

The Optimal Degree of Monetary-Discretion in a New Keynesian Model with Private Information*

Yuichiro Waki[†]

Richard Dennis

University of Queensland

University of Glasgow

Ippei Fujiwara

Keio University/ANU

January 14, 2015

Abstract

This paper considers the optimal degree of discretion in monetary policy when the central bank conducts policy based on its private information about the state of the economy and is unable to commit. Society seeks to maximize social welfare by imposing restrictions on the central bank's actions over time, and the central bank takes these restrictions and the New Keynesian Phillips curve as constraints. By solving a dynamic mechanism design problem we find that it is optimal to grant "constrained discretion" to the central bank by imposing both upper and lower bounds on permissible inflation, and that these bounds must be set in a history-dependent way. The optimal degree of discretion varies over time with the severity of the time-inconsistency problem, and, although no discretion is optimal when the time-inconsistency problem is very severe, our numerical experiment suggests that no-discretion is a transient phenomenon, and that some discretion is granted eventually.

*This study is conducted as a part of the Project "Social Security, Taxation, and Public Finance" undertaken at Research Institute of Economy, Trade and Industry (RIETI). We thank Ctirad Slavik, Ryo Jinnai, Giuseppe Moscarini, Begona Dominguez, and seminar/conference participants at Kyoto University, the Reserve Bank of New Zealand, the University of Auckland, Australian National University, the CIGS Conference on Macroeconomic Theory and Policy 2014, the Federal Reserve Bank of Philadelphia, the Econometric Society North American Summer Meeting 2014, Hitotsubashi University, and RIETI for their useful comments.

[†]Address for correspondence: Rm 613 Colin Clark Building, School of Economics, University of Queensland, Brisbane, Australia; email: y.waki@uq.edu.au

1 Introduction

How much flexibility should society allow a central bank in its conduct of monetary policy? At the center of the case for flexibility is the argument that central bankers have private information (Canzoneri, 1985) or information that are difficult to be encoded in a rule, perhaps about the economy's state or structure, or perhaps about the distributional costs of inflation arising through heterogeneous preferences (Sleet, 2004). If central banks have flexibility over policy decisions, then this gives them the ability to use for the public's benefit any private or non-contractible information that they have. However, if central banks face a time-inconsistency problem (Kydland and Prescott, 1977), then it may be beneficial to limit their flexibility. Institutionally, many countries have balanced these competing concerns by delegating monetary policy to an independent central bank that is required to keep inflation outcomes low and stable, often within a stipulated range, but that is otherwise given the freedom to conduct policy without interference. Inflation targeting is often characterized as "constrained discretion" (Bernanke and Mishkin, 1997) precisely because it endeavors to combine flexibility with rule-like behavior.

This paper examines the optimal degree of discretion in a monetary-policy delegation problem when the central bank has private information on the state of the economy and is unable to commit. We take the legislative approach of Canzoneri (1985) and Athey, Atkeson, and Kehoe (2005) (AAK, hereafter). Specifically, society imposes restrictions on the central bank's actions, and the benevolent central bank conducts policy subject to these restrictions and to a Phillips curve. Society cannot achieve the first-best because of the central bank's private information, but some restrictions on the central bank can ameliorate its inability to commit and are therefore beneficial. We solve a dynamic mechanism design problem to examine how much discretion society should grant to the central bank and to reveal the form of the optimal constrained discretion policy.

Unlike in AAK, we find that the optimal mechanism is history-dependent. A key aspect of our analysis is that inflation outcomes are governed by a forward-looking New Keynesian Phillips curve. This Phillips curve relates inflation outcomes to the output gap and to expected *future* inflation and allows policy-makers to deliver better outcomes today by tailoring future policy according to the current state of the economy, thereby giving a crucial role to policy

promises. We show that the optimal direct mechanism can be expressed as a function of last period's promised inflation and of the central bank's current private information (its type).

For each value of last period's promised inflation, the optimal mechanism has the interpretation of an "interval delegation," where society specifies an interval for permissible inflation and the central bank chooses from that interval. In general this interval does not serve as a binding constraint for some types, and we interpret that these types have *discretion*. Importantly, this interval, and hence the number of types that have discretion, varies with last period's promised inflation so that the central bank is incentivized to deliver inflation that is, on average, consistent with last period's promised inflation. There are, as a result, only three types of discretionary outcomes — no discretion when this interval constrains all types, full discretion when the interval does not constrain any type, and bounded discretion when the interval constrains only a subset of types.

How does the optimal degree of discretion vary with last period's promised inflation? There is one value of promised inflation at which full discretion is granted, and social welfare is maximized at that value. The further last period's inflation promise departs from this value, the less degree of discretion is granted, and in extreme cases no discretion is granted. This pattern is naturally explained by the severity of the time-inconsistency problem. For the central bank the gain from reneging on last period's inflation promise crucially depends on the value of promised inflation. At the welfare-maximizing value of promised inflation, promised inflation is delivered even if society lets the central bank conduct policy without restriction, and granting full-discretion is optimal. The gain from reneging increases as promised inflation departs from its welfare-maximizing value, making the time-inconsistency problem more severe, and society must impose tighter restrictions on the central bank's actions in order to deliver the promised inflation, reducing the central bank's degree of discretion.

The optimal mechanism also exhibits an interesting, limited form of history-dependence — history as encoded in the state variable is disregarded for types that have discretion. We find that for each type of the central bank there is an interval of inflation promise in which that type has discretion. In such an interval, inflation, the output gap, and the continuation mechanism depend on the history only through the current value of private information, and the history-dependence is disposed of. This property resembles the "amnesia" property that

Kocherlakota (1996) finds in a full-information limited-commitment model of risk-sharing.

How can we implement the second-best with a non-direct mechanism? We propose a history-dependent inflation targeting scheme which stipulates a band of permissible inflation that varies with inflation promise announced by the central bank last period. This scheme allows the central bank to constrain its future-self through its choice of inflation promise, mitigating the time-inconsistency problem. We show that when designed appropriately this scheme implements the outcome of the optimal direct mechanism. Importantly, unlike in AAK, there are situations in which a lower limit of inflation imposes a binding constraint on the central bank's choice.

Finally, we examine how the optimal degree of discretion changes over time, using a numerical example. We find that some discretion are always granted in the ergodic set of inflation promise. This implies that, even if we impose a hypothetical initial inflation promise made in period -1 , no-discretion is at most a short-run, transient phenomenon, and some discretion are eventually granted. Interestingly, no discretion is given only when the initial inflation promise is sufficiently far from the value that maximizes social welfare. The ergodic set contains the peak of the social welfare function, from which the fully optimal mechanism starts off. Therefore, some discretion is always granted in the fully optimal mechanism.

The remainder of this paper is organized as follows. Section 2 reviews related literatures. Section 3 describes the set-up and illustrates how private information enters the model. Section 4 formulates an optimal (direct) mechanism design problem, along with two benchmark policies, the full-information policy and the optimal discretionary policy, that serve as counterpoints to the optimal private-information policy. In Section 5 we discuss theoretical results. Section 6 presents the numerical results that emerge from the benchmark policies and from the private-information policy. Section 7 offers concluding comments. Appendices contain technical material, including proofs of theoretical results and complete descriptions of how the various solutions were computed.

2 Related literature

We build on the literature of monetary policy with private information, which includes Canzoneri (1985), Sleet (2001), and AAK. Like ourselves, they study models in which the central bank receives a private signal about the state of the economy and conducts policy subject to a

Phillips curve. Their settings are distinct from ours in that they use a static Phillips curve, containing contemporaneous rather than forward-looking inflation expectation, which severs the connection between time-inconsistency and history dependence.¹ By using a forward-looking Phillips curve we show that the optimal degree of discretion should vary over time in a history-dependent manner.

Our work is also related to the vast literature on policy making in New Keynesian models with symmetric information (Woodford, 2003; Galí, 2008). This literature has generally focused on settings in which the society cannot directly constrain the central bank's action set and in which granting some discretion to the central bank is simply suboptimal.² Our paper differs from this literature in that it introduces private information on the side of the central bank and uses the legislative approach to examine the optimal balance between rules and discretion. By focusing on constrained discretion, our work is related to the literature on inflation targeting, as summarized in, for example, Bernanke, Laubach, Mishkin, and Posen (2001).

Finally, our paper is related to the literatures on optimal delegation and dynamic contracting. While static problems are typically considered in the optimal-delegation literature, our problem is dynamic and we show how the optimality of interval delegation generalizes to dynamic settings.³ Using an approach akin to Athey, Bagwell, and Sanchirico's (2004), we show that our problem can be formulated recursively as a function-valued dynamic programming problem in which last period's inflation promise serves as the state variable. This formulation not only enables us to characterize theoretically the optimal mechanism, but also reduces significantly the computational burden, compared to a set-valued dynamic programming approach (Abreu, Pearce, and Stacchetti, 1990), which is common in the dynamic contracting literature.⁴

¹Canzoneri (1985) analyzes the effects of several specific rules that are incentive-compatible but not necessarily optimal. Sleet (2001) considers an optimal incentive-compatible mechanism in a full-fledged general equilibrium model with two types, and AAK does the same in a reduced-form model with a continuum of types.

²In Kurozumi (2008), private agents behave strategically, and they may be able to deter the central bank from taking undesirable actions on the equilibrium path. In some studies it is assumed that society can assign a loss function to a central bank and that the central bank is required to minimize it (e.g. Jensen (2002)). Neither approach allows society to remove certain actions from the central bank's choice set.

³For static delegation problems, Holmström (1984), Alonso and Matouschek (2008), and Amador and Bagwell (2013) give some sufficient conditions for interval delegation to be optimal. Atkeson (1991), Sleet (2004), and Amador, Werning, and Angeletos (2006) essentially consider dynamic delegation problems. Athey, Atkeson, and Kehoe (2005) consider a repeated delegation problem, but, as the optimal mechanism is shown to be static, their problem in the end reduces to a static delegation problem.

⁴A similar result is obtained in Atkeson (1991) in a hidden action model of an optimal international lending, and in Sleet (2004) in a two-type hidden information model of optimal taxation. We consider a monetary policy model with hidden information and a continuum of types. A common feature of Atkeson (1991), Sleet (2004), and this paper is the assumption that the objectives of the mechanism designer and the agent coincide, which is not very common in the literature of dynamic contract.

