Optimal Contracts and Inflation Targets Revisited*

Torsten Persson† and Guido Tabellini‡

Draft, May 27, 2024

Abstract

We revisit the optimal-contract approach to the design of monetary institutions, in the light of the Zero Lower Bound (ZLB) on interest rates and the resort to Quantitative Easing (QE) in recent years. Four of our lessons have not yet been incorporated in the practices of inflation targeting central banks. First, the optimal contract and the implied inflation-targeting regime should condition on being at the ZLB or out of it. Second – as already argued by others – the optimal inflation target should be raised to deal with the possibility of being at the ZLB, and more so the greater the risk of being there. But this qualitative lesson does not appear to warrant major quantitative changes of inflation targets. Third, the relevance of the ZLB suggests that it may be desirable to expand central-bank mandates to encompass financial stability, broadly defined, besides price and output stability. Fourth, accountability for inflation performance is a central mechanism in a successful monetary-policy framework. How exactly to change those mechanisms in practice is a new and difficult challenge, which is at least as important as the search for optimal policy rules that has attracted so much recent attention.

*Paper presented at the Riksbank Conference on "The quest for nominal stability. Lessons from three decades with inflation targeting", May 23-24, 2024, Stockholm. We gratefully acknowledge financial support from the Swedish Research Council for financial support, and research assistance by Anton Arbman Hansing. Dmytro Sergeyev, our discussants Don Kohn and Carl Walsh, and Michael Woodford as well as other conference participants, provided helpful comments, while Stefan Laséen Jesper Lindé, Dmytro Sergeyev, and Silvana Tenreyro gave us help and advice on data. The views expressed in this working paper are our own responsibility, and should not be interpreted as reflecting the views of the Executive Board of Sveriges Riksbank.

†IIES, Stockholm University

‡Department of Economics and IGIER, Bocconi University

1 Introduction

By the 1980s, many countries had suffered through two decades of high inflation, following the breakdown of the Bretton-Woods system. Gradually, academics as well as policymakers came to interpret the high inflation as an equilibrium outcome, since low-inflation objectives faced a major credibility problem with political-institutional roots (Kydland and Prescott 1977, Calvo 1978, Barro and Gordon 1983). The search for a new monetary-policy anchor converged to a regime in which targeting of low inflation was delegated to an independent central bank.

With five countries adopting inflation targeting over a mere three years – New Zealand (1990), Canada (1991), the UK (1992), Sweden (1993), and Australia (1993) – this regime caught on quickly and broadly from the 1990s. Until the recent bout of inflation, many developed countries had thus experienced almost three decades of low inflation supported by inflation targeting.

Given this real-world success, the attention of central bankers, as well as of academics, turned from political/institutional issues towards more technical issues, such as the design of Taylor rules (Taylor 1993), inflation-forecast targeting (Svensson 1997a), and forward guidance (Woodford 2005).

However, monetary policymakers have also had to grapple with new striking policy challenges during the inflation-targeting era. The most dramatic event was the great financial crisis 15 years ago, followed by a prolonged depression of demand. In the financial crisis and its aftermath, most central banks ran into a Zero Lower Bound (ZLB) – or effective lower bound – on their policy interest rates, although Japan had already faced it for a good decade. Many observers have attributed the ZLB to two drivers: declining actual and expected inflation rates and falling natural (Wicksellian) real interest rates.

As pointed out by Krugman (1997), the ZLB raised a new credibility problem in the opposite direction from the high inflation regimes: at the ZLB, the central bank would like to raise expected inflation, but it cannot do so. Policymakers approached this challenge by using new policy tools, chiefly expansions of central-bank balance sheets, known as Quantitative Easing (QE). Such tools came to even more aggressive use to fend off diving demands in the early COVID-19 pandemic. After the pandemic and during the Ukraine war, high inflation reappeared and by now it may or may not have come to an end.

Figure 1 shows the evolution of (monthly) inflation rates, policy rates, and natural rates at the monthly frequency over the last 40 years in four countries – the US, UK, Euro Area and Sweden. The graphs clearly illustrate three features of the data. (1) Inflation rates went down from their high levels in the 1970s and 1980s, seemingly stabilized around 2%, but shot up again in the last few years. (2) Policy rates were close to the ZLB during much of the time after the financial crisis. (3) At the same time, measures of the natural rate went down considerably and even turned negative.

Given this changing and more volatile environment, monetary policymakers have to cope with two credibility problems: central banks want to keep expected inflation anchored despite temporary inflationary shocks, and at the same time they want to keep up inflation expectations to stay out of the ZLB.

In this paper, we revisit the issues around the credibility of monetary policy and their institutional underpinnings in light of this dual challenge. Questions about optimal monetary policy when the ZLB sometimes binds have been extensively addressed. However, few contributions ask how to make optimal policies credible through the design of monetary institutions, when expected inflation is either too high or too low. This is precisely what we aim to do. Our goal is thus to reconsider the main lessons from earlier research on central-bank credibility, incentives, and inflation targeting (Rogoff 1985, Persson and Tabellini 1993, Walsh 1995) in light of the challenges raised by the ZLB. In addition, we touch upon how to analyze credibility issues around alternative policies like Quantitative Easing (QE). Afrouzi et al (2024) also bring back credibility and political-economy perspectives on monetary policy into the analysis of inflation, but do not more than mention the design of central-banking institutions.

More precisely, we study a setting where monetary policy faces two credibility problems: the traditional inflation bias associated with an expectations-augmented supply side (as in Kydland and Prescott 1977, Calvo 1978, Barro and Gordon 1983) and the deflation bias associated with controlling demand at the ZLB (as in Krugman 1997). Under commitment, the optimal policy rule finds a compromise between excessive inflation when the economy is off the ZLB and higher inflationary expectations when it is on the ZLB. Under discretion, equilibrium inflation can be too high or too low, depending on which credibility problem is more severe, and (unlike in the traditional literature on the inflation bias), equilibrium output can be driven away from the natural rate.

In this setting, we ask how a central bank that operates under discretion can be motivated to follow the optimal policy rule. An optimal state-contingent contract for central-bank leaders entails an inflation "tax" or "subsidy," depending upon circumstances, and implements the optimal policy. A more realistic non-fully state-contingent optimal contract resembles an inflation-targeting regime, but has three distinct features. First, it targets inflation only when the economy is out of the ZLB. Second, the inflation target in the contract is higher than socially optimal inflation. Third, the tolerance for inflation deviations – literally, the penalty on the central bank for missing the target – could be asymmetric in either direction.

The paper is organized as follows. In Section 2, we present our basic framework for monetary policy with a ZLB, and derive the optimal policy rule under commitment as well as the equilibrium policy under discretion. Section 3 is about the design of optimal contracts with the purpose of incentivizing a central bank under discretion to implement the optimal monetary-policy rule. In Section 4, we take stock of the findings and compare them to results and ideas in the existing literature as well as to central-bank practice. In the context of this discussion, we also sketch how to extend our basic model with alternative policy instruments like QE. Section 5 concludes. Formal proofs of all our results are relegated to an Appendix.

2 Two Credibility Problems in Monetary Policy

In this section, we propose a simple framework that entails two credibility problems in monetary policy. The first is the traditional problem to control expected inflation when aggregate supply responds to surprise inflation, as highlighted in the classical credibility literature (Kydland and Prescott 1977, Barro and Gordon 1983). The second is to control aggregate demand when the ZLB binds, as highlighted in a more recent literature (Krugman 1998, Eggertson and Woodford 2003). We use this framework to analyze the equilibrium policy outcomes for central banks that operate under commitment and discretion, respectively.

2.1 A Simple Model

We use Occam's razor and formulate the simplest possible model that allows for these two problems and permits us to derive analytical results. The supply side is captured by

$$x^s = \theta + (\pi - \pi^e) - \varepsilon \ , \tag{1}$$

where x is the level of output, π is realized inflation, and π^e expected inflation. Random variable ε is an *iid* supply shock with mean 0 and variance υ_ε, while θ is the natural rate of output, which is subject to iid shocks with mean $\bar{\theta}$ and variance υ_θ.

Aggregate demand is given by an IS-like equation, namely

$$x^d = \theta - \sigma(i - \pi^e - \rho) \ . \tag{2}$$

Here, i is the nominal interest rate, controlled by the central bank, $\sigma > 0$ is a parameter, while ρ is the Wicksellian natural – or equilibrium – real interest rate, a random variable which is subject to *iid* demand shocks.

For given inflationary expectations π^e, this model is recursive in that output is determined by (2) and inflation by (1). However, in a rational-expectations equilibrium, π^e will have to be consistent with a joint solution for inflation and output. We omit time periods, because we deal with a stationary, stochastic environment with iid shocks.

Inflation expectations and policy We define expected inflation of private agents – whether they set wages (prices) or express aggregate demand – by $\pi^e = E(\pi \mid I)$. The information set I always includes θ, but not ε and ρ.

This static model is equivalent to a stationary dynamic model under two assumptions. First, when forming expectations about future inflation, private agents do not observe the current value of ρ[1]. Second, the same value of expected inflation enters both aggregate supply and aggregate demand. This is satisfied,

[1] If not, the optimal policy rule for what to do in period t would depend on the realization of ρ , as in Eggertson and Woodford (2003).

if π^e reflects forward-looking price setting in (1), and forward-looking expected real rate of interest on saving or borrowing in (2).[2]

Private rational expectations notwithstanding, policymakers still have an information advantage and can stabilize fluctuations due to shocks ε and ρ, but not due to θ. A central bank that can commit to a policy rule may thus take into account how the rule affects expectations π^e, while a central bank that acts under discretion must take these expectations as given.

Social losses Although the model relations can be given microeconomic underpinnings, we take a shortcut and formulate society's loss function in any period directly over macroeconomic outcomes

$$E[L(\pi, x)] = \frac{1}{2}E[(\pi - \overline{\pi})^2 + \lambda(x - \overline{x})^2] \ . \tag{3}$$

Here, the expectation refers to all random variables $\{\theta, \varepsilon, \rho\}$. The policy targets $\overline{\pi} > 0$ and $\overline{x} \geq \theta$ are socially optimal levels of inflation and output.

Parameter $\lambda > 0$ is the relative weight on the deviations from these policy targets. In line with the analysis in Halac and Yared (2020) and Afrouzi et al (2024), we could also introduce fluctuations in λ and interpret these as "political shocks." If the value of λ would enter the private sector's information set, as we already assume, then such shocks would have similar effects as shocks to θ. Subsections 3.3 and 4.4 include further remarks on this prospective extension.

Assumptions on parameters and shocks To keep the analysis simple, we make four (sets of) additional assumptions. First, we assume $i \geq 0$, namely at 0 the central bank faces a hard constraint and cannot set negative interest rates. Thus, at the ZLB monetary policy has no tool to control aggregate demand. In Subsection 4.3, we discuss how to relax this assumption. There, we assume that, at $i = 0$, policymakers have access to an additional instrument, which is associated with some social costs (unlike a non-negative interest rate). This additional instrument can be thought of either as QE or as a negative interest rate.

Second, the random variable ρ that captures natural-rate shocks on the demand side takes on only two values: $\rho = r < 0$ with probability q, and $\rho = R > 0$ with probability $(1-q)$. Thus, its expected value is $E(\rho) = qr + (1-q)R > 0$.

Third, we also impose two parameter restrictions

$$-r > \overline{\pi} + \lambda(\overline{x} - \theta) > 0 \tag{\textbf{Assumption 1}}$$

$$q < \frac{1}{1 + \sigma(1 + \lambda)} \ . \tag{\textbf{Assumption 2}}$$

[2] It is also satisfied if $\pi^e(\theta)$ in (1) reflects sticky nominal wages set in the previous period and θ is constant over time, as this implies constant inflation expectations both under commitment and discretion.

Assumption 1 can be weakened, but is needed for the ZLB to be binding both under commitment and discretion. **Assumption 2** ensures that we have a well-defined equilibrium under discretion.

Fourth, the supports of random shocks ε and θ are sufficiently narrow that, in equilibrium, the ZLB binds if and only if aggregate demand is low ($\rho = r$).

With these assumptions, $i > 0$ if $\rho = R$ and $i = 0$ if $\rho = r$, for all realizations of ε and θ. We can thus think about q as the probability that the ZLB binds, which only hinges on the realization of natural-rate (demand) shocks. Later in the paper, we discuss the consequences of allowing (some) supply shocks ε to affect the probability of being at the ZLB.

State-dependent macroeconomic outcomes Let x^ρ and π^ρ denote equilibrium output and inflation in state $\rho = R, r$. Under our assumption that expectations are formed with information about θ, but not about ε and ρ, we can thus define expected inflation as a weighted average of its possible outcomes

$$\pi^e = qE(\pi^r|\theta) + (1-q)E(\pi^R|\theta) \ . \tag{4}$$

Consider first state $\rho = r$. Setting $i = 0$ and using (1) and (1), we get

$$\begin{aligned} x^r &= \theta + \sigma r + \sigma \pi^e & (5) \\ \pi^r &= (1+\sigma)\pi^e + \sigma r + \varepsilon \ . & (6) \end{aligned}$$

With **Assumption 1**, the low demand associated with $r < 0$ cuts realized inflation. Here, $r < 0$ brings down aggregate demand and pushes output below the natural rate, θ. Once we reach the point $i = 0$, monetary policy can only shape aggregate outcomes indirectly: expectations π^e do reflect policy in the good state $\rho = R$.

Consider that state and set $x^d = x^s$ to obtain.

$$\begin{aligned} x^R &= \theta - \sigma(i - R) + \sigma \pi^e & (7) \\ \pi^R &= (1+\sigma)\pi^e - \sigma(i - R) + \varepsilon \ . & (8) \end{aligned}$$

In this traditional setting, the central bank directly affects both inflation and output via its control of the policy rate i.

In the next two subsections, we discuss the full equilibrium outcomes under commitment and discretion, respectively. In each regime, the central bank sets its policy rate in an optimal way – though under different constraints – and the private sector forms its expectations accordingly. However, before going there we briefly compare our framework to two existing literatures.

Comparison to earlier research Our simple framework of unconstrained monetary policy (in state $\rho = R$), more or less coincides with that in the early credibility literature (Kydland and Prescott 1977, Barro and Gordon 1983) and its later extensions to institution design, as inflation targeting sustained by contracts for central bankers (Persson and Tabellini 1993, Walsh 1995, Svensson 1997b).

As in that literature, a central bank acting under discretion has to accept the social cost of higher expected and actual inflation due to its ex post inflation incentives. But as mentioned in the introduction, the institution-design approach has not yet been extended to settings where the ZLB sometimes binds (in state $\rho = r$). Our doing so is one way to describe the contribution of this paper.

The framework is also reminiscent of earlier modeling of monetary policy and the ZLB (Eggertson and Woodford 2003, Eggertson and Giannoni 2013, and Kiley and Roberts 2017). In their settings, like in ours, the central bank would like to raise expected inflation under the ZLB. This is because a higher π^e would cut the state-r real rate $i - \pi^e = -\pi^e$, and thus raise both π^r and x^r.

