_1)(1+\hat{\alpha})-\hat{\alpha}^2\beta} > 0$. *Inflation is given by* $\frac{1}{\zeta}\pi_t = \nu_t - \frac{\hat{\alpha}\beta+v_1}{\hat{\alpha}\beta}\nu_{t-1} + v_0 - v_1\theta_t$.

To illustrate the economic forces at play, we again start by analyzing the risky steady state of the economy under the dynamic inflation target, which corresponds formally to the limit under a sufficiently long shock realization with $\theta_t = \varrho_t = 0$. In this limit, the target level converges to

$$\tau_t \to \tau = \zeta\left(v_0 + \delta_0\right) + \zeta\left(\delta_1 - \frac{\hat{\alpha}\beta + v_1}{\hat{\alpha}\beta}\right)\nu > 0,$$

which is positive. Unlike in Section 4.1, the risk of a lower bound spell in this economy leads to a positive long-run target level. It is well understood that proximity to an occasionally-binding effective lower bound provides an economic motive for a higher inflation target level. Proposition 5 formalizes this logic, showing that the dynamic inflation target features a positive target level in the risky steady state, even in the presence of persistent private information. If we shut off the risk of hitting the occasionally-binding effective lower bound, i.e., take the limit as $\lambda_0, \lambda_1 \to 0$, we have $\tau \to 0$ as in Section 4.1.[29]

While the implications of the effective lower bound for the optimal inflation target *level* may be well appreciated, our qualitatively novel result is that optimal long-run target *flexibility* is also

[29] Even in this limit, we study an economy at a non-distorted steady state in this subsection, whereas we considered a distorted steady state in Section 4.1. While the time consistency problem that stems from this distorted affects target flexibility in the risky steady state, this discussion highlights that the target level is unaffected. The commitment Ramsey allocation features 0 long-run inflation for the same economic rationale as we discussed in Section 4.1.

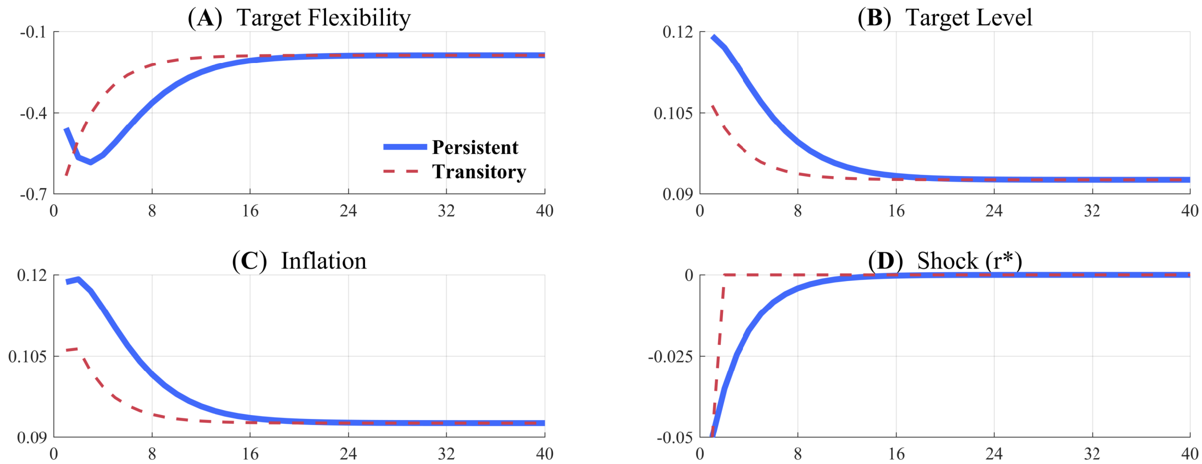

Figure 2. Impulse Responses: r^*

Note. Figure 2 plots the impulse responses of inflation and the dynamic inflation target after a decline in the natural rate of interest. Panels (A) through (D) show target flexibility, target level, inflation, and the shock, respectively. We target a quarterly calibration, staying as close as possible to Galí (2015), setting $\beta = 0.99$, $\alpha = 0.75$, and $\kappa = \frac{(1-\alpha)(1-\alpha\beta)}{\alpha}$. The blue solid line corresponds to a persistent shock ($\rho = 0.6$) and the red dashed line to a transitory shock ($\rho = 0$). In each case, we initialize the economy at the risky steady state and consider a shock at time 0.

affected by the risk of lower bound spells. In particular, in the risky steady state we have

$$\nu_t \to \nu = \frac{1}{1-\delta_1}\delta_0 = -\frac{\delta_1}{1-\delta_1}\frac{1+\alpha(1-\beta)}{\alpha(1-\beta\delta_1)}v_0 < 0.$$

A surprising result is that target flexibility is *negative* in the risky steady state: the central bank is *rewarded* for higher inflation. Intuitively, this application abstracts from steady state distortions, so the central bank lacks the usual inflationary bias for raising output that results from a distorted steady state. Without the effective lower bound we would therefore have $\nu_t = 0$, giving from to the divine coincidence.[30] With an effective lower bound, the central bank could still implement the allocation $\pi_t = y_t = 0$, but doing so requires negative nominal interest rates in response to large demand shocks. This generates first-order welfare losses from negative interest rates, while welfare losses from allowing inflation and output gaps around this allocation would be of second order. Hence, optimal policy pushes away from the divine coincidence allocation by raising inflation expectations, which is implemented by increased target flexibility in the risky steady state, $\nu < 0$. Intuitively in this case, the time consistency problem is a tendency towards *too little* inflation as a result of the effective lower bound.

To characterize the inflation target dynamics in response to a fall in r^*, we initialize the

[30] That is to say, demand shocks would simply alter the welfare-irrelevant nominal interest rate.

economy at its risky steady state and consider a shock $\theta_0 < 0$. Figure 2 plots the associated impulse responses. A decline in r^* prompts an *increase* in target flexibility (i.e., a fall in ν_t). Intuitively, a lower natural rate pushes the economy towards the ELB, raising the risk of lower bound spells, and thus making it valuable to raise inflation expectations. Under a dynamic inflation target mechanism, the central bank achieves this by raising its target's flexibility. And the resulting rise in inflation expectations is again reflected in an increase in the inflation target level, τ_t. A decline in r^* therefore requires an adjustment in both target parameters. When the fall in the natural rate is persistent, $\rho > 0$, the dynamic response of target flexibility, ν_t, is again hump-shaped, following the same economic logic as we discussed in Section 4.1.

Target flexibility with ELB and declining r^*. Two surprising insights emerge from this application. First, in the presence of an occasionally-binding lower bound constraint, the standard time consistency problem is offset by the benefit of higher inflation expectations. In this application, the optimal long-run inflation target features $\nu < 0$, thus rewarding the central bank for higher inflation. This result is stark because we abstract from the steady state distortion that motivated the penalty $\nu > 0$ in Section 4.1. More generally, two opposing forces shape the optimal level of target flexibility: The standard time consistency problem leads to inflationary bias if unaddressed and requires an inflation penalty. The risk of lower bound spells, however, makes higher inflation socially valuable, pushing optimal target flexibility in the opposite direction. In fact, the closer the economy is to the ELB, the stronger this force becomes, as highlighted by the impulse response in Panel (a) of Figure 2. The long-run level of the natural rate is therefore a critical determinant of optimal inflation target flexibility.

Second, both target parameters optimally respond to a decline in r^*. Academic and policy discourse has largely focused on implications for the optimal inflation target *level*. Our theory suggests that appropriate adjustments in target *flexibility* are just as important in order for the inflation target mechanism to implement the efficient Ramsey allocation. While our results echo previous arguments for an increase in the target level, they emphasize that a persistently lower natural rate of interest also calls for increased target flexibility.

4.3 Flattening Phillips Curve

The flattening of the U.S. Phillips curve has garnered much attention since the Great Recession (Blanchard, 2016; Galí and Gambetti, 2019; Rubbo, 2020; Del Negro et al., 2020). This debate has also engulfed monetary policy discourse in recent years (Brainard, 2015). In this application, we analyze the implications of a persistent flattening of the Phillips curve for the central bank's dynamic inflation target. We associate θ_t with a persistent shock to the social benefit of stimulating output, which corresponds to time variation in the slope of the NKPC.

Social welfare is characterized by a New Keynesian loss function around a distorted steady state, $\mathcal{U}_t(\pi_t, y_t, \theta_t) = -\frac{1}{2}\pi_t^2 - \frac{1}{2}\alpha y_t^2 + \theta_t y_t$. For tractability, we set $\alpha = 0$. In this case, internalizing the NKPC yields reduced-form utility

$$U(\pi_t, \mathbb{E}_t\pi_{t+1}, \theta_t) = -\frac{1}{2}\pi_t^2 + \frac{1}{\kappa/\theta_t}\left(\pi_t - \beta\mathbb{E}_t\pi_{t+1}\right). \tag{13}$$

Note that an increase in θ_t corresponds to an effective reduction in the slope $\kappa_t = \kappa/\theta_t$ of the Phillips curve. We assume that $\mathbb{E}_t\theta_{t+1} = 1 - \rho + \rho\theta_t$ with $0 \leq \rho \leq 1$, reflecting reversion of the slope of the Phillips curve towards κ over time. We now characterize the optimal inflation target response to a persistent change in the slope of the Phillips curve.

Proposition 6. *The dynamic inflation target that implements the full-information Ramsey allocation is*

$$\nu_t = \frac{1}{\kappa/\theta_t}$$

$$\tau_t = (1-\rho)\left(\frac{1}{\kappa} - \frac{1}{\kappa/\theta_t}\right)$$

Optimal inflation is $\pi_t = \frac{1}{\kappa/\theta_t} - \frac{1}{\kappa/\theta_{t-1}}$.

We again start our discussion of the economic forces by studying the risky steady state that obtains after a sufficiently long realization of $\theta_t = 1$. In this limit, the slope of the Phillips curve is constant and given by κ. It is therefore easy to verify that $\nu_t \to \nu = \frac{1}{\kappa}$ and $\tau_t \to \tau = 0$. Intuitively, there is a constant inflationary bias that arises from the assumed steady state distortion. The target flexibility $\nu = \frac{1}{\kappa}$ offsets the central bank's incentive to use inflation to stimulate output precisely because the slope is constant: the marginal value of stimulating output today is equal to the marginal cost of reducing output yesterday, resulting in zero expected inflation being optimal. This justifies a long-run target level of $\tau_t = 0$ and steady state inflation of $\pi_t = 0$.

We now initialize the economy in the risky steady state and consider a shock $\theta_0 > 0$ that flattens the Phillips curve. The target adjustment dynamics are displayed in Panels (a) and (b) of Figure 3. A flattening of the Philips curve *reduces* optimal target flexibility, implying larger punishments for high inflation. Intuitively, a flatter Phillips curve implies a larger marginal benefit from stimulating output and a larger marginal cost of expected future inflation. This exacerbates the central bank's time consistency problem. The optimal target adjustment is therefore to reduce flexibility. As the central bank faces larger penalties for positive inflation, firms' inflation expectations fall. The optimal inflation target adjustment therefore features both a fall in the level and a reduction of flexibility in response to a flattening of the Phillips curve. As in previous applications, persistent shocks lead to more persistent dynamics.

One interesting observation from Proposition 6 is that the on-impact response of the target level is larger when the shock is transitory. A transitory flattening of the Phillips curve implies that stimulating output is valuable today relative to the future. The Ramsey allocation exchanges

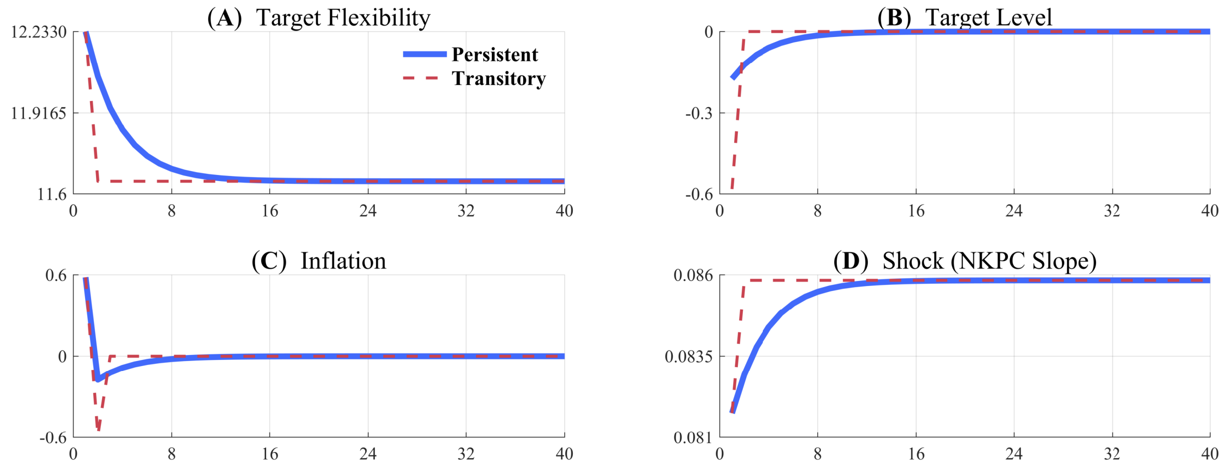

Figure 3. Impulse Responses: Flattening Phillips Curve

Note. Figure 3 plots the impulse responses of inflation and the dynamic inflation target after a flattening of the Phillips curve. Panels (A) through (D) show target flexibility, target level, inflation, and the shock, respectively. We target a quarterly calibration, staying as close as possible to Galí (2015), setting $\beta = 0.99$, $\alpha = 0.75$, and $\kappa = (1-\alpha)(1-\alpha\beta)/\alpha$. The blue solid line corresponds to a persistent shock ($\rho = 0.6$) and the red dashed line to a transitory shock ($\rho = 0$). In each case, we initialize the economy at the risky steady state and consider a shock at time 0.

an output boost today for a future output contraction. Inflation expectations and thus the target level fall sharply on impact. By contrast, a persistent flattening implies that stimulating output is valuable over a longer horizon. This tempers incentives to stimulate current output and dampens the impact on inflation expectations. The on-impact response of target flexibility, on the other hand, is unaffected by shock persistence because the time consistency problem is governed by the contemporaneous Phillips curve slope. In the limit $\rho \to 1$, the central bank adopts a permanently less flexible target while keeping the target level at $\tau_t = 0$.

Flattening Phillips curve vs. declining r^*. The U.S. has recently experienced two empirically prominent structural changes—the decline in r^* and the flattening Phillips curve. Much policy discourse seems to have stressed the benefits of raising the target level and allowing for more flexibility. In fact, the Federal Reserve has adopted an average inflation target in August 2020, which arguably reflects an increase in target flexibility. We have shown in Section 4.2 that these target adjustments are indeed optimal in response to a decline in r^*. Surprisingly, however, a flattening of the Phillips curve pushes in the opposite direction in both dimensions: The optimal policy response is to *lower* the target level and to *remove* the central bank's flexibility around the target because of an exacerbated time consistency problem. These results have important policy implications if the flattening of the Phillips curve proves persistent.

5 Duration and Persistence of Time Inconsistency

A dynamic inflation target implements the Ramsey allocation in an economy with persistent shocks and persistent private information. Our mechanism delegates to the central bank the authority to adjust its own target, as long as it does so *one period in advance*. To consider the implications of our result for policy design in practice, a natural question emerges: How long is a period? We now generalize our theory in the necessary dimensions to tackle this question and characterize the determinants of the optimal target adjustment horizon.

We introduce a longer-horizon time consistency problem in Section 5.1 and define its *duration* and *persistence*. Our main result in Section 5.2 shows that a generalized dynamic inflation target still implements the efficient Ramsey allocation in the presence of persistent private information. A *commitment curve* now summarizes the size of commitments the central bank makes at various horizons. We develop our main policy application in Section 5.3, where we discuss how a fixed-horizon review process as practiced by the Bank of Canada can yield a good approximation of a dynamic inflation target in practice.

5.1 Long-Horizon Time Consistency Problems

The Phillips curve of Section 2 features one-period-ahead inflation expectations. It gives rise to a time consistency problem that has a *duration* of one period: Under discretion, the central bank fails to internalize that policy decisions at time t affect inflation expectations formed at time $t-1$. To study longer-horizon time consistency problems, we introduce a generalized Phillips curve, under which output depends on $K \geq 1$ periods of inflation expectations,

$$y_t = F_t\Big(\pi_t,\, \mathbb{E}_t[\pi_{t+1} \,|\, \tilde{\theta}_t],\, \ldots,\, \mathbb{E}_t[\pi_{t+K} \,|\, \tilde{\theta}_t]\Big), \tag{14}$$

where $\mathbb{E}_t[\pi_{t+k} \,|\, \tilde{\theta}_t]$ denotes firms' k-period-ahead inflation expectation. Implementability conditions like (14) emerge naturally in many settings.[31] We leave the model of Section 2 otherwise unchanged. Substituting into social preferences $\mathcal{U}_t(\pi_t, y_t, \theta_t)$ yields the lifetime social welfare of the government

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t U_t\Big(\pi_t,\, \mathbb{E}_t[\pi_{t+1} \,|\, \tilde{\theta}_t],\, \ldots\,,\, \mathbb{E}_t[\pi_{t+K} \,|\, \tilde{\theta}_t],\, \theta_t\Big). \tag{15}$$

The case $K = 1$ corresponds to the baseline model.

[31] For example, the non-linear pricing equation that emerges in time-dependent rational expectations models of nominal rigidites features an infinite sequence of expectation terms (Calvo, 1983; Galí, 2015). It is only when linearizing around a 0-inflation steady state that the standard NKPC (10) with a single expectation term emerges. In Section 5.3, we study a generalized NKPC by linearizing the standard Calvo model around a steady state with positive inflation, which is an important and policy-relevant benchmark. Many other prominent models of nominal rigidities yield pricing equations of the form (14). Starting with Fischer (1977) and Taylor (1980), multi-period staggered wage and price contracts have become a popular model of nominal rigidities. An influential paper in this tradition is Chari et al. (2000). More recently, Werning (2022) studies the pass-through of inflation expectations and considers Phillips curves with generalized beliefs that also take a form similar to (14).

We again consider the full-information Ramsey allocation as an efficiency benchmark, which we then seek to implement through an incentive compatible mechanism.

Proposition 7. *The full-information Ramsey allocation is characterized by*

$$\frac{\partial U_t}{\partial \pi_t} = \sum_{k=0}^{K-1} \nu_{t-1,k}\,, \qquad \text{where } \nu_{t-1,k} = \begin{cases} -\frac{1}{\beta^{1+k}} \frac{\partial U_{t-1-k}}{\partial \mathbb{E}_{t-1-k}[\pi_t \mid \theta_{t-1-k}]} & \textit{if } t-1-k \geq 0 \\ 0 & \textit{if } t-1-k < 0 \end{cases} \tag{16}$$

Proposition 7 generalizes Proposition 1. Inflation at date t affects flow utility not only in period $t-1$ but also in $t-2$ through $t-K$. Facing the implementability condition (14), a Ramsey planner under commitment therefore finds it valuable to make promises about inflation K periods into the future. Such promises affect output y_t contemporaneously through firms' inflation expectations. Such commitments are valuable because they improve the contemporaneous inflation-output tradeoff: When there is a tradeoff between output and inflation stabilization in period t (i.e., Divine Coincidence does not hold), then being able to backload inflation adjustments—for a given desired output gap—into periods $t+1$ through $t+K$ *smooths* the cost of inflation across these periods.[32] This intuition is directly reflected in the optimality condition (16) that characterizes the full-information Ramsey allocation.

These promises that the Ramsey planner finds it valuable to make are time inconsistent, in the sense that a planner reoptimizing in period $t+s$ would have an incentive to reneg on them. It is in this sense that implementability condition (14) leads to a long-horizon time consistency problem. The *duration of time inconsistency* is then naturally defined as the largest integer S such that expectations about period $t+S$ appear in forward-looking implementability conditions. The duration of time inconsistency implied by (14) is K. And as the duration K increases, the first-order condition (16) becomes increasingly backward looking. Intuitively, the duration tells us for how many periods into the future the planner finds it valuable to make promises in order to improve the contemporaneous inflation-output tradeoff.

Proposition 7 defines $\nu_{t-1,k}$ analogously to ν_{t-1} in the baseline model as a date $t-1-k$ adapted constant. In this environment, we can think of

$$\bar{\nu}_{t-1} \equiv \sum_{k=0}^{K-1} \nu_{t-1,k}$$

[32] This intuition is true even in the standard model, where the NKPC features a single expectation term. Optimal policy under commitment in this benchmark is history-dependent: The planner makes promises for all dates into the future. But this is not because promises arbitrarily far into the future improve the contemporaneous inflation-output tradeoff the planner faces in period t. Instead, the planner smooths the cost of inflation adjustments between periods t and $t+1$ initially, which is possible due to firm expectations, but then finds it valuable to again smooth the promised inflation adjustment between periods $t+1$ and $t+2$, and so forth. Under the generalized Phillips curve (14), promises K periods into the future *directly* improve the contemporaneous inflation-output tradeoff.

as the total time consistency problem—or *total inflationary bias*—that needs to be corrected at date t in order to implement the full-information Ramsey allocation. Total inflationary bias represents the sum of K time inconsistent promises $\nu_{t-1,k}$. Not all of these promises are created equal, however. The planner will find it valuable to make stronger promises about future inflation in some periods and weaker promises for other periods. We introduce the notion of *persistence of time inconsistency* to capture this: Formally, the persistence of time inconsistency in period t is the $K \times 1$ vector whose kth element $\nu_{t-1,k-1}$ captures the strength of promises made k periods ago for date t.

5.2 Dynamic Inflation Targets with Long-Horizon Time Inconsistency

We now develop the main result of this section: a *K-horizon dynamic inflation target* implements the full-information Ramsey allocation.

Definition 8 (K-horizon Dynamic Inflation Target)**.** A *K-horizon dynamic inflation target* is an affine transfer rule mechanism, $T_t = -b_{t-1}(\pi_t - \tau_{t-1})$, whose *target level* equals a weighted average of the past K inflation forecasts,

$$\tau_{t-1} = \sum_{k=0}^{K-1} \omega_{t-1,k}\, \mathbb{E}_{t-1-k}[\pi_t \mid \tilde{\theta}_{t-1-k}],$$

for some weights $\omega_{t-1,k}$, and whose *target flexiblity* is the slope b_{t-1}.

The K-horizon dynamic inflation target reverts to the dynamic inflation target of Section 3 when $K = 1$. When $K > 1$, however, the target level τ_{t-1} is based on the last K forecasts for date t inflation, which is reminiscent of results on inflation forecast targeting (Svensson, 1997a; Svensson and Woodford, 2004). We are now ready to prove the following generalization of our main result.

Proposition 9. *A K-horizon dynamic inflation target implements the full-information Ramsey allocation in a locally incentive compatible mechanism. The weights for the target level are* $\omega_{t-1,k} = \frac{\nu_{t-1,k}}{\bar{\nu}_{t-1}}$*, and the target flexibility is* $b_{t-1} = \bar{\nu}_{t-1}$.

Proposition 9 generalizes our main result to longer-horizon time consistency problems. Intuitively, the mechanism still serves the same two roles emphasized in Section 3: correcting the time consistency problem in the central bank's contemporaneous inflation choice, and providing incentives for correctly updating the target. The target flexibility, b_{t-1}, is again set to address the former. When $K = 1$, inflationary bias is simply captured by $\bar{\nu}_{t-1} = \nu_{t-1,0}$ as in Section 3, i.e., the impact of current inflation on last period's output. When $K > 1$, the total time consistency problem $\bar{\nu}_{t-1}$ summarizes the cumulative impact of current inflation on output over the last K periods.