3 The set-up

Our set-up is similar to the canonical setting that is used in the New Keynesian policy literature to analyze the optimal policy without commitment (see e.g. Woodford, 2003). We consider an infinite horizon economy that has a central bank and the private sector. Time is discrete and goes from $t = 0$ to infinity. Each period, the central bank conducts monetary policy subject to the New Keynesian Phillips curve, but because it is unable to commit to its future actions it takes the private sector's inflation expectation as given. However, there are two important differences. First, in our set-up the central bank privately observes shocks that hit the economy every period, and, for this reason, we incorporate a communication stage wherein the central bank sends a message to society. Second, society can, each period, limit the central bank's action by specifying the set of acceptable actions from which the central bank must choose.

Policy is conducted each period, and at the beginning of each period, society specifies a compact set of acceptable pairs of inflation and the output gap from which the central bank *must* choose. We call this set a delegation set and denote it by D . A delegation set D must be a subset of $\Pi \times X \subset \mathbb{R}^2$, where $\Pi := [\underline{\pi}, \bar{\pi}]$ and $X = [\underline{x}, \bar{x}]$ are (large) compact intervals in \mathbb{R} that contain all the available inflation and the output gap choices, respectively. Essentially a delegation set is a menu of alternatives that society offers to the central bank, and it determines the central bank's degree of flexibility: the set D may consist of only one option, forcing the central bank to choose that action, or it may contain a number of options, giving the central bank some flexibility over its action.

After receiving D , the central bank privately observes the state of the economy, θ , which is drawn from a compact interval $\Theta := [\underline{\theta}, \bar{\theta}] \subset \mathbb{R}$, according to an i.i.d. density p . The density is strictly positive everywhere, i.e. $p(\theta) > 0$ for all $\theta \in \Theta$, and its cumulative distribution function is denoted by P . Society and the private sector never observes the state, θ .

After the central bank observes θ , public communication takes place. Specifically, the central bank sends a message $m \in M$, where M is a message space, to society, which is observed also by the private sector. The private sector then forms its *one-period-ahead* inflation expectation, $\pi^e \in \Pi$. It is crucial that this is expected, next period's inflation. We assume that the private sector does not act strategically, and that its sole objective is to form rational expectation regarding next period's inflation.

Once π^e is formed, the central bank chooses inflation, π , and the output gap, x , from D , taking π^e as given. Both π and x are publicly observable, and the central bank must choose them so that they satisfy the New Keynesian Phillips curve (NKPC):

$$\pi = \kappa x + \beta \pi^e, \quad (1)$$

where $\kappa > 0$ and $\beta \in (0, 1)$ are parameters that do not vary over time or with the state, θ .⁵ Given a delegation set D and an inflation expectation, π^e , the state does not affect the central bank's set of feasible actions. Equation (1) is a standard log-linear NKPC without a cost-push shock, and is forward-looking in that it involves expected future inflation.⁶ We assume that X contains both $(\underline{\pi} - \beta \bar{\pi})/\kappa$ and $(\bar{\pi} - \beta \underline{\pi})/\kappa$.

Social welfare is time-separable with the discount factor β . The momentary social welfare function, $R(\pi, x, \theta)$, depends on inflation, the output gap, and the state of the economy.⁷ The central bank is benevolent and $R(\pi, x, \theta)$ also equals its momentary payoff. The return function, R , is continuous in (π, x, θ) , and is strictly concave in (π, x) . We allow R to depend on θ to reflect the time-varying welfare costs of inflation and the output gap, a dependence that can arise, for example, if the re-distributional effects of inflation are time varying.⁸ An example of R is the following quadratic specification:

$$R(\pi, x, \theta) = -\frac{1}{2}(\pi - \theta)^2 - \frac{1}{2}bx^2, \quad b > 0, \quad (2)$$

where θ represents the inflation rate that minimizes the welfare loss from inflation.

We assume that, although the central bank is unable to commit, society is able to com-

⁵We follow a standard practice in the New Keynesian policy literature when assuming that the central bank directly chooses inflation and the output gap subject to NKPC (see e.g. Gali, 2008). This assumption is based on the idea that the central bank can only implement policies that are consistent with private-sector incentives. The NKPC constrains the central bank because the central bank has a first-mover advantage relative to the private sector within the period.

⁶The NKPC constitutes an equilibrium condition in many New Keynesian models, and it can be derived from various costly price adjustment models, including time-dependent pricing specifications, such as Calvo-style pricing (Calvo, 1983) and quadratic price adjustment costs (Rotemberg, 1982), as well as some state-dependent pricing specifications, such as Gertler and Leahy (2008).

⁷This is not inconsistent with the unobservability of θ . We can interpret θ as a private signal for an observable shock, s , and $R(\pi, x, \theta)$ as the expected social welfare conditional on θ , $E[r(\pi, x, s)|\theta]$, where $r(\pi, x, s)$ is realized social welfare. As long as the conditional distribution of s given θ has full support, true value of θ is never revealed even if s or $r(\pi, x, s)$ is observable.

⁸AAK make the same assumption and interpret it as follows: “[i]ndividual agents in the economy have either heterogeneous preferences or heterogeneous information regarding the optimal inflation rate, and the monetary authority sees an aggregate of that information that the private agents do not see.”

mit. Due to its inability to commit, the central bank is unable to manage the private sector’s inflation expectation by committing to a certain inflation choice in the next period. Society can improve welfare, because the way it sets a future delegation set affects the future central bank’s inflation choice, thereby influencing the private sector’s inflation expectation. However, an overly restrictive delegation set can prevent the central bank from utilizing its private information and may be undesirable. The question we ask can be framed as, how should society design delegation sets in order to maximize social welfare?

3.1 Discussion

Our set-up shares much in common with AAK. The distinct feature of our set-up is the forward-looking NKPC. In contrast, AAK’s benchmark example assumes a static Phillips curve,

$$\pi_t = \pi_t^e - (u_t - u^n), \tag{3}$$

where u is the unemployment rate, π^e is expected *contemporaneous* inflation, rather than expected future inflation, and u^n is the natural rate of unemployment. Equation (3) implies that the set of pairs of inflation and the output gap that the central bank can choose is independent of future policy. We view the forward-looking Phillips curve in equation (1) as more relevant, because it is a center-piece of many New Keynesian models and is widely used in central banks. Moreover, it captures an important channel for policy, allowing central banks to use forward-guidance to manage inflation expectations. In addition, the forward-looking NKPC curve gives rise to a “stabilization bias” (see e.g. Svensson, 1997), which differs from the “inflation bias” (Barro and Gordon, 1983), present in AAK and Sleet (2001). Our set-up enables us to examine what implications this difference has on policy.

As is usual in the delegation literature, the central bank is not allowed to choose inflation and the output gap that are not contained in the delegation set, D . We have to assume this, because the principal (society) cannot directly influence social welfare and it lacks a tool to punish such observable deviations. An alternative approach could be to have the private sector set inflation strategically (while the central bank sets only the output gap), and to consider the best sustainable equilibrium in that game, as Kurozumi (2008) does in a full information setting. Our set-up provides a useful benchmark for analyses of that kind.

4 Optimal mechanism design problem

In light of the Revelation Principle, we consider a direct revelation game in which the message space M equals Θ and the central bank reports its private information each period. We focus on public strategies that depend on history only through the central bank's report history. Because society is able to commit, society's problem is to choose its strategy so that the best equilibrium given that strategy yields the highest (period-0) social welfare. We formulate this problem as a dynamic mechanism design problem in which society specifies a direct mechanism, instead of delegation sets, that maps a history of reports into inflation, the output gap, and inflation expectation. Although this potentially expands society's set of tools, we later show that society can implement the optimal direct mechanism through appropriately specified delegation sets.

A mechanism is a sequence of measurable functions $\{(\pi_t, x_t, \pi_t^e)\}_{t=0}^\infty$ such that, for all t , $(\pi_t, x_t, \pi_t^e) : \Theta^{t+1} \rightarrow \Pi \times X \times \Pi$ are functions of report history. Society must choose a mechanism that satisfies the NKPC and is consistent with rational expectations: for all t and report history $\theta^t := (\theta_0, \theta_1, \dots, \theta_t) \in \Theta^{t+1}$,

$$\pi_t(\theta^t) = \kappa x_t(\theta^t) + \beta \pi_t^e(\theta^t), \quad (4)$$

and

$$\pi_t^e(\theta^t) = \int_{\Theta} \pi_{t+1}(\theta^t, \theta_{t+1}) p(\theta_{t+1}) d\theta_{t+1}. \quad (5)$$

For simplicity, we refer to equations (4) and (5) as the *feasibility constraint*.

The central bank chooses how to report its type over time. A reporting strategy is a sequence of measurable functions $\sigma := \{\sigma_t\}_{t=0}^\infty$ with $\sigma_t : \Theta^{t+1} \rightarrow \Theta$ for all t . The truth-telling strategy is a reporting strategy with $\sigma_t(\theta^t) = \theta_t$ for all t and θ^t . A mechanism $\{(\pi_t, x_t, \pi_t^e)\}_{t=0}^\infty$ is said to be *incentive-compatible* if and only if, for any report history θ^{t-1} , for any current type θ_t and for any reporting strategy σ ,

$$\begin{aligned} & R(\pi_t(\theta^{t-1}, \theta_t), x_t(\theta^{t-1}, \theta_t), \theta_t) + \beta \sum_{s=t+1}^{\infty} \int_{\Theta^{s-t}} \beta^{s-t-1} R(\pi_s(\theta^{t-1}, \theta_t^s), x_s(\theta^{t-1}, \theta_t^s), \theta_s) \mu^{s-t}(d\theta_{t+1}^s) \\ & \geq R(\pi_t(\theta^{t-1}, \sigma_0(\theta_t)), x_t(\theta^{t-1}, \sigma_0(\theta_t)), \theta_t) \\ & \quad + \beta \sum_{s=t+1}^{\infty} \int_{\Theta^{s-t}} \beta^{s-t-1} R(\pi_s(\theta^{t-1}, \sigma^{s-t}(\theta_t^s)), x_s(\theta^{t-1}, \sigma^{s-t}(\theta_t^s)), \theta_s) \mu^{s-t}(d\theta_{t+1}^s), \end{aligned} \quad (6)$$

where, for $s \geq t+1$, $\theta_{t+1}^s := (\theta_{t+1}, \theta_{t+2}, \dots, \theta_s) \in \Theta^{s-t}$ is a history of states from $t+1$ to s , μ^{s-t}

is the product measure that is consistent with density p , and $\sigma^{s-t}(\theta_t^s)$ is the report history from period t to period s when the central bank uses the reporting strategy σ from period t onward.⁹ The set of these inequalities is referred to as the *incentive-compatibility constraint*. In words, a mechanism is incentive-compatible if and only if, after any report history, the central bank finds it optimal to follow the truth-telling strategy.