That research on the ZLB differs from our model in that the role of monetary policy in the stabilization of supply shocks is typically not considered.[3] Moreover, the models studied in that literature are dynamic: the evolution of ρ reflects realizations of a Markov stochastic process rather than stationary *iid* shocks. This may be realistic, as periods of low demand at the ZLB have been stretched out in time, at least until the early COVID-19 pandemic. However, as ρ becomes a state variable, the results rely on numerical solutions. Our stationarity assumption allows for analytical solutions and thus for more transparent theoretical results.

More importantly, existing work on the ZLB normatively evaluates the performance of different policy rules, but does not ask how different institutions shape their incentive compatibility. It is thus an open question whether and how desirable rules can be implemented. Our tackling of that question is another way to describe the contribution of this paper.

2.2 Optimal Policy under Commitment

In this subsection, we consider optimal policies under commitment, when the ZLB sometimes binds.

A rule for interest rates By assumption, the central bank can only use the interest rate in (good) state $\rho = R$. If it can indeed commit to an interest-rate rule, we may derive the optimal one. Without loss of generality, we write this policy rule as

$$i^R = E(i^R) + \iota^{R,\varepsilon}\varepsilon \ ,$$

where the linear form reflects the linear-quadratic (LQ) structure of the model and the policy problem. The first component captures the rule's average interest rate and the second component its response to the random, mean-zero supply shock ε. Note that this rule is state-contingent in two ways: it is conditional on high realized demand (natural-interest) shock $\rho = R$ and, given that, on the realized supply shock ε.

[3] However, see Eggertson (2012) who does study supply shocks at the ZLB.

When discussing the optimal policy, it is sometimes convenient to write the resulting macroeconomic outcomes in state $\rho = R, r$ on a similar form, namely

$$\pi^{\rho} = E(\pi^{\rho}) + \pi^{\rho,\varepsilon}\varepsilon \qquad \text{and} \qquad x^{\rho} = E(x^{\rho}) + x^{\rho,\varepsilon}\varepsilon \ .$$

The first component reflects the rule's average interest rate and the second its response to the mean-zero supply shock ε. The former term is thus governed by the expected interest rate $E(i^R)$ (as well as the privately observed shock θ). Why is there such a term in state r when the ZLB binds? As mentioned in the Subsection 3.1, average state-R policy indirectly affects macroeconomic outcomes at the ZLB – expected policy alters $E(\pi^R)$ and hence $\pi^e = qE(\pi^r) + (1-q)E(\pi^R)$, which changes the real interest rate $i^r - \pi^e = -\pi^e$. Internalizing the effects of state-R policies on state-r outcomes is, indeed, a key concern when designing the optimal rule.

Equilibrium outcomes with the optimal policy rule We now derive the optimal rule for the interest rate under commitment, when the ZLB binds with probability q, and compute the resulting macroeconomic outcomes. While the details are relegated to the Appendix, the bottom-line results are described in

Proposition 1 *The optimal policy rule under commitment gives outcomes*

$$\pi^{C,R}(q,r,\varepsilon) = \bar{\pi} - \delta^R(q)(\bar{\pi}+r) + \frac{\lambda}{1+\lambda}\varepsilon$$

$$\pi^{C,r}(q,r,\varepsilon) = \bar{\pi} - \delta^r(q)(\bar{\pi}+r) + \varepsilon$$

$$\pi^{C,e}(q,r) = \bar{\pi} - [\delta^R(q)(1-q) + q\delta^r(q)](\bar{\pi}+r) \geq \bar{\pi}$$

$$x^{C,R}(q,r,\varepsilon) = \theta - q[\delta^R(q) - \delta^r(q)](\bar{\pi}+r) - \frac{1}{1+\lambda}\varepsilon$$

$$x^{C,r}(q,r,\varepsilon) = \theta + (1-q)[\delta^R(q) - \delta^r(q)](\bar{\pi}+r) - \varepsilon,$$

where $\delta^R(q)$, $\delta^r(q)$ *and* $[\delta^R(q)(1-q) + q\delta^r(q)]$ *are all strictly increasing in* q *and satisfy* $\delta^r(q) < \delta^R(q)$, $\delta^r(0) < \delta^R(0) = 0$ *and* $0 \leq \delta^R(q) < 1$.

To describe the results, we begin with the special case when q – the chance of a negative demand shock $\rho = r$ – is zero. Then, **Proposition 1** says that average policy implies $E(\pi^{C,R}(0,r,\varepsilon)) = \bar{\pi} = \pi^{C,e}(q,r)$ and $E(x^{C,R}(0,r,\varepsilon)) = \theta$. Thus, the optimal rule anchors actual and expected inflation at its socially optimal rate, which keeps average output at its natural rate.

But when $q > 0$, optimal state-R inflation on average exceeds the target, $E(\pi^{C,R}(q,r,\varepsilon)) > \bar{\pi}$, and average output is above the natural rate, $E(x^{C,R}(q,r,\varepsilon)) > \theta$ (recall that by **Assumption 1,** $-r > \bar{\pi} + \lambda(\overline{x} - \theta) > 0$**.** These deviations carry a cost for society in state R. However, they buy a lower real rate in state r where the ZLB binds, which brings social benefits by raising average inflation π^r and average output x^r in the right direction. Importantly, expected inflation always exceeds the target, $\pi^{C,e}(q,r) > \overline{\pi}$ if $q > 0$. If q and the absolute value of r rise, so does the optimal inflation rate in state R. This drives up expected inflation, which makes the ZLB less distorting in the bad state. Moreover, these effects are interactive: the optimal inflation rate in state R and expected inflation respond more forcefully to a higher q the deeper is the recession with the bad demand shock – i.e., the more negative is parameter r.[4]

Finally, the optimal response to supply shocks ε is familiar from the existing literature (see e.g., Persson and Tabellini 2000, Ch.15). Thus, in state R, the optimal policy rule trades off the stabilization of output and inflation around their desired values at a rate that depends on their relative weight λ in the social-loss function. In the bad state r, nothing can be done to stabilize supply shocks**.**

Relation to existing work Our result that the ZLB makes it optimal to raise expected inflation – and hence to raise actual inflation when outside of the ZLB – is not novel. It has been stressed by Krugman (1997), Eggertson and Woodford (2003), Coibion et al. (2017), Kiley and Roberts (2017) among others. In the dynamic model considered in the ZLB literature, the optimal policy under commitment is also history-dependent: it should be more expansionary (and hence raise expected inflation) immediately after the exit from the ZLB. This dynamic implication is missing in our stationary model. We return to a comparison of our results and existing work in Section 4.

2.3 Equilibrium Policy under Discretion

In this subsection, we turn to equilibrium policies when the central bank can no longer commit to an optimal policy rule, but instead sets the interest rate under discretion.

Policymaking under discretion As before, the central bank freely alters its policy rate in the good state $\rho = R$. However, as the bank acts under discretion, it minimizes the loss function only in that state and for a particular realized supply shock ε, always taking expected inflation π^e as given. Thus, the central bank acts in an ex post optimal fashion, as in the classical credibility literature. The difference is that this failure to internalize the policy effects on expected inflation spills over to macro outcomes at the ZLB, which binds with probability q.

Outcomes under discretion As in the Subsection 3.2, we just state the equilibrium results, relegating the proof to the Appendix. We use a similar notation and format as in **Proposition 1**, except that we index all variables with a D, rather than a C, superscript. The results are

Proposition 2 *The equilibrium policy under discretion gives outcomes*

$$\pi^{D,R}(q,r,\varepsilon) = \bar{\pi} + \lambda(\bar{x} - \theta) + \alpha^R(q)[\bar{\pi} + \lambda(\bar{x} - \theta) + r] + \frac{\lambda}{1+\lambda}\varepsilon$$

[4] Formally, all derivatives above are multiplicative in r, so their cross derivatives with r have the same sign – implying e.g., $\pi_q^{C,R}(q, r + dr, \varepsilon) > \pi_q^{C,r}(q, r, \varepsilon)$ for $dr < 0$.

$$\pi^{D,r}(q,r,\varepsilon) = \bar{\pi} + \lambda(\bar{x}-\theta) + \alpha^r(q)[\bar{\pi} + \lambda(\bar{x}-\theta) + r] + \varepsilon$$

$$\pi^{D,e} = \bar{\pi} + \lambda(\bar{x}-\theta) + [(1-q)\alpha^R(q) + q\alpha^r(q)][\bar{\pi} + \lambda(\bar{x}-\theta) + r]$$

$$x^{D,R}(q,r,\varepsilon) = \theta + q[\alpha^R(q) - \alpha^r(q)][\bar{\pi} + \lambda(\bar{x}-\theta) + r] - \frac{1}{1+\lambda}\varepsilon$$

$$x^{D,r}(q,r,\varepsilon) = \theta - (1-q)[\alpha^R(q) - \alpha^r(q)][\bar{\pi} + \lambda(\bar{x}-\theta) + r] - \varepsilon,$$

where $\alpha^R(q)$,$\alpha^r(q)$ *and* $[(1-q)\alpha^R(q) + q\alpha^r(q)]$ *are all strictly increasing in* q, *and* $0 \leq \alpha^R(q) < \alpha^r(q)$, *with* $\alpha^R(0) = 0$ *and* $\alpha^R(q) > 0$ *if* $q > 0$.

Let us compare these with the results under commitment. As we did then, consider first the case when $q = 0$. Because $\alpha^R(0) = 0$, **Proposition 2** says that average discretionary policy gives $E(\pi^{D,R}(0,r,\varepsilon)) = \bar{\pi} + \lambda(\overline{x} - \theta)$ and $E(x^{D,R}(0,r,\varepsilon)) = \theta$. This is the familiar inflation-bias result in monetary policy. Inflation is higher than under commitment, by $\lambda(\bar{x}-\theta)$, which reflects the ex post incentive to inflate. As this incentive is anticipated by the private sector, inflation has no systematic effect on output.

Let us instead suppose that $q > 0$. Then, average inflation in good state R is unambiguously below $\bar{\pi} + \lambda(\bar{x}-\theta)$, because $\alpha^R(q) > 0 > \bar{\pi} + \lambda(\bar{x}-\theta) + r$, and it is decreasing in q and in the absolute value of r. Thus, the inflation bias shrinks with a higher q and a lower r. Since $\alpha^r(q) > \alpha^R(q)$, inflation is even lower in bad state r. Correspondingly, average output is above (below) the natural rate in state R (state r), due to, a positive (negative) inflation surprise. Formally, we have $E(x^{D,R}(q,r,\varepsilon)) > \theta > E(x^{D,r}(q,r,\varepsilon))$.

Importantly, expected inflation falls unambiguously short of $\bar{\pi} + \lambda(\bar{x}-\theta)$, but could be either below or above $\bar{\pi}$ depending on parameter values. The Appendix shows that $\pi^{D,e}$ is more likely below $\bar{\pi}$ the greater is q and the larger is the absolute value of r, while $\pi^{D,e}$ is more likely above $\bar{\pi}$ the greater is $\lambda(\bar{x}-\theta)$.

Intuitively, the two credibility problems push policy in opposite directions. Under discretion, the central bank fails to internalize the consequences of its policy for inflation expectations. One consequence is higher inflation in the good state, which in itself imposes a social cost by driving up inflationary expectations. But higher expected inflation also implies a social benefit in the bad state where the central bank is constrained. If the ZLB binds more often, this benefit weighs more heavily and the gap shrinks between socially optimal and discretionary state-R inflation. Whether $\pi^{D,e}$ is ultimately above or below $\bar{\pi}$, depends on the relative weight of these two incentive problems.

Another way to see this result appeals to the theory of the second best (Lipsey and Lancaster 1956). Specifically, as the state-R discretionary equilibrium is "distorted" by one credibility problem, and the state-r ZLB equilibrium by another, the two distortions combined may yield a better outcome.

Finally, let us note that monetary policy under discretion stabilizes supply shocks in an optimal fashion. Formally, state-dependent inflation and output depend on ε in the same way in Propositions 1 and 2. Nevertheless, output is more volatile under discretion than under commitment, because monetary policy stabilize demand shocks due to ρ in a less effective way under discretion (this can be seen by noting that $x^{D,R}(q,r,\varepsilon) - x^{D,r}(q,r,\varepsilon) > x^{C,R}(q,r,\varepsilon) - x^{C,r}(q,r,\varepsilon) > 0$).

In the next section, we ask how to design institutions that give policymakers under discretion appropriate incentives to implement the optimal interest-rate rule under commitment.

3 Optimal Institution Design

To solve the classical credibility problem, Rogoff (1985) proposed that monetary policy be delegated to a "conservative central banker", interpreted as an independent central bank with a mandate stressing inflation control. But he did not study how to optimally hold this banker accountable for her actions. Subsequent research by Persson and Tabellini (1993) and Walsh (1995) took this step by studying an outright institution-design problem: how to provide the right incentives in an optimal contract. In this section, we ask how the earlier institution-design results hold up in the presence of the ZLB. We begin by stating the general contracting problem, and then spell out its solution under two alternative constraints on how granular incentives can be designed.

3.1 Central-bank Contracts

Suppose the central bank has a mandate to pursue society's goals as captured by loss function, $L(\pi, x)$. But it also has a "performance contract" $T(\pi) \gtreqless 0$, a reward/penalty scheme, which satisfies the central-bank leadership's participation constraint and depends, positively or negatively, on realized inflation. Moreover, the central bank operates under discretion and full information: having observed all shocks, it sets i to minimize $L(\pi, x) + T(\pi)$. Hence, if $T(\pi) > 0$, this corresponds to a negative utility transfer to the central bank.

Take the case when T is increasing in π. Then, higher realized inflation corresponds to a larger expected utility loss to the central bank. In material terms, this can be seen as the result of a lower probability of receiving a bonus payment, or to be reappointed to the central-bank leadership. Taken less literally, it can be interpreted as an intrinsic loss – such as a lower social or personal reputation – e.g., due to public criticism after a careful outside ex post evaluation of policymaking. Section 4 provides a further discussion of real-world accountability. Here, we study the formal problem of picking the optimal form for $T(\pi)$ under different assumptions about constraints on its state contingency.

The formal problem Let the central bank face a set of state-contingent contractual payments $T(\pi^{\rho}, \theta, \varepsilon)$, a set of "taxes" or "subsidies" for realized inflation. In the general case, we allow these payments to be contingent on realized inflation under the two demand-side shocks, as well as the realized values of the supply-side shocks θ and ε.

In our model, the bank shapes all aggregate outcomes – directly or indirectly – solely by its policy stance in good state R (see the proofs of **Propositions 1** and **2**). This makes it natural to focus on state-R inflation. Hence, we study a contract, which is contingent just on realized state-R inflation $T^R(\pi^R, \theta, \varepsilon)$. Specifically, we write that contract in the following form

$$T^R(\pi^R, \theta, \varepsilon) = T_0^R + T_\pi^R(\theta, \varepsilon)\pi^R + T_{\pi\pi}^R(\theta, \varepsilon)(\pi^R)^2 + \ ...,$$

where the intercept T_0^R can be set freely, to satisfy any participation constraint for central-bank leaders. We ask how to optimally design the terms of this contract, so that a central bank operating under discretion chooses an inflation rate $\pi^R(q, r, \varepsilon)$ as close as possible to the ex-ante social optimum. Ideally, the optimal contract implements the optimal policy rule under commitment stated in **Proposition 1**.