The target's level, τ_{t-1}, is again used to overcome the model's core informational frictions: the central bank's incentives to manipulate firm and government beliefs. The mechanism sets the target level equal to a weighted average of inflation forecasts for date t made over the last K periods. Intuitively, the weight $\omega_{t-1,k+1} = \frac{\nu_{t-1,k+1}}{\bar{\nu}_{t-1}}$ assigned to the inflation forecast k periods ago is the fraction of the total time consistency problem, $\bar{\nu}_{t-1}$, that originates from the impact of inflation on output k periods ago, $\nu_{t-1,k+1}$. Large weights are assigned to past dates with large time consistency problems. The K-horizon dynamic inflation target again achieves *informational divine coincidence*: The central bank's incentive to bias firm inflation expectations *downward* is exactly offset by its incentive to bias government transfers *upward* at the equivalent horizon. It holds precisely because the slope $\nu_{t+k,k}$ for future transfers is equal to the marginal utility cost of future inflation through changes in current output.

Evolution of commitments and commitment curve. As in the baseline model, the central bank only needs to know the target parameters $(\bar{\nu}_{t-1}, \tau_{t-1})$, and not the history that gave rise to those parameters, to set optimal inflation policy for date t. To update the target for the next period, however, the central bank must respect cumulative past commitments made for period $t+1$. We now show how to summarize this information in two $K \times 1$ vectors, $\boldsymbol{V}_t = \{\boldsymbol{V}_{t,1}, \ldots, \boldsymbol{V}_{t,K}\}$ and $\boldsymbol{T}_t = \{\boldsymbol{T}_{t,1}, \ldots, \boldsymbol{T}_{t,K}\}$, that serve as sufficient statistics for the target adjustment process.

We define $\boldsymbol{V}_{t,k}$ as cumulative promises made for target flexbility $k-1$ periods ahead, i.e., in period $t-1+k$. Under this convention, $\boldsymbol{V}_{t,1} = \bar{\nu}_{t-1}$ corresponds to target flexibility at date t and summarizes all commitments made over the past K periods. By contrast, $\boldsymbol{V}_{t,k}$ for $k > 1$ reflects the cumulative *partial commitments* the central bank has made so far. We refer to these as partial commitments precisely because they can still be updated at date t. We can track the evolution of partial commitments using the recursion

$$\boldsymbol{V}_{t+1,k} = \boldsymbol{V}_{t,k+1} + \nu_{t+k-1,k}$$

where $\boldsymbol{V}_{t,K+1} \equiv 0$ and $\nu_{t+k-1,k}$ reflects the new promise made at date t for target flexibility k periods ahead. To illustrate, note that $\boldsymbol{V}_{t+1,1} = \boldsymbol{V}_{t,2} + \nu_{t,1} = \bar{\nu}_t$: target flexibility for period $t+1$ results from adding a new partial commitment made in period t, $\nu_{t,1}$, to our measure of cumulative promises made in the past, $\boldsymbol{V}_{t,2}$. Vector $\boldsymbol{V}_t$ thus summarizes all relevant information for updating target flexiblity.

To update the target level τ_t, the central bank must compute a weighted average of historical inflation forecasts. The evolution of this weighted average of forecasts satisfies the recursion

$$\begin{aligned}\tau_t &= \frac{\nu_{t,0}}{\bar{\nu}_t}\mathbb{E}_t[\pi_{t+1}|\tilde{\theta}_t] + \sum_{k=1}^{K}\frac{\nu_{t,k}}{\bar{\nu}_t}\mathbb{E}_{t-k}[\pi_{t+1}|\tilde{\theta}_{t-k}] \\ &= \frac{\nu_{t,0}}{\nu_{t,0}+\boldsymbol{V}_{t,2}}\mathbb{E}_t[\pi_{t+1}|\tilde{\theta}_t] + \frac{\boldsymbol{V}_{t,2}}{\nu_{t,0}+\boldsymbol{V}_{t,2}}\underbrace{\sum_{k=1}^{K}\frac{\nu_{t,k}}{\boldsymbol{V}_{t,2}}\mathbb{E}_{t-k}[\pi_{t+1}|\tilde{\theta}_{t-k}]}_{\equiv \boldsymbol{T}_{t,2}},\end{aligned}$$

where the first line expresses τ_t as an average of current and historical inflation forecasts with weights directly taken from Proposition 9. We introduce $\boldsymbol{T}_t$ to track the evolution of average forecasts and summarize the information needed by the central bank to update its target level. Its first element reflects the current target level, $\boldsymbol{T}_{t,1} = \tau_{t-1}$, which is taken as given at date t. For $k > 1$, $\boldsymbol{T}_{t,k}$ summarizes the cumulative weighted average of historical forecasts for inflation in period $t-1+k$. Its evolution satisfies the recursion

$$\boldsymbol{T}_{t+1,k} = \frac{\boldsymbol{V}_{t,k+1}}{\boldsymbol{V}_{t,k+1}+\nu_{t+k-1,k}}\boldsymbol{T}_{t,k+1} + \frac{\nu_{t+k-1,k}}{\boldsymbol{V}_{t,k+1}+\nu_{t+k-1,k}}\mathbb{E}_t[\pi_{t+k}|\tilde{\theta}_t].$$

To implement the K-horizon dynamic inflation target, the central bank must therefore keep track of $(\boldsymbol{V}_t, \boldsymbol{T}_t)$. Intuitively, these two vectors encode a notion of forward guidance in the form of partial commitments for what the central bank will do for the next K periods. At date t, the central bank takes as given its target for the current date, $\tau_{t-1} = \boldsymbol{T}_{t,1}$ and $b_{t-1} = \boldsymbol{V}_{t,1}$, and lacks any ability to update this target. The central bank has partial ability to update its target for periods $t+k$, for $1 \leq k < K$, taking as given its prior commitments that are encoded in $\boldsymbol{V}_{t,k-1}$ and $\boldsymbol{T}_{t,k-1}$. Finally, the central bank has no prior commitment over inflation at date $t+K$, and so makes its first partial commitment for this period at date t. This provides a generalized notion of the iterated one-period commitments of the baseline model: The central bank here makes iterated K-period *partial* commitments.

A useful representation of this partial commitment process is the *commitment curve*. It encodes the size of the partial commitment made by the central bank at date t for date $t+k$, which is precisely $\nu_{t+k,k}$.

Definition 10 (Commitment Curve)**.** The *commitment curve* at date t is the curve $(k, \nu_{t+k,k})$ of commitments made at date t for all $k \geq 1$.

The commitment curve provides a natural representation of the persistence of time inconsistency and commitments under the K-horizon dynamic inflation target. Intuitively, its shape conveys how long the horizon of commitments made by the central bank truly is: A sharply downward-sloping curve means the central bank is only making large commitments for the near term, while a flat curve means the central bank is making large commitments over a long horizon. The commitment curve provides an instructive conceptual framework for characterizing the optimal horizon of target adjustments and answering the policy-relevant question of *how long is a period*.

5.3 Practical Policy Implications

This section develops our main policy application, leveraging the commitment curve introduced above to characterize the determinants of the optimal target adjustment horizon. We do so in the context of a generalized New Keynesian Phillips curve (GNKPC) that emerges when linearizing the standard Calvo model around positive steady state or trend inflation, denoted $\gamma = 1 + \bar{\pi}$ (Ascari, 2004; Ascari and Sbordone, 2014).[33]

Following closely Ascari and Ropele (2007), we study a linearized New Keynesian model that comprises a standard dynamic IS equation of the form (11), with EIS $\sigma = 1$, and a GNKPC, together given by

$$y_t = \mathbb{E}_t y_{t+1} - (i_t - \mathbb{E}_t \pi_{t+1}) \tag{17}$$

$$\pi_t = \kappa y_t + (\beta\gamma + \tilde{\beta})\mathbb{E}_t \pi_{t+1} + \tilde{\beta}\mathbb{E}_t \Big[\sum_{s=1}^{\infty} \tilde{\delta}^s \pi_{t+1+s} \Big]. \tag{18}$$

where $\tilde{\beta} = (\gamma - 1)\beta(1 - \alpha\gamma^{\epsilon-1})(\epsilon - 1)$ and $\tilde{\delta} = \alpha\beta\gamma^{\epsilon-1}$. The slope of the GNKPC is $\kappa = \frac{(1-\alpha\gamma^{\epsilon-1})(1-\alpha\beta\gamma^{\epsilon})}{\alpha\gamma^{\epsilon-1}}$. We denote by $(1 - \alpha)$ the probability that a firm can reset its price each period and by ϵ the elasticity of substitution between intermediate inputs. Note that in the case with no trend inflation, $\gamma = 1$ and $\tilde{\beta} = 0$, we recover the standard New Keynesian Phillips curve (10).[34] Unlike in previous sections, we now denote by π_t and y_t the percent deviations from a deterministic steady state with trend inflation γ.

Given a preference function $\mathcal{U}(\pi_t, y_t, \theta_t)$, we can sharply characterize the shape of the commitment curve associated with the GNKPC (18).

---

[33] The linearized model with trend inflation is an important and policy-relevant benchmark. Ascari and Sbordone (2014) argue that "*the conduct of monetary policy should be analyzed by appropriately accounting for the positive trend inflation targeted by policymakers.*" However, many other models also yield generalized Phillips curves of the form (14)—see also Footnote 31.

[34] Ascari and Ropele (2007) represent the GNKPC in terms of an auxiliary variable

$$\pi_t = \kappa y_t + \beta\gamma\mathbb{E}_t \pi_{t+1} + (\gamma - 1)\beta(1 - \alpha\gamma^{\epsilon-1})\mathbb{E}_t \Big[(\epsilon - 1)\pi_{t+1} + \phi_{t+1} \Big]$$

$$\phi_t = \alpha\beta\gamma^{\epsilon-1}\mathbb{E}_t \Big[(\epsilon - 1)\pi_{t+1} + \phi_{t+1} \Big]$$

where we have already set the EIS to $\sigma = 1$. Defining $\tilde{\beta}$ and $\tilde{\delta}$ as above, then dividing through the second equation by $\epsilon - 1$, defining $\varphi_t = \frac{1}{\epsilon-1}$, and solving forward we have $\varphi_t = \sum_{s=1}^{\infty} \tilde{\delta}^s \mathbb{E}_t \pi_{t+s}$. Substituting into the first equation and reallocating terms, we get $\pi_t = \kappa y_t + (\beta\gamma + \tilde{\beta})\mathbb{E}_t \pi_{t+1} + \tilde{\beta}\mathbb{E}_t \sum_{s=1}^{\infty} \tilde{\delta}^s \pi_{t+1+s}$.

Proposition 11. *The commitment curve has a quasi-hyperbolic shape, that is*

$$\nu_{t+k,k} = \beta^* \delta^{*(k-1)} \nu_{t+1,1}$$

where $\beta^* = \frac{\tilde{\beta}}{\tilde{\beta}+\beta\gamma}$ *and* $\delta^* = \frac{\tilde{\delta}}{\beta}$.

Proposition 11 reveals that the commitment curve in the GNKPC model has a shape associated with quasi-hyperbolic discounting (Laibson, 1997). This implies a large and discrete drop, $\beta^*\delta^*$, in the commitment curve between $k = 1$ and $k = 2$. For $k \geq 2$, the curve is governed by exponential discounting, with $\nu_{t+k+1,k+1} = \delta^* \nu_{t+k,k}$.

The quasi-hyperbolic shape emerges because long-horizon inflation expectations have a lower pass-through to current inflation. According to the GNKPC (18), the relative effect of inflation expectations on current output at different horizons is given by

$$\frac{\partial y_t / \partial \mathbb{E}_t \pi_{t+k}}{\partial y_t / \partial \mathbb{E}_t \pi_{t+1}} = \frac{-\frac{1}{\kappa}\tilde{\beta}\tilde{\delta}^{k-1}}{-\frac{1}{\kappa}(\beta\gamma + \tilde{\beta})} = \beta^* \tilde{\delta}^{k-1} < 1.$$

The pass-through of long-horizon relative to one-period-ahead inflation expectations is muted. The smaller β^* and $\tilde{\delta}$, the more quickly the effects of long-horizon inflation expectations decay.

The long-horizon time consistency problem that emerges from the Phillips curve (18) is governed precisely by the sensitivity of current output to long-horizon inflation expectations. Likewise, the commitment curve of Proposition 11 is therefore shaped by the parameters β^* and $\tilde{\delta}$ that also govern the relative pass-through of k-horizon inflation expectations. Intuitively, the quasi-hyperbolic shape implies that promises made for date t in the previous period $t-1$ tend to be larger by a factor β^* than partial commitments made in earlier periods. Here, β^* reflects the disproportionate impact that one-period-ahead inflation expectations have on output, and it therefore governs the relative importance of short- and long-horizon commitments. δ^*, on the other hand, determines how quickly the importance of partial commitments decays at longer horizons.

Notice that the standard NKPC (10) features $\tilde{\beta} = 1$ and $\beta^* = 0$. The standard New Keynesian model therefore corresponds to an extreme case of quasi-hyperbolic discounting, where only the first point on the commitment curve is nonzero.

Bank of Canada mechanism. The duration of the time consistency problem implied by the Phillips curve (18) is $K = +\infty$. Only an infinite-horizon dynamic inflation target could therefore implement the Ramsey allocation, requiring an infinite sequence of forward-looking partial commitments that are updated every period. The central bank would have to continuously update infinite-horizon target commitments, which is impractical.

Proposition 11 underscores, however, that not all commitments are created equal. The quasi-hyperbolic shape of the commitment curve has two implications. First, long-horizon commitments

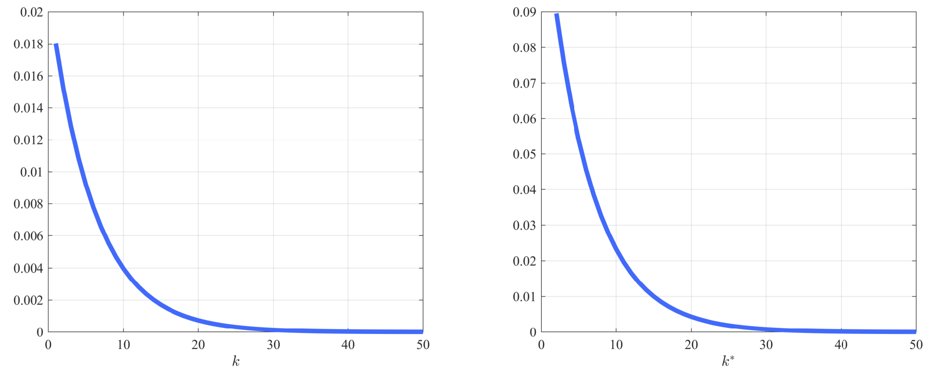

Figure 4. Commitment Curve

Note. Figure 4 plots the commitment curve $(k, \nu_{t+k,k})$ for $k > 1$ relative to $k = 1$ in Panel (a) and the measure $\boldsymbol{V}_{t-1,k^*}$ in Panel (b). We target a quarterly calibration and stay close to Ascari and Ropele (2007), setting $\beta = 0.99$, $\alpha = 0.75$, $\epsilon = 11$ and $\gamma = 1.01$.

become increasingly less relevant because the severity of the time consistency problem at longer horizons decays exponentially. Second, commitments for the very near term are disproportionately important because of the quasi-hyperbolic discount β^*. Together, these observations suggest that approximating the optimal infinite-horizon mechanism with an appropriate finite-horizon one may not generate large welfare losses.

The Bank of Canada follows a regular 5-year review process of its inflation target. In our framework, this implies that the target level and flexibility (τ, ν) are optimally updated every 5 years, but held constant between reviews. Our theory provides a natural framework to ask what the optimal horizon of such a review process is. Under the Bank of Canada mechanism, the choice of adjustment horizon K involves a tradeoff between short- and long-horizon commitments. Intuitively, increasing K is valuable to address the long-horizon time consistency problem implied by (18). Since the target is held fixed between reviews, however, a larger K is costly because it implies less flexibility to respond to persistent shocks in the short run. Proposition 11 allows us to compare the relative importance of short- and long-horizon commitments. In the limiting case as $\beta^* \to 0$, no long-horizon commitments are made and we revert to the standard New Keynesian Phillips curve. In this limit, the Bank of Canada mechanism with a review horizon of $K = 1$ period in fact implements the Ramsey allocation.

By contrast, as β^* and δ^* become larger, the commitment curve flattens. To characterize the relative importance of long-horizon commitments, we can ask what share of the current total commitment, $\bar{\nu}_{t-1}$, is accounted for by partial commitments made more than k^* periods ago. This share is precisely captured by $\boldsymbol{V}_{t-1,k^*}$, which we can compute in closed form around the economy's risky steady state, where $\nu_{t+1,1} = \nu$. We have

$$\boldsymbol{V}_{t-1,k^*} = \frac{\beta^* \sum_{k=k^*}^{\infty} (\delta^*)^{k-1}}{1 + \beta^* \sum_{k=2}^{\infty} (\delta^*)^{k-1}} = \frac{\beta^* \frac{\delta^*}{1-\delta^*}}{1 + \beta^* \frac{\delta^*}{1-\delta^*}} \delta^{*(k^*-2)}.$$

Figure 4 plots both the commitment curve and $\boldsymbol{V}_{t-1,k^*}$ for a standard quarterly calibration in Panels (a) and (b), respectively. Panel (b) suggests that 90% of the total is accounted for by one-quarter-ahead commitments, with 10% of the total coming from long-horizon commitments. The vast majority of commitments are made over a 10-quarter horizon, and virtually no commitment comes from more than 30 quarters—roughly 7 years—into the future. This reveals two important takeaways. First, finite-horizon dynamic inflation targets approximate the optimal mechanism in this environment. They likely incur only small welfare losses even though the duration of time inconsistency is $K = +\infty$. Second, commitments over the short term are particularly important and account for a large share of the total. The Bank of Canada's 5-year review horizon is therefore long enough to capture nearly all long-horizon commitments. Conversely, Figure 4 suggests that at a 5-year horizon, the marginal impact of long-horizon commitments is likely small. However, the short-run commitments that would be made if the central bank could adjust its target continuously might be large. Our results therefore suggest that a shorter adjustment horizon may strike a better balance between flexibility to respond to persistent shocks in the short run and commitment to address long-horizon time inconsistency.

6 Extensions

In Section 3, we showed that a dynamic inflation target can implement the full-information Ramsey allocation. This constitutes an optimal mechanism under three conditions: (i) the informational divine coincidence holds, that is firms and the government have the same information sets; (ii) enforcement is costless to the government; (iii) the government and central bank have the same preferences. We study (i) and (ii) in this Section, and (iii) in Appendix C.2.

6.1 The Importance of Information

Our first extension relaxes the assumption that firms and the government have the same information sets. We assume that a fraction of firms are *informed* and directly observe the state θ_t. We show that the optimal mechanism is a dynamic inflation target with a *penalized* adjustment process. Intuitively, penalized adjustments are required to compensate the central bank for information rents earned from informed firms. This extension demonstrates the robustness of the dynamic inflation target framework to different information structures.

Let a fraction $\gamma \in [0, 1]$ of firms directly observe the state θ_t. The remaining firms are uninformed and learn the state from central bank reports. Average inflation expectations are thus given by $\mathbb{E}_t^{\text{avg}}\pi_{t+1} = \gamma\mathbb{E}_t[\pi_{t+1} \mid \theta_t] + (1-\gamma)\mathbb{E}_t[\pi_{t+1} \mid \tilde{\theta}_t]$. We now write reduced-form preferences over average inflation expectations as $U_t(\pi_t,\, \mathbb{E}_t^{\text{avg}}\pi_{t+1},\, \theta_t)$. The full-information Ramsey allocation, including π_t and ν_{t-1}, is as in Proposition 1.

Following the same steps as in the proof of Proposition 3, we obtain the new envelope condition for incentive compatibility,

$$\frac{\partial \mathcal{W}_t(\theta^t)}{\partial \theta_t} = \overbrace{\frac{\partial U_t\left(\pi_t, \mathbb{E}_t\left[\pi_{t+1}|\theta_t\right], \theta_t\right)}{\partial \theta_t} + \beta\mathbb{E}_t\left[\mathcal{W}_{t+1}(\theta^{t+1})\frac{\partial f(\theta_{t+1}|\theta_t)/\partial\theta_t}{f(\theta_{t+1}|\theta_t)}\bigg|\theta_t\right]}^{\text{Previous Terms}} \\ \underbrace{+\,\gamma\frac{\partial U_t\left(\pi_t, \mathbb{E}_t\left[\pi_{t+1}|\theta_t\right], \theta_t\right)}{\partial \mathbb{E}_t\left[\pi_{t+1}|\theta_t\right]}\mathbb{E}_t\left[\pi_{t+1}\frac{\partial f(\theta_{t+1}|\theta_t)/\partial\theta_t}{f(\theta_{t+1}|\theta_t)}\bigg|\theta_t\right]}_{\text{Information Rent from Informed Firms}}. \tag{19}$$

Notice that inflation expectations of informed and uninformed firms coincide under the truth-telling mechanism. The first line of (19) captures the same information rents as before in equation (6). The second line, however, now reflects a new source of central bank information rents, earned from informed firms. This new force represents the key departure from the baseline model. It reflects how information about the state affects informed firms' inflation expectations.

Equation (19) reveals that a simple dynamic inflation target is no longer incentive compatible because it neglects this new information rent. We need to augment the mechanism accordingly. Denote the negative of the new information rent (omitting γ) at the Ramsey allocation by $\omega_t = \beta\nu_t\mathbb{E}_t\left[\pi_{t+1}\frac{\partial f(\theta_{t+1}\,|\,\theta_t)/\partial\theta_t}{f(\theta_{t+1}\,|\,\theta_t)} \mid \theta_t\right]$. We now define a *penalized dynamic inflation target* as an affine transfer rule with an additional penalty P_t for target adjustments at date t, which we will associate with the new information rent ω_t,

$$T_t = -b_{t-1}(\pi_t - \mathbb{E}_t\pi_{t-1}) - \gamma P_t.$$

Finally, we define the lifetime expected penalty as $\overline{P}_t = P_t + \mathbb{E}_t[\sum_{k=1}^{\infty}\beta^k P_{t+k}|\theta_t]$, which admits a recursive representation $\overline{P}_t = P_t + \beta\mathbb{E}_t\overline{P}_{t+1}$.

Proposition 12. *A penalized dynamic inflation target implements the full-information Ramsey allocation in a locally incentive compatible mechanism, with target flexibility $b_{t-1} = \nu_{t-1}$. The lifetime penalty function $\overline{P}$ is given in recursive form by*[35]

$$\overline{P}_t(\theta^t) = \int_{\underline{\theta}}^{\theta_t}\omega_t(\theta^{t-1}, x_t)dx_t + \int_{\underline{\theta}}^{\theta_t}\beta\mathbb{E}_t\left[\overline{P}_{t+1}\frac{\partial f(\theta_{t+1}|x_t)/\partial x_t}{f(\theta_{t+1}|x_t)} \;\Big|\; x_t\right]dx_t.$$

Proposition 12 generalizes our main result to environments with informed firms. It demonstrates that dynamic inflation targets are robust to alternative information structures—in this case requiring an additional penalty P_t.

[35] Note that the static penalty, P_t, can be obtained by combining this equation with the recursive representation above.

The lifetime penalty has a "static" and a "dynamic" component. The marginal static penalty is ω_t, which is the information rent from the central bank's private information about informed firm expectations. The information rent depends on how firms' inflation expectations covary with the shock structure. If high types θ_t signal high future types θ_{t+1} (monotone likelihood) and high future types signal high inflation π_{t+1}, then $\omega_t > 0$, that is there is a penalty for upwards target adjustments.[36] Intuitively, the unpenalized dynamic inflation target gives too much surplus to high θ types, and the penalization process deters lower types from deviating upwards. The marginal dynamic penalty, $\mathbb{E}_t[\overline{P}_{t+1}\Lambda_{t+1}|\theta_t]$, reflects that once a penalized adjustment process is in place, the central bank also possesses persistent private information about the distribution of future penalties.