Society's objective is to maximize social welfare. The (time-0) social welfare from a mechanism $\{(\pi_t, x_t, \pi_t^e)\}_{t=0}^\infty$ is the expected discounted sum of future returns:

$$\sum_{t=0}^{\infty} \int_{\Theta^{t+1}} \beta^t R(\pi_t(\theta^t), x_t(\theta^t), \theta_t) \mu^t(d\theta^t), \quad (7)$$

where, for each t , μ^t is the product measure that is consistent with the density p .

It is worth noting that, because $\pi_{t+1}(\theta^t, \theta_{t+1})$ is weighted by $p(\theta_{t+1})$ in equation (5), rational expectation is required to hold only when the central bank tells the truth in period $t+1$. The fact that its deviation from truth-telling in period $t+1$ may, ex-post, violate the period- t rational expectation condition captures our assumption that the central bank is unable to commit.

The problem we consider is to choose a mechanism $\{(\pi_t, x_t, \pi_t^e)\}_{t=0}^\infty$ to maximize social welfare (equation (7)) subject to the feasibility constraint (equations (4) and (5)), and the incentive-compatibility constraint (equation (6)).¹⁰

4.1 Two benchmarks

We compare the solution to the problem where θ is private to two benchmark alternatives. The first alternative is the full-information solution the second alternative is the “optimal discretionary policy”. For more detail see Clarida, Gali and Gertler (1999) or Woodford (2003).

4.1.1 Full-information solution

The full-information problem is to choose a mechanism $\{(\pi_t, x_t, \pi_t^e)\}_{t=0}^\infty$ to maximize social welfare (equation (7)) subject to the feasibility constraint (equations (4) and (5)). The solution

⁹This history is recursively defined: $\sigma^0(\theta_t^t) := \sigma_0(\theta_t)$, and $\sigma^{s-t}(\theta_t^s) = (\sigma^{s-1-t}(\theta_t^{s-1}), \sigma_{s-t}(\sigma^{s-1-t}(\theta_t^{s-1}), \theta_s))$ for any $s \geq t+1$.

¹⁰Because the central bank at time 0 before observing θ_0 has the same preference as society, we may interpret society as the time-0 central bank. Then the mechanism design problem here can be interpreted as the central bank's optimal commitment problem without self-control in Amador, Werning, and Angeletos (2006). In Amador, Werning, and Angeletos (2006) decision maker's preference itself is time-inconsistent, while in our problem time-inconsistency arises from the New Keynesian Phillips curve.

corresponds to the optimal policy under commitment in the New Keynesian policy literature.

4.1.2 Optimal discretionary policy

The optimal discretionary policy concerns a situation in which society does not impose any restrictions on the central bank's choice, i.e. $D = \Pi \times X$. Following the New Keynesian policy literature, we focus on a Markov perfect equilibrium. A *Markov perfect equilibrium under "no restriction"* consists of (i) the policy function, $(\pi^{MP}, x^{MP}) : \Theta \rightarrow \Pi \times X$, (ii) the inflation expectation, $\pi^{e,MP} \in \Pi$, and (iii) the value, $W^{MP} \in \mathbb{R}$, such that

1. For all $\theta \in \Theta$,

$$(\pi^{MP}(\theta), x^{MP}(\theta)) \in \arg \max_{\pi, x} R(\pi, x, \theta) + \beta W^{MP}$$

subject to $\pi = \kappa x + \beta \pi^{e,MP}$,

2. $\pi^{e,MP} = \int \pi^{MP}(\theta) p(\theta) d\theta$, and
3. $W^{MP} = (1 - \beta)^{-1} \int R(\pi^{MP}(\theta), x^{MP}(\theta), \theta) p(\theta) d\theta$.

The best Markov perfect equilibrium is referred to as the optimal discretionary policy.¹¹

4.2 Recursive formulation

The optimal mechanism design problem and the full-information problem have at least one forward-looking constraint. We rewrite these constraints to obtain recursive formulations.

4.2.1 Feasibility

First, we argue that the feasibility constraint (equations (4) and (5)) implies that last period's inflation expectation serves as a state variable. Observe that, by treating (x_t, π_t, π_t^e) as choice variables in period t , equation (4) amounts to a static constraint in period t . In period $t + 1$, the previously chosen π_t^e imposes a constraint, reflected in equation (5), on the current choice for inflation, π_{t+1} , i.e. π_t^e is a state variable in period $t + 1$. To put it differently, in every period, the mechanism promises an expected level of inflation in the next period, while delivering (on average) the inflation promised in the previous period. We therefore refer to π^e as *inflation*

¹¹When R takes the quadratic form in equation (2), a Markov perfect equilibrium is unique.

promise. The full-information problem also has this constraint, and thus it too has a recursive formulation in which last period's inflation promise serves as a state variable.¹²

4.2.2 Incentive compatibility

Second, the incentive compatibility constraint (6) can be written recursively, by adding the agent's continuation, or promised, utility as a choice variable (e.g.). Let

$$U = \left[\frac{\int_{\Theta} \{\min_{x,\pi} R(\pi, x, \theta)\} p(\theta) d\theta}{1 - \beta}, \frac{\int_{\Theta} \{\max_{x,\pi} R(\pi, x, \theta)\} p(\theta) d\theta}{1 - \beta} \right],$$

then the expected discounted value of future returns always lies in this compact interval. As is standard in the dynamic contracting literature (e.g. Green, 1987; Thomas and Worrall, 1990), it can be shown that a mechanism $\{(x_t, \pi_t, \pi_t^e)\}_{t=0}^{\infty}$ is incentive compatible if and only if

1. There exists a sequence of measurable functions $\{W_t\}_{t=-1}^{\infty}$ with $W_t : \Theta^t \rightarrow U$ for all $t \geq -1$, such that for all $t \geq 0$ and θ^t ,

$$W_{t-1}(\theta^{t-1}) = \int_{\Theta} [R(\pi_t(\theta^t), x_t(\theta^t), \theta_t) + \beta W_t(\theta^t)] p(\theta_t) d\theta_t. \quad (8)$$

2. For all t , θ^{t-1} , θ_t , and $\theta' \neq \theta_t$,

$$R(\pi_t(\theta^t), x_t(\theta^t), \theta_t) + \beta W_t(\theta^t) \geq R(\pi_t(\theta^{t-1}, \theta'), x_t(\theta^{t-1}, \theta'), \theta_t) + \beta W_t(\theta^{t-1}, \theta'). \quad (9)$$

4.2.3 Interim problem

The mechanism design problem is then equivalent to the problem of choosing W_{-1} and the sequence of measurable functions $\{(x_t, \pi_t, \pi_t^e, W_t)\}_{t=0}^{\infty}$ to maximize social welfare (equation (7)) subject to constraints (4), (5), (8), and (9). However, because period-0 inflation choice is not subject to a constraint like (5), there is asymmetry between period 0 and all other periods, and the problem is not fully recursive.

¹²In the literature of optimal monetary policy in New Keynesian models with symmetric information, it is a common practice to set up a linear-quadratic regulator problem and solve a sequence problem by Lagrangian method. We relate this approach to ours in Appendix B, using the quadratic social welfare function in (2).

We therefore consider the *interim problem* with the following auxiliary initial condition:

$$\pi_{-1}^e = \int_{\Theta} \pi_0(\theta) p(\theta), \quad (10)$$

where π_{-1}^e is a given number in Π and represents the inflation promise made in period -1 . The interim problem has, as shown in the next section, a recursive formulation and therefore enables us to obtain a clear characterization of its solution. We refer to a solution to the interim problem as an optimal interim mechanism, or, when not confusing, simply as an optimal mechanism, because a solution to the original problem is obtained from an optimal interim mechanism by choosing the best initial condition π_{-1}^e . For any $\pi_{-1}^e \in \Pi$, we say that $\{(x_t, \pi_t, \pi_t^e, W_t)\}_{t=0}^{\infty}$ is *feasible from π_{-1}^e* if and only if it satisfies equations (4), (5), and (10), and that it is *incentive-feasible from π_{-1}^e* if and only if it is feasible from π_{-1}^e and satisfies equations (8) and (9).¹³

5 Theoretical results

In this section, we first establish that, under certain conditions, a solution to the interim problem can be obtained by solving a function-valued dynamic programming problem with promised inflation as the state variable. Then we characterize its properties. Depending on last period's promised inflation, the optimal degree of discretion is shown to take one of three forms: full-discretion, no-discretion, or bounded-discretion. It is also shown that the optimal mechanism features amnesia — history is forgotten for types that have discretion. Finally, we propose an inflation targeting rule that achieves the same outcome as the optimal mechanism.

5.1 Dynamic Programming

To facilitate characterization, we follow AAK and restrict our attention to allocations that satisfy the following:

Assumption 1 *For all t and θ^{t-1} , $\pi_t(\theta^{t-1}, \cdot) : \Theta \rightarrow \Pi$ is a piecewise C^1 function.*

¹³For any $\pi_{-1}^e \in \Pi$, it is straightforward to prove that the constraint set is non-empty. For a given $\pi_{-1}^e \in \Pi$, consider a mechanism such that $\pi_t(\theta^t) = \pi_{-1}^e$ and $x_t(\theta^t) = (1 - \beta)\pi_{-1}^e/\kappa$ for all t and θ^t . These functions are clearly measurable and satisfy the auxiliary initial condition. As we assume that X is an interval that contains both $(\underline{\pi} - \beta\bar{\pi})/\kappa$ and $(\bar{\pi} - \beta\underline{\pi})/\kappa$, this allocation satisfies the NKPC after any history. Since this allocation is independent of history, it is incentive-compatible.