3.2 State-contingent Contracts

In this subsection, we suppose that the bank contract can be fully state-contingent. That is, the contractual terms can depend on the realized demand shock ρ, as well as supply shocks θ and ε.

As we showed in Section 2, stabilization – the policy response to ε – is the same under commitment and discretion. Thus, there is no need to incentivize the bank's response to supply shocks ε, if the contract implements the optimal policy rule. Below, we show that this is indeed the case. Therefore, we do not make the contract terms contingent on ε and focus on contractual payments $T^R(\pi^R, \theta)$.

Next, we observe that, by **Propositions 1** and **2**, equilibrium inflation rates in state R under commitment and discretion differ by constant terms, which however depend on the realization of θ and other known parameters, such as q, r and $\bar{\pi}$. A *linear* contract over state-R inflation contingent on θ (and other parameters), $T^R(\pi^R, \theta) = T_0^R + T_\pi^R(\theta)$, should therefore induce a central bank under discretion to implement the optimal policy under commitment. The following proposition shows that this is indeed the case:

Proposition 3 *The optimal state-contingent contract has the form* $T(\pi^R, \theta) = T_0^R + T_\pi^R(\theta)\pi^R$, *where*

$$T_\pi^R(\theta) = \lambda(\bar{x} - \theta) + \beta^R(q)(\bar{\pi} + r), \tag{9}$$

and $\beta^R(q) = \alpha^R(q) + \delta^R(q) \geq 0$ *is strictly increasing in* q *with* $\beta^R(0) = 0$ *and* $\beta^R(q) > 0$ *if* $q > 0$, *and where* $\delta^R(q)$ *and* $\alpha^R(q)$ *are defined in Propositions 1 and 2. This contract induces a central bank operating under discretion to implement the optimal policy under commitment.*

As in Persson and Tabellini (1993) and Walsh (1995), the optimal state-contingent contract is linear in inflation. But this contract is state-dependent: it holds the central bank accountable for its performance, only if $\rho = R$. When $\rho = r$, the central bank is at the ZLB and inflation is beyond its control. The next section discusses the situation when the central bank can use an alternative (costly) policy instrument, like QE, to control aggregate demand at the ZLB.

Intuition and new insights To build intuition for **Proposition 3**, consider first the special case $q = 0$. Then, the proposition says that (as $\beta^R(0) = 0$) only the first term appears in the optimal contract $T_\pi^R(\theta) = \lambda(\overline{x} - \theta)$. Unsurprisingly, this coincides with the contract derived in Persson and Tabellini (1993) and Walsh (1995) in the absence of the ZLB. However, for positive q, the second term of $T_\pi^R(\theta)$ is negative, so – for larger q – it is ambiguous whether the central bank should face a corrective, "Pigouvian" inflation tax or a subsidy.

The novel results – compared to the previous literature – are two. First, the optimal contract is state contingent in two ways. As in Persson and Tabellini (1993) and Walsh (1995), it depends on the natural-rate of output shock θ, as that shock shapes the classical credibility problem. But the contract is also contingent on the natural-rate of interest shock ρ, as that shock shapes the new credibility problem under the ZLB.

Second, the power of the contract – i.e., its slope in realized state-R inflation – reflects the interplay of the two distortions under discretion that we discussed in Subsection 2.3. In particular, it may entail either a tax or a subsidy on inflation. A subsidy is more likely the less serious is the classical credibility problem tied to boosting output in state R. This is captured by the first contract term $\lambda(\bar{x} - \theta) > 0$. But a subsidy is also more likely the more serious is the credibility problem tied to boosting demand in state r (the ZLB). This is captured by the second term $\beta^R(q)(\bar{\pi} + r) < 0$. Note that $\beta^R(q) = \alpha^R(q) + \delta^R(q) > 0$, where $\alpha^R(q)$ is the dependence of state-R inflation on $(\bar{\pi} + r)$ under discretion and $-\delta^R(q)$ the same dependence under the optimal rule – this is how the optimal contract resolves the second credibility problem.

Depending on which of the two problems is more severe, the net effect is a tax or a subsidy for inflation. Since the relevance of the ZLB increases with q and with the absolute value of r, the optimal contract is more likely to entail an inflation subsidy if q is larger and the absolute value of r is larger.

Finally, like in Persson and Tabellini (1993) and Walsh (1995), the optimal contract is linear. Therefore, upward and downward deviations cancel each other and the contract does not distort the response to the supply shock ε over which the central bank has an information advantage. The contract also eliminates the inefficiency in how the discretionary equilibrium stabilizes demand shocks ρ. All in all, the optimal state-contingent contract in Proposition 3 implements the optimal policy under commitment. The next subsection shows how both linearity and full implementation can fail.

3.3 Non-state-contingent Contracts

We now impose the plausible constraint that not all relevant contingencies may be incorporated in an inflation contract. If we think about writing a literal contract, all contingencies may not be observable or verifiable. This applies in particular to the natural rate of output (growth) and the supply shock – θ and ε in our model – which may only be meaningfully defined in an empirical model. But models are not unique and the values of θ or ε may not be verifiable, while in principle observable. If we think about a more general central-bank institution, it will have to be designed under a veil of uncertainty about these shocks. Conditioning on realizations of θ and ε thus seems doubtful.

For simplicity, we continue to assume that the realization of ρ is verifiable, so the performance contract is only defined on inflation in state R, when monetary policy has a role to play. One justification for this assumption is that whether $i = 0$ or not is easily verifiable. Moreover, under the contract the central bank has no incentive to pretend that the ZLB binds when in fact it does not.

With this motivation – and a similar logic as in the previous subsection – we derive the optimal contract for holding the central bank accountable for state-R inflation, π^R, when the incentives cannot be contingent on the realizations of θ and ε. Since the loss function is quadratic in inflation, we allow both linear and quadratic terms in the contract. As in the preceding subsection, the central bank chooses its policy rate under discretion and full information, so as to minimize $L(\pi, x) + T(\pi)$.

The main result The Appendix proves that if **Assumptions 1** and **2** are satisfied and the variance of θ, υ_θ, is sufficiently large relative to the variance of ε, υ_ε, then

Proposition 4 *The optimal non-state-contingent contract for inflation away from the ZLB is*

$$T(\pi) = T_0^R + T_\pi^R(\bar{\theta})\pi^R + \frac{T_{\pi\pi}^R(\upsilon_\theta, \upsilon_\varepsilon)}{2}(\pi^R - \bar{\pi}^R(q))^2,$$

where $\bar{\pi}^R = E(\pi^{C,R}(q, r, \varepsilon)) = \bar{\pi} - \delta(q)(\bar{\pi} + r) \geq \bar{\pi}$ *coincides with expected state-*R *inflation in Proposition 1, where* $T_\pi^R(\bar{\theta}) \gtreqless 0$ *has the same form as in Proposition 3 evaluated at expected value* $\bar{\theta}$, *and where* $T_{\pi\pi}^R(\upsilon_\theta, \upsilon_\varepsilon) > 0$ *is increasing in* υ_θ *and decreasing in* υ_ε. *This contract does not implement the optimal policy under commitment.*

By **Proposition 4**, the optimal non-state-contingent contract resembles an inflation target: its last term penalizes deviations from the optimal average inflation rate under commitment $\bar{\pi}^R > \bar{\pi}$. The inflation target coincides with the social optimum $\bar{\pi}$ only if the probability of being at the ZLB is zero. Otherwise, $\bar{\pi}^R$ exceeds socially optimal inflation, and does so by more the higher is q. As in **Proposition 3**, the additional penalty (the second term) for missing the inflation target could be asymmetric in either direction – i.e., $T_\pi^R(\bar{\theta}) \gtreqless 0$. Again, this reflects whether the credibility problem tied to the ex post inflation incentives is larger or smaller than the credibility problem tied to raising demand at the ZLB, except that $T_\pi^R(\bar{\theta})$ is now defined by the expected value of θ rather than its specific realization.[5]

Different from a state-contingent contract, an optimal non-state-contingent contract can no longer incentivize the central bank to implement the optimal policy rule. In particular, the non-linear inflation contract distorts the responses to two types of shocks and has to make a compromise between them. The penalty $T_{\pi\pi}^R(\upsilon_\theta, \upsilon_\varepsilon)$ for deviations from $\bar{\pi}^R(q)$ is more severe – $T_{\pi\pi}^R$ larger – the

[5] In the extension discussed in Subsection 2.1 with political shocks to λ, and if λ were observable but non-verifiable, the contract would include $\overline{\lambda}$ – the average value of λ.

larger the variance of the natural-rate shocks θ that cause the inflation bias, and less severe – $T^R_{\pi\pi}$ smaller – the larger the variance of the supply shocks ε that the central bank ought to stabilize.

4 Discussion

This section relates our results on optimal contracts to the existing literature on monetary policy, as well as to accountability mechanisms for real-world central banks. In addition to this discussion of *how* policy should be delegated, we also bring up the old question *which tasks* should be delegated to central banks. Finally, we discuss how one could relax some of our simplifying assumptions and extend the analysis in other ways.

As a premise, we have assumed that – absent an optimal contract – the central bank fully and correctly internalizes society's loss function as captured by $L(\pi, x)$. The purpose of the contract is to counteract the specific incentive problems that result from sequential decision making and the resulting lack of commitment. An inflation contract should thus not be confused with the general objectives assigned to the central bank, such as the US Federal Reserve's dual mandate of maximum employment and price stability, or the ECB's primary objective of price stability (and without prejudice to price stability to support the general economic policies of the EU). In our conceptual framework, these general goals can be interpreted as the form of society's loss function, $L(\pi, x)$. The optimal contract $T(\pi)$ goes beyond these general statements, and is meant to hold the central bank accountable for realized inflation, according to pre-specified procedures and contingencies.

Our discussion does not pretend to be exhaustive. We single out some issues that we think are particularly relevant today in designing appropriate monetary policy institutions. However, we do not address a number of important operational questions, such as which measure of inflation to target, how large deviations to tolerate around the stated target, how fast to return after target deviations, how to communicate policy decisions and future intentions, or how to forecast inflation.

4.1 Comparison with Existing Research and Practice

We begin by pointing out a few upshots of the optimal contracts we have derived, comparing them to other proposals in the research on monetary policy and to policy strategies recently adopted by central banks.

A higher inflation target? As in Persson and Tabellini (1993), it is natural to interpret the optimal non-state-contingent contract in **Proposition 4** as an inflation-targeting structure. By its last term, this contract holds the central bank accountable for keeping inflation close to a pre-specified target. If the ZLB is expected to bind, the target inflation rate should be higher than the socially optimal rate of inflation: $\bar{\pi}^R > \bar{\pi}$. Qualitatively, this is consistent with proposals in the aftermath of the financial crisis (Blanchard et al. 2010, Krugman 2014) to raise the inflation target from about 2% (the benchmark in most countries) towards 4%.[6]

Even though our model is very simple, we present some back-of-the-envelope calculations to judge its quantitative normative implications. Suppose that $\bar{\pi} = 0.02$, $\sigma = 1$, $\lambda = 0.25$ (cf. Kiley and Roberts 2017) and $r = -0.03$. Then, equation (A34) for $\delta^R(q)$ in the Appendix says that the inflation target to be applied off the ZLB should be related to q as

$$\bar{\pi}^R = \bar{\pi} - \frac{[1 + (1 + \lambda)\sigma]\sigma q}{1 - q + (1 + \lambda)\sigma^2 q}(\bar{\pi} + r) = 2 + \frac{2,25}{0,25 + 1/q} .$$

Given the incidence of the ZLB in the last 20 years (Kiley and Roberts 2017), it is not unreasonable to set $q = 0.25$. This would imply $\bar{\pi}^R \approx 2,5\%$. To reach $\bar{\pi}^R = 3\%$, we have to assume $q > 0.4$ or $r < -0.033$. The upshot that the ZLB does not call for an inflation target much higher than 3% is also supported by the micro-founded welfare simulations in Coibion et al. (2012). However, allowing for additional sources of aggregate demand shocks, or for correlated supply and demand shocks, might raise the frequency of being at the ZLB, which would raise the implied inflation target.

State-dependent inflation targets? A second – perhaps more subtle – aspect distinguishes our optimal contract from the prescription of raising the inflation target to 3 or 4%. In our contracts, the central bank is held accountable for inflation *only* in states when the ZLB does not bind – thus, the inflation-targeting framework only applies if $\rho = R$. The proposals mentioned above do not make this distinction between being constrained by the ZLB, or not. Since average inflation is lower in state r than in state R, this suggests further caution in calling for a target much above 2% for average inflation – i.e., irrespective of whether the ZLB binds. Kiley and Roberts (2017) do acknowledge that it is important to distinguish macro outcomes by whether the ZLB binds, but they do not analyze how to design central bank incentives in order to make state-contingent policy rules credible.

Asymmetric inflation targets? At their introduction, some inflation-targeting frameworks were asymmetric, stating that realized inflation above the target was worse than downward deviations. For example, the ECB described its initial strategy as the goal of keeping medium-run inflation close to or below 2%. These asymmetries have since been removed.[7]

[6] More recently, Blanchard called for an inflation target of 3% (Blanchard 2022).

[7] In 2021, the ECB stated that : "The Governing Council considers that price stability is best maintained by aiming for 2% inflation over the medium term. The Governing Council's commitment to this target is symmetric. Symmetry means that the Governing Council considers negative and positive deviations from this target as equally undesirable." cf. https://www.ecb.europa.eu/home/search/review/html/ecb.strategyreview_monpol_strategy_statement.en.html

To gauge if and how the optimal contract should be asymmetric – formally, if the second term $T^R_\pi(\bar{\theta})$ in **Proposition 4** is positive or negative – requires an estimate for the average inflation bias $\lambda(\bar{x} - \bar{\theta})$. Data from the 1960s and 1970s would produce very high estimates for the inflation bias, and hence a penalty for upward deviations ($T^R_\pi(\bar{\theta}) > 0$). One may argue that, by now, the public and monetary authorities may have understood that chasing very low output gaps could backfire into high inflation – in terms of our model, the desired level of output $\bar{x}$ is very close to the average natural rate $\bar{\theta}$. But we note the recent arguments in Afrouzi et al. (2024) that political-economy and credibility problems may well make a comeback as inflation drivers.

In analogy with our earlier numeric discussion of the target level, we can ask how high the inflation bias must be to generate a symmetric optimal contract – i.e., $T^R_\pi(\bar{\theta}) = 0$. Consider the same parameter values as above, including $q = 0.25$. By expression (A20) for $\beta(q)$ in the Appendix, the optimal contract is symmetric for an inflation bias $\lambda(\bar{x} - \bar{\theta})$ about 0,6%. With $\lambda = 0.25$, this value would apply if the central bank systematically tried to keep output 1.5% above its natural rate. A higher (lower) inflation bias, or a value of q lower (higher) than 0.25 would imply an optimal inflation tax (subsidy).