Proposition 12 yields important insights on the design of central bank inflation targets. The immediate consequence is that information heterogeneity necessitates a penalized target adjustment process. Penalties play the intuitive role of ensuring that a central bank that should implement low inflation is not incentivized towards excessive upward adjustments. A more nuanced perspective is that this suggests a complexity-based argument for central banks to be responsible for collecting and disseminating information about the structural state of the economy to firms. When all firms are uninformed and learn from the central bank, an unpenalized dynamic inflation target implements the Ramsey allocation. By contrast when some or all firms are informed, a dynamic inflation target requires a penalization process to control target adjustments.

6.2 Costly Enforcement

Our second extension allows for costly mechanism enforcement. While the optimal mechanism no longer implements the Ramsey allocation, we show the optimal allocation rule is similar to that under a dynamic inflation target. Moreover, the optimal mechanism reverts to a dynamic inflation target at the extremes of the shock distribution. In Appendix B.1, we revisit the applications of Section 4 under costly enforcement.

To capture the social cost of implementing and enforcing a monetary policy mechanism, we assume social preferences are now

$$\mathbb{E}\left[\sum_{t=0}^{\infty}\beta^t\left(U_t(\pi_t, \mathbb{E}_t[\pi_{t+1}|\tilde{\theta}_t], \theta_t) - \kappa T_t\right)\right], \tag{20}$$

where transfers are socially costly in proportion to $\kappa \geq 0$.[37] In conjunction, we introduce the central bank participation constraint, given by $\mathcal{W}_0 \geq 0$, normalizing the outside option to 0 without loss of generality.[38] Recall that a mechanism is a mapping $(\pi_t, T_t) : \Theta^t \to \mathbb{R}^2$ that must be incentive compatible, as defined in Section 2.3. We again solve for the optimal relaxed mechanism that

[36] This means the information rent is *negative*.

[37] This corresponds to a standard (quasilinear) transferable utility model. As usual, T_t may also correspond to non-quasilinear utilities, provided they are transferable in this form.

[38] At the end of the proof of Proposition 13, we show that a dynamic inflation target is an optimal mechanism when there is instead an average participation constraint, $\mathbb{E}\mathcal{W}_0 \geq 0$.

enforces the envelope characterization of local incentive compatibility (6).

Proposition 13. *The solution to an optimal allocation rule of the relaxed problem is given by the first-order conditions*

$$\frac{\partial U_t}{\partial \pi_t} - K\Gamma_t \frac{\partial^2 U_t}{\partial \theta_t \partial \pi_t} = \lambda_{t-1} \tag{21}$$

where $K = \frac{\kappa}{1+\kappa}$*, where*

$$\lambda_{t-1} = \begin{cases} -\frac{1}{\beta}\frac{\partial U_{t-1}}{\partial \mathbb{E}_{t-1}\pi_t} + K\Gamma_{t-1}\frac{1}{\beta}\frac{\partial^2 U_{t-1}}{\partial \theta_{t-1}\partial \mathbb{E}_{t-1}\pi_t} & \text{for } t \geq 1 \\ 0 & \text{for } t = 0 \end{cases}$$

and where $\Gamma_t(\theta^t)$ *is given recursively by*

$$\Gamma_t(\theta^t) = \Gamma_{t-1}(\theta^{t-1})h^{-1}(\theta_t|\theta_{t-1})\mathbb{E}_{t-1}\left[\Lambda(s_t|\theta_{t-1})\middle| s_t \geq \theta_t\right] \tag{22}$$

where $\Gamma_0(\theta^0) = h^{-1}(\theta_0)$*, where* $h^{-1}(\theta_t|\theta_{t-1}) = \frac{1-F(\theta_t|\theta_{t-1})}{f(\theta_t|\theta_{t-1})}$ *is the inverse hazard rate, and where* $\Lambda(s_t|\theta_{t-1}) = \frac{\partial f(s_t|\theta_{t-1})/\partial\theta_{t-1}}{f(s_t|\theta_{t-1})}$ *is the derivative of the likelihood ratio.*

Proposition 13 characterizes the allocation rule under the optimal mechanism with costly enforcement.[39] If transfers are not costly, $\kappa = K = 0$, the optimal mechanism reverts to a dynamic inflation target that implements the Ramsey allocation. Two new economic forces emerge when transfers are socially costly.

First is the classical information rent earned by the central bank (agent), manifesting in the term $K\Gamma_t \frac{\partial U_t}{\partial \theta_t \pi_t}$ on the LHS of (21). Intuitively, this reflects the surplus that the central bank receives from revealing its persistent private information to the government. This surplus, manifesting as larger transfers for a given allocation, is costly to the government in proportion to the enforcement costs $K > 0$. Thus when transfers are not costly, information rents earned by the central bank have no cost to the government, and this term drops out. This information rent parallels the usual information rent in models with persistent private information (Pavan et al., 2014): it is higher when an increase in inflation yields a larger increase in marginal utility for higher types, that is $\partial^2 U_t/\partial\theta_t\partial\pi_t > 0$, and when the information signaled about the current type from past types, Γ_t, is higher.

Second is an information rent due to time inconsistency, i.e., the forward-looking Phillips curve, reflected by the term $K\Gamma_{t-1}\frac{1}{\beta}\frac{\partial^2 U_{t-1}}{\partial\theta_{t-1}\partial\mathbb{E}_{t-1}\pi_t}$. Much as an increase in contemporaneous inflation can disproportionately affect higher current types, the historical information rent matters to the extent that increases in *past inflation expectations* may disproportionately affect higher past types

[39] It should be noted that the government can still use a dynamic inflation target to implement the Ramsey allocation, but this mechanism is no longer optimal due to enforcement costs.

θ_{t-1}. This intuition is encoded in $\frac{\partial^2 U_{t-1}}{\partial\theta_{t-1}\partial\mathbb{E}_{t-1}\pi_t}$. Suppose that higher expected inflation lowers the information rent by worsening the previous period's inflation-output trade-off. Then this effect in fact calls for a *higher* inflation rate at date t than under allocative efficiency. Intuitively, the higher inflation rate improves planner welfare by lowering the central bank's information rents in the prior period, even though it worsens social surplus.

Proposition 13 highlights the importance of shock persistence to the optimal mechanism with costly enforcement. If shocks were not persistent, then $\Gamma_0 = \frac{1-F(\theta_0)}{f(\theta_0)}$ but $\Gamma_t = 0$ for all $t \geq 1$, since $\Lambda = 0$ (past shocks convey no information about the current shock). This implies that the optimal allocation satisfies $\frac{\partial U_t}{\partial \pi_t} = -\frac{1}{\beta}\frac{\partial U_{t-1}}{\partial\mathbb{E}_{t-1}\pi_t}$ for all $t \geq 2$: the optimal mechanism reverts to a dynamic inflation target along any history for all dates $t \geq 2$. This reflects a variant of the standard intuition that absent persistent shock, the principal extracts all surplus ex ante by promising the agent her optimal allocation after the initial date. Our result differs in two ways due to the time consistency problem from the Phillips curve. First, the (undistorted) optimal allocation is the Ramsey allocation, rather than the discretion allocation, which would be optimal absent time inconsistency. Second, when extracting surplus at date 0 the government internalizes the impact of date 1 inflation on date 0 information rents through the Phillips curve. The reversion to the Ramsey allocation only occurs at date 2 as a result.

Despite costly enforcement, the optimal allocation rule bears important similarities to that under a dynamic inflation target. The marginal impact of inflation on flow utility net of information rents today, $\frac{\partial U_t}{\partial \pi_t} - K\Gamma_t\frac{\partial U_t}{\partial\theta_t\partial\pi_t}$, is equated with the marginal impact of inflation today on flow utility net of information rents the prior period, λ_{t-1}. This historical impact is represented by the single statistic λ_{t-1}. Thus, the history dependence of the mechanism can be encoded in the triple $(\lambda_{t-1}, \Gamma_{t-1}, \theta_{t-1})$. λ_{t-1} encodes the total time consistency problem, while $(\Gamma_{t-1}, \theta_{t-1})$ encodes the persistence of information rents (used to determine Γ_t). This triple is a sufficient statistic at date t for characterizing the allocation and transfer rule to implement the optimum of Proposition 13. In this respect, a key qualitative insight of the dynamic inflation target that carries over is that there is a simple sufficiently statistic, λ_{t-1}, that summarizes the consequences of time inconsistency for the evolution of the optimal mechanism. Unlike in the baseline model, however, this variable encapsulates not only the impact on allocative efficiency, but also the impact on past information rents.

Multiplicative taste shocks. A canonical case in principal-agent frameworks is multiplicative taste shocks, $U_t(\pi_t, \mathbb{E}_t\pi_{t+1}, \theta_t) = \theta_t u_t(\pi_t, \mathbb{E}_t\pi_{t+1})$. In this case, the optimal allocation rule reduces to

$$\vartheta_t \frac{\partial u_t}{\partial \pi_t} = \vartheta_{t-1}\frac{-1}{\beta}\frac{\partial u_{t-1}}{\partial\mathbb{E}_{t-1}\pi_t}, \tag{23}$$

where $\vartheta_t = \theta_t - K\Gamma_t$ is the principal's *virtual value*, and where $\vartheta_0 = \theta_0 - \frac{1-F(\theta_0)}{f(\theta_0)}$ is the canonical virtual value from static frameworks. Absent a time consistency problem, equation (23) reduces to the usual problem of maximizing flow utility u_t when the virtual value is positive, and minimizing promised utility u_t when the virtual value is negative.[40] With the time consistency problem, the planner trades off marginal utility at date t, weighted by virtual value ϑ_t, against marginal utility at date $t-1$, weighted by virtual value ϑ_{t-1}. Thus, the allocation rule is the Ramsey allocation of a planner whose type is the virtual value ϑ. This tells us that the direction of distortion relative to the Ramsey allocation depends on the relative distortion of the virtual value relative to the true type. In particular, the central bank promotes *more* inflation on the margin when $\frac{\vartheta_t}{\theta_t} > \frac{\vartheta_{t-1}}{\theta_{t-1}}$, i.e., when the virtual value of the central bank at date t is higher relative to the true type than at date $t-1$.

Reversion to dynamic inflation target. Proposition 13 implies that the optimal mechanism reverts to a dynamic inflation target at both extremes of the shock distribution. The following corollary formalizes these no-top- *and* no-bottom-distortion results.

Corollary 14. If $\theta_t \in \{\underline{\theta}, \overline{\theta}\}$, then the optimal allocation at dates $t+1+s$ $(s \geq 0)$ can be implemented by a dynamic inflation target.

Corollary 14 parallels no-top- and no-bottom-distortion results that arise absent time consistency problems (Pavan et al., 2014). Since there are no central bank types above $\overline{\theta}$, no types above $\overline{\theta}$ earn information rents from the allocation of type $\overline{\theta}$. There is consequently no reason to distort that allocation. Persistent private information furthermore implies a no-distortion at the bottom result because the lowest type earns no rents from revealing information about the distribution of future types. In our model, the time consistency problem implies that the optimal allocation rule we revert to is the full-information Ramsey allocation. As a result, the optimal mechanism reverts to the dynamic inflation target at the limits of the distribution.

7 Conclusion

We develop a theory of how a central bank should update its inflation target in the presence of persistent economic shocks that are private information of the central bank. We show that a dynamic inflation targeting mechanism can implement the Ramsey allocation. The dynamic inflation target corrects not only the central bank's time consistency problem but also its strategic incentives to reveal information to firms about the persistent underlying state. The target's level and flexibility are both adjusted over time, and adjustments must be made *one period in advance*. We ntroduce the commitment curve, which summarizes the size of commitments the central bank makes for the future and helps inform the persistence of commitment to the current target. Our results suggest that a mechanism of adjustment at restricted points in time—for example every five years as practiced by the Bank of Canada—could be a desirable adjustment method.

[40] In canonical buyer-seller frameworks, this corresponds to selling the good only if the virtual value is positive.

References

Amador, Manuel, Iván Werning, and George Marios Angeletos. 2006. Commitment vs. flexibility. *Econometrica*, 74(2):365–396.

Andrade, Philippe, Jordi Galí, Hervé Le Bihan, and Julien Matheron. 2018. The optimal inflation target and the natural rate of interest. Technical report, National Bureau of Economic Research.

Angeletos, George-Marios, Christian Hellwig, and Alessandro Pavan. 2006. Signaling in a global game: Coordination and policy traps. *Journal of Political economy*, 114(3):452–484.

Ascari, Guido. 2004. Staggered prices and trend inflation: some nuisances. *Review of Economic dynamics*, 7(3):642–667.

Ascari, Guido and Tiziano Ropele. 2007. Optimal monetary policy under low trend inflation. *Journal of monetary Economics*, 54(8):2568–2583.

Ascari, Guido and Argia M Sbordone. 2014. The macroeconomics of trend inflation. *Journal of Economic Literature*, 52(3):679–739.

Athey, Susan, Andrew Atkeson, and Patrick J Kehoe. 2005. The optimal degree of discretion in monetary policy. *Econometrica*, 73(5):1431–1475.

Backus, David and John Driffill. 1985. Inflation and reputation. *The American Economic Review*, 75(3):530–538.

Ball, Laurence M. 2014. *The case for a long-run inflation target of four percent*. International Monetary Fund.

Barro, Robert J and David B Gordon. 1983. Rules, discretion and reputation in a model of monetary policy. *Journal of Monetary Economics*, 12(1):101–121.

Beshears, John, James Choi, Christopher Clayton, Christopher Harris, David Laibson, and Brigitte Madrian. 2020. Optimal Illiquidity.

Bhattarai, Saroj, Gauti B Eggertsson, and Bulat Gafarov. 2019. Time Consistency and Duration of Government Debt: A Model of Quantitative Easing. *The Review of Economic Studies*.

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

Blanchard, Olivier, Giovanni Dell'Ariccia, and Paolo Mauro. 2010. Rethinking macroeconomic policy. *Journal of Money, Credit and Banking*, 42:199–215.

Brainard, Lael. 2015. No Title. Technical report, Board of Governors of the Federal Reserve System. Speech delivered at the 57th National Association for Business Economics Annual Meeting.

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

Campbell, Jeffrey R, Charles L Evans, Jonas D M Fisher, Alejandro Justiniano, Charles W Calomiris, and Michael Woodford. 2012. Macroeconomic effects of federal reserve forward guidance [with comments and discussion]. *Brookings papers on economic activity*, pp. 1–80.

Canzoneri, Matthew B. 1985. Monetary policy games and the role of private information. *American Economic Review*, 75(5):1056–1070.

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

Clarida, Richard H. 2019. The Federal Reserve's Review of Its Monetary Policy Strategy, Tools, and Communication Practices. Technical report, Board of Governors of the Federal Reserve System. Speech delivered at the 2019 U.S. Monetary Policy Forum, sponsored by the Initiative on Global Markets at the University of Chicago Booth School of Business, held in New York, February 22.

Coibion, Olivier, Yuriy Gorodnichenko, and Johannes Wieland. 2012. The optimal inflation rate in New Keynesian models: should central banks raise their inflation targets in light of the zero lower bound? *Review of Economic Studies*, 79(4):1371–1406.

Cukierman, Alex and Allan H Meltzer. 1986. A theory of ambiguity, credibility, and inflation under discretion and asymmetric information. *Econometrica: journal of the econometric society*, pp. 1099–1128.

Dávila, Eduardo and Andreas Schaab. 2022. Optimal monetary policy with heterogeneous agents: A timeless ramsey approach. *Available at SSRN 4102028*.

Del Negro, Marco, Michele Lenza, Giorgio E Primiceri, and Andrea Tambalotti. 2020. What's up with the Phillips Curve? Technical report, National Bureau of Economic Research.

DellaVigna, Stefano and Ulrike Malmendier. 2004. Contract Design and Self-Control: Theory and Evidence. *Quarterly Journal of Economics*, 119(2):353–402.

Eggertsson, Gauti B, Neil R Mehrotra, and Jacob A Robbins. 2019. A model of secular stagnation: Theory and quantitative evaluation. *American Economic Journal: Macroeconomics*, 11(1):1–48.

Eggertsson, Gauti B and Michael Woodford. 2003. Zero bound on interest rates and optimal monetary policy. *Brookings papers on economic activity*, 2003(1):139–233.

Farhi, Emmanuel and Iván Werning. 2013. Insurance and taxation over the life cycle. *Review of Economic Studies*, 80(2):596–635.

Felix Hüfner. 2004. *Foreign Exchange Intervention as a Monetary Policy Instrument: Evidence for Inflation Targeting Countries*. Heidelberg: Physica.

Fischer, Stanley. 1977. Long-term contracts, rational expectations, and the optimal money supply rule. *Journal of political economy*, 85(1):191–205.

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

Galí, Jordi and Luca Gambetti. 2019. Has the US wage Phillips curve flattened? A semi-structural exploration. Technical report, National Bureau of Economic Research.

Galperti, Simone. 2015. Commitment, Flexibility, and Optimal Screening of Time Inconsistency. *Econometrica*, 83(4):1425–1465.

Gürkaynak, Refet S, Brian Sack, and Eric Swanson. 2005. The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models. *American economic review*, 95(1):425–436.

Halac, Marina and Pierre Yared. 2014. Fiscal Rules and Discretion Under Persistent Shocks. *Econometrica*, 82(5):1557–1614.

—. 2018. Fiscal Rules and Discretion in a World Economy. *American Economic Review*, 108(8):2305–2334.

—. 2022. Instrument-Based vs. Target-Based Rules. *Review of Economic Studies*, 89(1):312–345.

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.

Krishnamurthy, Arvind and Annette Vissing-Jorgensen. 2012. The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy. *Brookings Papers on Economic Activity Fall 2011*, p. 215.

Krugman, Paul. 2014. Inflation targets reconsidered. In *Draft paper for ECB Sintra conference*.

Kuttner, Kenneth N. 2001. Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of monetary economics*, 47(3):523–544.

Kydland, Finn E and Edward C Prescott. 1977. Rules Rather than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy*, 85(3):473–491.

Laibson, David. 1997. Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2):443–478.

Laubach, Thomas and John C Williams. 2016. Measuring the natural rate of interest redux. *Business Economics*, 51(2):57–67.

L'Huillier, Jean-Paul and Raphael Schoenle. 2019. Raising the Target: How Much Extra Room Does It Really Give?

Lucca, David O and Emanuel Moench. 2015. The pre-FOMC announcement drift. *The Journal of finance*, 70(1):329–371.

Marcet, Albert and Ramon Marimon. 2019. Recursive contracts. *Econometrica*, 87(5):1589–1631.

McDermott, John and Rebecca Williams. 2018. Inflation Targeting in New Zealand: an experience in evolution. Technical report.

Mele, Antonio. 2014. Repeated moral hazard and recursive lagrangeans. *Journal of Economic Dynamics and Control*, 42:69–85.

Messner, Matthias, Nicola Pavoni, and Christopher Sleet. 2012. Recursive methods for incentive problems. *Review of Economic Dynamics*, 15(4):501–525.

Pavan, Alessandro, Ilya Segal, and Juuso Toikka. 2014. Dynamic Mechanism Design: A Myersonian Approach. *Econometrica*, 82(2):601–653.

Pavoni, Nicola, Christopher Sleet, and Matthias Messner. 2018. The Dual Approach to Recursive Optimization: Theory and Examples. *Econometrica*, 86(1):133–172.

Persson, Torsten and Guido Tabellini. 1993. Designing institutions for monetary stability. *Carnegie-Rochester Confer. Series on Public Policy*, 39(C):53–84.

Powell, Jerome H. 2019. Monetary policy in the post-crisis era. Speech at "Bretton Woods: 75 Years Later - Thinking about the Next 75".

Powell, Jerome H. 2020. New Economic Challenges and the Fed's Monetary Policy Review. Technical report, Board of Governors of the Federal Reserve System. Speech delivered at "Navigating the Decade Ahead: Implications for Monetary Policy," an economic policy symposium sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming.

Rogoff, Kenneth. 1985. The Optimal Degree of Commitment to an Intermediate Monetary Target. *The Quarterly Journal of Economics*, 100(4):1169–1189.

Romer, Christina D and David H Romer. 2000. Federal Reserve Information and the Behavior of Interest Rates. *American Economic Review*, 90(3):429–457.

Rubbo, Elisa. 2020. Networks, Phillips Curves and Monetary Policy.

Sargent, Thomas J and Neil Wallace. 1975. " Rational" expectations, the optimal monetary instrument, and the optimal money supply rule. *Journal of political economy*, 83(2):241–254.

Schmitt-Grohé, Stephanie and Martin Uribe. 2010. The optimal rate of inflation. In *Handbook of monetary economics*, volume 3, pp. 653–722. Elsevier.

Sleet, Christopher. 2001. On credible monetary policy and private government information. *Journal of Economic Theory*, 99(1-2):338–376.

Sublet, Guillaume. 2022. Fiscal Rules with Discretion for an Economic Union.

Svensson, Lars E O. 1997a. Inflation forecast targeting: Implementing and monitoring inflation targets. *European economic review*, 41(6):1111–1146.

—. 1997b. Optimal Inflation Targets, "Conservative" Central Banks, and Linear Inflation Contracts. *The American Economic Review*, 87(1):98–114.

Svensson, Lars E O and Michael Woodford. 2004. Implementing optimal policy through inflation-forecast targeting. In *The inflation-targeting debate*, pp. 19–92. University of Chicago Press.

Svensson, Lars E.O. 2010. Inflation targeting. *Handbook of Monetary Economics*, 3(C):1237–1302.

Taylor, John B. 1980. Aggregate dynamics and staggered contracts. *Journal of political economy*, 88(1):1–23.

Waki, Yuichiro, Richard Dennis, and Ippei Fujiwara. 2018. The optimal degree of monetary discretion in a new Keynesian model with private information. *Theoretical Economics*, 13(3):1319–1367.

Walsh, Carl E. 1995. Optimal contracts for central bankers. *American Economic Review*, 85(1):150–167.

Werning, Ivan. 2011. Managing a liquidity trap: Monetary and fiscal policy. Technical report, National Bureau of Economic Research.

Werning, Iván. 2022. Expectations and the Rate of Inflation. Technical report, National Bureau of Economic Research.

Online Appendix

A Proofs

A.1 Proof of Proposition 1

Under full information, the objective function of the government is

$$\sup_{\pi_t} E_0 \sum_{t=0}^{\infty} \beta^t U_t \left(\pi_t, E_t \left[\pi_{t+1}|\theta_t\right], \theta_t\right).$$

Taking the FOC in π_t, we have

$$0 = \beta^{t-1} \frac{\partial U_{t-1}}{\partial \mathbb{E}_{t-1}\pi_t} \frac{\partial \mathbb{E}_{t-1}\pi_t}{\partial \pi_t(\theta^t)} f(\theta^{t-1}) + \beta^t \frac{\partial U_t}{\partial \pi_t} f(\theta^t)$$

From here, we have $\frac{\partial \mathbb{E}_{t-1}\pi_t}{\partial \pi_t(\theta^t)} = f(\theta_t|\theta_{t-1})$, so that we have

$$0 = \frac{\partial U_{t-1}}{\partial \mathbb{E}_{t-1}\pi_t} + \beta \frac{\partial U_t}{\partial \pi_t}$$

from which the result follows.