Let Ω be the set of (π_{-1}^e, W_{-1}) 's such that there exists a sequence of measurable functions $\{(x_t, \pi_t, \pi_t^e, W_t)\}_{t=0}^\infty$ that satisfies Assumption 1 and equations (4), (5), (8), (9), and (10). Because the interim problem has a non-empty constraint set for all $\pi_{-1}^e \in \Pi$, the projection of Ω into Π is simply Π . For all $\pi_-^e \in \Pi$, the maximized social welfare given π_-^e is given by

$$\bar{W}(\pi_-^e) = \sup_{W_- \text{ s.t. } (\pi_-^e, W_-) \in \Omega} W_-, \quad (11)$$

which implies that if we obtain Ω , we also obtain the maximized social welfare.

This set Ω is, however, difficult to characterize in our setting. In settings with discrete types or discrete action spaces, Ω can be characterized as the largest fixed point of some set operator *à la* Abreu, Pearce, and Stacchetti (1990). In our set-up, there are a continuum of types and continuous action spaces, and the measurability restriction is difficult to impose in the APS type set operator. We instead use an approach akin to that in Athey, Bagwell, and Sanchirico (2004) to characterize directly the function \bar{W} , rather than the whole set Ω .

To characterize \bar{W} , we first consider the *factored problem*: for each $\pi_-^e \in \Pi$,

$$\bar{V}(\pi_-^e) = \sup_{\pi, x, \pi^e, W} \int_{\Theta} \{R(\pi(\theta), x(\theta), \theta) + \beta W(\theta)\} p(\theta) d\theta, \quad (12)$$

subject to

$$\pi_-^e = \int_{\underline{\theta}}^{\bar{\theta}} \pi(\theta) p(\theta) d\theta, \quad (13)$$

$$\pi(\theta) = \kappa x(\theta) + \beta \pi^e(\theta), \forall \theta, \quad (14)$$

$$R(\pi(\theta), x(\theta), \theta) + \beta W(\theta) \geq R(\pi(\theta'), x(\theta'), \theta) + \beta W(\theta'), \forall \theta, \theta' \neq \theta, \quad (15)$$

$$\pi \text{ is a piecewise } C^1 \text{ function}, \quad (16)$$

$$(\pi^e(\theta), W(\theta)) \in \Omega, \forall \theta, \quad (17)$$

and (π, x, π^e, W) are measurable. It follows that $\bar{V} \geq \bar{W}$. They may not be identical because the factored problem relaxes the measurability restriction.

To show that $\bar{V} = \bar{W}$, we formulate a relaxed problem by replacing equation (17) with the weaker constraint

$$\forall \theta, \quad \pi^e(\theta) \in \Pi \quad \text{and} \quad W(\theta) \leq \bar{W}(\pi^e(\theta)).$$

In words, society can set the continuation utility to any level that is lower than \bar{W} . It is convenient to define a Bellman operator associated with this relaxed problem. Let $B(\Pi)$ be the space of bounded functions on Π . Then $(B(\Pi), \|\cdot\|)$ where $\|\cdot\|$ is the sup norm is a Banach space. Define a Bellman operator $\mathbb{T} : B(\Pi) \rightarrow B(\Pi)$ as follows: for all $F \in B(\Pi)$, for all $\pi_-^e \in \Pi$,

$$\mathbb{T}F(\pi_-^e) = \sup_{\pi(\cdot), x(\cdot), \pi^e(\cdot), W(\cdot)} \int_{\underline{\theta}}^{\bar{\theta}} \left\{ R(\pi(\theta), x(\theta), \theta) + \beta W(\theta) \right\} p(\theta) d\theta, \quad (18)$$

subject to constraints (13), (14), (15), (16), and

$$\forall \theta, \quad \pi^e(\theta) \in \Pi \quad \text{and} \quad W(\theta) \leq F(\pi^e(\theta)). \quad (19)$$

For any $F \in B(\Pi)$ and $\pi_-^e \in \Pi$, we refer to the maximization problem in equation (18) as the $\mathbb{T}F(\pi_-^e)$ -problem. The value of this relaxed problem is given by $\mathbb{T}\bar{W}$, and, because we are considering a relaxed problem, it follows that $\mathbb{T}\bar{W} \geq \bar{V}$. This implies $\mathbb{T}\bar{W} \geq \bar{W}$.

Proposition 1 \mathbb{T} is a β -contraction mapping.

Proof. Blackwell's sufficient condition is satisfied. (See e.g. Stokey et al., 1989.) ■

Let \tilde{W} be the fixed point of \mathbb{T} . Because \mathbb{T} is monotone, it follows that $\tilde{W} \geq \mathbb{T}\bar{W} \geq \bar{V} \geq \bar{W}$. Therefore, $\tilde{W} = \bar{W} = \bar{V}$ is implied if it can be established that $\bar{W} \geq \tilde{W}$. To show this inequality, we show that, under certain assumptions, there is a quadruple of measurable functions $(\pi_*, x_*, \pi_*^e, W_*)$ of (θ, π_-^e) such that, at each π_-^e , $(\pi_*(\cdot, \pi_-^e), x_*(\cdot, \pi_-^e), \pi_*^e(\cdot, \pi_-^e), W_*(\cdot, \pi_-^e))$ attains the maximum of the $\mathbb{T}\tilde{W}(\pi_-^e)$ -problem, and that $W_*(\cdot, \pi_-^e) = \tilde{W}(\pi_*^e(\cdot, \pi_-^e))$. By iterating forward, such a quadruple generates for each $\pi_-^e \in \bar{\Pi}$ a mechanism that is incentive-feasible from π_-^e .¹⁴ This implies $\bar{W} \geq \tilde{W}$.

Several assumptions are now in order. First we make the following assumptions on the return function, R , and the density function, p .

Assumption 2

$$R(\pi, x, \theta) = A(\pi) + B(x) + \pi \times \theta$$

where A and B are strictly concave C^2 functions. The first derivative of A , $A'(\pi)$, goes to ∞ ($-\infty$) as $\pi \rightarrow -\infty$ (∞ , respectively). Also $A''(\pi) \leq \bar{A}'' < 0$ and $B''(x) \leq \bar{B}'' < 0$ for some

¹⁴See Chapter 9 in Stokey et al. (1989).

constant \bar{A}'' and \bar{B}'' .¹⁵

Assumption 3 (i) $(1 - P(\theta))/p(\theta)$ is strictly decreasing and $P(\theta)/p(\theta)$ is strictly increasing in θ . (ii) The density function p is continuous and strictly positive for all $\theta \in [\underline{\theta}, \bar{\theta}]$.

Under Assumption 2, social welfare is separable in x and (π, θ) , and the private information governs the marginal social cost of inflation: the higher is θ , the lower is the marginal social cost of inflation. The quadratic specification in equation (2) satisfies Assumption 2, and is used later in our numerical experiments. Assumption 3 is the monotone hazard condition.

Now we show that the Bellman operator \mathbb{T} preserves certain properties of a function, and that if F has these properties the equation (19) is always satisfied with equality. Let $\mathcal{V}(\bar{\Pi})$ be the set of strictly concave C^1 functions on a compact subinterval $\bar{\Pi} \subset \Pi$ whose first derivative is C^1 except on a finite set of points and both the right- and the left-derivatives of the first derivative exist everywhere, including $-\infty$. We make two additional assumptions:

Assumption 4 There is a compact interval, $\bar{\Pi} \subset \Pi$, such that, for any $F \in \mathcal{V}(\bar{\Pi})$ and any $\pi_-^e \in \bar{\Pi}$, a solution to the $\mathbb{T}F(\pi_-^e)$ -problem satisfies $(\pi(\theta), x(\theta), \pi^e(\theta)) \in \text{int}(\Pi) \times \text{int}(X) \times \text{int}(\bar{\Pi})$ for all θ .

Assumption 5 For all $\pi_-^e \in \bar{\Pi}$, a maximum is attained in the $\mathbb{T}F(\pi_-^e)$ -problem when $F \in \mathcal{V}(\bar{\Pi})$.

The following proposition shows that \mathbb{T} maps $\mathcal{V}(\bar{\Pi})$ into itself, and that the equation (19) is always satisfied with equality in the $\mathbb{T}F(\pi_-^e)$ -problem for all $\pi_-^e \in \bar{\Pi}$ when $F \in \mathcal{V}(\bar{\Pi})$.

Proposition 2 Suppose $F \in \mathcal{V}(\bar{\Pi})$. Then under Assumptions 2, 3, 4, and 5, (i) $\mathbb{T}F \in \mathcal{V}$, (ii) there is a quadruple of continuous functions $(\pi, x, \pi^e, W) : \Theta \times \bar{\Pi} \rightarrow \Pi \times X \times \bar{\Pi} \times U$ such that, for each $\pi_-^e \in \bar{\Pi}$, $(\pi(\cdot, \pi_-^e), x(\cdot, \pi_-^e), \pi^e(\cdot, \pi_-^e), W(\cdot, \pi_-^e))$ attains the maximum for the $\mathbb{T}F(\pi_-^e)$ -problem, that (iii) $\pi(\cdot, \pi_-^e)$ is piecewise C^1 for each $\pi_-^e \in \bar{\Pi}$, and that (iv) $W(\theta, \pi_-^e) = F(\pi^e(\theta, \pi_-^e))$ for all (θ, π_-^e) .

¹⁵We could include an additively separable term $d(\theta)$ that is non-linear in θ , but we set $d(\theta) = 0$ without loss of generality. The assumption that the term $\pi\theta$ has a positive unit coefficient is not restrictive, as one can scale up or down θ when the coefficient is not one, and can re-define $-\theta$ as the type when the coefficient is negative. This form also allows us to normalize $E[\theta] = 0$ without loss of generality.

The proof is in Appendix A. Proposition 2 implies that the fixed point \tilde{W} is a continuous, weakly concave function, but it is not guaranteed to be in $\mathcal{V}(\bar{\Pi})$. To characterize the solution further, we assume:

Assumption 6 $\tilde{W} \in \mathcal{V}(\bar{\Pi})$.

Under Assumption 6, we can apply Proposition 2 to $F = \tilde{W}$, and it follows that $\bar{V} = \bar{W} = \tilde{W}$.¹⁶

5.2 Optimal degree of discretion

The previous results allow us to characterize the optimal degree of discretion using the policy functions $(\pi_*, x_*, \pi_*^e, W_*)$, which solve the $\mathbb{T}\tilde{W}$ -problem. To quantify the degree of discretion, we consider the inflation that the central bank would choose if it were given a certain form of policy-flexibility, and compare it to the inflation prescribed by the optimal mechanism. This approach is analogous to AAK: they define the “static best response” of the central bank — the inflation choice that maximizes the momentary social welfare for a given inflation expectation — and compare it to the optimal mechanism. In our setting, expected inflation may vary with the central bank’s message, implying that the optimal mechanism is dynamic, and that the static best response doesn’t provide a useful benchmark for comparison. Instead we introduce the notion of the “one-shot discretionary best response.”