Once again, these numbers – in our simple model – suggest that the ZLB does not radically change the properties of the basic optimal-inflation contract, even though some marginal changes may be desirable.

Price-level targets or inflation targets? Some scholars have proposed price-level targeting as a preferred way to implement stable inflation (Svensson 1999b). How will such a regime work with a ZLB? In the model of Eggertson and Woodford (2003), the optimal monetary policy under commitment can be approximated by a price-level targeting rule around a rising trend. Intuitively, such a rule calls for higher actual (and expected) inflation after a sequence of low inflation rates – the likely outcome after some time at the ZLB. If the natural interest rate follows a Markov process (or has positive serial correlation), then this raises expected inflation when needed.

However, the welfare consequences – and hence the credibility – of such a policy rule become questionable in more general settings. For instance, it may be very costly in terms of output losses to return to a prescribed price-level path after a large supply shock, especially if inflation is backward-looking and hence sticky. Price-level targeting would also rule out using inflation as a shock absorber in the wake of major fiscal shocks (Lucas and Stokey 1983).

Moreover, targeting a rising price level may be more complicated to explain to the public than targeting inflation. This, in turn, would make it more difficult to hold the central bank accountable, possibly undermining the credibility of monetary policy. For these reasons, price-level targeting has not been adopted in the real world, except by Sweden that targeted a constant price level in the 1930s (Fisher 1934, Ch. X).

Average inflation targets? In August 2020, the US Federal Reserve announced it would target average inflation, for a period longer than a year. Doing so, the Fed implicitly introduced elements of history dependence into its strategy: as yearly inflation had been below 2% for a while, the Fed would now accept a period with inflation above 2%. If believed, this announcement would have raised expected inflation when most needed, like a (credible) price-level target.[8]

However, the drawback of this strategy became evident a couple of years later, when supply shocks raised the inflation rate. Since the strategy and its underpinnings had not been explicitly spelled out, market participants found it difficult to form expectations. Was the Fed "behind the curve" and reacting too slowly to inflationary shocks – implementing a procyclical monetary policy and raising output and inflation volatility – maybe because it did not forecast inflation correctly? Or was it just following its pre-announced strategy of targeting inflation for a longer period? And over exactly what period was average inflation measured? Without clarity on these questions, the benefits of an inflation-targeting framework for policy credibility seemed lost.

This obliquity contrasts with the transparency of an optimal-contract approach, as articulated in Section 3. The underlying economic model may certainly be criticized for being too simple, in particular for being static. But a contract built on simple logic is more transparent. If appropriate accountability mechanisms foster contract compliance, transparency promotes credibility, which facilitates expectations formation.

4.2 Real-world Implementation

In this section, we briefly discuss what the results in Section 3 tell us about how to delegate monetary policy to central banks and to hold them accountable, comparing these implications to real-world practices (Svensson 1999a discusses how actual inflation targeting was implemented in its early days).

Our analysis suggests that the objective and horizon of the delegation – including any conditions in terms of state dependence – should be publicly spelled out ex ante by the bank's principal. As the contracts we have derived are based on realized inflation, performance should be evaluated ex post. To set incentives right, the central-bank leadership should know about this evaluation and expect that they will be held accountable for their policy performance with explicit reference to the delegation terms.

How to delegate? Some countries on inflation targeting have codified the delegation of monetary policy in an official document. Thus, as stipulated in a new law from the year before, the first of many Policy Target Agreements (PTA) was signed in March 1990 between New Zealand Finance Minister David Caygill

[8] The Fed 2020 framework was also asymmetric: it called for keeping inflation "moderately above target for some time" after persistent outcomes below target, but no such makeups for overshoots. Moreover, policy ought to consider shortfalls from "maximum employment" but not overshoots.

and Reserve Bank Governor Don Brash. This public document spelled out the objective (0-2%), how it were to be measured and at what horizon (2 years). Signing such a PTA was mandatory before any appointment, or reappointment, of the governor who was singled out as ultimately responsible.

Similarly, under the Bank of England Act, the Chancellor of the Exchequer writes an annual public remit to the members of the Bank of England's MPC, which spells out a 2% target for the annual inflation rate with an obligation for the bank to explain any deviation of more than 1% in a public letter. This remit regularly makes reference to other goals, as well as to economic shocks, but it has not made the inflation target state dependent. In particular, the 2% target has never been altered despite long periods at the ZLB.

Other central banks – such as the US Fed, the ECB, and the Riksbank – that enjoy legislated instrument independence and are obliged by law to pursue price stability – have instead practiced some goal independence by formulating their own inflation targets so as to clarify their monetary-policy strategies.

Accountability mechanisms? As for exercising accountability, we also observe a variety of practices. Up to 2019, the New Zealand PTA designated the Bank 's signatory, namely the Governor, as ultimately responsible for fulfilling the target. But nowadays the MPC is collectively responsible, as is the MPC of the Bank of England. Some countries have developed a practice of periodic backward-looking evaluations by outsiders, sometimes led by leading international experts. Thus, the Riksbank's inflation-targeting policy has undergone four long-term evaluations commissioned by the fiscal committee of parliament – the parliament (rather than the government) being the Riksbank's principal.[9]. In Norway, annual policy evaluations – Norges Bank Watch – have been carried out by the independent Center for Monetary Economics for more than 20 years on behalf of the Ministry of Finance. The delivery of these outside reports are associated with public questioning of bank representatives arranged by the contracting authority. These periodic evaluations often have the dual purpose of assessing central-bank performance and suggesting possible improvements to the targeting framework.

Other central banks regularly submit their own reports to their principals, including the ECB, the US Fed, and the Bank of England – e.g., by the Federal Reserve Act, the US Fed reports to Congress every year, and by the Bank of England Act, the bank publishes quarterly Monetary Policy Reports (Inflation Reports up to 2019). However, external periodic evaluations of the Fed and the Bank of England have been called for by a group at the Boston Fed (Fuhrer et al. 2018), and by the Upper House Economic Affairs Committee (House of Lords 2023), respectively. Given the results in Section 3, an outside and independent evaluation of contract fulfilment does indeed appear to be an important mechanism for exercising accountability.

[9] The first evaluation was carried out by Francesco Giavazzi and Fredric Mishkin (for 1995-2005), and subsequent ones by Charles Goodhart and Jean-Charles Rochet (2005-10), by Marvin Goodfriend and Mervyn King (2010-15), and by Karnit Flug and Patrick Honohan (2015-20).

These real-world procedures to exercise accountability entail both extrinsic and intrinsic rewards or punishments. The (earlier) proviso that a New Zealand Bank Governor can only be appointed and reappointed after signing a PTA is a clear example of the former. The latent reputational losses of Riksbank MPC members who are criticized for policy failures by renowned international academics, or by domestic politicians – and are exposed to this criticism in highly publicized hearings – is a clear example of the latter.

Bottom line The spirit of our contracting analysis suggests that the principal, rather than the agent, ought to design the details of the delegation. Accountability will be more effective, and incentives stronger, when the target and the procedure are assigned by the principal rather than being self-declared policy targets. But an explicit assignment of a target, by itself, may not be sufficient to alter central-bank incentives. It needs to be complemented by a specific procedure to hold the central bank accountable, assessing its policy performance against the delegated target.

Combining the theory and practice of inflation targeting suggests that an effective accountability procedure would be to require the central bank to hold yearly inflation close to target on average over a pre-defined period (say, 4 years). Significant yearly deviations from target should be clearly explained as attempts to reach other goals specified in the central bank mandate. Periodic external reviews of central-bank performance in light of its mandate, and reviews of the targeting framework should also be part of the procedure. Existing theory has little to say on whether the target should be specified as a point or as a range.

In existing practices, accountability mechanisms are not always as strong and effective as they could be. Not surprisingly, central banks do not demand stronger accountability, while their political principals may lack the expertise to design more effective procedures.

4.3 Alternative Instruments

So far, our analysis has relied on two simplifying assumptions. (1) The central bank lacks other policy instruments when the ZLB excludes further cuts in interest rates. (2) Policymakers and society only worry about the effects of central-bank policy on inflation and output. But many central banks relied on QE, in one form or another, when they intervened in the aftermath of the financial crisis and in the early COVID-pandemic. And even though such measures may decrease the short-run risk of financial crises, boosting liquidity and asset prices by large asset purchases may have harmful side effects. These may include more pressure on future inflation, spillovers on fiscal policy, or a higher risk of future financial crises (Acharya and Rajan 2023, House of Lords 2023). Some central banks, including the ECB and the Riksbank, have also resorted to negative policy interest rates, appealing to an "effective lower bound" rather than a ZLB. But the effectiveness and desirability of negative interest rates is controversial (see, e.g., Brunnermeier et al 2023, or Eggertson et al 2024).

In this section, we sketch how to expand our simple model when the central bank can use a proxy for QE, or negative interest rates, under the ZLB. We also assume that such use has an additional social cost, which we will initially interpret as reflecting a higher risk of future financial instability. We don't pretend to discuss optimal policies at the ZLB – a lot has been written on that already. Rather, we discuss the additional incentive problems that may arise when attempting to steer the economy by other means at the ZLB. As we shall see, this raises new questions about institution design – not just "how" to delegate, but "what" to delegate.

A simple model of QE Consider a minimal extension of our model. Let us replace (2), our previous equation for aggregate demand, by

$$x^d = \theta - \sigma(i - \pi^e - \rho) + g \ , \tag{10}$$

where g is a measure of QE (or of negative interest rates).[10] Suppose further that society's loss function is

$$L = \frac{1}{2}E[(\pi^\rho - \bar{\pi})^2 + \lambda(x^\rho - \bar{x})^2] + \kappa g \ . \tag{11}$$

for $\rho = R, r$. Parameter $\kappa > 0$ captures the additional social costs associated with this alternative policy instrument, due to higher risks of future financial instability. While this is an extremely simple reduced form, such an additional cost could be derived from a more primitive model. Suppose that an expansion of liquidity via higher central-bank asset purchases raises the probability of a large future financial crisis. Suppose further that we model the consequences of a large financial crisis as a major economic disaster, along the lines of Barro (2006). Then, the last term in (11) could capture (a first-order approximation of) the higher expected cost due to a higher risk of a major disaster – maybe better so if g denoted cumulated asset purchases by the CB (or the cost imposed on the banking system by negative interest rates.

Finally, we assume that a central bank operating under discretion (taking expectations as given) minimizes a loss function of the same form as (11), except that it internalizes only a fraction $\gamma \leq 1$ of the true social cost of using instrument g – meaning that we get the objective of the central bank by replacing the last term on the RHS of (11) with $\gamma\kappa g$. We make this assumption to illustrate the consequences of today's arrangement that the main responsibility for most central banks is to maintain macroeconomic stability. But the costs (or benefits) that QE impose on society may largely be associated with other policy-related outcomes outside of a narrow mandate, costs which may therefore not be fully internalized by the central bank.

[10] If g refers to negative inetrest rates, we may write $g = \eta\sigma i$, where $\eta < 0$ is a parameter that captures the fact that negative interest rates may transmit differently to demand than positive ones.

Differences between discretion and commitment How does the equilibrium under the optimal policy rule with commitment differ from the equilibrium under discretion, given these assumptions? The Appendix proves:

Proposition 5 *The difference between the equilibrium under discretion and the optimum under commitment, also taking into account the difference in the minimized loss functions, is*

$$
\begin{aligned}
\pi^{D,R} - \pi^{C,R} &= \lambda(\bar{x} - \theta) - q(1+\sigma)\kappa + \frac{(1-\gamma)\lambda q\kappa}{(1+\lambda)} \\
\pi^{D,r} - \pi^{C,r} &= \lambda(\bar{x} - \theta) - q(1+\sigma)\kappa + \frac{(1-\gamma)(1+q\lambda)\kappa}{(1+\lambda)} \\
\pi^{D,e} - \pi^{C,e} &= \lambda(\bar{x} - \theta) - q(1+\sigma)\kappa + (1-\gamma)q\kappa \\
x^{D,R} - x^{C,R} &= -\frac{q(1-\gamma)\kappa}{(1+\lambda)} \\
x^{D,r} - x^{C,r} &= \frac{(1-q)(1-\gamma)\kappa}{(1+\lambda)}.
\end{aligned}
$$

Compared to the optimal policy under commitment, discretionary policy entails two distortions. The first arises because a central bank under discretion neglects the effects of both policy instruments on expected inflation. It is apparent if we set $\gamma = 1$ in the expressions in **Proposition 5**. Under discretion, inflation in both states $\rho = R, r$ could be too high or too low, compared to the optimum under commitment. This depends on whether the conventional inflation bias $\lambda(\bar{x} - \theta)$ – the first term in $\pi^{D,R} - \pi^{C,R}$ and $\pi^{D,r} - \pi^{C,r}$ – is higher or lower than the neglected effect of expected inflation on aggregate demand at the ZLB – the second term $q(1+\sigma)\kappa$ in these expressions. Because using QE is costly, higher expected inflation would reduce the need to resort to this costly alternative instrument. As in Sections 2 and 3, this deflation bias is higher the larger is q, the probability of being at the ZLB.

This difference between the discretionary equilibrium and the optimal policy is qualitatively similar to the difference highlighted by **Propositions 1** and **2**, and it could be corrected with a suitable state-contingent linear inflation contract. Unlike in the two previous sections though, the inflation difference between discretion and commitment is the same in both states, $\rho = R, r$. Hence, an optimal state-contingent inflation contract would *not* need to distinguish between these two states. Since the distortion is the same in the two states, the same inflation tax or subsidy would remove it. But this result hinges on the simple linear social cost of using g and state-dependence might thus reappear with a more general cost function.

However, a second distortion arises if the central bank does not correctly internalize the social costs (or benefits) of using alternative instrument g – i.e., when parameter $\gamma \neq 1$. Suppose in particular that $\gamma < 1$, so that the central banks does not fully internalize the adverse effects of QE on asset markets, or of negative interest rates on the banking system, and thus on financial stability. Then, any contract with a state-ρ independent inflation target – as in existing real-world regimes – worsens the second distortion, as it induces the central bank to rely too much on QE or negative interest rates at the ZLB. As is evident by comparing the third terms in $\pi^{D,R} - \pi^{C,R}$ and $\pi^{D,r} - \pi^{C,r}$, the second distortion is higher in state r than in state R (for positive q and $\gamma < 1$). In principle, this distortion could be remedied by a state-contingent inflation contract. The less the central bank internalizes the social cost of g (the lower is γ), the more it over-uses it to stimulate the economy at the ZLB, and the higher is realized inflation in state r. To remedy this, we should raise the inflation tax (or reduce the inflation subsidy) in both states, $\rho = R, r$, but more so in state r, and the more so the lower is γ. In other words, our result from Section 3 that the optimal inflation contract should be state-contingent on ρ reappears, even with a linear social cost of g. Similar considerations (with opposite implications) would arise if the central bank exaggerates the social costs of using g (i.e., if $\gamma > 1$). This would be the case, e.g., if it does not correctly take into account that QE could prevent financial crisis in circumstances when liquidity is scarce, as argued by Allen et al. (2024).