A.2 Proof of Proposition 3

The proof strategy is as follows. First, we derive the relevant envelope condition associated with local incentive compatibility, which defines necessary conditions on the value function associated with an incentive compatible mechanism.[41] We then show that the value function generated by our proposed mechanism satisfies this envelope condition.

Envelope condition. Suppose that the central bank has a history $\tilde{\theta}^{t-1}$ of reports and a history θ^t of true types at date t. Given a mechanism with transfer rule $T_t(\tilde{\theta}_t)$ and allocation rule $\pi_t(\tilde{\theta}_t)$, the value function of a central bank that has truthfully reported in the past, assuming truthful reporting in the future, is given by

$$\mathcal{W}_t(\theta^t) = \max_{\tilde{\theta}_t} T_t + U_t\left(\pi_t, \mathbb{E}_t\left[\pi_{t+1}|\tilde{\theta}_t\right], \theta_t\right) + \beta \mathbb{E}_t\left[\mathcal{W}_{t+1}(\theta^t, \tilde{\theta}_t, \theta_{t+1})\middle|\theta_t\right]$$

Notice that the Phillips curve expectation $\mathbb{E}_t\left[\pi_{t+1}|\tilde{\theta}_t\right]$ is based on the date t reported type, not the date t true type. Furthermore, notice that $\mathcal{W}_{t+1}$ depends on the reported type $\tilde{\theta}_t$, but not on the true type θ_t. This is because flow utility at dates $t+s$ ($s \geq 0$) do not depend on past true types and

[41] This portion of the argument follows from the arguments in Farhi and Werning (2013), or more generally from Pavan et al. (2014), but we state it out for completeness and for clarity.

because the shock structure is Markov. This implies that we can in fact write $\mathcal{W}_{t+1}(\theta^{t-1}, \tilde{\theta}_t, \theta_{t+1})$. As a result, the Envelope Condition in the true type θ_t, evaluated at truthful reporting $\tilde{\theta}_t = \theta_t$, is

$$\frac{\partial \mathcal{W}_t(\theta^t)}{\partial \theta_t} = \frac{\partial U_t\left(\pi_t, \mathbb{E}_t\left[\pi_{t+1}|\tilde{\theta}_t\right], \theta_t\right)}{\partial \theta_t} + \beta \frac{\partial \mathbb{E}_t\left[\mathcal{W}_{t+1}(\theta^{t-1}, \tilde{\theta}_t, \theta_{t+1})\middle|\theta_t\right]}{\partial \theta_t}$$

where we have

$$\begin{aligned}\frac{\partial \mathbb{E}_t\left[\mathcal{W}_{t+1}(\theta^{t-1}, \tilde{\theta}_t, \theta_{t+1})\middle|\theta_t\right]}{\partial \theta_t} &= \frac{\partial}{\partial \theta_t}\int_{\underline{\theta}}^{\overline{\theta}} \mathcal{W}_{t+1}(\theta^{t-1}, \tilde{\theta}_t, \theta_{t+1}) f(\theta_{t+1}|\theta_t) d\theta_{t+1} \\ &= \mathbb{E}_t\left[\mathcal{W}_{t+1}(\theta^{t-1}, \tilde{\theta}_t, \theta_{t+1}) \frac{\partial f(\theta_{t+1}|\theta_t)/\partial \theta_t}{f(\theta_{t+1}|\theta_t)}\middle|\theta_t\right]\end{aligned}$$

Substituting in and evaluating at truthful reporting, we obtain

$$\frac{\partial \mathcal{W}_t(\theta^t)}{\partial \theta_t} = \frac{\partial U_t\left(\pi_t, \mathbb{E}_t\left[\pi_{t+1}|\theta_t\right], \theta_t\right)}{\partial \theta_t} + \beta \mathbb{E}_t\left[\mathcal{W}_{t+1}(\theta^{t+1}) \frac{\partial f(\theta_{t+1}|\theta_t)/\partial \theta_t}{f(\theta_{t+1}|\theta_t)}\middle|\theta_t\right]$$

which provides a conventional envelope condition for incentive compatibility. For clarity, note that $\frac{\partial U_t(\pi_t, \mathbb{E}_t[\pi_{t+1}|\theta_t], \theta_t)}{\partial \theta_t}$ is the derivative of U_t in the direct type θ_t, but *not* including the Phillips curve expectation, which is the derivative in the reported type.

Verifying the envelope condition. We now verify the value function under our mechanism satisfies the envelope condition. Our mechanism has a transfer rule $T_t(\theta^t) = -\nu_{t-1}(\theta^{t-1})\Big(\pi_t(\theta^t) - \mathbb{E}_{t-1}[\pi_t|\theta_{t-1}]\Big)$ and an allocation rule given by the constrained efficient allocation of Proposition 1. The value function associated with this mechanism is

$$\mathcal{W}_t(\theta^t) = -\nu_{t-1}\Big(\pi_t - \mathbb{E}_{t-1}[\pi_t|\theta_{t-1}]\Big) + U_t\left(\pi_t, \mathbb{E}_t\left[\pi_{t+1}|\theta_t\right], \theta_t\right) + \beta \mathbb{E}_t\left[\mathcal{W}_{t+1}(\theta^{t+1})\middle|\theta_t\right]$$

where $\nu_{t-1}, \pi_t, \mathbb{E}_{t-1}[\pi_t|\theta_{t-1}])$ are the constrained efficient values associated with Proposition 1, given the realized shock history. From here, recall that ν_{t-1} and $\mathbb{E}_{t-1}[\pi_t|\theta_{t-1}]$ are only functions of θ^{t-1}. Therefore, $\frac{\partial \nu_{t-1}}{\partial \theta_t} = \frac{\partial \mathbb{E}_{t-1}[\pi_t|\theta_{t-1}]}{\partial \theta_t} = 0$. Thus differentiating the value function in θ_t, we have

$$\begin{aligned}\frac{\partial \mathcal{W}_t(\theta^t)}{\partial \theta_t} =& \frac{\partial U_t}{\partial \theta_t} + \beta \mathbb{E}_t\left[\mathcal{W}_{t+1}(\theta^{t+1}) \frac{\partial f(\theta_{t+1}|\theta_t)/\partial \theta_t}{f(\theta_{t+1}|\theta_t)}\middle|\theta_t\right] \\ & - \nu_{t-1}\frac{\partial \pi_t}{\partial \theta_t} + \frac{\partial U_t}{\partial \pi_t}\frac{\partial \pi_t}{\partial \theta_t} + \frac{\partial U_t}{\partial \mathbb{E}_t\left[\pi_{t+1}|\theta_t\right]}\frac{d\mathbb{E}_t\left[\pi_{t+1}|\theta_t\right]}{d\theta_t} + \beta \mathbb{E}_t\left[\frac{\partial \mathcal{W}_{t+1}(\theta^{t+1})}{\partial \theta_t}\middle|\theta_t\right]\end{aligned}$$

The first line on the RHS are the terms associated with the envelope condition. The second line are derivatives that arise because in equilibrium, the reported type equals the true type, and we have evaluated the value function given truthful reporting. It therefore remains to show that the second line sums to zero and hence our mechanism satisfies the required envelope condition.

It is helpful to write out the continuation value function $\mathcal{W}_{t+1}$ in sequence notation. Iterating forward, we obtain

$$\begin{aligned}\mathcal{W}_{t+1}(\theta^{t+1}) = & -\nu_t\Big(\pi_{t+1} - \mathbb{E}_t[\pi_{t+1}|\theta_t]\Big) \\ & - \mathbb{E}_{t+1}\Bigg[\sum_{s=1}^{\infty}\beta^s\nu_{t+s}\Big(\pi_{t+1+s} - \mathbb{E}_{t+s}[\pi_{t+1+s}|\theta_{t+s}]\Big)\Bigg|\theta_{t+1}\Bigg] \\ & + \mathbb{E}_{t+1}\Bigg[\sum_{s=0}^{\infty}\beta^s U_{t+1+s}\left(\pi_{t+1+s}, \mathbb{E}_{t+1+s}\left[\pi_{t+2+s}|\theta_{t+1+s}\right], \theta_{t+1+s}\right)\Bigg|\theta_{t+1}\Bigg]\end{aligned}$$

The first two lines on the RHS are total expected discounted value arising from transfers. The third line on the RHS is total expected discounted value arising from flow utility.

Notice from here that the second line is equal to zero. To see this, applying Law of Iterated Expectations, when $s \geq 1$ we have

$$\mathbb{E}_{t+1}\Bigg[\nu_{t+s}\pi_{t+1+s}|\theta_{t+1}\Bigg] = \mathbb{E}_{t+1}\Bigg[\mathbb{E}_{t+s}\Bigg[\nu_{t+s}\pi_{t+1+s}\Bigg|\theta_{t+s}\Bigg]|\theta_{t+1}\Bigg] = \mathbb{E}_{t+1}\Bigg[\nu_{t+s}\mathbb{E}_{t+s}\Bigg[\pi_{t+1+s}\Bigg|\theta_{t+s}\Bigg]|\theta_{t+1}\Bigg]$$

since ν_{t+s} is a function only of θ^{t+s}, and so is known at date $t+s$. As a result, the second line is zero, and we can write

$$\begin{aligned}\mathcal{W}_{t+1}(\theta^{t+1}) = & -\nu_t\Big(\pi_{t+1} - \mathbb{E}_t[\pi_{t+1}|\theta_t]\Big) \\ & + \mathbb{E}_{t+1}\Bigg[\sum_{s=0}^{\infty}\beta^s U_{t+1+s}\left(\pi_{t+1+s}, \mathbb{E}_{t+1+s}\left[\pi_{t+2+s}|\theta_{t+1+s}\right], \theta_{t+1+s}\right)\Bigg|\theta_{t+1}\Bigg]\end{aligned}$$

From here, we differentiate the continuation value $\mathcal{W}_{t+1}(\theta^{t+1})$ in the date t type θ_t, yielding

$$\begin{aligned}\frac{\partial\mathcal{W}_{t+1}(\theta^{t+1})}{\partial\theta_t} = & -\frac{\partial\nu_t}{\partial\theta_t}\Big(\pi_{t+1} - \mathbb{E}_t[\pi_{t+1}|\theta_t]\Big) - \nu_t\left(\frac{\partial\pi_{t+1}}{\partial\theta_t} - \frac{d\mathbb{E}_t[\pi_{t+1}|\theta_t]}{d\theta_t}\right) \\ & + \mathbb{E}_{t+1}\Bigg[\sum_{s=0}^{\infty}\beta^s\left(\frac{\partial U_{t+1+s}}{\partial\pi_{t+1+s}}\frac{\partial\pi_{t+1+s}}{\partial\theta_t} + \frac{\partial U_{t+1+s}}{\partial\mathbb{E}_{t+1+s}\left[\pi_{t+2+s}|\theta_{t+1+s}\right]}\mathbb{E}_{t+1+s}\left[\frac{\partial\pi_{t+2+s}}{\partial\theta_t}|\theta_{t+1+s}\right]\right)\Bigg|\theta_{t+1}\Bigg]\end{aligned}$$

Notice in the above derivation that only the first line includes a total derivative of firm expectations, $\frac{d\mathbb{E}_t[\pi_{t+1}|\theta_t]}{d\theta_t}$, which accounts for the changes in probability density. All later lines only include the direct change in inflation policy. This is because conditional expectations at date $t+1$ are taken with respect to θ_{t+1}, not θ_t.

We now rearrange the first term on the second line as follows. In particular, we write

$$\sum_{s=0}^{\infty} \beta^s \frac{\partial U_{t+1+s}}{\partial \pi_{t+1+s}} \frac{\partial \pi_{t+1+s}}{\partial \theta_t} = \frac{\partial U_{t+1}}{\partial \pi_{t+1}} \frac{\partial \pi_{t+1}}{\partial \theta_t} + \sum_{s=0}^{\infty} \beta^{s+1} \frac{\partial U_{t+2+s}}{\partial \pi_{t+2+s}} \frac{\partial \pi_{t+2+s}}{\partial \theta_t}$$

which extracts the first element of the sum, and relabels the remainder of the sum to continue to start from $s = 0$. Substituting back in, we obtain

$$\begin{aligned}\frac{\partial \mathcal{W}_{t+1}(\theta^{t+1})}{\partial \theta_t} = &- \frac{\partial \nu_t}{\partial \theta_t} \Big(\pi_{t+1} - \mathbb{E}_t[\pi_{t+1}|\theta_t] \Big) - \nu_t \Big(\frac{\partial \pi_{t+1}}{\partial \theta_t} - \frac{d\mathbb{E}_t[\pi_{t+1}|\theta_t]}{d\theta_t} \Big) + \frac{\partial U_{t+1}}{\partial \pi_{t+1}} \frac{\partial \pi_{t+1}}{\partial \theta_t} \\ &+ \mathbb{E}_{t+1} \Bigg[\sum_{s=0}^{\infty} \beta^{s+1} \Bigg(\frac{\partial U_{t+2+s}}{\partial \pi_{t+2+s}} \frac{\partial \pi_{t+2+s}}{\partial \theta_t} \\ &\qquad + \frac{1}{\beta} \frac{\partial U_{t+1+s}}{\partial \mathbb{E}_{t+1+s}\left[\pi_{t+2+s}|\theta_{t+1+s}\right]} \mathbb{E}_{t+1+s}\left[\frac{\partial \pi_{t+2+s}}{\partial \theta_t} |\theta_{t+1+s} \right] \Bigg) \Bigg| \theta_{t+1} \Bigg].\end{aligned}$$

By definition, we have $\nu_{t+s+1} = -\frac{1}{\beta} \frac{\partial U_{t+1+s}}{\partial \mathbb{E}_{t+1+s}[\pi_{t+2+s}|\theta_{t+1+s}]}$, given the allocation rule we are using in constructing the value function is the constrained efficient allocation rule. By Proposition 1, we also have $\frac{\partial U_{t+2+s}}{\partial \pi_{t+2+s}} = \nu_{t+s+1}$ for the same reason. Therefore, we can write

$$\begin{aligned}&\mathbb{E}_{t+1} \Bigg[\sum_{s=0}^{\infty} \beta^{s+1} \Bigg(\frac{\partial U_{t+2+s}}{\partial \pi_{t+2+s}} \frac{\partial \pi_{t+2+s}}{\partial \theta_t} + \frac{1}{\beta} \frac{\partial U_{t+1+s}}{\partial \mathbb{E}_{t+1+s}\left[\pi_{t+2+s}|\theta_{t+1+s}\right]} \mathbb{E}_{t+1+s}\left[\frac{\partial \pi_{t+2+s}}{\partial \theta_t} |\theta_{t+1+s} \right] \Bigg) \Bigg| \theta_{t+1} \Bigg] \\ &= \mathbb{E}_{t+1} \Bigg[\sum_{s=0}^{\infty} \beta^{s+1} \Bigg(\nu_{t+1+s} \frac{\partial \pi_{t+2+s}}{\partial \theta_t} - \nu_{t+1+s} \mathbb{E}_{t+1+s}\left[\frac{\partial \pi_{t+2+s}}{\partial \theta_t} |\theta_{t+1+s} \right] \Bigg) \Bigg| \theta_{t+1} \Bigg] \\ &= 0\end{aligned}$$

where the last line follows by Law of Iterated expectations,

$$\begin{aligned}\mathbb{E}_{t+1} \left[\nu_{t+1+s} \mathbb{E}_{t+1+s}\left[\frac{\partial \pi_{t+2+s}}{\partial \theta_t} |\theta_{t+1+s} \right] \Bigg| \theta_{t+1} \right] &= \mathbb{E}_{t+1} \left[\mathbb{E}_{t+1+s} \left[\nu_{t+1+s} \frac{\partial \pi_{t+2+s}}{\partial \theta_t} |\theta_{t+1+s} \right] \Bigg| \theta_{t+1} \right] \\ &= \mathbb{E}_{t+1} \left[\nu_{t+1+s} \frac{\partial \pi_{t+2+s}}{\partial \theta_t} \Bigg| \theta_{t+1} \right].\end{aligned}$$

Therefore, we obtain

$$\frac{\partial \mathcal{W}_{t+1}(\theta^{t+1})}{\partial \theta_t} = - \frac{\partial \nu_t}{\partial \theta_t} \Big(\pi_{t+1} - \mathbb{E}_t[\pi_{t+1}|\theta_t] \Big) - \nu_t \Big(\frac{\partial \pi_{t+1}}{\partial \theta_t} - \frac{d\mathbb{E}_t[\pi_{t+1}|\theta_t]}{d\theta_t} \Big) + \frac{\partial U_{t+1}}{\partial \pi_{t+1}} \frac{\partial \pi_{t+1}}{\partial \theta_t}.$$

Finally, notice that as before, by Proposition 1 we have $\nu_t = \frac{\partial U_{t+1}}{\partial \pi_{t+1}}$, and therefore we can write

$$\frac{\partial \mathcal{W}_{t+1}(\theta^{t+1})}{\partial \theta_t} = - \frac{\partial \nu_t}{\partial \theta_t} \Big(\pi_{t+1} - \mathbb{E}_t[\pi_{t+1}|\theta_t] \Big) + \nu_t \frac{d\mathbb{E}_t[\pi_{t+1}|\theta_t]}{d\theta_t}.$$

We are now ready to substitute back in to the expression for $\frac{\partial \mathcal{W}_t}{\partial \theta_t}$. Substituting back in, we have

$$\begin{aligned}\frac{\partial \mathcal{W}_t(\theta^t)}{\partial \theta_t} =& \frac{\partial U_t}{\partial \theta_t} + \beta \mathbb{E}_t \left[\mathcal{W}_{t+1}(\theta^{t+1}) \frac{\partial f(\theta_{t+1}|\theta_t)/\partial \theta_t}{f(\theta_{t+1}|\theta_t)} \middle| \theta_t \right] \\ & - \nu_{t-1}\frac{\partial \pi_t}{\partial \theta_t} + \frac{\partial U_t}{\partial \pi_t}\frac{\partial \pi_t}{\partial \theta_t} + \frac{\partial U_t}{\partial \mathbb{E}_t\left[\pi_{t+1}|\theta_t\right]}\frac{d\mathbb{E}_t\left[\pi_{t+1}|\theta_t\right]}{d\theta_t} \\ & + \beta \mathbb{E}_t\left[-\frac{\partial \nu_t}{\partial \theta_t}\Big(\pi_{t+1} - \mathbb{E}_t[\pi_{t+1}|\theta_t]\Big) + \nu_t \frac{d\mathbb{E}_t[\pi_{t+1}|\theta_t]}{d\theta_t} \middle| \theta_t \right]\end{aligned}$$

The arguments now are familiar. The first term on the third line is zero, since

$$\mathbb{E}_t\left[-\frac{\partial \nu_t}{\partial \theta_t}\Big(\pi_{t+1} - \mathbb{E}_t[\pi_{t+1}|\theta_t]\Big) \middle| \theta_t \right] = -\frac{\partial \nu_t}{\partial \theta_t}\mathbb{E}_t\left[\pi_{t+1} - \mathbb{E}_t[\pi_{t+1}|\theta_t] \middle| \theta_t \right] = 0.$$

From here, we can rearrange terms to get

$$\begin{aligned}\frac{\partial \mathcal{W}_t(\theta^t)}{\partial \theta_t} =& \frac{\partial U_t}{\partial \theta_t} + \beta \mathbb{E}_t \left[\mathcal{W}_{t+1}(\theta^{t+1}) \frac{\partial f(\theta_{t+1}|\theta_t)/\partial \theta_t}{f(\theta_{t+1}|\theta_t)} \middle| \theta_t \right] \\ & + \left[-\nu_{t-1} + \frac{\partial U_t}{\partial \pi_t}\right]\frac{\partial \pi_t}{\partial \theta_t} + \frac{\partial U_t}{\partial \mathbb{E}_t\left[\pi_{t+1}|\theta_t\right]}\frac{d\mathbb{E}_t\left[\pi_{t+1}|\theta_t\right]}{d\theta_t} + \beta \mathbb{E}_t\left[\nu_t \frac{d\mathbb{E}_t[\pi_{t+1}|\theta_t]}{d\theta_t} \middle| \theta_t \right]\end{aligned}$$

By Proposition 1, we have $-\nu_{t-1} + \frac{\partial U_t}{\partial \pi_t} = 0$.[42] Likewise from the definition of ν_t, we have $\frac{\partial U_t}{\partial \mathbb{E}_t[\pi_{t+1}|\theta_t]} = -\beta \nu_t$. Therefore, we also have

$$\frac{\partial U_t}{\partial \mathbb{E}_t\left[\pi_{t+1}|\theta_t\right]}\frac{d\mathbb{E}_t\left[\pi_{t+1}|\theta_t\right]}{d\theta_t} + \beta \mathbb{E}_t\left[\nu_t \frac{d\mathbb{E}_t[\pi_{t+1}|\theta_t]}{d\theta_t} \middle| \theta_t \right] = -\beta\nu_t \frac{d\mathbb{E}_t\left[\pi_{t+1}|\theta_t\right]}{d\theta_t} + \beta\nu_t \frac{d\mathbb{E}_t[\pi_{t+1}|\theta_t]}{d\theta_t} = 0.$$

Thus, the entire second line is zero, and we are left with

$$\frac{\partial \mathcal{W}_t(\theta^t)}{\partial \theta_t} = \frac{\partial U_t}{\partial \theta_t} + \beta \mathbb{E}_t \left[\mathcal{W}_{t+1}(\theta^{t+1}) \frac{\partial f(\theta_{t+1}|\theta_t)/\partial \theta_t}{f(\theta_{t+1}|\theta_t)} \middle| \theta_t \right]$$

which is the required envelope condition. This concludes the proof.