Imagine that the central bank is allowed to *choose any* (π, x, π^e) for one period, subject only to the NKPC and, in particular, not subject to the constraint $\pi_-^e = E\pi$, but faces the optimal mechanism in all subsequent periods. The one-shot discretionary best response is the optimal inflation that would be chosen by the central bank in this hypothetical situation.

Definition 1 *The one-shot discretionary best response is a function $\pi_D : \Theta \rightarrow \Pi$ that, for each θ , solves*

$$\max_{\pi} \left\{ A(\pi) + \theta\pi + \max_{(x, \pi^e): \pi = \kappa x + \beta\pi^e} \{ B(x) + \beta\tilde{W}(\pi^e) \} \right\}. \quad (20)$$

The one-shot discretionary best response is well-behaved:

Lemma 1 *$\pi_D(\cdot)$ is a strictly increasing, continuous, piecewise C^1 function.*

¹⁶We can relax Assumptions 4 and 6 and replace them with the following assumption: there exist a compact sub-interval $\bar{\Pi} \subset \Pi$ and $F_0 \in \mathcal{V}(\bar{\Pi})$ such that $F_0 \leq \mathbb{T}F_0$, and that for all $n \geq 1$ and $\pi_-^e \in \bar{\Pi}$, the bound constraints are not binding for the $(\mathbb{T})^n F_0$ -problem. (In words, the bound constraints never bind during value function iteration starting from F_0 .) Then we can show that a solution to the $\mathbb{T}\tilde{W}$ -problem satisfies the properties (ii) and (iv) in Proposition 2 and the properties shown in the next section, but (iii) is not guaranteed.

It is natural to interpret that, for a given π_-^e , a type θ has discretion when $\pi_*(\theta; \pi_-^e) = \pi_D(\theta)$. The degree of discretion at π_-^e is naturally defined as the probability of the event $\{\pi_*(\theta; \pi_-^e) = \pi_D(\theta)\}$. We say that the central bank has *full-discretion* at π_-^e if the degree of discretion is one. The next proposition shows that the central bank has full-discretion at only one value for π_-^e , π_-^{e*} , and that π_-^{e*} is the expected value of the one-shot discretionary best response.

Proposition 3 *\tilde{W} is uniquely maximized at $\pi_-^e = \pi_-^{e*} = E[\pi_D]$, and the policy function satisfies $\pi_*(\cdot; \pi_-^{e*}) = \pi_D(\cdot)$.*

From Proposition 3 we can think of π_-^{e*} as the most desirable initial condition: if $\pi_{-1}^e = \pi_-^{e*}$, then social welfare from time zero onward is maximized. Furthermore, if the central bank were allowed to renege on previously promised inflation expectation, then it would behave as if expected inflation were π_-^{e*} . This leads us to the following definition of time-inconsistency:

Definition 2 *The policy function is time-consistent at π_-^e if and only if $\tilde{W}(\pi_-^e) \geq \tilde{W}(\tilde{\pi}_-^e)$ for any $\tilde{\pi}_-^e$. The severity of the time-inconsistency problem at π_-^e is measured by $\max_{\tilde{\pi}_-^e} \tilde{W}(\tilde{\pi}_-^e) - \tilde{W}(\pi_-^e)$.*

Note that this definition of time-inconsistency is non-standard, as we define it point-wise. An immediate implication of Proposition 3 is that the policy function is time-consistent only at $\pi_-^e = \pi_-^{e*}$. Since \tilde{W} is strictly concave and has a peak at π_-^{e*} , the severity of the time-inconsistency problem increases as π_-^e moves away from π_-^{e*} .

To characterize less than full discretion, it is convenient to define two types of discretion: *no discretion* and *bounded discretion*.

Definition 3 *The optimal policy has no-discretion at π_-^e if $\pi_*(\cdot; \pi_-^e)$ is constant. It has bounded discretion at π_-^e if $\pi_*(\cdot; \pi_-^e)$ is not constant and either $\pi_*(\cdot; \pi_-^e) = \max\{\pi', \pi_D(\cdot)\}$ or $\pi_*(\cdot; \pi_-^e) = \min\{\pi', \pi_D(\cdot)\}$ for some constant π' .*

Note that bounded discretion is equivalent to a cut-off property: there is a threshold value for θ such that the optimal inflation is constant either above or below that threshold. It turns out that, for any π_-^e , the optimal mechanism takes a rather simple form — full discretion, bounded discretion, or no discretion — and that the optimal degree of discretion is linked closely to the severity of the time-inconsistency problem. This result is summarized by the next proposition.




Figure 1: Policy function $\pi_*(\theta, \pi_-^e)$ as a function of θ , for different values of π_-^e

Proposition 4 *There exist two strictly increasing, continuous, piecewise C^1 threshold functions $T_1 : (\pi_D(\underline{\theta}), \pi_-^{e*}) \rightarrow \Theta$ and $T_2 : (\pi_-^{e*}, \pi_D(\bar{\theta})) \rightarrow \Theta$ such that the policy function for inflation, π_* , features*

1. Full discretion, $\pi_*(\theta; \pi_-^e) = \pi_D(\theta)$ for all θ , if $\pi_-^e = \pi_-^{e*}$; or,
2. No discretion, $\pi_*(\theta; \pi_-^e) = \pi_-^e$ for all θ , if $\pi_-^e \leq \pi_D(\underline{\theta})$ or $\pi_-^e \geq \pi_D(\bar{\theta})$; or,
3. Bounded discretion, if $\pi_-^e \in (\pi_D(\underline{\theta}), \pi_-^{e*})$,

$$\pi_*(\theta; \pi_-^e) = \begin{cases} \pi_D(\theta), & \forall \theta \in [\underline{\theta}, T_1(\pi_-^e)] \\ \pi_D(T_1(\pi_-^e)), & \forall \theta \in [T_1(\pi_-^e), \bar{\theta}] \end{cases},$$

or if $\pi_-^e \in (\pi_-^{e*}, \pi_D(\bar{\theta}))$, then

$$\pi_*(\theta; \pi_-^e) = \begin{cases} \pi_D(T_2(\pi_-^e)), & \forall \theta \in [\underline{\theta}, T_2(\pi_-^e)] \\ \pi_D(\theta), & \forall \theta \in (T_2(\pi_-^e), \bar{\theta}] \end{cases}.$$

Proposition 4 reveals three important properties. First, when π_-^e is sufficiently far from π_-^{e*} (either above or below), then there is no discretion. This is depicted in panels (i) and (ii) in Figure 1. Red dashed lines represent π_D and blue solid lines represent $\pi_*(\cdot; \pi_-^e)$ for a given π_-^e . Second, when π_-^e is not too far from π_-^{e*} , the policy function exhibits bounded discretion, or a cut-off property. When π_-^e is less than π_-^{e*} , but not too low, inflation rates for low- θ types are the same as their one-shot discretionary best response (in this sense they are unconstrained) while high- θ types are constrained to a single level of inflation (panel (iii) in Figure 1). Similarly, when π_-^e is higher than π_-^{e*} , but not too high, inflation rates for high- θ types are the same as their one-shot discretionary best response while low- θ types are constrained to a single level of inflation (panel (iv) in Figure 1). As π_-^e moves away from π_-^{e*} , the time-inconsistency problem becomes more severe, the degree of discretion becomes smaller, and eventually no discretion is permitted.

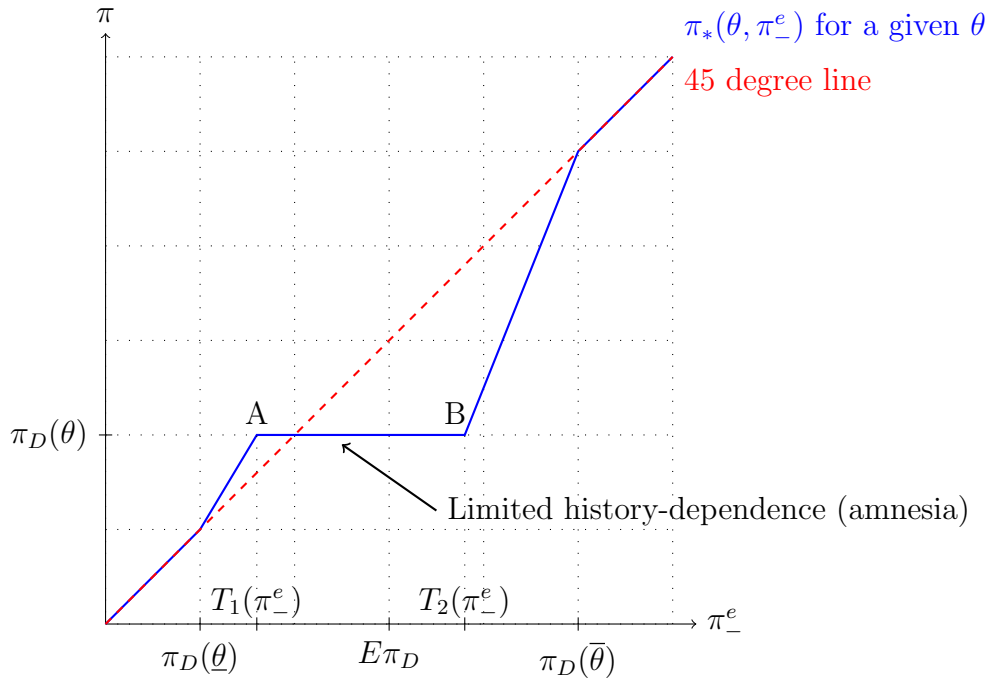


Figure 2: Policy function $\pi_*(\theta; \pi_-^e)$ as a function of π_-^e , given θ

Third, when we view the policy function for a given θ as a function of π_-^e , it is strictly

increasing up to $\pi_-^e = T_1^{-1}(\theta)$, is then flat up to $\pi_-^e = T_2^{-1}(\theta)$, and is strictly increasing after that. This property is depicted in Figure 2. Note that between points A ($\pi_-^e = T_1^{-1}(\theta)$) and B ($\pi_-^e = T_2^{-1}(\theta)$), the policy function is flat and its value equals $\pi_D(\theta)$. Importantly, the fact that the policy function is flat on an interval implies that the history-dependence, as encoded in the state variable, is disposed of on this interval. If the state variable in period t reside within such an interval for given θ_t , then the continuation mechanism from period $t + 1$ onward does not depend on θ^{t-1} . For types that have discretion, the optimal mechanism therefore features amnesia in the sense of Kocherlakota (1996).