More far-reaching implications Nonetheless, there is a simpler and more transparent way to address the distortion (incentive problem) that the central bank does not fully internalize the social cost imposed by this alternative instrument. Our discussion here is particularly relevant with regard to QE. If the effects (positive or negative) of QE mostly concern costs taking the form of imposing higher future financial risks – or benefits taking the form of preventing current liquidity crises – then the most natural incentive contract would add an auxiliary clause to make the central bank internalize these effects. In other words, the common logic of "addressing the distortion directly at its source" (Bhagwati 1969) would suggest that the central-bank mandate be expanded to cover financial stability, on top of price-*cum*-output stability. If the incentive contract manages to induce the central bank to correctly internalize the social costs and benefits of QE in the financial-stability domain, then – by the discussion above – the incentive-contract clause over inflation must no longer be state-contingent (given the linear social cost).

Here, financial stability could have a broader interpretation than avoiding banking and financial crises, which also includes financial amplifications of business cycles. As discussed by Borio et al. (2022), central banks around the world are developing new frameworks for macroprudential measures. Integrating new policy instruments in a unified framework for price stability and financial stability raises operational issues. Holding the central bank accountable for broader objectives also raises new challenges. Not only do policy horizons differ across goals, but new financial innovations keep altering the constraints for central-bank operations – finding appropriate policy instruments is thus like "aiming at a moving target." To address these challenges may be the next crucial frontier in the design of monetary institutions.

Finally, the additional effects of QE may go beyond the risk of financial instability, and involve spillovers to the domain of fiscal policy. This might call for more explicit coordination between fiscal and monetary policy when the central bank is at the ZLB. For instance, to preserve central-bank independence, mounting capital losses on large asset holdings in the central bank's balance sheet must eventually be absorbed by the Treasury. In fact, such absorption was institutionally backed after the financial crisis in the UK – though not without frictions – by the creation of the so-called Asset Purchase Facility.

Anyway, thinking about the ZLB – and the temptation to use instruments like QE – alters our perspective not only on how to delegate a certain objective to the central bank. Our extended model also sheds new light on which objectives should be delegated in the first place, and hints at the desirability of coordination with other policymakers.

4.4 What's Missing?

Wrapping up this section, we briefly point to a few obvious omissions in our analysis and a few aspects of our approach that one could fruitfully develop further.

Relax drastic assumptions To facilitate closed-form – rather than numerical – solutions, we work with a stationary framework, where the probability of being at the ZLB is constant across time. This flies in the face of the evidence from Japan and from the aftermath of the financial crisis, which strongly suggest some path dependence. We also assume that the same expectations of inflation affect both the demand and the supply side of the economy. Deriving optimal contracts in a dynamic setting where the ZLB is history-dependent is a vital but complex task.

To simplify the analysis, we subsume all demand fluctuations in a binary shock ρ to the natural rate of interest, one realization of which always triggers the ZLB. Similarly, we assume that supply shocks ε (and θ) are uncorrelated with ρ and have a narrow enough support that they do not affect the probability of being at the ZLB. More general distributions of demand and supply shocks would be more realistic and allow for a richer and more interesting analysis of when the ZLB binds.

Related to this point, the optimal contract that we have studied in the previous section is contingent on the realization of ρ, an exogenous event that we have assumed to be verifiable. A more realistic assumption is that the verifiable event is whether $i = 0$. However, this contingency is not the same as a shock to the natural real rate, as i is a choice variable of the central bank. An optimal sustainable contract contingent on whether $i > 0$ or $i = 0$ would need to have an additional property: under the contract the central bank would have incentives to tell the truth about the realization of ρ. We don't know if such a truth-telling constraint would require additional restrictions on parameter values in our simple model.

Include policy spillovers As it stands, our model only studies the traditional macro outcomes of inflation and output. But in the preceding subsection, we argued that QE – and other instruments at the ZLB – may have stronger spillovers on other policy-related outcomes than does conventional monetary (interest-rate) policy. This raises old issues of assignment of policy instruments and coordination among policymakers (Tinbergen 1956, Mundell 1962). Taking such spillovers explicitly into account, would be an important and interesting extension of the appropriate institutional framework for monetary (and fiscal) policy.

Study richer policymaking incentives More broadly, our model maintains the simple contracting approach of earlier research (Persson and Tabellini 1993, Walsh 1995), which directs our thinking to *material* motives. However, in monetary policymaking by experts, extrinsic incentives are probably much less important than the *intrinsic* incentives tied to decision-maker reputations (Benabou and Tirole 2003). To study how alternative accountability mechanisms motivate central-bank leaders, who are modeled as motivated agents (Besley and Ghatak 2005), may thus enrich our analysis.

Further, we do not distinguish central banks with *single* and *multiple* decision-makers. That distinction may be vital for information aggregation and decisions in genuinely new situations, such as the ZLB. Extending our model in that direction would be worthwhile, as most central banks delegate decisions to an MPC, which is more or less collectively accountable depending on the precise institutional setting.

Introduce political shocks In Section 2, we mentioned that our model could be extended with political shocks to λ, the relative weight on inflation and output, as in Halac and Yared (2020) and Afrouzi et al (2024). In Section 3, we conjectured that the optimal non-state contingent central-bank contract in such a setting would entail the average value of λ. It is natural to suppose that the political shocks are tied to uncertain election outcomes whereby two political parties – with different weights λ – alternate in power.

It would be interesting to study the political feasibility of delegating an optimal central-bank contract in this partisan setting. It is conceivable that the parties would prefer such a compromise under a veil of ignorance about future political power. The incentives for an ex ante agreement would be analogous to the incentives for parties to find a way to compromise in Alesina (1987) and Lagunoff (2001).[11]

5 Concluding Remarks

Inflation targeting has been an effective innovation in the conduct of monetary policy, because it has altered policymaking incentives. The fact that monetary

[11] Alesina (1987) considered a compromise via a reputational equilibrium in the wake of partisan disagreement over how to stabilize macroeconomic outcomes, while Lagunoff (2001) considered a compromise via a constitutional rule in the wake of partisan disgreement over how much to extend civil rights.

authorities have gained credibility and become increasingly able to influence expectations reflects changes in institutions, much more than novel communication strategies.

But inflation-targeting frameworks were developed to deal with high and volatile inflation, at a time when central banks did not encounter feasibility constraints in setting interest rates. Since then the world has changed. Thanks to the success in bringing down inflation, and a downward trend in the natural interest rate, the ZLB has become a recurring restriction on interest-rate policy.

In this paper, we ask whether and how the inflation-targeting framework ought to be adapted to this new contingency. We do so by revisiting the optimal-contract approach to the design of monetary institutions, along the lines of Persson and Tabellini (1993) and Walsh (1995). Four lessons stand out.

First, the optimal contract and the implied inflation-targeting regime should condition on being at the ZLB or out of it. Since monetary policy operates in a different manner in these circumstances, which are observable and verifiable, the distinction between being at the ZLB or out of it should be reflected in central-bank targets and how the central bank is held accountable for its performance. This simple lesson has not been incorporated in existing central-bank practices.

Second, and as already argued by others, the optimal inflation target should be raised to deal with the possibility of being at the ZLB, and more so the greater the risk of being there. Moreover, upward deviations of inflation from target should not be discouraged more than downward deviations. But these qualitative lessons do not appear to warrant major quantitative changes of inflation targets, at least based on our suggestive computations.

Third, the relevance of the ZLB suggests that it may be desirable to expand central-bank mandates to encompass financial stability, broadly defined, besides price and output stability. At the ZLB, the central bank has to rely on non-standard policy tools, which likely alter how financial markets function and influence business cycles as well as reshape the risks of financial crises. The lesson from our approach is that one cannot take it for granted that, absent appropriate institution design and explicit responsibility, monetary authorities have the right incentives to handle this broader set of policy instruments.

Last, but not least, accountability for inflation performance is a central ingredient in a successful monetary-policy framework. As we just argued, inflation targeting mattered mainly by changing policymaker incentives. The lessons in our contracting results on new features of effective inflation targeting thus ought to be reflected in modified accountability mechanisms. How exactly to change those mechanisms in practice is a new and difficult challenge, which goes beyond the scope of this paper. Addressing that challenge is an urgent task for future research, at least as important as the search for optimal policy rules that has attracted so much recent attention.

6 References

Acharya, Viral A. and Raghuram Rajan 2023. "Liquidity, Liquidity Everywhere, not a Drop to Use: Why Flooding Banks with Central Bank Reserves May Not Expand Liquidity," University of Chicago, BFI Working Paper 2022-19.

Afrouzi, Hassan, Marina Halac, Kenneth Rogoff, and Pierre Yared 2024. "Changing Central Bank Pressures and Inflation," Brookings Papers of Economic Activity, forthcoming.

Alesina, Alberto 1987. "Macroeconomic Policy in a Two-party System as a Repeated Game," Quarterly Journal of Economics 102, 651-678.

Allen, Franklin, Jae Hyoung Kim, and Ansgar Walther 2024. "Inflation Targeting and Financial Stability," mimeo.

Bhagwati, Jagdish N. 1969. "The Generalized Theory of Distortions and Welfare," MIT, Economics Department Working Paper, No. 39.

Barro, Robert J. 2006. "Rare Disasters and Asset Markets in the Twentieth Century," Quarterly Journal of Economics 121, 823-866.

Barro, Robert J. and David B. Gordon 1983. "Rules, Discretion and Reputation in a Model of Monetary Policy," Journal of Monetary Economics 12, 101-121.

Bénabou Roland and Jean Tirole 2003. "Intrinsic and Extrinsic Motivation," Review of Economic Studies 70, 489-520.

Besley, Timothy and Maitreesh Ghatak 2005. "Competition and Incentives with Motivated Agents," American Economic Review 95, 616-636.

Blanchard, Olivier, Giovanni Dell'Ariccia, and Paolo Mauro 2010. "Rethinking Macroeconomic Policy," IMF Staff Position Note.

Blanchard, Olivier 2022. "It is Time to Revisit the 2% Inflation Target," Financial Times, November 28.

Borio, Claudio, Ilhyock Shim, and Hyun Song Shin 2022. "Macro-financial Stability Frameworks: Experience and Challenges," BIS Working Papers, No. 1057.

Brunnermeier, Marcus K., Joseph Abadi, and Yann Koby 2013. "The Reversal Interest Rate," American Economic Review 113, 2084-2120.

Calvo, Guillermo A. 1978. "On the Time Consistency of Optimal Policy in a Monetary Economy," Econometrica 46, 1411–1428.

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, 1371-1406.

Davis, Josh, Cristian Fuenzalida, Leon Huetsch, Benjamin Mills, and Alan M. Taylor 2024. "Global Natural Rates in the Long Run: Postwar Macro Trends and the Market-implied r* in 10 Advanced Economies," Journal of International Economics, 103919.

Eggertson, Gauti B. 2012. "Was the New Deal Contractionary?" American Economic Review 102, 524-55.

Eggertsson, Gauti B., Ragnar E. Juelsrud, Lawrence H. Summers, and Ella Getz Wold (2024). "Negative Nominal Interest rates and the Bank Lending

Channel," forthcoming in Review of Economic Studies.

Eggertsson, Gauti B. and Marc P. Giannoni 2013. "The Inflation Output Trade-Off Revisited," Federal Reserve Bank of New York Staff Report No 608.

Eggertsson, Gauti B. and Michael Woodford 2003. "The Zero Bound on Interest Rates and Optimal Monetary Policy," Brookings Papers on Economic Activity 34, 139-235.

Fisher, Irving 1934. Stabilised Money, London, George Allen & Unwin Ltd.

Fuhrer, Jeff, Giovanni P. Olivei, Eric S. Rosenberg, and Geoffrey M.B. Tootell 2018. "Should the Federal Reserve Regularly Evaluate its Monetary Policy Framework?," Brookings Papers on Economic Activity, 443–97.

Halac, Marina and Pierre Yared 2020. "Inflation Targeting under Political Pressure," in Independence, Credibility and Communication of Central Banking, Santiago, Chile.

Holston, Kathryn, Thomas Laubach, and John C. Williams 2023. "Measuring the Natural Rate of Interest after COVID-19," Federal Reserve Bank of New York Staff Reports, No. 1063.

House of Lords 2023. "Making an Independent Bank of England Work Better," Economic Affairs Committee, 1st Report of Session 2023–24.

Kiley, Michael T. and John M. Roberts 2017. "Monetary Policy in a Low Interest Rate World," Brookings Papers on Economic Activity 48, 317-396.

Kydland, Finn E. and Edward C. Prescott 1977. "Rules Rather Than Discretion: The Inconsistency of Optimal Plans," Journal of Political Economy 85, 473-491.

Krugman, Paul R. 1997. "It's Baack: Japan's Slump and the Return of the Liquidity Trap," Brookings Papers on Economic Activity 2, 137-205.

Krugman, Paul R. 2014. "Inflation Targets Reconsidered," Paper presented for the European Forum on Central Banking.

Lagunoff, Roger 2001. "A Theory of Constitutional Standards and Civil Liberties, Review of Economic Studies 68, 109-132.

Lipsey, Richard. G. and Kevin Lancaster 1956. "The General Theory of Second Best," Review of Economic Studies 24, 11–32.

Lucas, Robert E. and Nancy Stokey 1983. "Optimal Fiscal and Monetary Policy in an Economy without Capital," Journal of Monetary Economics 12, 55-93.

Mundell, Robert A. 1962. "The Appropriate Use of Monetary and Fiscal Policy for Internal and External Stability," IMF Staff Papers 9, 70-79.

Persson, Torsten and Guido Tabellini 1993. "Designing Institutions for Monetary Stability," Carnegie-Rochester Conference Series on Public Policy 39, 53-84.

Persson Torsten and Guido Tabellini 2000. Political Economics, MIT Press, Cambridge

Rogoff, Kenneth 1985. "The Optimal Degree of Commitment to an Intermediate Monetary Target," Quarterly Journal of Economics 100, 1169–89.

Svensson, Lars E.O. 1997a. "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets," European Economic Review 41, 1111-1146.

Svensson, Lars E.O. 1997b. "Optimal Inflation Targets, 'Conservative' Central Banks, and Linear Inflation Contracts," American Economic Review 87, 98-114.

Svensson, Lars E.O. 1999a. "Inflation Targeting as a Monetary Policy Rule," Journal of Monetary Economics 43, 607-654.

Svensson, Lars E.O. 1999b. "Price-level Targeting versus Inflation Targeting: A Free Lunch?," Journal of Money Credit and Banking 31, 277-295.

Taylor, John B. 1993. "Discretion vs. Policy Rules in Practice," Carnegie-Rochester Conference Series on Public Policy 39, 195-214.

Tinbergen, Jan 1956. Economic Policy: Principles and Design, Amsterdam, North-Holland Publishing.

Walsh, Carl. E. 1995. "Optimal Contracts for Central Bankers," American Economic Review 85, 150-167.

Woodford, Michael 2005. "Central-bank Communication and Policy Effectiveness," in The Greenspan Era: Lessons for the Future, Kansas City: Federal Reserve Bank of Kansas City.