A.3 Proof of Proposition 4

Using reduced form preferences, our two key equations are

$$\nu_{t-1} = -\pi_t - \hat{\alpha}\Big(\pi_t - \beta \mathbb{E}_t \pi_{t+1}\Big) + \hat{\lambda}$$

[42] For completeness, note that when considering the date 0 value function, we have $\nu_{-1} = 0$ and have $\frac{\partial U_t}{\partial \pi_t} = 0$ by Proposition 1.

$$\nu_t = -\hat{\alpha}\Big(\pi_t - \beta\mathbb{E}_t\pi_{t+1}\Big) + \hat{\lambda} - \frac{1}{\beta}\theta_t$$

Summing the two equations, we get

$$\nu_t = \nu_{t-1} + \pi_t - \frac{1}{\beta}\theta_t$$

Now, we guess and verify a linear solution of the form,

$$\nu_t = \gamma_0 + \gamma_1\theta_t + \gamma_2\nu_{t-1}$$

Using our key equation, we get

$$\pi_t = \nu_t - \nu_{t-1} + \frac{1}{\beta}\theta_t$$

Leading one period and taking expectations,

$$\mathbb{E}_t\pi_{t+1} = \gamma_0 + (\gamma_2 - 1)\nu_t + \Big(\gamma_1 + \frac{1}{\beta}\Big)\rho\theta_t$$

Now, substituting back in to the equation for ν_t and rearranging,

$$\Big(1 + \hat{\alpha} - \hat{\alpha}\beta(\gamma_2 - 1)\Big)\nu_t = \hat{\alpha}\beta\gamma_0 + \hat{\lambda} + \Big[\hat{\alpha}\beta\Big(\gamma_1 + \frac{1}{\beta}\Big)\rho - \frac{1+\hat{\alpha}}{\beta}\Big]\theta_t + \hat{\alpha}\nu_{t-1}$$

Now, we solve by coefficient matching. Coefficient matching on γ_2, we have

$$\Big(1 + \hat{\alpha} - \hat{\alpha}\beta(\gamma_2 - 1)\Big)\gamma_2 = \hat{\alpha}$$

$$0 = \hat{\alpha}\beta\gamma_2^2 - \Big(1 + \hat{\alpha} + \hat{\alpha}\beta\Big)\gamma_2 + \hat{\alpha}$$

and so the non-explosive root is

$$\gamma_2 = \frac{1 + \hat{\alpha} + \hat{\alpha}\beta - \sqrt{\Big(1 + \hat{\alpha} + \hat{\alpha}\beta\Big)^2 - 4\hat{\alpha}^2\beta}}{2\hat{\alpha}\beta}$$

Now, we can coefficient match on the constant,

$$\gamma_0 = \frac{\hat{\alpha}}{1 + \hat{\alpha} - \hat{\alpha}\beta(\gamma_2 - 1)}\frac{\hat{\alpha}\beta\gamma_0 + \hat{\lambda}}{\hat{\alpha}}$$

$$\gamma_0 = \frac{\gamma_2}{1 - \beta\gamma_2}\frac{\hat{\lambda}}{\hat{\alpha}}$$

Finally, coefficient mathcing on γ_1,

$$\gamma_1 = \frac{\hat{\alpha}}{1+\hat{\alpha}-\hat{\alpha}\beta(\gamma_2-1)} \frac{\left[\hat{\alpha}\beta\left(\gamma_1+\frac{1}{\beta}\right)\rho - \frac{1+\hat{\alpha}}{\beta}\right]}{\hat{\alpha}}$$

$$\gamma_1 = \frac{\gamma_2}{1-\gamma_2\beta\rho}\left[\rho - \frac{1+\hat{\alpha}}{\hat{\alpha}}\frac{1}{\beta}\right]$$

A.4 Proof of Proposition 5

Consider reduced-form preferences,

$$U_t(\pi_t, \mathbb{E}_t\pi_{t+1}, \theta_t) = -\frac{1}{2}\pi_t^2 - \frac{1}{2}\alpha\left(\pi_t - \beta\mathbb{E}_t\pi_{t+1}\right)^2 + v(\mathbb{E}_t\pi_{t+1}+\theta_t)$$

where for notational convenience we use α in place of $\hat{\alpha}$ in the derivations (and then simply replace α with $\hat{\alpha}$ at the end). Thus, we have derivatives

$$\frac{\partial U_t}{\partial \pi_t} = -\pi_t - \alpha\left(\pi_t - \beta\mathbb{E}_t\pi_{t+1}\right)$$

$$\frac{\partial U_t}{\partial \mathbb{E}_t\pi_{t+1}} = \alpha\beta\left(\pi_t - \beta\mathbb{E}_t\pi_t\right) + v'(i_t^*)$$

Under the usual definitions of ν_t, we then have

$$\nu_{t-1} = -\pi_t - \alpha\left(\pi_t - \beta\mathbb{E}_t\pi_{t+1}\right) \tag{24}$$

$$\nu_t = -\alpha\left(\pi_t - \beta\mathbb{E}_t\pi_{t+1}\right) - v_0 + v_1\mathbb{E}_t\pi_{t+1} + v_1\theta_t \tag{25}$$

where we have used $v'(i_t) = \beta v_0 - \beta v_1 i_t$ and $i_t^* = \mathbb{E}_t\pi_{t+1} + \theta_t$.

We now guess and verify a linear solution of the form

$$\nu_t = \gamma_0 + \gamma_1\nu_{t-1} + \gamma_2\theta_t.$$

Rearranging equation (24), we get

$$\beta\mathbb{E}_t\pi_{t+1} = \frac{1}{\alpha}\nu_{t-1} + \frac{1+\alpha}{\alpha}\pi_t, \tag{26}$$

and substituting into equation (25) we get

$$\nu_t = -v_0 + \frac{(\alpha\beta+v_1)(1+\alpha)-\alpha^2\beta}{\alpha\beta}\pi_t + \frac{\alpha\beta+v_1}{\alpha\beta}\nu_{t-1} + v_1\theta_t.$$

From here, we denote $\frac{1}{\zeta} \equiv \frac{(\alpha\beta+v_1)(1+\alpha)-\alpha^2\beta}{\alpha\beta} > 0$. Thus rearranging the above equation, we have

$$\frac{1}{\zeta}\pi_t = \nu_t + v_0 - \frac{\alpha\beta + v_1}{\alpha\beta}\nu_{t-1} - v_1\theta_t \tag{27}$$

We now lead this equation forward one period and take expectations,

$$\frac{1}{\zeta}\mathbb{E}_t\pi_{t+1} = \mathbb{E}_t\nu_{t+1} + v_0 - \frac{\alpha\beta + v_1}{\alpha\beta}\nu_t - v_1\mathbb{E}_t\theta_{t+1}$$

and now, we can use the guess for ν_t along with the property $\mathbb{E}_t\theta_{t+1} = \rho\theta_t$ to obtain

$$\frac{1}{\zeta}\mathbb{E}_t\pi_{t+1} = \gamma_0 + v_0 + \left(\gamma_1 - \frac{\alpha\beta + v_1}{\alpha\beta}\right)\nu_t + (\gamma_2 - v_1)\rho\theta_t.$$

Now, equations (26) and (27) jointly imply

$$\frac{1}{\zeta}\mathbb{E}_t\pi_{t+1} = \frac{1}{\zeta}\frac{1}{\alpha\beta}\nu_{t-1} + \frac{1+\alpha}{\alpha\beta}\left(\nu_t + v_0 - \frac{\alpha\beta + v_1}{\alpha\beta}v_{t-1} - v_1\theta_t\right)$$

and so substituting in, we obtain

$$\gamma_0 + v_0 + \left(\gamma_1 - \frac{\alpha\beta + v_1}{\alpha\beta}\right)\nu_t + (\gamma_2 - v_1)\rho\theta_t = \frac{1}{\zeta}\frac{1}{\alpha\beta}\nu_{t-1} + \frac{1+\alpha}{\alpha\beta}\left(\nu_t + v_0 - \frac{\alpha\beta + v_1}{\alpha\beta}v_{t-1} - v_1\theta_t\right)$$

which rearranges and simplifies to

$$\left(\gamma_1 - \frac{1+\alpha+\alpha\beta+v_1}{\alpha\beta}\right)\nu_t = \left(\frac{1+\alpha-\alpha\beta}{\alpha\beta}v_0 - \gamma_0\right) - \frac{1}{\beta}v_{t-1} - \left(\frac{1+\alpha-\alpha\beta\rho}{\alpha\beta}v_1 + \gamma_2\rho\right)\theta_t.$$

The LHS is linear, so using our guess $\nu_t = \gamma_0 + \gamma_1\nu_{t-1} + \gamma_2\theta_t$ and coefficient matching, we have the system

$$\gamma_0 = \frac{\frac{1+\alpha(1-\beta)}{\alpha\beta}v_0 - \gamma_0}{\gamma_1 - \frac{1+\alpha+\alpha\beta+v_1}{\alpha\beta}}$$

$$\gamma_1 = -\frac{1}{\beta}\frac{1}{\gamma_1 - \frac{1+\alpha+\alpha\beta+v_1}{\alpha\beta}}$$

$$\gamma_2 = \frac{-\left(\frac{1+\alpha(1-\beta\rho)}{\alpha\beta}v_1 + \gamma_2\rho\right)}{\gamma_1 - \frac{1+\alpha+\alpha\beta+v_1}{\alpha\beta}}$$

The second equation rearranges to a quadratic $\beta\gamma_1^2 - \frac{1+\alpha+\alpha\beta+v_1}{\alpha}\gamma_1 + 1 = 0$ in γ_1. We choose the non-explosive lower root to maintain consistency with the transversality condition, which yields

$$\gamma_1 = \frac{1 + \alpha(1+\beta) + v_1 - \sqrt{\left(1 + \alpha(1+\beta) + v_1\right)^2 - 4\alpha^2\beta}}{2\alpha\beta}$$

From here, the equation for γ_0 can be rewritten as $\gamma_0 = -\beta\gamma_1\left(\frac{1+\alpha(1-\beta)}{\alpha\beta}v_0 - \gamma_0\right)$, and rearranging yields

$$\gamma_0 = -\gamma_1 \frac{1+\alpha(1-\beta)}{\alpha(1-\beta\gamma_1)} v_0$$

Similarly, the eequation for γ_2 is rewritten as $\gamma_2 = \beta\gamma_1\left(\frac{1+\alpha(1-\beta\rho)}{\alpha\beta}v_1 + \gamma_2\rho\right)$, which rearranges to

$$\gamma_2 = \frac{1}{\alpha}\frac{1+\alpha(1-\beta\rho)}{1-\beta\gamma_1\rho}\gamma_1 v_1$$

Thus, we have our solution.

Inflation is given by

$$\frac{1}{\zeta}\pi_t = \nu_t - \frac{\alpha\beta + v_1}{\alpha\beta}\nu_{t-1} + v_0 - v_1\theta_t$$

A.5 Proof of Proposition 6

Given reduced form preferences $U_t = -\frac{1}{2}\pi_t^2 + \theta_t \frac{\pi_t - \beta\mathbb{E}_t\pi_{t+1}}{\kappa}$, then we have

$$\frac{\partial U_t}{\partial \pi_t} = -\pi_t + \frac{1}{\kappa}\theta_t$$

$$\frac{\partial U_{t-1}}{\partial \mathbb{E}_{t-1}\pi_t} = -\frac{\beta}{\kappa}\theta_{t-1}$$

Thus substituting in the definitions,

$$\nu_{t-1} = -\pi_t + \frac{1}{\kappa/\theta_t}$$

$$\nu_{t-1} = \frac{1}{\kappa/\theta_{t-1}}$$

Thus putting them together, we get $\pi_t = \frac{1}{\kappa/\theta_t} - \frac{1}{\kappa/\theta_{t-1}}$. Finally, using $\mathbb{E}_t\pi_{t+1} = 1 - \rho + \rho\theta_t$ we get

$$\mathbb{E}_t\pi_{t+1} = \frac{\mathbb{E}_t\theta_{t+1} - \theta_t}{\kappa} = (1-\rho)\frac{1}{\kappa} - (1-\rho)\frac{\theta_t}{\kappa}$$

which gives the result.

A.6 Proof of Proposition 7

Consider the Ramsey problem,

$$\max_{\pi} \sum_{t=0}^{\infty} \beta^t U_t(\pi_t, \mathbb{E}_t[\pi_{t+1}|\tilde{\theta}_t], ..., \mathbb{E}_t[\pi_{t+K}|\tilde{\theta}_t], \theta_t)$$

It is expositionally helpful to extend the sum to include $U_{-1}, ..., U_{-K} = 0$. Under this extended sum, differentiating in $\pi_t(\theta^t)$ for $t \geq 0$, we have

$$0 = \sum_{s=t-K}^{t-1} \beta^s \frac{\partial U_s}{\partial \mathbb{E}_s[\pi_t|\theta_s]} \frac{\partial \mathbb{E}_s[\pi_t|\theta_s]}{\partial \pi_t(\theta^t)} f(\theta^s) + \beta^t \frac{\partial U_t}{\partial \pi_t} f(\theta^t).$$

From here, note that we have

$$\frac{\partial \mathbb{E}_s[\pi_t|\theta_s]}{\partial \pi_t(\theta^t)} f(\theta^s) = f(\theta^t|\theta^s) f(\theta^s) = f(\theta^t)$$

Thus rearranging and dividing through, we have

$$\frac{\partial U_t}{\partial \pi_t} = -\sum_{s=t-K}^{t-1} \beta^{s-t} \frac{\partial U_s}{\partial \mathbb{E}_s[\pi_t|\theta_s]}.$$

Substituting in the definition of $\nu_{t,k}$ gives the result.

A.7 Proof of Proposition 9

The proof strategy is as follows. First, we derive the relevant envelope condition associated with local incentive compatibility, which defines necessary conditions on the value function associated with an incentive compatible mechanism. We then show that the value function generated by our proposed mechanism satisfies this envelope condition.

Envelope Condition. Suppose that the central bank has a history $\tilde{\theta}^{t-1}$ of reports and a history θ^t of true types at date t. Given a mechanism with transfer rule $T_t(\tilde{\theta}_t)$ and allocation rule $\pi_t(\tilde{\theta}_t)$, the value function of a central bank that has truthfully reported in the past, assuming truthful reporting in the future, is given by

$$\mathcal{W}_t(\theta^t) = \max_{\tilde{\theta}_t} T_t + U_t(\pi_t, \mathbb{E}_t[\pi_{t+1}|\tilde{\theta}_t], ..., \mathbb{E}_t[\pi_{t+K}|\tilde{\theta}_t], \theta_t) + \beta \mathbb{E}_t \left[\mathcal{W}_{t+1}(\theta^t, \tilde{\theta}_t, \theta_{t+1}) \middle| \theta_t \right]$$

Notice that the expectations at date t are based on the date t reported type, not the date t true type. Furthermore, notice that $\mathcal{W}_{t+1}$ depends on the reported type $\tilde{\theta}_t$, but not on the true type θ_t. This is because flow utility at dates $t+s$ ($s \geq 0$) do not depend on past true types and because the shock structure is Markov. This implies that we can in fact write $\mathcal{W}_{t+1}(\theta^{t-1}, \tilde{\theta}_t, \theta_{t+1})$. As a result, the

Envelope Condition in the true type θ_t, evaluated at truthful reporting $\tilde{\theta}_t = \theta_t$, is

$$\frac{\partial \mathcal{W}_t(\theta^t)}{\partial \theta_t} = \frac{\partial U_t}{\partial \theta_t} + \beta \frac{\partial \mathbb{E}_t \left[\mathcal{W}_{t+1}(\theta^{t-1}, \tilde{\theta}_t, \theta_{t+1}) \middle| \theta_t \right]}{\partial \theta_t}$$

where we have

$$\begin{aligned}
\frac{\partial \mathbb{E}_t \left[\mathcal{W}_{t+1}(\theta^{t-1}, \tilde{\theta}_t, \theta_{t+1}) \middle| \theta_t \right]}{\partial \theta_t} &= \frac{\partial}{\partial \theta_t} \int_{\underline{\theta}}^{\overline{\theta}} \mathcal{W}_{t+1}(\theta^{t-1}, \tilde{\theta}_t, \theta_{t+1}) f(\theta_{t+1}|\theta_t) d\theta_{t+1} \\
&= \mathbb{E}_t \left[\mathcal{W}_{t+1}(\theta^{t-1}, \tilde{\theta}_t, \theta_{t+1}) \frac{\partial f(\theta_{t+1}|\theta_t)/\partial \theta_t}{f(\theta_{t+1}|\theta_t)} \middle| \theta_t \right]
\end{aligned}$$

Substituting in and evaluating at truthful reporting, we obtain

$$\frac{\partial \mathcal{W}_t(\theta^t)}{\partial \theta_t} = \frac{\partial U_t \left(\pi_t, \mathbb{E}_t \left[\pi_{t+1}|\theta_t \right], ..., \mathbb{E}_t[\pi_{t+K}|\theta_t], \theta_t \right)}{\partial \theta_t} + \beta \mathbb{E}_t \left[\mathcal{W}_{t+1}(\theta^{t+1}) \frac{\partial f(\theta_{t+1}|\theta_t)/\partial \theta_t}{f(\theta_{t+1}|\theta_t)} \middle| \theta_t \right]$$

which provides a conventional envelope condition for incentive compatibility. For clarity, note that $\frac{\partial U_t}{\partial \theta_t}$ is the derivative of U_t in the direct type θ_t, but *not* including the Phillips curve expectation, which is the derivative in the reported type.

Verifying the Envelope Condition. We now verify the value function under our mechanism satisfies the envelope condition. Our mechanism has a transfer rule

$$T_t = -\sum_{k=1}^{K} \nu_{t,k}(\pi_t - \mathbb{E}_{t-k}\pi_t)$$

and an allocation rule given by the constrained efficient allocation of Proposition 7. It will at times be helpful to define

$$v_{t-1} = \sum_{k=1}^{K} \nu_{t,k}.$$

The value function associated with this mechanism is

$$\begin{aligned}
\mathcal{W}_t(\theta^t) = &-\sum_{k=1}^{K} \nu_{t,k}(\pi_t - \mathbb{E}_{t-k}\pi_t) \\
&+ U_t \left(\pi_t, \mathbb{E}_t \pi_{t+1}, ..., \mathbb{E}_t \pi_{t+K}, \theta_t \right) + \beta \mathbb{E}_t \left[\mathcal{W}_{t+1}(\theta^{t+1}) \middle| \theta_t \right]
\end{aligned}$$

where all objects are evaluated at their constrained efficient values associated with Proposition 7, given the realized shock history. Differentiating the value function in θ_t, we have

$$\begin{aligned}\frac{\partial \mathcal{W}_t(\theta^t)}{\partial \theta_t} =& \frac{\partial U_t}{\partial \theta_t} + \beta \mathbb{E}_t\left[\mathcal{W}_{t+1}(\theta^{t+1})\frac{\partial f(\theta_{t+1}|\theta_t)/\partial\theta_t}{f(\theta_{t+1}|\theta_t)}\middle|\theta_t\right] \\ & - v_{t-1}\frac{\partial \pi_t}{\partial \theta_t} + \frac{\partial U_t}{\partial \pi_t}\frac{\partial \pi_t}{\partial \theta_t} + \sum_{k=1}^{K}\frac{\partial U_t}{\partial \mathbb{E}_t\pi_{t+k}}\frac{d\mathbb{E}_t\pi_{t+k}}{d\theta_t} + \beta\mathbb{E}_t\left[\frac{\partial \mathcal{W}_{t+1}(\theta^{t+1})}{\partial \theta_t}\middle|\theta_t\right]\end{aligned}$$

The first line on the RHS are the terms associated with the envelope condition. The second line are derivatives that arise because in equilibrium, the reported type equals the true type, and we have evaluated the value function given truthful reporting. It therefore remains to show that the second line sums to zero and hence our mechanism satisfies the required envelope condition.

We begin by noting that the first two terms on the second line sum to zero, that is

$$-v_{t-1}\frac{\partial \pi_t}{\partial \theta_t} + \frac{\partial U_t}{\partial \pi_t}\frac{\partial \pi_t}{\partial \theta_t} = 0.$$

This follows immediately from Proposition 7 given the definition of v_{t-1}. We are therefore left to study the final two terms, and so we write

$$\begin{aligned}\frac{\partial \mathcal{W}_t(\theta^t)}{\partial \theta_t} =& \frac{\partial U_t}{\partial \theta_t} + \beta \mathbb{E}_t\left[\mathcal{W}_{t+1}(\theta^{t+1})\frac{\partial f(\theta_{t+1}|\theta_t)/\partial\theta_t}{f(\theta_{t+1}|\theta_t)}\middle|\theta_t\right] \\ & + \sum_{k=1}^{K}\frac{\partial U_t}{\partial \mathbb{E}_t\pi_{t+k}}\frac{d\mathbb{E}_t\pi_{t+k}}{d\theta_t} + \beta\mathbb{E}_t\left[\frac{\partial \mathcal{W}_{t+1}(\theta^{t+1})}{\partial \theta_t}\middle|\theta_t\right]\end{aligned}$$

It is helpful to write out the continuation value function $\mathcal{W}_{t+1}$ in sequence notation. Iterating forward, we obtain

$$\mathcal{W}_{t+1}(\theta^{t+1}) = \mathbb{E}_{t+1}\sum_{s=0}^{\infty}\beta^s\left[-\sum_{k=1}^{K}\nu_{t+1+s,k}(\pi_{t+1+s} - \mathbb{E}_{t+1+s-k}\pi_{t+1+s}) + U_{t+1+s}\right]$$

Now, we differentiate in θ_t. Here, we obtain

$$\begin{aligned}\frac{\partial \mathcal{W}_{t+1}(\theta^{t+1})}{\partial \theta_t} =& \mathbb{E}_{t+1}\sum_{s=0}^{\infty}\beta^s\left[-\sum_{k=1}^{K}\frac{\partial \nu_{t+1+s,k}}{\partial\theta_t}(\pi_{t+1+s} - \mathbb{E}_{t+1+s-k}\pi_{t+1+s})\right] \\ & + \mathbb{E}_{t+1}\sum_{s=0}^{\infty}\beta^s\left[\sum_{k=1}^{K}\nu_{t+1+s,k}\frac{d\mathbb{E}_{t+1+s-k}\pi_{t+1+s}}{d\theta_t}\right] \\ & + \mathbb{E}_{t+1}\sum_{s=0}^{\infty}\beta^s\left[-\sum_{k=1}^{K}\nu_{t+1+s,k}\frac{\partial \pi_{t+1+s}}{\partial\theta_t} + \frac{\partial U_{t+1+s}}{\partial \pi_{t+1+s}}\frac{\partial \pi_{t+1+s}}{\partial \theta_t}\right] \\ & + \mathbb{E}_{t+1}\sum_{s=0}^{\infty}\beta^s\left[\sum_{k=1}^{K}\frac{\partial U_{t+1+s}}{\partial \mathbb{E}_{t+1+s}\pi_{t+1+s+k}}\mathbb{E}_{t+1+s}\frac{\partial \pi_{t+1+s+k}}{\partial\theta_t}\right]\end{aligned}$$

To begin with, note that the third line is zero, from Proposition 7. Thus we can write,

$$\begin{aligned}\frac{\partial \mathcal{W}_{t+1}(\theta^{t+1})}{\partial \theta_t} =& \mathbb{E}_{t+1}\sum_{s=0}^{\infty}\beta^s\Big[-\sum_{k=1}^{K}\frac{\partial \nu_{t+1+s,k}}{\partial \theta_t}(\pi_{t+1+s}-\mathbb{E}_{t+1+s-k}\pi_{t+1+s})\Big] \\ &+\mathbb{E}_{t+1}\sum_{s=0}^{\infty}\beta^s\Big[\sum_{k=1}^{K}\nu_{t+1+s,k}\frac{d\mathbb{E}_{t+1+s-k}\pi_{t+1+s}}{d\theta_t}\Big] \\ &+\mathbb{E}_{t+1}\sum_{s=0}^{\infty}\beta^s\Big[\sum_{k=1}^{K}\frac{\partial U_{t+1+s}}{\partial \mathbb{E}_{t+1+s}\pi_{t+1+s+k}}\mathbb{E}_{t+1+s}\frac{\partial \pi_{t+1+s+k}}{\partial \theta_t}\Big]\end{aligned}$$

Next, recall that we can write

$$\nu_{t,k} = -\frac{1}{\beta^k}\frac{\partial U_{t-k}}{\partial \mathbb{E}_{t-k}\pi_t}$$

Therefore, we can equivalently write

$$\frac{\partial U_{t+1+s}}{\partial \mathbb{E}_{t+1+s}\pi_{t+1+s+k}} = -\beta^k\nu_{t+1+s+k,k}$$

Thus substituting into the third line,

$$\begin{aligned}\frac{\partial \mathcal{W}_{t+1}(\theta^{t+1})}{\partial \theta_t} =& \mathbb{E}_{t+1}\sum_{s=0}^{\infty}\beta^s\Big[-\sum_{k=1}^{K}\frac{\partial \nu_{t+1+s,k}}{\partial \theta_t}(\pi_{t+1+s}-\mathbb{E}_{t+1+s-k}\pi_{t+1+s})\Big] \\ &+\mathbb{E}_{t+1}\sum_{s=0}^{\infty}\beta^s\Big[\sum_{k=1}^{K}\nu_{t+1+s,k}\frac{d\mathbb{E}_{t+1+s-k}\pi_{t+1+s}}{d\theta_t}\Big] \\ &+\mathbb{E}_{t+1}\sum_{s=0}^{\infty}\beta^s\Big[\sum_{k=1}^{K}-\beta^k\nu_{t+1+s+k,k}\mathbb{E}_{t+1+s}\frac{\partial \pi_{t+1+s+k}}{\partial \theta_t}\Big]\end{aligned}$$