Finally, Proposition 5 characterizes the policy functions for the output gap and inflation promise, x_* and π_*^e , respectively.

Proposition 5 *Let $(x_S, \pi_S^e) : \Pi \rightarrow X \times \bar{\Pi}$ be a pair of functions such that, for any $\pi \in \Pi$, $(x_S(\pi), \pi_S^e(\pi))$ maximizes $B(s) + \beta \tilde{W}(\pi^e)$ subject to $\pi = \kappa x + \beta \pi^e$. Then x_S and π_S^e are strictly increasing and continuous, and the policy functions for the output gap and promised inflation satisfy $x_*(\theta; \pi_-^e) = x_S(\pi_*(\theta; \pi_-^e))$ and $\pi_*^e(\theta; \pi_-^e) = \pi_S^e(\pi_*(\theta; \pi_-^e))$ for all (θ, π_-^e) .*

Recall that some discretion is given when $\pi_-^e \in (\pi_D(\underline{\theta}), \pi_D(\bar{\theta}))$. Because Proposition 5 implies that π_S^e is strictly increasing, it follows that $\pi_*^e(\theta; \pi_-^e)$ lies between $\pi_S^e(\pi_D(\underline{\theta}))$ and $\pi_S^e(\pi_D(\bar{\theta}))$. Imagine a situation in which this interval is contained in $(\pi_D(\underline{\theta}), \pi_D(\bar{\theta}))$. This implies that, once some discretion is given, the continuation mechanism always prescribes some discretion afterwards, and the probability of visiting the no-discretion region from the some-discretion region is zero. Moreover, because the fully optimal mechanism starts from the optimal initial condition π_-^{e*} , which is contained in $(\pi_D(\underline{\theta}), \pi_D(\bar{\theta}))$, it never visits the no-discretion region, and therefore some discretion is always given. This is indeed the case in our numerical experiment.

Another implication of Proposition 5 is that, if society can ensure that the central bank's inflation choice is the same as what the optimal mechanism prescribes and that the value of choosing π^e is $\tilde{W}(\pi^e)$, then the central bank finds it optimal (subject to NKPC) to choose the output gap and inflation promise that are prescribed by the optimal mechanism. In what follows, we show that society can ensure these two things by a particular inflation targeting rule, and, therefore, that society does not need to restrict the central bank's choice of the output gap and inflation promise.

5.3 Implementation by inflation targeting

Here we show that the optimal allocation can be implemented by a simple *inflation targeting rule*. Consider the game described in Section 3 where the central bank at the communication stage announces next period's inflation promise, i.e. $M = \bar{\Pi}$. Society chooses a delegation set in a Markovian way, based on last period's inflation promise made by the central bank. In particular, society chooses delegation sets of the form of $\Gamma(m_-) \times X$ where $\Gamma : \bar{\Pi} \rightrightarrows \Pi$ is a correspondence and m_- denotes last period's message, or inflation promise, sent by the central bank. Delegation sets of this form do not impose any constraint on the central bank's output gap choice, and thus can be interpreted as inflation targeting.

The equilibrium concept we use is a Markov perfect equilibrium, the same as that used for the optimal discretionary policy. A *Markov perfect equilibrium under an inflation targeting rule* Γ consists of (i) the central bank's policy function, $(\pi_{IT}, x_{IT}, m_{IT}) : \Theta \times \bar{\Pi} \rightarrow \Pi \times X \times \bar{\Pi}$, (ii) the private sector's policy function, $\pi_{IT}^e : \bar{\Pi} \times \bar{\Pi}$, and (iii) the central bank's value function $W^{IT} : \bar{\Pi} \rightarrow \mathbb{R}$ such that

1. For all $m_- \in \bar{\Pi}$ and any $m \in \bar{\Pi}$,

$$\pi_{IT}^e(m; m_-) = \int_{\underline{\theta}}^{\bar{\theta}} \pi_{IT}(\theta; m) p(\theta) d\theta,$$

2. For all $(\theta, m_-) \in \Theta \times \bar{\Pi}$,

$$(\pi_{IT}(\theta, m_-), x_{IT}(\theta, m_-), m_{IT}(\theta, m_-)) \in \arg \max_{\pi, x, m} R(\pi, x, \theta) + \beta W^{IT}(m)$$

subject to $\pi = \kappa x + \beta \pi_{IT}^e(m, m_-)$ and $\pi \in \Gamma(m_-)$, and

3. $W^{IT} : \bar{\Pi} \rightarrow \mathbb{R}$ satisfies a recursion: for all $m_- \in \bar{\Pi}$,

$$W^{IT}(m_-) = \int_{\underline{\theta}}^{\bar{\theta}} \{R(\pi_{IT}(\theta, m_-), x_{IT}(\theta, m_-), \theta) + \beta W^{IT}(m_{IT}(\theta, m_-))\} p(\theta) d\theta.$$

The second and third conditions state that, given the private sector's policy function, the central bank's policy and value functions satisfy optimality condition. Similarly, the first condition states that the private sector's policy function is consistent with rational expectation, given the

central bank's policy function.

We say that Γ implements the optimal policy if $(\pi_{IT}, x_{IT}, m_{IT}) = (\pi_*, x_*, \pi_*^e)$, $\pi^e = \pi_*^e$, and $W_{IT} = \tilde{W}$ constitute a Markov perfect equilibrium under Γ . Note that the condition $m_{IT} = \pi_*^e$ requires that the central bank find it optimal to tell the truth in that its announcement of promised inflation equals next period's expected inflation.

We propose the following inflation targeting rule, denoted by $\bar{\Gamma}$.

$$\bar{\Gamma}(m_-) = \begin{cases} \Pi \cap (-\infty, m_-] & \text{if } m_- \leq \pi_D(\underline{\theta}) \\ \Pi \cap (-\infty, \pi_D(T_1(m_-))] & \text{if } m_- \in (\pi_D(\underline{\theta}), \pi_-^{e*}) \\ \Pi \cap (-\infty, \infty) & \text{if } m_- = \pi_-^{e*} \\ \Pi \cap [\pi_D(T_2(m_-)), \infty) & \text{if } m_- \in (\pi_-^{e*}, \pi_D(\bar{\theta})) \\ \Pi \cap [m_-, \infty) & \text{if } m_- \geq \pi_D(\bar{\theta}) \end{cases}$$

Proposition 6 $\bar{\Gamma}$ implements the optimal policy.

Proof is in Appendix.

There are many inflation range targeting rules other than $\bar{\Gamma}$ that also implement the optimal policy. $\bar{\Gamma}$ is the largest among them. The smallest correspondence is

$$\underline{\Gamma}(m_-) = \{\pi \in \Pi \mid \pi = \pi_*(\theta; m_-) \text{ for some } \theta\}.$$

It is straightforward to see that a necessary and sufficient condition for Γ to implement the optimal policy is $\underline{\Gamma}(m_-) \subset \Gamma(m_-) \subset \bar{\Gamma}(m_-)$ for all $m_- \in \bar{\Pi}$.

This condition highlights the necessity of imposing an upper limit on inflation for $m_- < \pi_-^{e*}$ and a lower limit for $m_- > \pi_-^{e*}$. This is in contrast to AAK's result that a constant inflation cap, or an upper-bound, can implement the optimal mechanism. This difference arises because, in AAK, the source of the time-inconsistency problem is the inflation bias embodied in the social welfare function, and because the severity of time-inconsistency problem is constant. In our setting, an inflation cap suffices only when $m_- < \pi_-^{e*}$. An inflation cap does not suffice generally because the source of the time-inconsistency problem is stabilization bias and the direction of the bias can go either way. When last period's inflation promise m_- is higher than π_-^{e*} , the central bank wants to renege on its promise and restart the economy with a lower initial

condition π_-^{e*} . Therefore the central bank in our setting has a deflation bias when $m_- > \pi^{e*}$, and a lower limit must be imposed to deliver the promised level of inflation.

This necessary and sufficient condition also implies that history-dependence of a delegation set is necessary to implement the optimal policy. When next period's delegation set depends on the current period inflation promise, the central bank can use promised inflation to restrict the action taken by its future-self, thereby mitigating the time-inconsistency problem. Because the severity of time-inconsistency problem depends on promised inflation, the optimal upper- and lower-limits for inflation must also vary with this promise. Inflation targeting with a fixed range is therefore unable to implement the optimal policy.

Note that, for a given m_- , the minimal delegation set $\underline{\Gamma}(m_-)$ takes the form of an interval. This result is closely related to the optimality of interval delegation obtained in many static delegation problems (Holmström, 1977 and 1984, Alonso and Matouschek, 2008, and Amador and Bagwell, 2013). One distinct feature of our result is that, to implement the optimal policy, we need a communication stage on top of a delegation set. This is because there is an information asymmetry between the central bank and the private sector. Society and the central bank, from the ex-ante point of view, both benefit from the introduction of a communication stage, and thus agree on the use of communication, because their objectives coincide ex-ante (Melumand and Shibano, 1991).

6 Numerical results

We have seen that the optimal degree of discretion is endogenous and depends upon π_-^e . But it is not yet clear from our theoretical results how the degree of discretion changes over time. In this section we use a numerical experiment to examine this issue. The numerical procedure is described in Appendix C.

6.1 Parameter values

Our parameterization of the model is largely standard. We assume that a period corresponds to a quarter in length and set the discount factor, β , to 0.99. With the Calvo-pricing model as our guide, we set the slope of the Phillips curve, κ , to 0.12875. This value for κ is supported by a Calvo-pricing parameter of 0.75, by an elasticity of substitution between goods, ϵ , of 5.00,

implying a 25 percent mark-up, and by a momentary utility function for the representative household of the form $\ln c_t - h_t^{1+\nu}/(1+\nu)$, where c_t denotes consumption, where h_t denotes hours worked, and where ν , the (inverse) Frisch labor-supply elasticity, is set to 0.50.