Figure 1 Monthly Inflation, Policy and Natural Rates since 1975 in Four Countries

Inflation rates

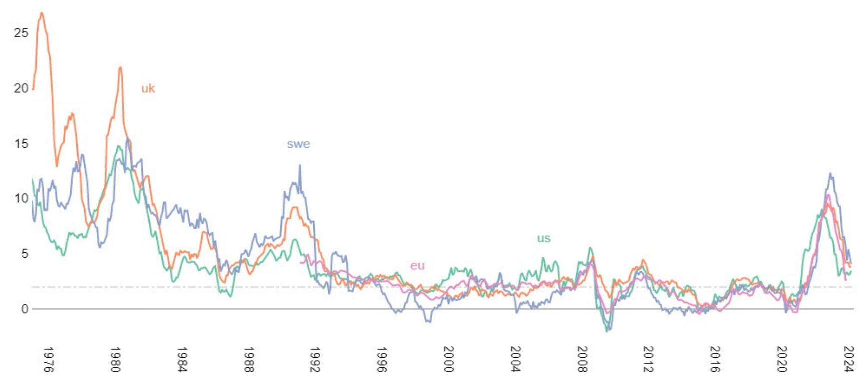

Policy rates

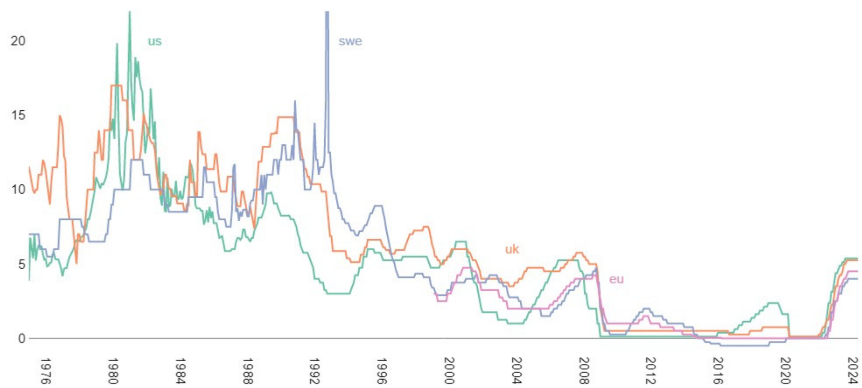

Natural rates

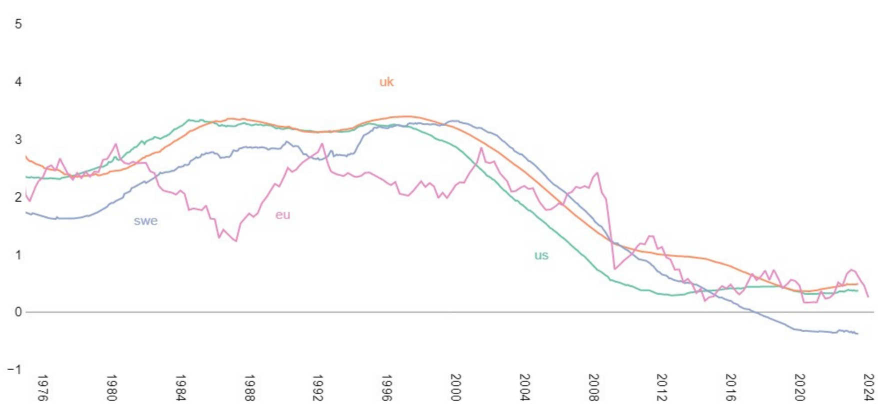

Note: The figure shows monthly data for (top to bottom) inflation rates, nominal (central-bank policy) interest rates, and real "natural" rates for the United States, United Kingdom, Sweden, and the Euro Area, going back to 1975 (or 1990). Inflation and policy rates are retrieved from Thomson Reuters for the Euro Area and from BIS, OECD, and the Riksbank, for the other countries. Natural rates are two-sided estimates from Davis et al. (2024), except for the Euro Area, where the estimates are from Holston, Laubach, and Williams (2023).

7 Appendix

7.1 Proof of Proposition 1

As explained in Section 2.2, the LQ structure of the problem allows us to study the optimal policy rule in two parts, one is the average (expected over ε) interest rate $E(i^R)$ and the other its (linear) response to supply shocks $\iota^{\varepsilon}\varepsilon$. Further, the one-to-one relation between i^R and π^R allows us to study this as a choice π^R, taking its direct and indirect effects on equilibrium macroeconomic outcomes, including the effects via inflationary expectations.

Average (expected) policy Thus, we first minimize loss function (3) with regard to expected policy $\pi^{C,Re}$, taking into account the effects on macroeconomic outcomes via (4)-(8). Using the latter five relations, we can re-express all these outcomes in terms of $\pi^{C,Re}$

$$x^{C,Re} = \theta - q\sigma\frac{(\pi^{C,Re}+r)}{1-q(1+\sigma)} \tag{A1}$$

$$\pi^{C,re} = \frac{(1-q)(1+\sigma)\pi^{C,Re}+\sigma r}{1-q(1+\sigma)} \tag{A2}$$

$$x^{C,re} = \theta + (1-q)\sigma\frac{(\pi^{C,Re}+r)}{1-q(1+\sigma)} \tag{A3}$$

$$\pi^{C,e} = \frac{(1-q)\pi^{C,Re}+q\sigma r)}{1-q(1+\sigma)}. \tag{A4}$$

Using (A1)-(A4) to rewrite the expected loss as a function of $x^{C,Re}$, the first-order condition for a minimum of expected loss $E[L(\pi,x)] = \frac{1}{2}E[(\pi-\overline{\pi})^2 + \lambda(x-\overline{x})^2$ with regard to π^{Re} can be written as

$$\begin{aligned} 0 \quad = \quad & \lambda q\frac{(1-q)\sigma}{1-q(1+\sigma)}\left(\theta + \frac{(1-q)\sigma}{1-q(1+\sigma)}(\pi^{C,Re}+r) - \overline{x}\right) \\ & -\lambda(1-q)\frac{q\sigma}{1-q(1+\sigma)}\left(\theta - \frac{q\sigma}{1-q(1+\sigma)}(\pi^{C,Re}+r) - \overline{x}\right) \\ & +q\frac{(1+\sigma)(1-q)}{1-q(1+\sigma)}\left(\frac{(1-q)(1+\sigma)\pi^{C,Re}+\sigma r}{1-q(1+\sigma)} - \overline{\pi}\right) \\ & +(1-q)(\pi^{C,Re}-\overline{\pi}). \end{aligned} \tag{A5}$$

Collecting terms, and simplifying the resulting expression, we obtain the following closed form solution for π^{Re}

$$\pi^{C,Re} = \overline{\pi} - \delta^R(q)(\bar{\pi}+r) \geq \bar{\pi}, \tag{A6}$$

where

$$1 > \delta^R(q) = \frac{[1+(1+\lambda)\sigma]\sigma q}{1-q+(1+\lambda)\sigma^2 q} > 0$$

with $\delta^R(0) = 0$ and

$$\frac{\partial \delta^R(q)}{\partial q} = \frac{[1 + (1 + \lambda)\sigma]\sigma}{[1 - q + (1 + \lambda)\sigma^2 q]^2} > 0.$$

Using this short-hand expression for π^{Re} as well as another short-hand function $\delta^r(q)$ defined by

$$\delta^r(q) = \delta^R(q) - \frac{\sigma}{1 + q((1 + \lambda)\sigma^2 - 1)} < \delta^R(q),$$

in equations (A2), (1) and (4), we obtain

$$\begin{aligned}
\pi^{C,re} &= \overline{\pi} - \delta^r(q)(\overline{\pi} + r) \\
\pi^{C,e} &= \overline{\pi} - [\delta^R(q)(1 - q) + q\delta^r(q)](\overline{\pi} + r) > \overline{\pi} \\
x^{C,Re} &= \theta - q[\delta^R(q) - \delta^r(q)](\overline{\pi} + r) \\
x^{C,re} &= \theta + [\delta^R(q) - \delta^{rR}(q)](1 - q)(\overline{\pi} + r).
\end{aligned}$$

Finally, note that $\delta^r(0) = -\sigma$, and that

$$\delta^R(q) - \delta^r(q) = \frac{\sigma}{1 + q((1 + \lambda)\sigma^2 - 1)} > 0$$

$$\begin{aligned}
\delta^R(q) - q[\delta^R(q) - \delta^r(q)] &= \frac{(1 + \lambda)\sigma^2 q}{1 - q + (1 + \lambda)\sigma^2 q} > 0 \\
\frac{\partial \left\{\delta^R(q) - q[\delta^R(q) - \delta^r(q)]\right\}}{\partial q} &= \frac{(1 + \lambda)\sigma^2}{[1 - q + (1 + \lambda)\sigma^2 q]^2} > 0.
\end{aligned}$$

These are the average macroeconomic outcomes that appear in **Proposition 1**.

Policy response to ε To fully characterize the macroeconomic effects of the best policy rule, it remains to derive the optimal interest-rate response to supply shocks $\iota^\varepsilon \varepsilon$. These are captured by the (private-sector) unexpected components of state-R inflation $\pi^{C,R\varepsilon} = \pi^{C,R} - \pi^{C,Re}$ and output $x^{C,R\varepsilon} = x^{C,R} - x^{C,Re}$.

The first-order condition when minimizing the loss with regard to ι^ε, taking the effect on $x^{C,R\varepsilon}$ and $\pi^{C,R\varepsilon}$ into account, is

$$\begin{aligned}
& E_\varepsilon[(1 - q)\lambda \frac{dx^{C,R\varepsilon}}{d\iota^\varepsilon}\varepsilon(x^R - \overline{x}) + (1 - q)\frac{d\pi^{C,R\varepsilon}}{d\iota^\varepsilon}\varepsilon(\pi^R - \overline{\pi})] \qquad \text{(A7)} \\
= \; & \sigma\varepsilon^2(1 - q)[-\sigma(1 + \lambda)\iota^\varepsilon + 1] = 0.
\end{aligned}$$

where E_ε denotes the expectations operator with regard to ε. Thus, we have $\iota^\varepsilon = \frac{1}{\sigma(1+\lambda)}$, which given (8) and (7) yields

$$\pi^{C,R\varepsilon} = \frac{\lambda}{1 + \lambda}\varepsilon \quad \text{and} \quad x^{C,R\varepsilon} = -\frac{1}{1 + \lambda}\varepsilon. \qquad \text{(A8)}$$

The expressions in (A8) are the random macroeconomic outcomes in state R under the optimal rule that appear in **Proposition 1**.

In state r, the central bank – by our assumption – operates under the ZLB for all realizations of ε. Because of this and because the only effect on state r outcomes of state-R policy runs via expected inflation, the central bank cannot use the real interest rate to stabilize any supply shocks in state r. Thus, we trivially have

$$\pi^{r\varepsilon}\varepsilon = \varepsilon \qquad \text{and} \qquad x^{r\varepsilon}\varepsilon = -\varepsilon,$$

the random macroeconomic outcomes in state r under the optimal rule according to **Proposition 1**. This completes the proof of the proposition. QED

7.2 Proof of Proposition 2

To derive the equilibrium macroeconomic outcomes under discretion that appear in **Proposition 2**, we use the same two-step procedure as in the proof of **Proposition 1**.

Average policy Thus, we begin by deriving the average outcomes. As discussed in Subsection 2.3, a central bank that cannot commit must optimize ex post and thus takes inflation expectations as given. Formally, we study which value of $\pi^{D,Re}$ minimizes loss function

$$L(\pi^\rho, x^\rho) = \frac{1}{2}[(\pi^\rho - \overline{\pi})^2 + \lambda(x^\rho - \overline{x})^2].$$

for given realizations of all random variables, ρ, θ and ε, and for a given value of π^e. Given the macroeconomic model, this minimizer must satisfy the first-order condition

$$[(\pi^{D,Re} - \overline{\pi}) + \lambda(\theta + \pi^{D,Re} - \pi^e - \overline{x})]\frac{d\pi^{D,Re}}{di^{Re}} = 0. \qquad \text{(A9)}$$

Denoting average inflation in state r by $\pi^{D,re}$, we can write expected inflation as

$$\pi^e = q\pi^{D,re} + (1-q)\pi^{D,Re}. \qquad \text{(A10)}$$

Substituting this expression into (A9), we obtain an expression for equilibrium inflation, namely

$$\pi^{D,Re} = \overline{\pi} + \lambda(\overline{x} - \theta) - q(\pi^{D,Re} - \pi^{D,re}). \qquad \text{(A11)}$$

But clearly, this is only an intermediate solution. Let us proceed by noting that the expression for $\pi^{D,re}$ is the same as under commitment, namely (A2) – although the precise values of $\pi^{D,re}$ and π^{re} will differ (as expected inflation differs across the two regimes). Together, (A2) and (A11) give

$$\pi^{D,Re} - \pi^{D,re} = \frac{\sigma}{(1-q(1+\sigma)}(\pi^{D,re} + r),$$

and

$$\pi^{D,Re} = \frac{(1-q(1+\sigma))(\overline{\pi}+\lambda(\overline{x}-\theta))+q\lambda\sigma r}{(1-q(1+\sigma(1+\lambda))}. \quad \text{(A12)}$$

Thus, we have parametric solutions for $\pi^{D,re}$ in (A2) and $\pi^{D,Re}$ in (A12), which are related as

$$\pi^{D,re} = \pi^{D,Re} + \frac{\sigma[\overline{\pi}+\lambda(\overline{x}-\theta)+r]}{(1-q(1+\sigma(1+\lambda))}, \quad \text{(A13)}$$

where **Assumption 1** ($\overline{\pi}+\lambda(\overline{x}-\theta)+r<0$) implies $\pi^{D,re}<\pi^{D,Re}$.

Next, we use (A12) to define a short-hand expression for $\pi^{D,re}$, viz.

$$\pi^{D,Re} = \bar{\pi}+\lambda(\bar{x}-\theta)+\alpha^R(q)[\bar{\pi}+\lambda(\bar{x}-\theta)+r], \quad \text{(A14)}$$

where

$$\alpha^R(q) = \frac{\lambda\sigma q}{1-q(1+\sigma)-\lambda\sigma q} > 0$$

is strictly increasing in q. By (A12) and (A13), we can also write

$$\pi^{D,re} = \bar{\pi}+\lambda(\bar{x}-\theta)+\alpha^r(q)[\bar{\pi}+\lambda(\bar{x}-\theta)+r], \quad \text{(A15)}$$

where

$$\alpha^r(q) = \frac{\lambda\sigma q+\sigma}{1-q(1+\sigma)-\lambda\sigma q} > \alpha^R(q)$$

is also strictly increasing in q. Moreover, $\alpha^r(0)>0=\alpha^R(0)$. To take the last steps, just use (A10) and (1) to obtain the expressions for π^e and for $x^{D,\mathrm{Re}}$ and for $x^{D,re}$ (i.e. for the average values of output in the two states) in **Proposition 2**.