From here, let us compare the second and third lines. When $s \geq k$, we know that $t+1+s-k \geq t+1$ and so we have

$$\frac{d\mathbb{E}_{t+1+s-k}\pi_{t+1+s}}{d\theta_t} = \mathbb{E}_{t+1+s-k}\frac{\partial \pi_{t+1+s}}{\partial \theta_t}.$$

We also know that all terms with $t+1+s-k < t$, that is $k > 1+s$, drop out of the second line (since they are date $t-1$ or lower adapted constants). What this leaves us with is that the second line cancels out with the third line except for the points where $t+1+s-k = t$, that is precisely the points where there is also a probability measure derivative. Put together and taking the expectation at date t, this gives us

$$\begin{aligned}\mathbb{E}_t\frac{\partial \mathcal{W}_{t+1}(\theta^{t+1})}{\partial \theta_t} =& \mathbb{E}_t\sum_{s=0}^{\infty}\beta^s\Big[-\sum_{k=1}^{K}\frac{\partial \nu_{t+1+s,k}}{\partial \theta_t}(\pi_{t+1+s}-\mathbb{E}_{t+1+s-k}\pi_{t+1+s})\Big] \\ &+\sum_{s=0}^{K-1}\beta^s\nu_{t+1+s,1+s}\frac{d\mathbb{E}_t\pi_{t+1+s}}{d\theta_t}\end{aligned}$$

where we note that all terms on the second line are t-adapated, so the expectation operator drops out. Now consider the first line. Here, we know that $\nu_{t+1+s,k}$ is date $t+1+s-k$ adapted. Therefore, it drops out for all $s<k$. When $s \geq k$, we know that $\mathbb{E}_{t+1+s-k}\pi_{t+1+s}$ is a date $t+1+s-k \geq t+1$ adapted constant, which is the same as $\nu_{t+1+s,k}$. Therefore by law of iterated expectations for $s \geq k$,

$$\mathbb{E}_t \frac{\partial \nu_{t+1+s,k}}{\partial \theta_t}(\pi_{t+1+s} - \mathbb{E}_{t+1+s-k}\pi_{t+1+s}) = \mathbb{E}_t \frac{\partial \nu_{t+1+s,k}}{\partial \theta_t}\mathbb{E}_{t+1+s-k}(\pi_{t+1+s} - \mathbb{E}_{t+1+s-k}\pi_{t+1+s}) = 0$$

Therefore, the entire first line is zero, and we are left with

$$\mathbb{E}_t \frac{\partial \mathcal{W}_{t+1}(\theta^{t+1})}{\partial \theta_t} = \sum_{s=0}^{K-1} \beta^s \nu_{t+1+s,1+s} \frac{d\mathbb{E}_t \pi_{t+1+s}}{d\theta_t}.$$

Finally, we can now go back and substitute in for our equation for the derivative of $\mathcal{W}_t$. Substituting in,

$$\begin{aligned}\frac{\partial \mathcal{W}_t(\theta^t)}{\partial \theta_t} =& \frac{\partial U_t}{\partial \theta_t} + \beta \mathbb{E}_t \left[\mathcal{W}_{t+1}(\theta^{t+1}) \frac{\partial f(\theta_{t+1}|\theta_t)/\partial \theta_t}{f(\theta_{t+1}|\theta_t)} \middle| \theta_t \right] \\ &+ \sum_{k=1}^{K} \frac{\partial U_t}{\partial \mathbb{E}_t \pi_{t+k}} \frac{d\mathbb{E}_t \pi_{t+k}}{d\theta_t} + \beta \sum_{s=0}^{K-1} \beta^s \nu_{t+1+s,1+s} \frac{d\mathbb{E}_t \pi_{t+1+s}}{d\theta_t}\end{aligned}$$

Substituting in $\nu_{t+1+s,1+s} = -\frac{1}{\beta^{1+s}} \frac{\partial U_t}{\partial \mathbb{E}_t \pi_{t+1+s}}$, we get

$$\begin{aligned}\frac{\partial \mathcal{W}_t(\theta^t)}{\partial \theta_t} =& \frac{\partial U_t}{\partial \theta_t} + \beta \mathbb{E}_t \left[\mathcal{W}_{t+1}(\theta^{t+1}) \frac{\partial f(\theta_{t+1}|\theta_t)/\partial \theta_t}{f(\theta_{t+1}|\theta_t)} \middle| \theta_t \right] \\ &+ \sum_{k=1}^{K} \frac{\partial U_t}{\partial \mathbb{E}_t \pi_{t+k}} \frac{d\mathbb{E}_t \pi_{t+k}}{d\theta_t} - \sum_{s=0}^{K-1} \frac{\partial U_t}{\partial \mathbb{E}_t \pi_{t+1+s}} \frac{d\mathbb{E}_t \pi_{t+1+s}}{d\theta_t}\end{aligned}$$

and the second line drops to zero (the two sums are equivalent replacing $k = 1+s$). Thus, we obtain

$$\frac{\partial \mathcal{W}_t(\theta^t)}{\partial \theta_t} = \frac{\partial U_t}{\partial \theta_t} + \beta \mathbb{E}_t \left[\mathcal{W}_{t+1}(\theta^{t+1}) \frac{\partial f(\theta_{t+1}|\theta_t)/\partial \theta_t}{f(\theta_{t+1}|\theta_t)} \middle| \theta_t \right]$$

which is the required envelope condition. This completes the proof.

A.8 Proof of Proposition 11

Recall that we have

$$\pi_t = \kappa y_t + (\beta\gamma + \tilde{\beta})\mathbb{E}_t \pi_{t+1} + \tilde{\beta}\mathbb{E}_t \Big[\sum_{s=1}^{\infty} \tilde{\delta}^s \pi_{t+1+s} \Big].$$

From Proposition 7 for $k \geq 1$,

$$\nu_{t+k,k} = -\frac{1}{\beta^k} \frac{\partial \mathcal{U}_t}{\partial y_t} \frac{\partial y_t}{\partial \mathbb{E}_t \pi_{t+k}}.$$

Thus, we can write for $k > 1$,

$$\begin{aligned}
\nu_{t+k,k} &= \frac{1}{\beta^{k-1}} \frac{\frac{\partial y_t}{\partial \mathbb{E}_t \pi_{t+k}}}{\frac{\partial y_t}{\partial \mathbb{E}_t \pi_{t+1}}} \nu_{t+1,1} \\
&= \frac{1}{\beta^{k-1}} \frac{\tilde{\beta}\tilde{\delta}^{k-1}}{\beta\gamma + \tilde{\beta}} \nu_{t+1,1} \\
&= \beta^* \delta^{*(k-1)} \nu_{t+1,1}
\end{aligned}$$

where $\delta^* = \frac{\tilde{\delta}}{\beta}$ and $\beta^* = \frac{\tilde{\beta}}{\beta\gamma + \tilde{\beta}}$, completing the proof.

A.9 Proof of Proposition 12

Replicating the proof of Proposition 3 and including a penalty function Ω_t, we have additional terms in our value function at date t,

$$-\gamma \frac{\partial P_t}{\partial \theta_t} - \gamma \mathbb{E}_t \left[\sum_{k=1}^{\infty} \beta^k \frac{\partial P_{t+k}}{\partial \theta_t} \middle| \theta_t \right].$$

For the envelope condition to be satisfied, these terms must equal the unaccounted for information rent, $-\gamma\omega_t$ from equation (19). Thus we must construct penalties satisfying

$$\frac{\partial P_t}{\partial \theta_t} + \mathbb{E}_t \left[\sum_{k=1}^{\infty} \beta^k \frac{\partial P_{t+k}}{\partial \theta_t} \middle| \theta_t \right] = \omega_t.$$

Now from here, we can totally differentiate the recursive formulation of $\overline{P}_t$ to write

$$\frac{\partial \overline{P}_t}{\partial \theta_t} = \frac{\partial P_t}{\partial \theta_t} + \beta \mathbb{E}_t [\frac{\partial \overline{P}_{t+1}}{\partial \theta_t} | \theta_t] + \beta \mathbb{E}_t [\overline{P}_{t+1} \frac{\partial f(\theta_{t+1}|\theta_t)/\partial \theta_t}{f(\theta_{t+1}|\theta_t)} | \theta_t].$$

Thus combining with the required condition above, we have

$$\frac{\partial \overline{P}_t}{\partial \theta_t} = \omega_t + \beta \mathbb{E}_t [\overline{P}_{t+1} \frac{\partial f(\theta_{t+1}|\theta_t)/\partial \theta_t}{f(\theta_{t+1}|\theta_t)} | \theta_t].$$

The final expression comes from integrating. Thus we have constructed the required penalty function to satisfy the envelope condition.

A.10 Proof of Proposition 13

Integrating the Envelope Condition (equation 6), we obtain integral incentive compatibility

$$\mathcal{W}_t(\theta^t) = \int_{\underline{\theta}}^{\theta_t} \frac{\partial U_t(\theta^{t-1}, s_t)}{\partial s_t} ds_t + \beta \int_{\underline{\theta}}^{\theta_t} \mathbb{E}_t \left[\mathcal{W}_{t+1} \frac{\partial f_t(\theta_{t+1}|s_t)/\partial s_t}{f_t(\theta_{t+1}|s_t)} | s_t \right] ds_t \tag{28}$$

Integral incentive compatibility relates the total date-t utility to the central bank to two information rents. Note that due to shock persistence, the central bank earns information rents not only due to the effect on current flow utility, but also on the conditional probability distribution.[43]

Integral incentive compatibility (28) gives a Bellman representation to the value function, in terms of only the allocation rule. We can re-express this Bellman equation in sequence form by iterating the Bellman equation forward. Doing so, we obtain the following result characterizing this sequence representation.

Lemma 15. *The value function $\mathcal{W}_t$ can be represented as*

$$\mathcal{W}_t(\theta^t) = \mathbb{E}_t\left[\sum_{s=0}^{\infty} \beta^s B_t^s(\theta^{t+s}) \middle| \theta_t\right] \qquad \forall t,$$

where B_t^s is given by

$$B_t^s(\theta^{t+s}) = \prod_{k=0}^{s-1} \frac{1}{f_{t+k}(\theta_{t+k+1}|\theta_{t+k})} \times$$
$$\int_{s_t \le \theta_t, \dots, s_{t+s} \le \theta_{t+s}} \frac{\partial U_{t+s}(\theta^{t-1}, s_t, \dots, s_{t+s})}{\partial s_{t+s}} \prod_{k=0}^{s-1} \frac{\partial f_{t+k}(\theta_{t+k+1}|s_{t+k})}{\partial s_{t+k}} ds_{t+s} \dots ds_t.$$

Proof. Suppose that we take the Bellman equation:

$$\mathcal{W}_t(\theta^t) = \int_{\underline{\theta}}^{\theta_t} \frac{\partial U_t(\theta^{t-1}, s_t)}{\partial s_t} ds_t + \beta \int_{\underline{\theta}}^{\theta_t} E_t\left[\mathcal{W}_{t+1} \frac{\partial f_t(\theta_{t+1}|s_t)/\partial s_t}{f_t(\theta_{t+1}|s_t)} | s_t\right]$$

And iterate it forward once. Iterating forward once, we obtain:

$$\mathcal{W}_t(\theta^t) = \int_{\underline{\theta}}^{\theta_t} E_t\left[\frac{\partial U_t(\theta^{t-1}, s_t)}{\partial s_t} ds_t + \frac{\partial f_t\left(\theta_{t+1}|s_t\right)/\partial s_t}{f_t\left(\theta_{t+1}|s_t\right)} \beta \left[\int_{\underline{\theta}}^{\theta_{t+1}} \frac{\partial U_t(\theta^{t-1}, s_t, s_{t+1})}{\partial s_{t+1}} + E_{t+1} \mathcal{W}_{t+2} \frac{f_{t+1}(\theta_{t+2}|s_{t+1})/\partial s_{t+1}}{f_{t+1}(\theta_{t+2}|s_{t+1})}\right.\right.$$

Iterating forward, suppose that we define the following recursive operator. In particular, we define:

$$\mathcal{B}_t^0(g, \theta) = \int_{\underline{\theta}}^{\theta} g ds_t$$

Note that for the function $g_t^0 = \frac{\partial U_t(\theta^{t-1}}{\partial s_t}$, we have that $\mathcal{B}_t^0$ is the first term in the infinite series defining $\mathcal{W}_t$.

And suppose we define next:

$$\mathcal{B}_t^1(g, \theta) = \int_{\underline{\theta}}^{\theta} E_t\left[\frac{\partial f_t(\theta_{t+1}|s_t)/\partial s_t}{f_t(\theta_{t+1}|s_t)} g \middle| s_t\right] ds_t$$

[43] Recall that we have normalized the date 0 outside option to zero.

Consider the function $g_t^1 = \int_{\underline{\theta}}^{\theta_{t+1}} \frac{\partial U_{t+1}(\theta^{t-1}, s_t, s_{t+1})}{\partial s_{t+1}} ds_{t+1}$. Taking the function $\mathcal{B}_t^1\left(g_t^1, \theta_t\right)$ and multiplying by β, we obtain the second term in the infinite series for $\mathcal{W}_t$.

From here, we define a recursive operator. Consider a function g_t^s that is a date $t+s$ adapted function. We define the operator:

$$\mathcal{B}_t^2\left(g_t^2, \theta_t\right) = \mathcal{B}_t^1\left(\mathcal{B}_{t+1}^1\left(g_t^2, \theta_{t+1}\right), \theta_t\right)$$

So that we have:

$$\mathcal{B}_t^2\left(g_t^2, \theta_t\right) = \int_{\underline{\theta}}^{\theta_t} E_t\left[\frac{\partial f_t(\theta_{t+1}|s_t)/\partial s_t}{f_t(\theta_{t+1}|s_t)} \int_{\underline{\theta}}^{\theta_{t+1}} E_{t+1}\left[\frac{\partial f_{t+1}(\theta_{t+2}|s_{t+1})/\partial s_{t+1}}{f_{t+1}(\theta_{t+2}|s_{t+1})} g_t^2\left(s_{t+1}, \theta_{t+2}\right)\middle| s_{t+1}\right] ds_{t+1}\middle| s_t\right] ds_t$$

Which, when $g_t^2(s_t, s_{t+1}, \theta_{t+2}) = \int_{\underline{\theta}}^{\theta_{t+2}} \frac{\partial U_{t+2}\left(\theta^{t-1}, s_t, s_{t+1}, s_{t+2}\right)}{\partial s_{t+2}} ds_{t+2}$ and multiplied by β^2, gives us the next term in the infinite series defining $\mathcal{W}_t$.

Continuosly defining these recursive operators as such, and defining functions $g_t^s\left(s_t, ..., s_{t+s-1}, \theta_{t+s}\right) = \int_{\underline{\theta}}^{\theta_{t+s}} \frac{\partial U_{t+s}\left(\theta^{t-1}, s_t, ..., s_{t+s}\right)}{\partial s_{t+s}}$, we obtain the infinite series that characterizes $\mathcal{W}_t$.

In other words, we can construct such recursive operators. From here, we look to simplify these operators. Let us start from the operator $\mathcal{B}_t^1\left(g, \theta_t\right)$. In particular, we have:

$$\begin{aligned}
\mathcal{B}_t^1\left(g, \theta_t\right) &= \int_{\underline{\theta}}^{\theta_t} E_t\left[\frac{\partial f_t(\theta_{t+1}|s_t)/\partial s_t}{f_t(\theta_{t+1}|s_t)} g\left(s_t, \theta_{t+1}\right)\middle| s_t\right] ds_t \\
&= \int_{\underline{\theta}}^{\theta_t} \int_{\theta_{t+1}} \frac{\partial f_t(\theta_{t+1}|s_t)}{\partial s_t} g(s_t, \theta_{t+1}) d\theta_{t+1} ds_t \\
&= \int_{\theta_{t+1}} \left[\int_{\underline{\theta}}^{\theta_t} \frac{\partial f_t(\theta_{t+1}|s_t)}{\partial s_t} g(s_t, \theta_{t+1}) ds_t\right] d\theta_{t+1} \\
&= \int_{\theta_{t+1}} \frac{\left[\int_{\underline{\theta}}^{\theta_t} \frac{\partial f_t(\theta_{t+1}|s_t)}{\partial s_t} g(s_t, \theta_{t+1}) ds_t\right]}{f_t(\theta_{t+1}|\theta_t)} f_t(\theta_{t+1}|\theta_t) d\theta_{t+1} \\
&= E_t\left[\frac{1}{f_t(\theta_{t+1}|\theta_t)}\left[\int_{\underline{\theta}}^{\theta_t} \frac{\partial f_t(\theta_{t+1}|s_t)}{\partial s_t} g(s_t, \theta_{t+1}) ds_t\right]\middle| \theta_t\right]
\end{aligned}$$

In particular, as applied to the function $g_t^1 = \int_{\underline{\theta}}^{\theta_{t+1}} \frac{\partial U_{t+1}(\theta^{t-1}, s_t, s_{t+1})}{\partial s_{t+1}} ds_{t+1}$, we obtain:

$$\mathcal{B}_t^1\left(g, \theta_t\right) = E_t\left[\frac{1}{f_t(\theta_{t+1}|\theta_t)}\left[\int_{\underline{\theta}}^{\theta_t} \int_{\underline{\theta}}^{\theta_{t+1}} \frac{\partial U_{t+1}(\theta^{t-1}, s_t, s_{t+1})}{\partial s_{t+1}} \frac{\partial f_t(\theta_{t+1}|s_t)}{\partial s_t} ds_{t+1} ds_t\right]\middle| \theta_t\right]$$

Which is of the form in the Lemma.

Now, let us consider the second operator. We have:

$$\mathcal{B}_t^2\left(g, \theta_t\right) = \mathcal{B}_t^1\left(\mathcal{B}_{t+1}^1\left(g, \theta_{t+1}\right), \theta_t\right)$$

Recall that the simplified operator above expresses:

$$\mathcal{B}_t^1(g,\theta_t) = E_t\left[\frac{1}{f_t(\theta_{t+1}|\theta_t)}\left[\int_{\underline{\theta}}^{\theta_t}\frac{\partial f_t(\theta_{t+1}|s_t)}{\partial s_t}g(s_t,\theta_{t+1})ds_t\right]\middle|\,\theta_t\right]$$

In other words, we have along history (θ^{t-1},s_t):

$$\mathcal{B}_{t+1}^1(g,\theta_{t+1}) = E_{t+1}\left[\frac{1}{f_{t+1}(\theta_{t+2}|\theta_{t+1})}\left[\int_{\underline{\theta}}^{\theta_{t+1}}\frac{\partial f_{t+1}(\theta_{t+2}|s_{t+1})}{\partial s_{t+1}}g(s_t,s_{t+1},\theta_{t+2})ds_{t+1}\right]\middle|\,\theta_{t+1}\right]$$

And applying this into the operator defining $\mathcal{B}_t^2$, we obtain:

$$\begin{aligned}
\mathcal{B}_t^2(g,\theta_t) &= E_t\left[\frac{1}{f_t(\theta_{t+1}|\theta_t)}\left[\int_{\underline{\theta}}^{\theta_t}\frac{\partial f_t(\theta_{t+1}|s_t)}{\partial s_t}\mathcal{B}_{t+1}^1(g,\theta_{t+1})\,ds_t\right]\middle|\,\theta_t\right]\\
&= E_t\left[\frac{1}{f_t(\theta_{t+1}|\theta_t)}\left[\int_{\underline{\theta}}^{\theta_t}\frac{\partial f_t(\theta_{t+1}|s_t)}{\partial s_t}E_{t+1}\left[\frac{1}{f_{t+1}(\theta_{t+2}|\theta_{t+1})}\left[\int_{\underline{\theta}}^{\theta_{t+1}}\frac{\partial f_{t+1}(\theta_{t+2}|s_{t+1})}{\partial s_{t+1}}g(s_t,s_{t+1},\theta_{t+2})ds_{t+1}\right]\middle|\,\theta_t\right.\right.\right.\\
&= E_tE_{t+1}\left[\frac{1}{f_t(\theta_{t+1}|\theta_t)}\left[\int_{\underline{\theta}}^{\theta_t}\frac{\partial f_t(\theta_{t+1}|s_t)}{\partial s_t}\left[\frac{1}{f_{t+1}(\theta_{t+2}|\theta_{t+1})}\left[\int_{\underline{\theta}}^{\theta_{t+1}}\frac{\partial f_{t+1}(\theta_{t+2}|s_{t+1})}{\partial s_{t+1}}g(s_t,s_{t+1},\theta_{t+2})ds_{t+1}\right]\middle|\,\theta_t\right.\right.\right.\\
&\overbrace{=}^{\text{LIE}} E_t\left[\frac{1}{f_t(\theta_{t+1}|\theta_t)}\frac{1}{f_{t+1}(\theta_{t+2}|\theta_{t+1})}\left[\int_{\underline{\theta}}^{\theta_t}\int_{\underline{\theta}}^{\theta_{t+1}}\frac{\partial f_t(\theta_{t+1}|s_t)}{\partial s_t}\frac{\partial f_{t+1}(\theta_{t+2}|s_{t+1})}{\partial s_{t+1}}g(s_t,s_{t+1},\theta_{t+2})ds_{t+1}ds_t\right]\middle|\,\theta_t\right]
\end{aligned}$$

And substituting in $g_t^2 = \int_{\underline{\theta}}^{\theta_{t+2}}\frac{\partial U_{t+2}\left(\theta^{t-1},s_t,s_{t+1},s_{t+2}\right)}{\partial s_{t+2}}ds_{t+2}$, we get the next expression from the Lemma. From here, the result follows from repeated iteration. ■

Lemma 15 allows us to represent the principal's optimization problem in a tractable way. Given an allocation rule for inflation, we use the characterization of the value function in Lemma 15 as well as the Bellman equation to characterize the transfer rule which implements the allocation,

$$T_t = \mathcal{W}_t - U_t - \beta\mathbb{E}_t[\mathcal{W}_{t+1}|\theta_t].$$

We can then substitute the implementing taxes into the government's utility function, and obtain the following result characterizing the relaxed social planning problem.

Lemma 16. *The relaxed social planning problem can be written as*

$$\max_{\{\pi_t\}}\mathbb{E}_{-1}\left[\sum_{t=0}^{\infty}\beta^t\left[-\frac{\kappa}{1+\kappa}B_0^t + U_t\right]\right],$$

where B_0^t is given as in Lemma 15. The implementing transfer rule is given by

$$T_t = \mathcal{W}_t - U_t - \beta \mathbb{E}_t[\mathcal{W}_{t+1}|\theta_t],$$

where $\mathcal{W}_t$ is given as a function of the allocation rule as in Lemma 15.

Proof. For any allocation rule, T_t provides the implementation. Recall that the government's welfare is given by:

$$\max E_{-1}\left[\sum_{t=0}^{\infty}\beta^t U_t - \kappa T_t\right],$$

Recall that bank welfare is given by:

$$\mathcal{W}_0 = E_0 \sum_{t=0}^{\infty}\left[\beta^t U_t + T_t\right]$$

In other words, we always have:

$$-E_0 \sum_{t=0}^{\infty} T_t = E_0 \sum_{t=0}^{\infty}\beta^t U_t - \mathcal{W}_0$$

Substituting in above, by Law of Iterated Expectations we obtain the planning problem:

$$\max E_{-1}\left[-\kappa \mathcal{W}_0 + \sum_{t=0}^{\infty}\beta^t(1+\kappa)U_t\right],$$

and where lastly, we use Lemma 4 substitute in for $\mathcal{W}_0$ to obtain the result. ■

Lemma 16 provides a characterization of the relaxed social planning problem, subject to integral incentive compatibility. We are now ready to characterize the optimal allocation in Proposition 13.[44]

Recall that our objective function for the second-best optimization problem was given by:

$$\max \int_{\theta_0}\left[\sum_{t=0}^{\infty}\beta^t\left[-\frac{\kappa}{1+\kappa}\mathcal{B}_0^s\left(g_0^t,\theta_0\right) + U_t\left(\pi_t,\pi_{t+1},\theta_t,\theta_t\right)\right]\right] dF_0(\theta_0)$$

Note that given the optimal mechanism implements truthful reporting, we may substitute in $\tilde{\theta}_t = \theta_t$.