We use the quadratic specification in equation (2) for the social welfare function, R . We set $b = \kappa/\epsilon$, on the basis that the second-order Taylor expansion of the representative agent's utility takes that form in the canonical new Keynesian model without private information (see e.g. Gali (2008)).

Turning to the type, θ , we assume that θ has a uniform probability density function on the interval $[-0.5\%, 0.5\%]$. These numbers are not annualized, and approximately correspond to $[-2\%, 2\%]$ per annum. For computational purposes, this continuous density is approximated using a uniform-grid containing 31 points.

6.2 Benchmark results

First we present the full-information solution and the optimal discretionary policy.

6.2.1 Full-information benchmark

We begin with the full information solution to understand how the state variable, π_-^e , and the shock, θ , affect the central bank's actions in the absence of private information. As we will see, some characteristics of the solution also hold for the case where information about θ is private. Usefully, the full-information problem is an example of an optimal linear-quadratic regulator problem, implying that the full-information solution is linear in π_-^e .

Figure 3 displays the policy functions for inflation promise, inflation, and the output gap, against the state variable, π_-^e . To make these plots visible, we report these policy functions for just 5 values of θ , including the lowest and the highest values. All numbers are expressed in terms of percentages, so $\pi = 1$ corresponds to an inflation rate of 1 percent per quarter.

Monotonicity in π_-^e : For each θ , the policy functions are increasing in last period's inflation promise. It is unsurprising that inflation is increasing in π_-^e because $\pi_-^e = E[\pi]$. It is also unsurprising that expected inflation, π^e is increasing in last period's inflation promise, because the policy trade-off between stabilizing inflation and the output gap requires inflation expectations to adjust gradually over time. The fact that the output gap is also increasing in π_-^e




Figure 3: Full-information Solution

may at first seem surprising, but it simply reflects the fact that commitment to its past policies requires the central bank to validate an increase in π_-^e with an increase in the output gap.

Monotonicity in θ : For each π_-^e , the policy functions are increasing in θ . To understand this result, note from the social welfare function that an increase in θ corresponds to a decline in the welfare cost of inflation, making it optimal for the central bank to raise inflation. In turn, the increase in current inflation leads to an increase in the output gap and in promised inflation.

Inflation promises in the long-run: The policy function for π^e is flatter than the 45 degree line, implying that expected inflation settles in a compact interval around zero in the long-run.

6.3 Optimal discretionary policy

Figure 4 depicts the optimal discretionary policy. The policy functions are depicted as a set of flat lines because they depend only on θ . Expected inflation is independent of θ too, and therefore only one flat line is shown in the left panel. Because $\pi^e = 0$ always, the NKPC implies that inflation and the output gap are proportional: $\pi = \kappa x$.




Figure 4: Optimal discretionary policy

6.4 Private information results

We now turn to our main results for the case where θ is private information. Considering inflation first, the left panel in Figure 5 displays the policy function for inflation as a function of the state for five values of θ . The middle and right panels then plot the differences between the private-information solution and the full-information solution and between the private-information solution and the optimal discretionary policy, respectively.

Looking at the policy functions shown in the left panel it is clear that the private information solution is nonlinear in the state variable, π_-^e , and that, despite the fact that we use a discrete-type model for computation, it exhibits properties that are largely consistent with the theoretical prediction of the continuous type model. When promised inflation is too low, the optimal mechanism prescribes no-discretion. As we move toward the right, we observe that the degree of discretion increases until π_-^e reaches zero, which is π_-^{e*} in this simulation. Moreover, the policy function for the lowest θ type becomes (almost) flat first, and then those for higher θ 's become (almost) flat. Once π_-^e exceeds zero, the degree of discretion decreases. Inflation for low- θ types starts increasing, while for high θ types' inflation stays constant.¹⁷

¹⁷Figure 5 reveals some behaviors that differ from the theoretical predictions of the continuous-type model.




Figure 5: Policy Function Comparison: π

Recall that the full-information solution is strictly increasing in π_-^e , for each θ . The fact that the private information solution features a flat interval implies that even when there is some degree of discretion the solution is distinct from the full-information solution. Despite these differences, it is nevertheless true that as the degree of discretion increases the private-information solution becomes closer to the full-information solution (middle panel). The distance between the private-information solution and the optimal discretionary policy is also small when the former is flat, but is non-negligible (right panel).

Figures 6 and 7 show the behavior of the output gap and expected inflation, respectively.¹⁸ While these figures reveal patterns that are qualitatively similar to those for inflation, it is interesting that the private-information solution and the optimal discretionary policy generates quantitatively very similar outcomes for the output gap when the private-information solution is flat, suggesting that the important differences between the two policies reside in their implications for inflation.

First, we observe full-discretion not only at one point ($\pi_-^e = \pi_-^{e*} = 0$), but on an interval around (though it is small). Second, on this interval, the policy function is not flat and slightly increasing. Third, the policy function for each θ decreases slightly before becoming virtually flat. These features are all likely to be the result of our use of a discrete-type model for the computations, and to the discreteness of our computational method. The interval of full-discretion indeed becomes small as we increase the number of types.

¹⁸Expected inflation is always zero under the optimal discretionary policy and is omitted from the figure.




Figure 6: Policy Function Comparison: x




Figure 7: Policy Function Comparison: π_e^e

Figure 8 plots the value function for the private-information policy, \tilde{W} , along with those for the two benchmark policies, W^{FI} and W^{MP} . Again, the optimal discretionary policy is independent of promised inflation and therefore its value is constant and depicted as a flat line. The optimal initial condition, π^{e*} , equals zero in this example, and therefore society finds it optimal for initial inflation to take this value. We evaluate the performance of the private-information solution using the measure,

$$\frac{\max_{\pi_-^e} \tilde{W}(\pi_-^e) - W^{MP}}{\max_{\pi_-^e} W^{FI}(\pi_-^e) - W^{MP}},$$

which, in this example, is 39.1%. Note that

$$\max_{\pi_-^e} W^{FI}(\pi_-^e) - W^{MP} = \left(\max_{\pi_-^e} W^{FI}(\pi_-^e) - \max_{\pi_-^e} \tilde{W}(\pi_-^e) \right) + \left(\max_{\pi_-^e} \tilde{W}(\pi_-^e) - W^{MP} \right).$$

The first term on the right-hand side represents the (absolute) loss from private information, and the second term represents the (absolute) loss associated with the sub-optimality of no-restriction. The latter is about 2/3 of the former in this example, which suggests that the optimal mechanism greatly improves upon no-restriction.¹⁹ Another way to measure the costs of private information is to calculate the inflation promise that, in the full-information setting, yields the same value as the maximized value in the private-information setting.²⁰ This is around 0.7% (or -0.7%), which implies that, taking the full-information setting as the status-quo, society is indifferent between the full-information solution with an inflation promise of 2.7% per year (while the optimal promise is 0%) and the private information solution. Similarly, we measure the cost of no-restriction by calculating the inflation promise that for the optimal interim mechanism yields the same value as the optimal discretionary policy. This inflation promise is around 0.5% (or -0.5%), implying that no restriction is as costly as using the optimal interim mechanism with a suboptimal initial inflation promise of 2% per year.

Another important property of the optimal interim mechanism is that no-discretion is at most a short-run phenomenon, and that promised inflation resides, in the long-run, within the

¹⁹We also compute the “expected” Ramsey policy defined in AAK, which solves the same mechanism design problem with an additional constraint that discretion is never given. We found that this policy performs worse than the optimal discretionary policy, regardless of the state π_-^e . This suggests that, at least in this example, gains from utilizing information are large.

²⁰This approach to quantifying the cost of discretion is analogous to the inflation equivalent measure of Jensen (2002) and Dennis and Soderstrom (2006).




Figure 8: Value Function Comparison

interval that prescribes some discretion. The left panel of Figure 9 depicts the policy function for π^e together with the 45 degree line (dashed). As for the full-information solution, promised inflation settles in an interval indicated by dashed lines, around zero, and once promised inflation enters this interval, it stays there forever. I.e. it is an ergodic set of promised inflation. Moreover, because the policy function is flatter for all θ than the 45 degree line outside this interval, promised inflation enters this interval in finite time, for any initial condition π_{-1}^e .

Note that this interval does not overlap at all the no-discretion regions, but instead contains only regions with some discretion. This implies that no-discretion is at most a short-run, transitional phenomenon in the optimal interim mechanism, and that at least some discretion is granted in the long-run. If society does not take π_{-1}^e as given and chooses the initial condition π_{-1}^e , then it chooses $\pi_{-1}^e = 0$, which is contained within the ergodic set. If π_{-1}^e can be chosen in the initial period, then we have the stronger result that the no-discretion region of the state space is never visited.

The right panel of Figure 9 displays the full-information solution for π^e alongside the private information solution for π^e and the 45 degree line. We see that the interval within which promised inflation settles is smaller for the private-information solution than for the full-information solution. Moreover, we generally observe that, for each π_{-1}^e , the difference be-