Finally, note that

$$\frac{\partial[(1-q)\alpha^R(q)+q\alpha^r(q)]}{\partial q} = \alpha^r(q)-\alpha^R(q)+(1-q)\frac{\partial\alpha^R(q)}{\partial q}+q\frac{\partial\alpha^r(q)}{\partial q} > 0.$$

Policy response to ε When it comes to the response to the random supply shocks ε under discretion, it is easy to show that the first-order conditions for a minimum loss coincide with those under the optimal rule under commitment – the proof is therefore omitted. This is why the random part of the macro outcomes in **Propositions 1** and **2** coincide. *QED*.

7.3 Proof of Proposition 3

As stated in the text, the equilibrium response to the ε shocks is the same under commitment and under discretion. Hence, the optimal state contingent contract only needs to address the distortion in expected inflation.

Simplify the central-bank first order condition under commitment with regard to π^{Re} by rewriting (A15) as

$$(1-q)[1-q(1+\sigma)](\pi^{C,Re}-\bar{\pi})-(1-q)\lambda q\sigma(x^{C,Re}-\bar{x})+ \quad \text{(A16)}$$

$$+q(1+\sigma)(1-q)(\pi^{C,re}-\bar{\pi})+q\lambda\sigma(1-q)(x^{C,re}-\bar{x})=0.$$

Under discretion (taking π^e as given) with the optimal contract, the first-order conditions with regard to π^R evaluated at $\varepsilon=0$ are more simply

$$(\pi^{Re}-\bar{\pi})+\lambda(x^{Re}-\bar{x})+T_\pi^R(q,\theta)=0. \tag{A17}$$

Hence, (A16) is equal to (A17) – i.e., the optimal contract implements the equilibrium under commitment – if

$$\begin{aligned} T_\pi^R(\theta) &= \{(1-q)[1-q(1+\sigma)]-1\}(\pi^{C,Re}-\bar{\pi})- \\ &\quad -\{(1-q)q\sigma+1\}\lambda(x^{C,Re}-\bar{x})+q\left\{(\pi^{C,re}-\bar{\pi})(1+\sigma)(1-q)+\lambda(x^{C,re}-\bar{x})\sigma(1-q)\right\}, \end{aligned}$$

which in turn implies that the contract has the form $T_0^R+T_\pi^R(\theta)\pi^R$. To derive a closed form expression for $T_\pi^R(\theta)$, note that in equilibrium $x^{Re}-x^{re}=\pi^{Re}-\pi^{re}$ and rewrite (9) as:

$$\begin{aligned} T_\pi^R(\theta) &= -(\pi^{C,Re}-\pi^{re})q(1+\sigma)(1-q)-q(\pi^{C,Re}-\bar{\pi})-q\lambda\sigma(1-q)(x^{C,Re}-x^{re})-\lambda(x^{C,Re}-\bar{x}) \\ &= (\pi^{C,Re}-\pi^{re})[-q(1+\sigma)(1-q)-q\lambda\sigma(1-q])-q(\pi^{C,Re}-\bar{\pi})-\lambda(x^{C,Re}-\bar{x}) \\ &= -q(1-q)(1+\sigma(1+\lambda))(\pi^{C,Re}-\pi^{C,re})-q(\pi^{C,Re}-\bar{\pi})-\lambda(x^{C,Re}-\bar{x}). \end{aligned} \tag{A18}$$

Now insert (A18) into the equilibrium expressions of the relevant variables under commitment, as stated in **Proposition 1**, to get

$$T_\pi^R(\theta)=\lambda(\bar{x}-\theta)+\beta(q)(\bar{\pi}+r), \tag{A19}$$

where

$$0\leq\beta(q)=\frac{q\sigma(1+\lambda)(1+\sigma)}{1-q+(1+\lambda)\sigma^2 q}<1, \tag{A20}$$

and where the second inequality follows from **Assumption 2.** The definition of $\beta(q)$ implies that

$$\frac{\partial\beta(q)}{\partial q}=\frac{\sigma(1+\lambda)(1+\sigma)}{[1-q+(1+\lambda)\sigma^2 q]^2}>0$$

with $\beta(0)=0$.

Finally, it follows from the definitions in (the proofs of) **Propositions 1** and **2** that $\beta(q)=\alpha^R(q)+\delta^R(q)$. *QED.*

7.4 Proof of Proposition 4

Consider a general quadratic contract on the form $T(\pi)=\tau_0+\tau_1\pi^R+\tau_2(\pi^R)^2/2$.

Equilibrium policy as a function of the contract parameters Under discretion and with full information in state R, the central bank chooses i to minimize

$$L=\frac{1}{2}[(\pi^R-\bar{\pi})^2+\lambda(x^R-\bar{x})^2]+\tau_1\pi^R+\frac{\tau_2}{2}(\pi^R)^2.$$

The first-order condition imply

$$(\pi^R - \bar{\pi}) + \lambda(x^R - \bar{x}) + \tau_1 + \tau_2 \pi^R = 0. \tag{A21}$$

Take the expectation with regard to ε and denote $\pi^{Re} = E_\varepsilon \pi^R$, $x^{Re} = E_\varepsilon x^R$. Then, (A21) implies

$$(\pi^{Re} - \bar{\pi}) + \lambda(x^{Re} - \bar{x}) + \tau_1 + \tau_2 \pi^{Re} = 0. \tag{A22}$$

Using the equlibrium relation between x^{Re} and π^{Re}, equation (A22) can also be written as:

$$(\pi^{Re} - \bar{\pi}) + \lambda(\theta - q\sigma \frac{\pi^{Re} + r}{1 - q(1+\sigma)} - \bar{x}) + \tau_1 + \tau_2 \pi^{Re} = 0.$$

This implies that the equilibrium inflation rate chosen by the central bank is (we denote equilibrium values with a * superscript):

$$\pi^{Re*} = \frac{\alpha_1[\bar{\pi} + \lambda(\bar{x} - \theta) - \tau_1] + \alpha_2 r}{\alpha_1(1+\tau_2) - \alpha_2}, \tag{A23}$$

where

$$\begin{aligned} \alpha_1 &= 1 - q(1+\sigma) \\ \alpha_2 &= \lambda\sigma q. \end{aligned}$$

Inserting (A23) in the remaining equilibrium expressions we also get:

$$x^{Re*} = \theta - \frac{q\sigma[\bar{\pi} + \lambda(\bar{x} - \theta) - \tau_1]}{\alpha_1(1+\tau_2) - \alpha_2} - \frac{q\sigma(1+\tau_2)r}{[\alpha_1(1+\tau_2) - \alpha_2]} \tag{A24}$$

$$x^{re*} = \theta + \frac{\sigma(1-q)[\bar{\pi} + \lambda(\bar{x} - \theta) - \tau_1]}{\alpha_1(1+\tau_2) - \alpha_2} + \frac{\sigma(1-q)(1+\tau_2)}{\alpha_1(1+\tau_2) - \alpha_2} r \tag{A25}$$

$$\pi^{re*} = \frac{(\alpha_1 + \sigma)[\bar{\pi} + \lambda(\bar{x} - \theta) - \tau_1] + [\alpha_2 + \sigma(1+\tau_2)]r}{\alpha_1(1+\tau_2) - \alpha_2}. \tag{A26}$$

Next, we have to compute equilibrium values of the relevant variables as a function of ε. Using (1), we rewrite (A21) as

$$(\pi^R - \bar{\pi}) + \lambda(\theta + \pi^R - \pi^e - \varepsilon \ - \bar{x}) + \tau_1 + \tau_2 \pi^R = 0.$$

Subtracting (A22) from this expression, we get:

$$(\pi^R - \pi^{Re})(1+\tau_2) + \lambda(\pi^R - \pi^{Re} - \varepsilon) = 0,$$

or

$$\pi^{R*} = \pi^{Re*} + \frac{\lambda}{(1+\lambda+\tau_2)}\varepsilon, \tag{A27}$$

where π^{Re*} is given by (A23). Using (1) and (2) we also get:

$$x^{R*} = x^{Re*} - \frac{1+\tau_2}{(1+\lambda+\tau_2)}\varepsilon \tag{A28}$$

$$x^{r*} = x^{re*} \tag{A29}$$

$$\pi^{r*} = \pi^{re*} + \varepsilon \tag{A30}$$

Optimal contract The expressions derived above fully describe the equilibrium under discretion as a function of the parameters in a quadratic contract. To characterize the optimal contract, we plug these expressions in the expected loss function and compute the optimal values of τ_1 and τ_2. That is, we now solve

$$Min_{\tau_1,\tau_2} \quad \frac{1}{2}E_\rho E_\theta E_\varepsilon[(\pi^{\rho*}-\bar{\pi})^2 + \lambda(x^{\rho*}-\bar{x})^2],$$

subject to the equilibrium expressions derived above.

Let $\pi_1^{\rho*} = \partial\pi^{\rho*}/\partial\tau_1$, $\pi_2^{\rho*} = \partial\pi^{\rho*}/\partial\tau_2$, and so on. The first-order conditions for τ_1 and τ_2 are

$$E_\theta E_\varepsilon\left\{(1-q)[(\pi^{R*}-\bar{\pi})\pi_1^{R*}] + (1-q)\lambda[(x^{R*}-\bar{x})x_1^{R*}] + q[(\pi^{r*}-\bar{\pi})\pi_1^{r*}] + q\lambda[(x^{r*}-\bar{x})x_1^{r*}]\right\} \tag{A31}$$

$$E_\theta E_\varepsilon\left\{(1-q)[(\pi^{R*}-\bar{\pi})\pi_2^{R*}] + (1-q)\lambda[(x^{R*}-\bar{x})x_2^{R*}] + q[(\pi^{r*}-\bar{\pi})\pi_2^{r*}] + q\lambda[(x^{r*}-\bar{x})x_2^{r*}]\right\} \tag{A32}$$

where

$$\begin{aligned}
\pi_1^{R*} &= \pi_1^{Re*} = -\frac{\alpha_1}{\alpha_1(1+\tau_2)-\alpha_2} \\
\pi_1^{r*} &= \pi_1^{re*} = -\frac{(1+\sigma)(1-q)}{\alpha_1(1+\tau_2)-\alpha_2} \\
x_1^{R*} &= x_1^{Re*} = \frac{q\sigma}{\alpha_1(1+\tau_2)-\alpha_2} \\
x_1^{r*} &= x_1^{re*} = -\frac{\sigma(1-q)}{\alpha_1(1+\tau_2)-\alpha_2}
\end{aligned}$$

$$\begin{aligned}
\pi_2^{R*} &= -\frac{\lambda}{(1+\lambda+\tau_2)^2}\varepsilon - \frac{\alpha_1\pi^{Re*}}{\alpha_1(1+\tau_2)-\alpha_2} \\
\pi_2^{r*} &= -\frac{\alpha_1\pi^{re*}}{\alpha_1(1+\tau_2)-\alpha_2} \\
x_2^{R*} &= -\frac{\lambda}{[(1+\lambda+\tau_2)]^2}\varepsilon - \frac{\alpha_1(x^{Re*}-\theta)+q\sigma r}{\alpha_1(1+\tau_2)-\alpha_2} \\
x_2^{r*} &= -\frac{\alpha_1(x^{re*}-\theta)-(1-q)\sigma r}{\alpha_1(1+\tau_2)-\alpha_2}.
\end{aligned}$$

Optimal value of τ_1 We begin by inserting the first four partial derivatives in (A31) using (A23)-(A26). Simplifying, taking expectations, and engaging in some tedious algebra, we get that the optimal value of τ_1

$$\tau_1^* = \lambda(\bar{x} - \bar{\theta}) + \beta(q)(\bar{\pi} + r) - \tau_2 \bar{\pi}^R, \tag{A33}$$

where $\beta(q)$ is given by (A20) in the proof of **Proposition 3.** Moreover,

$$\bar{\pi}^R = \bar{\pi} - \delta(q)(\bar{\pi} + r) > \bar{\pi}$$

with

$$1 > \delta(q) = \frac{[1 + (1 + \lambda)\sigma]\sigma q}{1 - q + (1 + \lambda)\sigma^2 q} = \beta(q) - \lambda\sigma q > 0 \tag{A34}$$

is the optimal expected rate of inflation under commitment. Thus, function $\delta(q)$

is the same function as the $\delta(q)$ derived in the proof of **Proposition 1.** Thus,

$$\frac{\partial \delta(q)}{\partial q} = \frac{[1 + (1 + \lambda)\sigma]\sigma}{[1 - q + (1 + \lambda)\sigma^2 q]^2} > 0$$

and $\delta(0) = 0$.

Optimal value of τ_2 Next, we insert the remaining partial derivatives in (A32). Using (A27)-(A30) and simplifying, we get

$$\begin{aligned}
&-(1-q)E_\theta E_\varepsilon \left\{ [\pi^{Re*} + \frac{\lambda}{(1 + \lambda + t_2)}\varepsilon - \bar{\pi}][\frac{\lambda}{(1 + \lambda + t_2)^2}\varepsilon + \frac{\alpha_1 \pi^{Re*}}{\alpha_1(1 + t_2) - \alpha_2}] \right\} - \\
&-(1-q)\lambda E_\theta E_\varepsilon \left\{ [x^{Re*} - \frac{1 + t_2}{(1 + \lambda + t_2)}\varepsilon - \bar{x}][\frac{\lambda}{[(1 + \lambda + t_2)]^2}\varepsilon + \frac{\alpha_1(x^{Re*} - \theta) + q\sigma r}{\alpha_1(1 + t_2) - \alpha_2}] \right\} - \\
&-qE_\theta E_\varepsilon \left\{ [\pi^{re*} + \varepsilon - \bar{\pi}][\frac{\alpha_1 \pi^{re*}}{\alpha_1(1 + t_2) - \alpha_2}] \right\} - \\
&-q\lambda E_\theta E_\varepsilon \left\{ (x^{re*} - \bar{x})[\frac{\alpha_1(x^{re*} - \theta) - (1 - q)\sigma r}{\alpha_1(1 + t_2) - \alpha_2}] \right\} = 0.
\end{aligned} \tag{A35}$$

In this expression, π^{Re*}, π^{re*}, x^{Re*}, π^{re*} are all linear functions of θ and τ_1 and non-linear functions of τ_2, as given by (A23)-(A26). Taking expectations with respect to θ and ε, using the expression for the optimal value of τ_1, τ_1^* in (A33), simplifying, and engaging in another round of tedious algebra, we get

$$\begin{aligned}
&-A(\tau_2)^2 \left\{ \frac{1 - q + q\sigma^2(1 + \lambda)}{\alpha_1} \right\} (1 - q) - \\
&-\frac{[1 - q + q\sigma^2(1 + \lambda)](1 - q)\alpha_1\lambda^2}{[\alpha_1(1 + \tau_2) - \alpha_2]^2} \upsilon_\theta - \\
&-A(\tau_2) \left\{ [2(1 + \sigma) + \lambda\sigma]\frac{\sigma q r}{\alpha_1} - \bar{\pi} \right\} (1 - q) - \\
&-q\sigma r \left\{ [\frac{\sigma r}{\alpha_1} - \bar{\pi}] \right\} + \frac{[\alpha_1(1 + \tau_2) - \alpha_2](1 - q)\lambda^2\tau_2}{(1 + \lambda + \tau_2)^3} \upsilon_\varepsilon = 0,
\end{aligned} \tag{A36}$$

where υ_θ and υ_ε denote the variances of θ and ε and where

$$A(\tau_2) = \frac{\alpha_1(1-\beta)\bar{\pi} + (\alpha_2 - \alpha_1\beta)r + \tau_2\bar{\pi}^R}{\alpha_1(1+\tau_2) - \alpha_2}.$$

Equation (A36) implicitly defines the optimal value of τ_2 as a function of the remaining parameters, and in particular of the two variances, $\tau_2^* = F(\upsilon_\theta, \upsilon_\varepsilon)$. Under **Assumption 2**, $\alpha_1 - \alpha_2 > 0$. If $\tau_2^* > 0$, the LHS of (A36) is thus increasing in υ_ε and decreasing in υ_θ. If the second-order condition for τ_2 is satisfied and if $\tau_2^* > 0$, it follows by the implicit function theorem that $\partial\tau_2^*/\partial\upsilon_\theta > 0 > \partial\tau_2^*/\partial\nu_\varepsilon$. To show that $\tau_2^* > 0$, consider the LHS of (A36) at the point $\tau_2 = 0$. It can be written as

$$\begin{aligned}
&(A(0))^2 \left\{\frac{1-q+q\sigma^2(1+\lambda)}{\alpha_1}\right\}(1-q) + \\
&+\frac{[1-q+q\sigma^2(1+\lambda)](1-q)\alpha_1\lambda^2}{[\alpha_1-\alpha_2]^2}\upsilon_\theta + \\
&+A(0)\left\{[2(1+\sigma)+\lambda\sigma]\frac{\sigma q r}{\alpha_1} - \bar{\pi}\right\}(1-q) + q\sigma r[\frac{\sigma r}{\alpha_1} - \bar{\pi}],
\end{aligned} \tag{A37}$$

which is certainly positive if υ_θ is large enough.