Recall further the simplified form of the operators:

$$\mathcal{B}_t^s = E_t\left[\prod_{k=0}^{s-1}\frac{1}{f_{t+k}(\theta_{t+k+1}|\theta_{t+k})}\int_{s_t\le\theta_t,\ldots,s_{t+s}\le\theta_{t+s}}\frac{\partial U_{t+s}(\theta^{t-1},s_t,\ldots,s_{t+s})}{\partial s_{t+s}}\prod_{k=0}^{s-1}\frac{\partial f_{t+k}(\theta_{t+k+1}|s_{t+k})}{\partial s_{t+k}}ds_{t+s}\ldots ds_t\middle|\theta_t\right]$$

[44] We characterize the optimal allocation assuming that π_t is interior.

Now, denote the *realized value* of the operator $\mathcal{B}_0^t$ by:

$$B_0^t(\theta^t) = \prod_{k=0}^{t-1} \frac{1}{f_k(\theta_{k+1}|\theta_k)} \int_{s_0 \leq \theta_0,...,s_t \leq \theta_t} \frac{\partial U_t(s_0, ..., s_t)}{\partial s_t} \prod_{k=0}^{t-1} \frac{\partial f_k(\theta_{k+1}|s_k)}{\partial s_k} ds_t ... ds_0$$

So that $B_0^t(\theta^t)$ is a random variable derived from the history θ^t of shocks. Given the definition of this random variable, denote E_{-1} to be the beginning-of-period-0 expectation, not conditional on the information θ_0. From here, we can rewrite the objective function of the government as:

$$\max E_{-1} \left[\sum_{t=0}^{\infty} \beta^t \left[-\frac{\kappa}{1+\kappa} B_0^t(\pi_t, \pi_{t+1}, \theta_t|\theta^{t-1}) + (1+\kappa) U_t(\pi_t, \pi_{t+1}, \theta_t|) \right] \right]$$

From here, consider the optimal choice of inflation $\pi_t(z^t)$, for a realized history $\theta^t = z^t$ of shocks. Note that the solution can be written in the form (for $t \geq 1$):

$$\frac{\partial U_{t-1}}{\partial \pi_t(z^t)} f(z^{t-1}) + \beta \frac{\partial U_t}{\partial \pi_t(z^t)} f(z^t) = \frac{\kappa}{1+\kappa} E_{-1} \sum_{s=t-1}^{t} \beta^{s-(t-1)} \frac{d}{d\pi_t(z^t)} B_0^s(\pi_s, \pi_{s+1}, \theta_s|\theta^s)$$

So that all that remains is to characterize the derivatives of B_0^s with respect to $\pi_t(z^t)$. When $s = t$, we have:

$$\frac{d}{dz^t} B_0^t(\theta^t) = \frac{d}{\pi_t(z^t)} \left[\prod_{k=0}^{t-1} \frac{1}{f_k(\theta_{k+1}|\theta_k)} \int_{s_0 \leq \theta_0,...,s_t \leq \theta_t} \frac{\partial U_t(s_0, ..., s_t)}{\partial s_t} \prod_{k=0}^{t-1} \frac{\partial f_k(\theta_{k+1}|s_k)}{\partial s_k} ds_t ... ds_0 \right]$$

Note that $\pi_t(z^t)$ appears in $\frac{\partial U_t(s_0,...,s_t)}{\partial s_t}$ only along the path given by $s_0 = z_0, s_1 = z_1, \ldots, s_t = z_t$. Essentially then, this derivative at a single point $\pi_t(z^t)$ comes down to extracting the derivative along that path under the integral. The derivative along that path is then given by:

$$\frac{d}{dz^t} B_0^t(\theta^t) = \mathbf{1}_{z_0 \leq \theta_0,...,z_t \leq \theta_t} \prod_{k=0}^{t-1} \frac{1}{f_k(\theta_{k+1}|\theta_k)} \frac{\partial^2 U_t}{\partial z_t \partial \pi_t(z^t)} \prod_{k=0}^{t-1} \frac{\partial f_k(\theta_{k+1}|z_k)}{\partial z_k}$$

Note the subtlety that the θ's are preserved, as the realization of the random history, whereas the s's are replaced by z's, as the path under the integrals that leads to the history z^t under the integrals. It is worth remembering then, when we substitute into the expectation, that θ_t is a random variable, and z^t is (fixed) the history being differentiated along, and so is not a random variable.

Note that by exactly the same logic, we obtain $\forall t \geq 2$

$$\frac{d}{dz^t} B_0^{t-1}(\theta^{t-1}) = \mathbf{1}_{z_0 \leq \theta_0,...,z_{t-1} \leq \theta_{t-1}} \prod_{k=0}^{t-2} \frac{1}{f_k(\theta_{k+1}|\theta_k)} \frac{\partial^2 U_{t-1}}{\partial z_{t-1} \partial \pi_t(z^t)} \prod_{k=0}^{t-2} \frac{\partial f_k(\theta_{k+1}|z_k)}{\partial z_k}$$

As a result, the right-hand side of the first-order condition becomes $\forall t \geq 2$

$$\begin{aligned}
\frac{1+\kappa}{\kappa}\text{RHS} &= E_{-1}\sum_{s=t-1}^{t}\frac{d}{d\pi_t(z^t)}B_0^s(\pi_s,\pi_{s+1},\theta_s|\theta^s) \\
&= E_{-1}\left[\mathbf{1}_{z_0\leq\theta_0,\ldots,z_{t-1}\leq\theta_{t-1}}\prod_{k=0}^{t-2}\frac{1}{f_k(\theta_{k+1}|\theta_k)}\frac{\partial^2 U_{t-1}}{\partial z_{t-1}\partial\pi_t(z^t)}\prod_{k=0}^{t-2}\frac{\partial f_k(\theta_{k+1}|z_k)}{\partial z_k}\right] \\
&\quad+\beta E_{-1}\left[\mathbf{1}_{z_0\leq\theta_0,\ldots,z_t\leq\theta_t}\prod_{k=0}^{t-1}\frac{1}{f_k(\theta_{k+1}|\theta_k)}\frac{\partial^2 U_t}{\partial z_t\partial\pi_t(z^t)}\prod_{k=0}^{t-1}\frac{\partial f_k(\theta_{k+1}|z_k)}{\partial z_k}\right] \\
&= \frac{\partial^2 U_{t-1}}{\partial z_{t-1}\partial\pi_t(z^t)}E_{-1}\left[\mathbf{1}_{z_0\leq\theta_0,\ldots,z_{t-1}\leq\theta_{t-1}}\prod_{k=0}^{t-2}\frac{1}{f_k(\theta_{k+1}|\theta_k)}\frac{\partial f_k(\theta_{k+1}|z_k)}{\partial z_k}\right] \\
&\quad+\frac{\partial^2 U_t}{\partial z_t\partial\pi_t(z^t)}\beta E_{-1}\left[\mathbf{1}_{z_0\leq\theta_0,\ldots,z_t\leq\theta_t}\prod_{k=0}^{t-1}\frac{1}{f_k(\theta_{k+1}|\theta_k)}\prod_{k=0}^{t-1}\frac{\partial f_k(\theta_{k+1}|z_k)}{\partial z_k}\right]
\end{aligned}$$

Where here, we applied the fact that we have chosen a specific history z^t, so that the cross-partials above are *not* random variables, but rather are specific realizations of those random variables. By contrast, the part inside the expectation corresponds to histories which contain these specific histories, and so are random variables.

Now, consider these two expectations. Now, we define $\Omega_t(z^t)$ by:

$$\begin{aligned}
\Omega_t(z^t) &\equiv E_{-1}\left[\mathbf{1}_{z_0\leq\theta_0,\ldots,z_t\leq\theta_t}\prod_{k=0}^{t-1}\frac{1}{f_k(\theta_{k+1}|\theta_k)}\prod_{k=0}^{t-1}\frac{\partial f_k(\theta_{k+1}|z_k)}{\partial z_k}\right] \\
&= \int_{z_t}^{\overline{\theta}}\int_{z_{t-1}}^{\overline{\theta}}\ldots\int_{z_0}^{\overline{\theta}}\prod_{k=0}^{t-1}\frac{\partial f_k(\theta_{k+1}|z_k)}{\partial z_k}f(\theta_0)d\theta_t\ldots d\theta_0 \\
&= \int_{z_t}^{\overline{\theta}}\frac{\partial f_k(\theta_t|z_{t-1})}{\partial z_k}\left[\int_{z_{t-1}}^{\overline{\theta}}\ldots\int_{z_0}^{\overline{\theta}}\prod_{k=0}^{t-1}\frac{\partial f_k(\theta_{k+1}|z_k)}{\partial z_k}f(\theta_0)d\theta_{t-1}\ldots d\theta_0\right]d\theta_t \\
&= \int_{z_t}^{\overline{\theta}}\frac{\partial f_k(\theta_t|z_{t-1})}{\partial z_{t-1}}\Omega_{t-1}(z^{t-1})d\theta_t \\
&= \Omega_{t-1}\left(z^{t-1}\right)\int_{z_t}^{\overline{\theta}}\frac{\partial f_k(\theta_t|z_{t-1})}{\partial z_{t-1}}d\theta_t
\end{aligned}$$

Which is well-defined for all $t \geq 1$. However, it requires an initial condition $\Omega_0(z^0)$. It is helpful to define this initial condition in the date 1 FOC. Note that at date 1, we have:

$$\mathcal{B}_0^{t-1}(\theta^{t-1}) = \mathcal{B}_0^0(\theta^0) = \int_{\underline{\theta}}^{\theta_0}\frac{\partial U_0}{\partial s_0}ds_0$$

So that we have $\frac{d}{d\pi_t(z^t)}\mathcal{B}_0^{t-1}(\theta^{t-1}) = \mathbf{1}_{z_0 \leq \theta_0}\frac{\partial U_0}{\partial \pi_1(z^1)}$. In particular then, the expectation is simply:

$$E_{-1}\left[\mathbf{1}_{z_0 \leq \theta_0}\right] = \int_{z_0}^{\overline{\theta}} f(\theta_0)d\theta_0 = 1 - F(z_0)$$

So that we have initial condition $\Omega_0(z^0) = 1 - F(z_0)$.

This gives us a state space reduction property, where we can fully determine Ω_t from Ω_{t-1} and z_{t-1} by a recursive sequence, where the initial value is $\Omega_0(z^0) = 1 - F(z_0)$.

From here, we can substitute back into the FOCs:

$$(1+\kappa)\left[\frac{\partial U_{t-1}}{\partial \pi_t(z^t)}f(z^{t-1}) + \beta\frac{\partial U_t}{\partial \pi_t(z^t)}f(z^t)\right] = \kappa\left[\Omega_{t-1}(z^{t-1})\frac{\partial^2 U_{t-1}}{\partial z_{t-1}\partial \pi_t(z^t)} + \beta\Omega_t(z^t)\frac{\partial^2 U_t}{\partial z_t \partial \pi_t(z^t)}\right]$$

From here, it is helpful to divide through by $f(z^{t-1})$:

$$(1+\kappa)\left[\frac{\partial U_{t-1}}{\partial \pi_t(z^t)} + \beta\frac{\partial U_t}{\partial \pi_t(z^t)}f(z_t|z_{t-1})\right] = \kappa\left[\frac{\Omega_{t-1}(z^{t-1})}{f(z^{t-1})}\frac{\partial^2 U_{t-1}}{\partial z_{t-1}\partial \pi_t(z^t)} + \beta\frac{\Omega_t(z^t)}{f(z^t)}\frac{\partial^2 U_t}{\partial z_t \partial \pi_t(z^t)}f(z_t|z_{t-1})\right]$$

And from here, we define $\Gamma_t(z^t) = \frac{\Omega_t(z^t)}{f(z^t)}$. Note that we have:

$$\Gamma_t(z^t) = \frac{\Omega_t(z^t)}{f(z^t)} = \frac{\Omega_{t-1}(z^{t-1})}{f(z^t)}\frac{\int_{z_t}^{\overline{\theta}}\frac{\partial f_k(\theta_t|z_{t-1})}{\partial z_k}d\theta_t}{f(z_t|z_{t-1})} = \Gamma_{t-1}(z^{t-1})\frac{\int_{z_t}^{\overline{\theta}}\frac{\partial f_k(\theta_t|z_{t-1})}{\partial z_k}d\theta_t}{f(z_t|z_{t-1})}$$

Giving us our key result for $t \geq 1$.

Note that the relevant initial condition is $\Gamma_0 = \frac{1-F(z_0)}{f(z_0)}$. This is the standard term in evaluating the virtual value in static mechanism design problems, and it is not surprising that it appears here. What it notable is that this term appears in the *date 1* optimality condition, in addition (as we will see) to the date-0 one. This is because of the time consistency problem.

Lastly, we can evaluate the FOC for π_0. In π_0, there is no time consistency element, and we are left with the simple tradeoff between current π and transfers. Repeating the steps from above, we obtain the simple condition

$$\frac{\partial U_0}{\partial \pi_0} = \frac{\kappa}{1+\kappa}\Gamma_0(z^0)\frac{\partial^2 U_0}{\partial z_0 \partial \pi_0}$$

which is a standard virtual value condition. This gives the full result.

This concludes the proof.

A.10.1 Second best with Average Transfers

In the baseline model, we impose the assumption that the outside option takes the form $\mathcal{W}_0(\theta^0) \geq 0$. We might alternatively have expressed this in the form

$$\int_{\theta_0} \mathcal{W}_0(\theta^0) f(\theta_0|\theta_{-1}) d\theta_0 \geq 0$$

The core difference between these two assumptions from a modeling perspective is on the timing of information arrival versus the participation decision. Under the baseline assumption, either θ_0 is already known to the central bank, or the central bank has the opportunity to revert to the outside option after learning θ_0. Under the second assumption, θ_0 is not known to the central bank, and the central bank does not have the option to revert to the outside option after learning it.

Under this alternative structure, the optimality of the dynamic inflation target returns. In particular, implementable allocations are still defined as in Lemma 15, while the transfer rule is $T_t(\theta^t) = \mathcal{W}_t - U_t - \beta \mathbb{E}_t \left[\mathcal{W}_{t+1}|\theta_t\right]$. The average participation constraint implies that we have

$$0 = E_{-1}\mathcal{W}_0 = \mathbb{E}_{-1} \sum_{t=0}^{\infty} \beta^t (U_t + T_t),$$

which is markedly different from the baseline model. In particular, substituting this expression into social welfare, we obtain the social optimization problem

$$\max_{\{\pi_t\}} \mathbb{E}_{-1} \sum_{t=0}^{\infty} \beta^t (1+\kappa) U_t$$

implying that the optimal allocation rule is constrained efficient. From here, we obtain the optimality of the dynamic inflation target.

Proposition 17. *Suppose that the participation constraint takes the form*

$$\int_{\theta_0} \mathcal{W}_0(\theta^0) f(\theta_0|\theta_{-1}) d\theta_0 \geq 0$$

Then, the optimal mechanism is a dynamic inflation target, and yields the constrained efficient allocation.

Proof. The proof follows immediately. The objective function is to maximize social welfare and hence the optimal allocation is the full-information Ramsey allocation. The mechanism that implements this is the dynamic inflation target, with a lump sum transfer at date 0 to achieve a binding participation constraint. ■

The intuition behind Proposition 17 is straight-forward: under the average constraint, the government can capture the full social surplus and simply reduce the average transfer to the central bank at date 0 to satisfy the participation constraint. This implies that the government chooses the mechanism and allocation that maximize social surplus, which is the dynamic inflation target.

A.11 Proof of Corollary 14

The proof follows immediately from the definition of Γ_t, which is equal to zero if $\theta_t \in \{\underline{\theta}, \overline{\theta}\}$. When $\Gamma_t = 0$, the allocation rule is constrained efficient for all Γ_{t+k}, $k \geq 1$, so the optimal mechanism reverts to constrained efficiency, which is implemented by the dynamic inflation target.

B Further Applications

In this Appendix, we develop several additional applications. Appendix B.1 revisits our main applications of Section 4 under costly enforcement. Appendix B.2 revisits the classical conservative central banker Rogoff (1985). We study persistent cost-push shocks in Appendix B.3 revisits the canonical New Keynesian consensus on inflation targets with persistent cost-push shocks. Appendix B.4 revisits persistent changes in the natural interest rate r_t^* (Section 4.2) and allows for arbitrary EIS, $\sigma > 0$.

B.1 Costly Enforcement: Main Applications Revisited

It is instructive to revisit how costly enforcement (Section 6.2) affects the optimal allocation rule in our main applications to lower bound spells (Section 4.1), changing natural interest rates r_t^* (Section 4.2), and changing slope of the Phillips curve (Section 4.3). We show that costly enforcement calls for *less* aggressive unconventional policies (e.g., forward guidance) when the economy experiences a lower bound spell, while it calls for *more* aggressive policies (e.g., raising the inflation target) in response to a decline in r^*. We document competing effects in the case of flattening Phillips curve that can call more more or less aggressive policies.

Lower bound spells. In the case of lower bound spells (Section 4.1), reduced-form preferences satisfy $\frac{\partial U_t}{\partial \pi_t \partial \theta_t} = 0$ and $\frac{\partial U_t}{\partial \mathbb{E}_t \pi_{t+1} \partial \theta_t} = c_0$ for a constant $c_0 > 0$. This reflects that high $\theta_t > 0$ corresponds to a binding lower bound and thus makes it valuable to promise more *future* inflation. However, because θ_t reflects a benefit of increasing the nominal rate and increasing inflation π_t does not directly increase the nominal rate, changes in the allocation rule π_t does not generate an information rent for the central bank at date t. This leads to an allocation rule given by

$$\frac{\partial U_t}{\partial \pi_t} = \nu_{t-1} + K\Gamma_{t-1} c_0,$$

where the RHS is λ_{t-1}.

Suppose that lower bound spells are persistent and higher current types signal higher future types (monotone likelihood). Then, $\Gamma_{t-1} > 0$, so that the optimal mechanism prescribes a marginal value of contemporaneous inflation that is *higher* under costly enforcement, all else equal. Intuitively, higher inflation expectations increase *past* information rents through by pushing the economy away from the lower bound. This leads the planner to prefer a less aggressive policy for promoting future inflation.

Declining r^*. In the case of changes in the natural rate $\theta_t = r^*_t$ (Section 4.2), reduced-form preferences satisfy $\frac{\partial U_t}{\partial \pi_t \partial \theta_t} = 0$ and $\frac{\partial U_t}{\partial \mathbb{E}_t \pi_{t+1} \partial \theta_t} = -c_1$ for a constant $c_1 > 0$. Intuitively, high θ_t corresponds to being further from the effective lower bound, which reduces the value of raising inflation expectations to get away from the ELB. The allocation rule under the optimal mechanism is given by

$$\frac{\partial U_t}{\partial \pi_t} = \nu_{t-1} - K\Gamma_{t-1}c_1,$$

where again the RHS is λ_{t-1}. The rule thus parallels the rule under lower bound spells, but in the opposite direction. This is because higher inflation expectations now *reduce* past information rents to the central bank, rather than raising them, by pushing the economy away from the ELB. This leads the planner to prefer a *more* aggressive policy for promoting future inflation.

These results highlight a surprising contrast between the two lower bound applications: costly enforcement calls for less aggressive unconventional policies in a lower bound spell, but for more aggressive policies in response to changing a natural rate. Intuitively once the economy is already in a lower bound spell, boosting inflation expectations raises central bank information rents by disproportionately benefiting central banks in worse conditions. By contrast if the economy has not yet hit the lower bound, boosting inflation expectations reduces central bank information rents by pushing all central banks away from the lower bound, reducing the value to the central bank of private information about r^*.

Flattening Phillips curve. In the case of a flattening Phillips curve (Section 4.3), reduced-form preferences satisfy $\frac{\partial U_t}{\partial \pi_t \partial \theta_t} = \frac{1}{\kappa}$ and $\frac{\partial U_t}{\partial \mathbb{E}_t \pi_{t+1} \partial \theta_t} = -\frac{\beta}{\kappa}$. This reflects that a flattening Phillips curve (higher θ_t) increases the value of stimulating current output through current inflation, but also increases the cost of higher inflation expectations that depress output. The optimal allocation rule is given by

$$\frac{\partial U_t}{\partial \pi_t} = \nu_{t-1} + \frac{K}{\kappa}\Delta\Gamma_t,$$

where again the RHS is λ_{t-1} and where $\Delta\Gamma_t \equiv \Gamma_t - \Gamma_{t-1}$. There are two competing effects from costly enforcement On the one hand, high θ_t means that the central bank's value of stimulating output rises, promoting higher current inflation. This increases information rents to the central bank and calls for lower inflation. On the other hand, high inflation also increases past inflation expectations, which reduces information rents to past central banks and calls for higher inflation (similarly to the r^* application). The relative magnitude of the two effects is determined by $\Delta\Gamma_t$, that is the change in the persistent portion of the information rent earned by the central bank between the two dates. From Proposition 13, we can write

$$\Delta\Gamma_t = \Gamma_{t-1}\left(h(\theta_t|\theta_{t-1})\mathbb{E}_t\left[\Lambda(s_t|\theta_{t-1})\middle|s_t \geq \theta_t\right] - 1\right).$$

where recall that $h^{-1}(\theta_t|\theta_{t-1}) = \frac{1-F(\theta_t|\theta_{t-1})}{f(\theta_t|\theta_{t-1})}$ is the inverse hazard rate and $\Lambda(s_t|\theta_{t-1}) = \frac{\partial f(s_t|\theta_{t-1})/\partial\theta_{t-1}}{f(\theta_t|\theta_{t-1})}$ is the derivative of the likelihood ratio. We know that the expected likelihood ratio derivative is zero at $\theta_t = \underline{\theta}$ while we know that the inverse hazard rate is zero at $\theta_t = \overline{\theta}$. Thus local to the two extremes of the shock distribution, we have $\Delta\Gamma_t < 0$ and hence the optimal mechanism promotes *higher* inflation. Interestingly, this suggests a tendency in this environment for the backward looking information rent to dominate the contemporaneous information rent, and hence generate a tendency to promote higher inflation to generate lower past information rents (at the expense of promoting higher current information rents). In the interior, two common assumptions are a nonincreasing inverse hazard rate and a monotone (increasing) likelihood ratio (higher past types signal high future types). These have competing effects on the response to a flattening Phillips curve. Intuitively, a lower inverse hazard rate reduces current virtual surplus whereas a higher likelihood ratio increases virtual surplus.

B.2 Revisiting Rogoff's Inflation-Conservative Central Banker

We ask whether dynamic inflation targets can be implemented by inflation-conservative central bankers in the spirit of Rogoff (1985). In particular, our inflation-conservative central banker places a greater penalty on inflation than the government. After appropriate intertemporal rearrangement of terms, we represent this by assuming central bank preferences equal to

$$V_t = U_t - c(\pi_t - \mathbb{E}_{t-1}[\pi_t|\tilde{\theta}_{t-1}]),$$

where as before U_t denotes the preferences of society and the government, and where c is the constant linear cost to the conservative central banker of inflation exceeding firm inflation expectations.[45] We obtain the following result.

Proposition 18. *With an inflation-conservative central banker, the full-information Ramsey allocation can then be implemented by a dynamic inflation target with* $b_{t-1} = \nu_{t-1} - c$.