To show that the second-order condition is satisfied, take the derivative of the LHS of (A36) with regard to τ_2, recalling that to derive (A36) we substituted away τ_1^* from (A33). This gives

$$\begin{aligned}
&2A(\tau_2)A_2\left\{\frac{1-q+q\sigma^2(1+\lambda)}{\alpha_1}\right\}(1-q) - \\
&-2\frac{[1-q+q\sigma^2(1+\lambda)](1-q)\alpha_1\lambda^2}{[\alpha_1(1+\tau_2)-\alpha_2]^3}\alpha_1\upsilon_\theta + A_2\left\{[2(1+\sigma)+\lambda\sigma]\frac{\sigma q r}{\alpha_1} - \bar{\pi}\right\}(1-q) - \\
&-\frac{[\alpha_1(1+\tau_2)-\alpha_2](1-q)\lambda^2}{(1+\lambda+\tau_2)^3}\upsilon_\varepsilon - \frac{\alpha_1(1-q)\lambda^2\tau_2}{(1+\lambda+\tau_2)^3}\upsilon_\varepsilon + 3\frac{[\alpha_1(1+\tau_2)-\alpha_2](1-q)\lambda^2\tau_2}{(1+\lambda+\tau_2)^4}\upsilon_\varepsilon,
\end{aligned} \tag{A38}$$

where

$$A_2 = \frac{\partial A(\tau_2)}{\partial \tau_2} = \frac{\bar{\pi}^R - \alpha_1 A(\tau_2)}{\alpha_1(1+\tau_2)-\alpha_2} \gtrless 0.$$

Hence, if υ_θ is large enough, the expression in (A38) is negative, implying that the second-order condition for τ_2 is satisfied.

Summary To summarize this lengthy and laborious proof, we have shown that if **Assumption 2** is satisfied and υ_θ is large enough, we can characterize the parameter values in the optimal contract $T(\pi) = \tau_0 + \tau_1\pi^R + \tau_2(\pi^R)^2/2$. Specifically, $\tau_2^* > 0$ is increasing in υ_θ and decreasing in υ_ε. We also get

$$\tau_1^* = \lambda(\bar{x} - \bar{\theta}) + \beta(q)(\bar{\pi} + r) - \tau_2^*\bar{\pi}^R \tag{A39}$$

with

$$\bar{\pi}^R = \bar{\pi} - \delta(q)(\bar{\pi} + r) > \bar{\pi},$$

where functions $\beta(q) > \delta(q)$ are both increasing in q with $\beta(0) = \delta(0) = 0$.

Exploiting (A39), we can then rewrite the optimal contract as

$$T(\pi) = T_0^R + T_\pi^R(\overline{\theta})\pi^R + \frac{T_{\pi\pi}^R(q)}{2}(\pi^R - \bar{\pi}^R)^2,$$

where

$$T_0^R = \tau_0 + \frac{T_{\pi\pi}^R(q)}{2}(\bar{\pi}^R)^2$$

$$T_\pi^R(\overline{\theta}) = \lambda(\bar{x} - \overline{\theta}) + \beta(q)(\bar{\pi} + r)$$

$$T_{\pi\pi}^R(q) = \tau_2^*,$$

and where τ_0 can be arbitrarily defined so as to meet the central bank participation constraint of the central bank's leaders. These are the properties stated in **Proposition 4.** *QED.*

7.5 Proof of Proposition 5

Discretion Under discretion, the central bank minimizes

$$L(\pi^\rho, x^\rho) = \frac{1}{2}[(\pi^\rho - \overline{\pi})^2 + \lambda(x^\rho - \overline{x})^2 + \gamma\kappa g] \ .$$

for a given realization of all random variables, $\rho, \theta, \varepsilon$, given (1) and (10), and taking expected inflation as given. At $\rho = R$, the central bank only uses i, since it is not costly. At $\rho = r$, it uses g as i is constrained by the ZLB. Using also the expression for aggregate supply, we get:

$$\begin{aligned} x^R &= \theta - \sigma(i - \pi^e - R) \\ x^r &= \theta + \sigma(\pi^e + r) + g \end{aligned}$$

$$\pi^R = (1+\sigma)\pi^e - \sigma(i - R) + \varepsilon$$

$$\pi^r = (1+\sigma)\pi^e + g + \sigma r + \varepsilon.$$

We employ the same notation as earlier, namely $\pi^{Re} = E_\varepsilon \pi^R$, $x^{Re} = E_\varepsilon x^R$ and so on. Then, we can write.

$$\pi^e = (1-q)\pi^{Re} + q\pi^{re} = \pi^{Re} - q(\pi^{Re} - \pi^{re}).$$

Let us first look at the optimality conditions with regard to i and g, taking π^e as given. They imply

$$(\pi^R - \bar{\pi}) + \lambda(x^R - \bar{x}) = 0 \tag{A40}$$

$$(\pi^r - \bar{\pi}) + \lambda(x^r - \bar{x}) + \gamma\kappa = 0. \tag{A41}$$

Subtract (A41) from (A40), and use the fact that from the supply curves $x^R - x^r = \pi^R - \pi^r$. This allows us to write

$$(\pi^R - \pi^r) = \gamma\kappa/(1+\lambda). \tag{A42}$$

Hence, expected inflation is

$$\pi^e = \pi^{Re} - \frac{\gamma q \kappa}{1+\lambda}. \tag{A43}$$

Next, we take expectation with regard to ε and use (A43) and the supply curve to rewrite (A40) as

$$\pi^{Re} - \bar{\pi} + \lambda(\theta + \frac{\gamma q \kappa}{1+\lambda} - \bar{x}) = 0,$$

such that

$$\pi^{Re} = \bar{\pi} + \lambda(\bar{x} - \theta) - \frac{\gamma q \lambda \kappa}{1+\lambda}. \tag{A44}$$

Insert (A44) in (A42) to get

$$\pi^{re} = \bar{\pi} + \lambda(\bar{x} - \theta) - \frac{(1+q\lambda)\gamma\kappa}{1+\lambda}. \tag{A45}$$

Expected inflation is thus

$$\pi^e = \bar{\pi} + \lambda(\bar{x} - \theta) - \gamma q \kappa.$$

To get the reaction to ε, write $\pi^\rho = \pi^{\rho e} + \pi^{\rho\varepsilon}\varepsilon$, with $\pi^{\rho\varepsilon}$ yet to be determined. We can write (A40) as

$$\begin{aligned}
\pi^{Re} + \pi^{\rho\varepsilon}\varepsilon - \bar{\pi} + \lambda(\theta + \pi^{Re} + \pi^{\rho\varepsilon}\varepsilon - \pi^e - \varepsilon - \bar{x}) &= 0 \\
\pi^{\rho\varepsilon}\varepsilon + \lambda(\pi^{\rho\varepsilon}\varepsilon - \varepsilon) &= 0 \\
(1+)\lambda\pi^{\rho\varepsilon}\varepsilon &= \lambda\varepsilon,
\end{aligned}$$

which gives

$$\pi^{\rho\varepsilon} = \frac{\lambda}{1+\lambda}.$$

Finally, we insert these results in the output equations, to get

$$\begin{aligned}
x^R &= \theta + \frac{\gamma q \kappa}{1+\lambda} - \frac{1}{1+\lambda}\varepsilon \\
x^r &= \theta + \gamma\kappa(q - \frac{(1+q\lambda)}{1+\lambda}) - \frac{1}{1+\lambda}\varepsilon \\
&= \theta - \frac{(1-q)\gamma\kappa}{1+\lambda} - \frac{1}{1+\lambda}\varepsilon.
\end{aligned}$$

Commitment Take the supply curves and the expression for expected inflation,

$$x^{\rho} = \theta + (\pi^{\rho} - \pi^{e}),$$

$$\pi^{e} = (1-q)\pi^{Re} + q\pi^{re} = \pi^{Re} - q(\pi^{Re} - \pi^{re}).$$

It follows that

$$\begin{aligned} x^{Re} &= \theta + q(\pi^{Re} - \pi^{re}) \\ x^{re} &= \theta + (1-q)(\pi^{re} - \pi^{Re}). \end{aligned}$$

Further, by the aggregate-demand function

$$\begin{aligned} g^{e} &= \pi^{re} - (1+\sigma)\pi^{e} - \sigma r \\ &= [1-(1+\sigma)q]\pi^{re} - (1+\sigma)(1-q)\pi^{Re} - \sigma r. \end{aligned}$$

The central bank minimizes by choice of π^{Re} and π^{re} :

$$\frac{1}{2}E_{\varepsilon}E_{\theta}(1-q)[(\pi^{R} - \overline{\pi})^{2} + \lambda(x^{R} - \overline{x})^{2}] + \frac{1}{2}E_{\varepsilon}E_{\theta}q[(\pi^{r} - \overline{\pi})^{2} + \lambda(x^{r} - \overline{x})^{2} + \kappa g]$$

Hence, the first order conditions are, respectively,

$$(1-q)(\pi^{Re} - \overline{\pi}) + \lambda q(1-q)(x^{Re} - \overline{x}) - q(1-q)\lambda(x^{re} - \overline{x}) - q(1+\sigma)(1-q)\kappa = 0$$

$$-(1-q)q\lambda(x^{Re} - \overline{x}) + q(\pi^{re} - \overline{\pi}) + q(1-q)\lambda(x^{re} - \overline{x}) + q[1-(1+\sigma)q]\kappa = 0.$$

Noting that $\pi^{Re} - \pi^{re} = x^{Re} - x^{re}$, rewrite the first-order conditions as

$$(1-q)(\pi^{Re} - \overline{\pi}) + \lambda q(1-q)(\pi^{Re} - \pi^{re}) - q(1+\sigma)(1-q)\kappa = 0$$

$$-(1-q)q\lambda(\pi^{Re} - \pi^{re}) + q(\pi^{re} - \overline{\pi}) + q[1-(1+\sigma)q]\kappa = 0$$

and add them together to get

$$(1-q)(\pi^{Re} - \overline{\pi}) + q(\pi^{re} - \overline{\pi}) = q\sigma\kappa$$

or

$$\pi^{e} = (1-q)\pi^{Re} + q\pi^{re} = \bar{\pi} + q\sigma\kappa > \bar{\pi}.$$

Moreover,

$$(\pi^{Re} - \pi^{re}) = [\pi^{Re} - \bar{\pi} - q\sigma\kappa]/q$$

$$(\pi^{Re} - \pi^{re}) = [\bar{\pi} + q\sigma\kappa - \pi^{re}]/(1-q).$$

Insert these expressions in the optimality conditions and simplify to obtain

$$\pi^{Re} = \overline{\pi} + \frac{[q(1+\sigma) + \lambda q\sigma]\kappa}{(1+\lambda)} > \overline{\pi}$$

$$\pi^{re} = \overline{\pi} - \frac{[1 - q[(1+\sigma(1+\lambda)]]\kappa}{(1+\lambda)} < \overline{\pi},$$

by **Assumption 2**.

Finally, using the expressions for output, we get

$$\begin{aligned} x^{Re} &= \theta + \frac{q\kappa}{(1+\lambda)} \\ x^{re} &= \theta - \frac{(1-q)\kappa}{(1+\lambda)}. \end{aligned}$$

We don't repeat the computations for the optimal stabilization policies, since they coincide with those under discretion. The proof is similar to that of Propositions 2 and 3.

Difference between discretion and commitment Taking the difference between these equilibrium expressions, we have:

$$\begin{aligned} \pi^{D,R} - \pi^{C,R} &= \lambda(\bar{x} - \theta) - \frac{[q(1+\sigma) + \lambda q(\sigma+\gamma)]\kappa}{(1+\lambda)} \\ &= \lambda(\bar{x} - \theta) + \frac{(1-\gamma)\lambda q\kappa}{(1+\lambda)} - q(1+\sigma)\kappa \\ \pi^{D,r} - \pi^{C,r} &= \lambda(\bar{x} - \theta) + \frac{[1-(1+q\lambda)\gamma - q[(1+\sigma(1+\lambda)]]\kappa}{(1+\lambda)} \\ &= \lambda(\bar{x} - \theta) + \frac{(1-\gamma)(1+q\lambda)\kappa}{(1+\lambda)} - q(1+\sigma)\kappa \end{aligned}$$

$$\begin{aligned} \pi^{D,e} - \pi^{C,e} &= \lambda(\bar{x} - \theta) + q(1-\gamma)\kappa - q(1+\sigma)\kappa \\ x^{D,R} - x^{C,R} &= -\frac{q(1-\gamma)\kappa}{(1+\lambda)} \\ x^{D,r} - x^{C,r} &= \frac{(1-q)(1-\gamma)\kappa}{(1+\lambda)}. \end{aligned}$$

Moreover, note that:

$$\begin{aligned} \frac{\partial(\pi^{D,Re} - \pi^{C,Re})}{\partial\gamma} &= -\frac{\lambda q\kappa}{(1+\lambda)} < 0 \\ \frac{\partial(\pi^{D,re} - \pi^{C,re})}{\partial\gamma} &= -\frac{(1+q\lambda)\kappa}{(1+\lambda)} < \frac{\partial(\pi^{D,Re} - \pi^{C,Re})}{\partial\gamma}. \end{aligned}$$

This completes the proof of **Proposition 5**. *QED*.