Proposition 18 demonstrates that the appointment of an inflation-conservative central banker does not obviate the fundamental need for a dynamic inflation target. Intuitively, the inflation-conservative central banker applies a constant penalty to inflation, given by c. In the presence of persistent shocks, the target flexibility ν_t of the dynamic inflation target changes over time.

[45] This is a special case of preference disagreement in Appendix C.2.

While an inflation-conservative central bank raises target flexibility on average, in the sense that $b_{t-1} = \nu_{t-1} - c < \nu_{t-1}$, the total implied inflation penalty $b_{t-1} + c$ is ν_{t-1} just as before. The inflation target mechanism that implements the full-information Ramsey allocation is still time-varying and responds to persistent shocks.

In the language of Svensson (1997b), however, appointing an inflation-conservative central banker can resolve *average* inflationary bias when c is set equal to the average value of ν_t in the stochastic steady state. When this average penalty is large (e.g., in the presence of a distorted steady state) but time variation in ν_t is small, approximating the dynamic inflation target with an inflation-conservative central bank may result in relatively small welfare losses.

Proposition 18 suggests that an alternative implementation of the dynamic inflation target might be to appoint new central bank chairs with appropriate inflation preferences in response to changes in ν_t. The inflation conservativeness of the central bank would then be time-varying and correspond to $c_t = \nu_{t-1}$. If in response to a shock at date $t-1$ the dynamic inflation target requires $\nu_{t-1} > \nu_{t-2}$, then a more dovish central banker at date $t-1$ should be replaced by a more hawkish central banker at t. Just as the dynamic inflation target must be updated one period in advance, the appointment of a new central banker would also be announced one period in advance.[46]

B.2.1 Proof of Proposition 18

The proof follows the same steps as in Proposition 3. The envelope condition is the same, given that the additional term $-c(\pi_t - \mathbb{E}_{t-1}[\pi_t|\tilde{\theta}_t])$ in V_t depends on reported types and not true types. From here, the value function at date t under our proposed mechanism given by

$$\begin{aligned}
\mathcal{W}_t(\theta^t) &= -c(\pi_t - \mathbb{E}_{t-1}\pi_t) + V_t + \beta\mathbb{E}_t\Big[\mathcal{W}_t(\theta^{t+1})|\theta_t\Big] \\
&= -(c + b_{t-1})(\pi_t - \mathbb{E}_{t-1}\pi_t) + U_t + \beta\mathbb{E}_t\Big[\mathcal{W}_t(\theta^{t+1})|\theta_t\Big] \\
&= -\nu_{t-1}(\pi_t - \mathbb{E}_{t-1}\pi_t) + U_t + \beta\mathbb{E}_t\Big[\mathcal{W}_t(\theta^{t+1})|\theta_t\Big]
\end{aligned}$$

which is the same value function as in the proof of Proposition 3 when evaluated at the constrained efficient allocation. Thus the result follows using the same proof as for Proposition 3.

[46] Importantly, just as a fixed central bank under the optimal mechanism was tasked with updating its own target, in an implementation with time varying conservativeness a central banker would be tasked with appointing her own replacement one period in advance (or at the least, would be responsible for naming her successor). However, this institutional arrangement is not typical (if used at all) in practice. For example, in the U.S. the president is tasked with appointing members of the Board of Governors, who must then be confirmed by the Senate.

B.3 Cost-Push Shocks, Flexible Inflation Targeting, and Price-Level Targeting

In this application, we study a persistent cost-push shock both with and without costly enforcement. This revisits the related full-information environment of Svensson and Woodford (2004) and studies the properties of the dynamic inflation target. Social welfare is characterized by a New Keynesian loss function around a non-distorted steady state, $\mathcal{U}_t(\pi_t, y_t, \theta_t) = -\frac{1}{2}\pi_t^2 - \frac{1}{2}\alpha(y_t - \theta_t)^2$. For simplicity, we set the slope of the Phillips curve to be $\kappa = 1$. Internalizing the NKPC (10) into the loss function yields reduced-form preferences

$$U(\pi_t, \mathbb{E}_t\pi_{t+1}, \theta_t) = -\frac{1}{2}\pi_t^2 - \frac{1}{2}\alpha(\pi_t - \beta E_t\pi_{t+1} - \theta_t)^2. \tag{29}$$

Note that θ_t is a cost-push shock in the usual sense: higher θ_t means higher current inflation is needed in order to maintain the same output loss. We assume the cost-push shock satisfies $\mathbb{E}_t\theta_{t+1} = \rho\theta_t$, where $0 \leq \rho \leq 1$ is its persistence. The following result characterizes the dynamic inflation target.

Proposition 19. *The dynamic inflation target that implements the full-information Ramsey allocation is*

$$\nu_t = \gamma_1\nu_{t-1} + \gamma_2\theta_t$$

$$\tau_t = -(1-\gamma_1)\gamma_1\nu_{t-1} + \gamma_2(\gamma_1 - 1 + \rho)\theta_t,$$

where $0 \leq \gamma_1 \leq 1$ *does not depend on* ρ*, and* $\gamma_2 \geq 0$ *increases in* ρ*. Optimal inflation sets* $\pi_t = \nu_t - \nu_{t-1}$*.*

Proposition 19 specializes the dynamic inflation target of Proposition 3 to the cost-push shock application. In response to a positive and persistent innovation in the shock, i.e., a high θ_t realization, the central bank updates both parameters of the target for the next period. First, the target flexibility *decreases* in the sense that ν_t rises. This happens because the cost-push shock leads to a larger output gap today, increasing the inflationary bias of the central bank.

Second, the response of the target level is ambiguous and depends on the shock persistence. When shocks are not persistent, a cost-push shock is followed by a *lower* target level. As shocks become more persistent, there is a critical level $\rho^* = 1 - \gamma_1$ after which the central bank raises the target level instead. This result reflects the common intuition of the cost-push shock model: The central bank would like to promise low future inflation to improve the contemporaneous inflation-output trade-off; as shocks become more persistent, however, it also wants to promise higher future inflation to mitigate future expected cost-push shocks.

The target also decreases as the *previous* period's target flexibility parameter ν_{t-1} rises. This reflects the history dependency: a high past inflationary bias leads to a desire for low inflation today, which in turn leads to a desire for low inflation tomorrow. This means that the increase in ν_t serves as a force for future deflationary pressures. Finally, contemporaneous inflation unambiguously rises in response to a positive cost-push shock. It is interesting to note that the target flexibility is *always* more responsive to a contemporaneous cost-push shock than its flexibility, since we have $-1 < \gamma_1 - 1 + \rho < 1$.

Costly enforcement. With costly enforcement, note that we have $\frac{\partial U_t}{\partial \pi_t \partial \theta_t} = \frac{1}{2}\alpha$ and $\frac{\partial U_t}{\partial \mathbb{E}_t \pi_{t+1} \partial \theta_t} = -\frac{1}{2}\alpha\beta$. The impacts are analogous to a flattening Phillips curve, and means we can write

$$\frac{\partial U_t}{\partial \pi_t} = \nu_{t-1} + \frac{1}{2}\frac{K}{\alpha}\Delta\Gamma_t$$

Thus relative to the Ramsey solution, the optimal mechanism adjusts the allocation trading off two effects on information rents. On the one hand, higher expected inflation reduces *past* information rents by increasing costs of inflation for central banks that experience large past cost push shocks. On the other hand, higher contemporaneous inflation increases *current* information rents by reducing costs of large contemporanous cost push shocks. The optimal allocation rule trades off these two effects. As once again $\Delta\Gamma_t < 0$ local to the boundaries of the shock distribution, particularly large or particularly small cost push shocks at date t lead past information rents to dominate, and calls for a *more* aggressive inflation response today in order to reduce historical information rents. Interestingly, this amplifies the response of inflation to a large cost push shock, pushing the allocation rule closer to the policy under discretion.

B.3.1 Proof of Proposition 19

Given reduced from preferences are

$$U(\pi_t, \mathbb{E}_t \pi_{t+1}, \theta_t) = -\frac{1}{2}\pi_t^2 - \frac{1}{2}\alpha(\pi_t - \beta E_t \pi_{t+1} - \theta_t)^2$$

then we have

$$\frac{\partial U_t}{\partial \pi_t} = -\pi_t - \alpha(\pi_t - \beta E_t \pi_{t+1} - \theta_t)$$

$$\frac{\partial U_{t-1}}{\partial \mathbb{E}_{t-1} \pi_t} = \beta\alpha(\pi_{t-1} - \beta E_{t-1}\pi_t - \theta_{t-1}).$$

By definition, we have

$$\nu_{t-1} = -\frac{1}{\beta}\frac{\partial U_{t-1}}{\partial \mathbb{E}_{t-1}\pi_t} = -\alpha(\pi_{t-1} - \beta E_{t-1}\pi_t - \theta_{t-1}).$$

Therefore, we can write the FOC for the full-information Ramsey allocation, $\frac{\partial U_t}{\pi_t} = \nu_{t-1}$, equivalently as

$$-\pi_t - \nu_t = \nu_{t-1}$$

or in other words, $\pi_t = \nu_t - \nu_{t-1}$. Combined with the definition of ν_{t-1} and the initial condition $\nu_{-1} = 0$, this gives us a complete system.

Suppose that $\mathbb{E}_t\theta_{t+1} = \rho\theta_t$, where $\rho = 1$ corresponds to full persistence. We thus think of cost push shocks as reverting towards zero. We guess and verify a linear solution

$$\nu_t = \gamma_1\nu_{t-1} + \gamma_2\theta_t.$$

Given this conjecture, we know from the FOC that

$$\pi_t = (\gamma_1 - 1)\nu_{t-1} + \gamma_2\theta_t.$$

Using the definition of ν_t,

$$\nu_t = -\alpha\pi_t + \alpha\beta\mathbb{E}_t\pi_{t+1} + \alpha\theta_t,$$

we substitute in the expression for π_t and our conjecture for ν_{t+1} to obtain

$$\nu_t = -\alpha\Big(\nu_t - \nu_{t-1}\Big) + \alpha\beta\Big((\gamma_1 - 1)\nu_t + \gamma_2\mathbb{E}_t\theta_{t+1}\Big) + \alpha\theta_t.$$

Now using that $\mathbb{E}_t\theta_{t+1} = \rho\theta_t$ and rearranging, we get

$$\nu_t = \frac{\alpha}{1 + \alpha + (1 - \gamma_1)\alpha\beta}\nu_{t-1} + \frac{\alpha\Big(\beta\gamma_2\rho + 1\Big)}{1 + \alpha + (1 - \gamma_1)\alpha\beta}\theta_t$$

Thus coefficient matching, we have the system of equations

$$\frac{\alpha}{1 + \alpha + (1 - \gamma_1)\alpha\beta} = \gamma_1$$

$$\frac{\alpha\Big(\beta\gamma_2\rho + 1\Big)}{1 + \alpha + (1 - \gamma_1)\alpha\beta} = \gamma_2$$

The first equation is defined solely in terms of γ_1. Thus taking it and rearranging, we obtain the quadratic

$$\alpha\beta\gamma_1^2 - \gamma_1(1 + \alpha + \alpha\beta) + \alpha = 0.$$

This quadratic has two roots, with the upper root being explosive since $\beta < 1$ implies $\gamma_1^+ > 1$. Thus selecting the non-explosive root gives $0 \leq \gamma_1 \leq 1$, where

$$\gamma_1 = \frac{1 + \alpha + \alpha\beta - \sqrt{(1 + \alpha + \alpha\beta)^2 - 4\alpha^2\beta}}{2\alpha\beta}.$$

Note that to see why this root lies between 0 and 1, the quadratic above equals $\alpha > 0$ for $\gamma_1 = 0$

and equals $-1 < 0$ when $\gamma_1 = 1$.

Given that $0 \leq \gamma_1 \leq 1$, we can solve for γ_2 using the second equation, which gives

$$\gamma_2 = \frac{\gamma_1}{1 - \beta\rho\gamma_1},$$

which is positive since $\beta\rho\gamma_1 \leq 1$. Thus we have our solution. Given this solution, the parameters of the target are

$$\nu_t = \gamma_1 \nu_{t-1} + \gamma_2 \theta_t$$

and

$$\begin{aligned}
\tau_t &= \mathbb{E}_t \pi_{t+1} \\
&= (\gamma_1 - 1)\nu_t + \gamma_2 \rho \theta_t \\
&= -(1 - \gamma_1)\gamma_1 \nu_{t-1} + \gamma_2(\gamma_1 - 1 + \rho)\theta_t
\end{aligned}$$

B.4 Application: r^* Revisited and the Commitment Curve

We revisit the application to persistent changes in the natural interest rate r_t^* (Section 4.2) but allow for $\sigma > 0$. The realized nominal interest rate is

$$i_t = \mathbb{E}_t \pi_{t+1} + \theta_t + \sigma \left[\mathbb{E}_t y_{t+1} - y_t \right] - \epsilon_t.$$

Intuitively, an expected rise in the output gap means household consumption is expected to rise, raising the nominal interest rate and pushing the central bank away from the ELB. Similar to Section 4.2, we can write $i_t = i_t^* - \epsilon_t$ and write the welfare losses $v(i_t^*)$ from the ELB. In this case with $\sigma > 0$, we have a change in the definition of i_t^* to

$$i_t^* = -\sigma \pi_t + \Big(1 + \sigma(1 + \beta) \Big) \mathbb{E}_t \pi_{t+1} - \sigma \beta \mathbb{E}_t \pi_{t+2} + \theta_t,$$

which reflects internalizing the NKPC to substitute out the output gap. Intuitively, higher inflation today, π_t, increases output today and so reduces the required nominal rate. Higher inflation π_{t+1} both directly increases the nominal rate and indirectly increases it by stimulating output y_{t+1}. Conversely, higher inflation π_{t+1} depresses output y_{t+1} and so reduces the nominal rate.

Consider the shape of the commitment curve. Recall that we have $U_t = -\frac{1}{2}\pi_t^2 - \frac{1}{2}\hat{\alpha}\Big(\pi_t - \beta \mathbb{E}_t \pi_{t+1}\Big)^2 + v(i_t^*)$. We can write

$$\nu_{t+1,1} = \nu_{t+1,1}^y + \nu_{t+1,1}^i$$

where $\nu_{t+1,1}^y = -\frac{1}{2}\hat{\alpha}\Big(\pi_t - \beta \mathbb{E}_t \pi_{t+1}\Big)$ is the usual output gap component, and where $\nu_{t+1,1}^i =$

$-\left(v_0 - \beta v_1 i_t^*\right)\left(1 + \sigma(1+\beta)\right) < 0$ is the component coming from the effective lower bound. From here, we can show that

$$\nu_{t+2,2} = -\beta^* \nu_{t+1,1}^i$$

where $\beta^* = \frac{\sigma}{1+\sigma(1+\beta)} < 1$ is increasing in σ.

Intuitively, in this case the commitment curve can be decomposed into two components. The first component is the output gap commitment curve, where we have $\nu_{t+1,1}^y > 0$ and $\nu_{t+k,k}^y = 0$ for all $k > 1$. This corresponds to the standard one period commitment to stabilize the output gap. The second component is the *effective lower bound commitment curve*, where $\nu_{t+1,1}^i < 0$ and $\nu_{t+2,2}^i = -\beta^* \nu_{t+1,1}^i > 0$. The effective lower bound commitment curve switches signs precisely because of the different effects of inflation at different horizons.

C Further Extensions

C.1 Welfare Gains from a Dynamic Inflation Target

We characterize the potential welfare gains under a dynamic inflation target. Suppose that the central bank adopts a permanent, static target (ν^*, τ^*) instead of the dynamic inflation target of Proposition 3.[47] The following proposition describes the first-order welfare gains from moving from the static target to a dynamic inflation target.

Proposition 20. *To first order, the welfare gains in allocative efficiency from moving from a static target (ν^*, τ^*) to the dynamic inflation target (ν_{t-1}, τ_{t-1}) of Proposition 3 are*

$$\mathbb{E} \sum_{t=1}^{\infty} \beta^t \Bigg[\underbrace{\nu_{t-1}^* - \nu^*}_{\text{Cost of Excess Inflation}} \Bigg] \Bigg[\underbrace{\mathbb{E}_{t-1}\pi_t^* - \tau_{t-1}}_{\text{Amount of Excess Inflation}} \Bigg].$$

The first order welfare gains available from moving to a dynamic inflation target depend on two forces. The first, $\nu_{t-1}^* - \nu^*$, is the intertemporal variation in the time consistency problem under the static target (where ν_{t-1}^* is the time consistency wedge evaluated at the allocation obtained under the static target). When $\nu_{t-1}^* > \nu^*$, the time consistency problem is more severe than the slope imposed ν^*, and hence inflation is too high relative to the efficient tradeoff. In other words, the first term reflects the cost of excess inflation. The second term, $\mathbb{E}_{t-1}\pi_t^* - \tau_{t-1}$, is the difference between inflation expectations under the static target and inflation expectations under the dynamic target. High welfare gains are therefore available when a large excess time consistency problem, $\nu_{t-1}^* - \nu^*$, coincides with substantial excess inflation, $\mathbb{E}_{t-1}\pi_t^* - \tau_{t-1}$, relative to the constrained

[47] To simplify analysis, we will characterize welfare under a static target with full information, even though the dynamic inflation target implements the Ramsey allocation under incomplete information. This streamlines analysis because under a static target absent full information, the central bank's reporting constraints would be nontrivial due to information effects.

efficient inflation level. The dynamic inflation target thus allows welfare gains not only by allowing for greater inflation when the static target would be too severe, but also by allowing for lower inflation when the static target would be too flexible.

C.1.1 Proof of Proposition 20

To first order, the welfare gains of an inflation perturbation from the static target is

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{\partial U_t}{\partial \pi_t} d\pi_t + \frac{\partial U_t}{\partial \mathbb{E}_t \pi_{t+1}} d\pi_{t+1} \right].$$

From here, the first order condition of the central bank is $\nu^* = \frac{\partial U_t}{\partial \pi_t}$, while by definition $\frac{\partial U_t}{\partial \mathbb{E}_t \pi_{t+1}} = -\beta \nu_t^*$. We have $\frac{\partial U_0}{\partial \pi_0} = 0$, so that we have

$$\mathbb{E}_0 \sum_{t=1}^{\infty} \beta^t \left[\nu^* - \nu_{t-1}^* \right] d\pi_t.$$

Finally, we have $\mathbb{E}_{t-1} d\pi_t = \tau_{t-1} - \mathbb{E}_{t-1} \pi_t^*$, giving the result.

C.2 Preference Differences

We extend the costly enforcement model (Section 6.2) to allow for preference disagreement. Formally, the central bank has utility U_t but the government has utility $V_t(\pi_t, \mathbb{E}_t[\pi_{t+1}|\tilde{\theta}_t], \theta_t)$. Social preferences of the government are now

$$\max \mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t \left(V_t(\pi_t, \mathbb{E}_t[\pi_{t+1}|\tilde{\theta}_t], \theta_t) - \kappa T_t \right) \right]. \tag{30}$$

As before there is a central bank participation constraint. Define $K = \frac{\kappa}{1+\kappa}$ as before, and define *weighted reduced form preferences* to be

$$Z_t = (1-K)V_t + KU_t.$$

Weighted reduced form preferences average the preferences of the government and central bank. A higher weight is assigned to central bank preferences the more costly enforcement is, that is K rises in κ. The optimal mechanism can be described as follows.

Proposition 21. *The solution to an optimal allocation rule of the relaxed problem is given by the first-order conditions*

$$\frac{\partial Z_t}{\partial \pi_t} - K\Gamma_t \frac{\partial U_t}{\partial \theta_t \partial \pi_t} = \lambda_{t-1}^*$$

where $\lambda_{t-1}^* = -\frac{1}{\beta} \frac{\partial Z_{t-1}}{\partial \mathbb{E}_{t-1} \pi_t} + K\Gamma_{t-1} \frac{1}{\beta} \frac{\partial^2 U_{t-1}}{\partial \theta_{t-1} \partial \mathbb{E}_{t-1} \pi_t}$ *and* Γ_t *is defined as in Proposition 13.*

The optimal allocation rule of Proposition 21 is similar to that of Proposition 13, but with one important difference: the weighted preference Z_t replaces the planner's utility. Intuitively, the government places value on the lifetime utility to the central bank because promising higher lifetime value allows the government to extract more surplus in the form of transfers. Counterveiling this force is information rents, which are analogous to before and only depend on central bank preferences U_t. Intuitively, these terms only depend on central bank preferences as information rents accrue based on central bank preferences. Otherwise, the intuitions of Section 6.2 carry over.

It is helpful to illustrate two dichotomous cases. If $K = 0$ and enforcement is costless, we have $Z_t = V_t$ and hence the optimal allocation is the *government*'s Ramsey allocation. This follows intuitively: the government has no cost to designing a scheme that incentives the central bank to choose the government's preferred allocation. At the other extreme, if $K = 1$ then $Z_t = U_t$, that is to first order the planner only values transfers. Interestingly, the optimal allocation collapses to that of Proposition 13. Intuitively when the principal only cares about transfers, the principal on the one hand wants to make utility as high as possible to the agent in order to relax the central bank's participation constraint and extract larger transfers ex ante. On the other hand, the principal also internalizes that higher agent utility increasess agent information rents. This leads to the same allocation rule as in the case where principal and agent preferences are aligned except for transfers.

At intermediate values of K, the optimal allocation rule trades off the two extremes. On the one hand, the planner wishes to push the allocation closer to her Ramsey allocation, which increases her direct utility from allocations. At the same time, the planner wishes to push the allocation closer to the central bank's Ramsey allocation in order to relax the participation constraint and extract greater transfers. This leads to a balancing act determined by K, which encodes a relative weight the principal assigns to the different motivations.

As in Corollary 14, following $\theta_t \in \{\underline{\theta}, \overline{\theta}\}$ the optimal allocation reverts to the Ramsey allocation associated with weighted reduced-form preferences Z_t. If $K = 1$, then this allocation coincides with that of the dynamic inflation target.

C.2.1 Proof of Proposition 21

Observe that the integral envelope condition (28) still holds and implies Lemma 15 characterizes the central bank's value function, given central bank preferences have not changed. Thus the transfer rule is still given by $T_t = \mathcal{W}_t - U_t - \beta\mathbb{E}_t[\mathcal{W}_{t+1}|\theta_t]$. Thus we still have

$$-\mathbb{E}\sum_{t=0}^{\infty} T_t = \mathbb{E}\sum_{t=0}^{\infty}\beta^t U_t - \mathcal{W}_0$$

where $\mathcal{W}_0 = \mathbb{E}_0\left[\sum_{s=0}^{\infty}\beta^s B_0^s(\theta^s)\middle|\theta_0\right]$ Given the change in preferences, the government's objective function is now

$$\mathbb{E}\left[\sum_{t=0}^{\infty}\beta^t V_t - \kappa T_t\right]$$

thus substituting in the transfer rule and definition of $\mathcal{W}_0$, the government's objective function is

$$\mathbb{E}\left[\sum_{t=0}^{\infty}\beta^t\left[V_t+\kappa U_t-B_0^t\right]\right]$$

Finally dividing through by $1+\kappa$ and defining $K=\frac{\kappa}{1+\kappa}$ $(1-K=\frac{1}{1+\kappa})$, we obtain

$$\mathbb{E}\left[\sum_{t=0}^{\infty}\beta^t\left[(1-K)V_t+KU_t-KB_0^t\right]\right]$$

Thus we simply define $Z_t=(1-K)V_t+KU_t$ and the derivation proceeds exactly the same as before with Z_t replacing U_t as the government's effective utility function. This recovers the first order condition given and completes the proof.