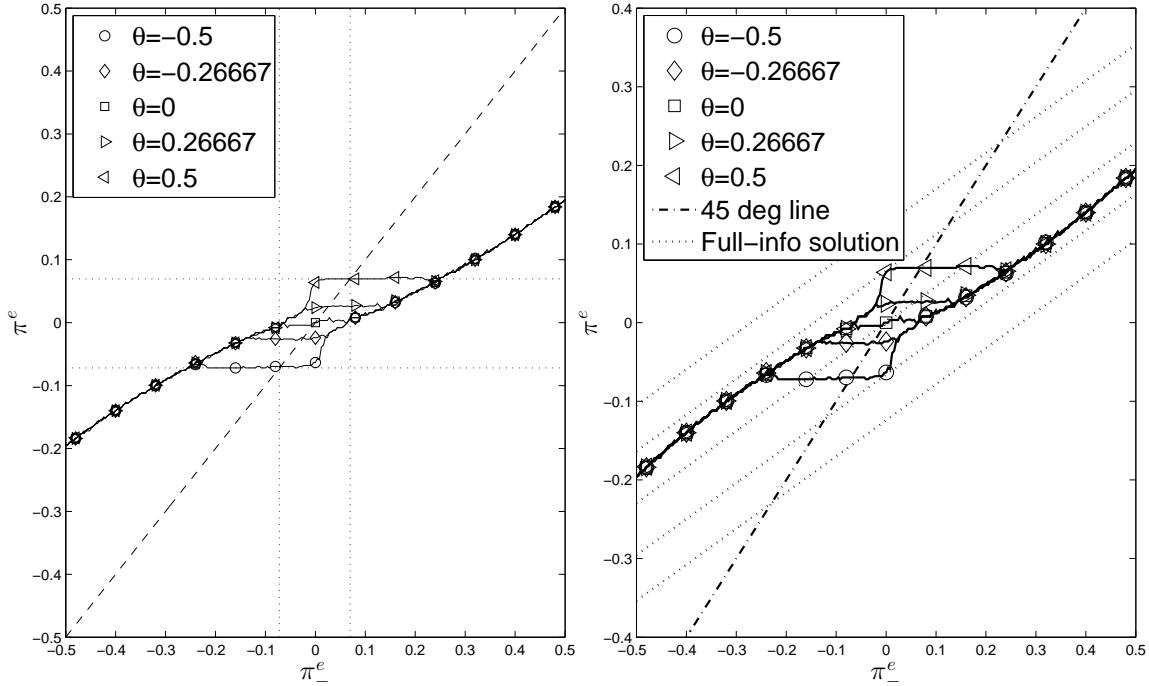


Figure 9: Long-run implication on π_-^e

tween the highest possible π^e and the lowest possible π^e is smaller for the private-information solution. Thus, one feature of the private-information solution is that it exhibits less volatility in expected inflation than the full-information solution, both along the transition path and in the ergodic set. This is natural. In the private-information setting, changing promised inflation from π_-^{e*} is more costly than in the full-information setting, because the mechanism imposes some restrictions on the central bank's ability to deliver such promises, which lowers social welfare. Accordingly, the private information solution makes less use of inflation promises.

7 Conclusion

In the context of the canonical New Keynesian model, we study the optimal degree of discretion that should be granted to a central bank when it has superior information about the welfare costs of inflation but is unable to commit. We show that the optimal mechanism depends on history through last period's inflation promise and that the optimal degree of discretion varies with this state variable. Although the central bank's ability to utilize its private information

should be curtailed, it is generally not optimal to grant the central bank either no-discretion or full-discretion. Full-discretion should be granted only when last period's inflation promise happens to maximize social welfare; no-discretion should be granted when last period's inflation promise is sufficiently far from the peak of social welfare; and some discretion should be granted for all intermediate values of last period's inflation promise. We demonstrate numerically that promised inflation must settle in a region within which some discretion is granted.

A practical implication of our analysis is that it is optimal to legislate an inflation-range targeting rule that specifies both upper and lower bounds on permissible inflation. It is essential to impose a lower bound, as the direction of the central bank's stabilization bias can be negative. Importantly, these bounds must be history-dependent to achieve the second-best, and a fixed-range targeting scheme is suboptimal. One way to encode history-dependence is to make the upper and lower bounds contingent upon promised inflation, a form of inflation target, announced by the central bank last period. Such history-dependence provides the central bank, which is unable to commit by itself, with a tool to restrict its future actions, and mitigates the stabilization bias.

Incorporating a persistent private shock would be an interesting extension of our analysis. In the full-information model, the gains from commitment do not change much when the shock persistence increases modestly, but decline sharply when the persistence becomes sufficiently high.²¹ We therefore conjecture that in the private information model it would still be optimal to limit the central bank's discretion to a similar extent as in the IID case when θ is moderately persistent, and that our results serve as a useful benchmark. When θ is highly persistent, the optimal degree of discretion can be much higher. A detailed analysis is warranted to examine the precise form of optimal delegation when θ is highly persistent.²² Also warranting a more detailed analysis is an environment in which the effects of private information can persist endogenously through inflation indexation or rule-of-thumb pricing. We leave this for future work.

Specifying a time-varying permissible inflation range has been seen in practice. Israel, for example, when adopting inflation targeting, did so by setting a sequence of decreasing tar-

²¹We confirmed this by computing the welfare difference between the full-information solution and the optimal discretionary policy for the quadratic specification with an AR(1) shock $\theta_t = \rho\theta_{t-1} + e_t$.

²²A recent paper by Halac and Yared (2014) considers the optimal, self-imposing fiscal rules when the government has a present-bias and persistent private information regarding the marginal value of public spending.

get ranges for the year-ahead inflation in an attempt to bring about disinflation (Bernanke, Laubach, Mishkin, and Posen, 2001). Our analysis suggests that fixed inflation targets, while practical, lack the sophistication needed to optimally trade off the gains and losses from discretion.

Appendix A: Proof of Propositions 2 — 5

Throughout, we assume that $F \in \mathcal{V}(\bar{\Pi})$.

7.1 Proofs of Propositions 2(iii), 2(iv), and 5.

Suppose that Assumptions 2, 3-(i), 4, and 5 hold.

We rewrite the $\mathbb{T}F$ -problem: define a slack variable $k(\theta) := W(\theta) - F(\pi^e(\theta))$ and replace the constraint $W(\theta) \leq F(\pi^e(\theta))$ with $k(\theta) \leq 0$. Then

$$\mathbb{T}F(\pi_-^e) = \max_{\pi(\cdot), x(\cdot), \pi^e(\cdot), k(\cdot)} \int_{\underline{\theta}}^{\bar{\theta}} \left\{ R(\pi(\theta), x(\theta), \theta) + \beta F(\pi^e(\theta)) + \beta k(\theta) \right\} p(\theta) d\theta,$$

subject to the constraints (13), (14), and

$$\begin{aligned} R(\pi(\theta), x(\theta), \theta) + \beta F(\pi^e(\theta)) + \beta k(\theta) &\geq R(\pi(\theta'), x(\theta'), \theta) + \beta F(\pi^e(\theta')) + \beta k(\theta'), \forall \theta, \theta' \neq \theta, \\ k(\theta) &\leq 0, \forall \theta. \end{aligned}$$

We call this problem (P1).

To relax (P1) further, define, for any $\pi \in \Pi$,

$$S(\pi; F) = \max_{(x, \pi^e) \in X \times \bar{\Pi}} B(x) + \beta F(\pi^e),$$

subject to $\pi = \kappa x + \beta \pi^e$. When obvious, we suppress the dependence of $S(\pi; F)$ on F . Then, for any triple (π, x, π^e) that satisfies the NKPC, $B(x) + \beta F(\pi^e) \leq S(\pi; F)$.

We now relax (P1), by introducing a slack variable $q(\theta) \leq 0$, replacing $B(x(\theta)) + \beta F(\pi^e(\theta))$ with $S(\pi(\theta)) + q(\theta)$, and dropping the NKPC. To simplify the notation, let

$$\tilde{R}(\pi, \theta) := A(\pi) + \pi\theta + S(\pi; F).$$

Then the following problem (P2) relaxes (P1):

$$\max_{\pi(\cdot), k(\cdot), q(\cdot)} \int_{\underline{\theta}}^{\bar{\theta}} \left[\tilde{R}(\pi(\theta); \theta) + \beta k(\theta) + q(\theta) \right] p(\theta) d\theta,$$

subject to the constraints (13),

$$\begin{aligned}\tilde{R}(\pi(\theta); \theta) + \beta k(\theta) + q(\theta) &\geq \tilde{R}(\pi(\theta'); \theta) + \beta k(\theta') + q(\theta'), \forall \theta, \theta' \neq \theta \\ k(\theta), q(\theta) &\leq 0, \forall \theta.\end{aligned}$$

The problem (P2) is equivalent to the following problem (P3):

$$\max_{\pi(\cdot), \delta(\cdot)} \int_{\underline{\theta}}^{\bar{\theta}} \left[\tilde{R}(\pi^e, \pi(\theta); \theta) + \delta(\theta) \right] p(\theta) d\theta,$$

subject to the constraints (13),

$$\begin{aligned}\tilde{R}(\pi(\theta); \theta) + \delta(\theta) &\geq \tilde{R}(\pi(\theta'); \theta) + \delta(\theta'), \forall \theta, \theta' \neq \theta, \\ \delta(\theta) &\leq 0, \forall \theta.\end{aligned}$$

Below we show that the solution to (P3) satisfies $\delta(\theta) = 0$, for all θ . This implies that the maximized value of (P3) is equal to that of (P2), and that (P2) has a solution with $k(\theta) = q(\theta) = 0$, for all θ . We then show that, from a solution to (P2) with $k(\theta) = q(\theta) = 0$, for all θ , one can recover a solution to (P1) with $k(\theta) = 0$, for all θ .

To show these results, we exploit the fact that, interpreting \tilde{R} as the return function in SWF and δ as the continuation value, (P3) has the same structure as the best payoff problem in AAK, except that average inflation, π^e , does not enter the return function and is exogenously fixed in this problem, whereas it enters the return function and is a choice variable in AAK.²³ Despite these differences, we argue that we can use AAK's results to prove that $\delta(\theta) = 0$, for all θ , is a property of the solution to (P3).

To apply AAK's results, their assumptions regarding \tilde{R} and p must be satisfied in our setting. We begin with "well-behavedness" of the return function – \tilde{R} is a strictly concave C^1 function with a piecewise C^1 derivative. We establish this by showing that the same properties hold for S . Consider the problem that defines the function S for a given $F \in \mathcal{V}(\bar{\Pi})$. Because F is strictly concave, it has a unique solution for each $\pi \in \Pi$. Let $(x_S(\cdot), \pi_S^e(\cdot)) : \Pi \rightarrow X \times \Pi$

²³The upper-bound for the continuation value in AAK is a constant that is not necessarily zero, but this difference is irrelevant.

be such that, for all $\pi \in \Pi$, $(x_S(\pi), \pi_S^e(\pi))$ is the solution to this maximization problem at π .²⁴ The following lemma establishes some properties of S and $(x_S(\cdot), \pi_S^e(\cdot))$.

Lemma 2 *S is a strictly concave C^1 function with a piecewise C^1 derivative on $\text{int}(\Pi)$. Both $x_S(\cdot)$ and $\pi_S^e(\cdot)$ are strictly increasing, continuous, piecewise C^1 functions.*

Proof. Because the graph of the constraint correspondence is convex and the objective function $B(x) + \beta F(\pi^e)$ is strictly concave by Assumption 2, S is strictly concave. The Benveniste-Scheinkman theorem implies that S is differentiable at any $\pi \in \text{int}(\Pi)$ with the derivative $B'(x_S(\pi))/\kappa$.

From the FONC of the maximization problem, $B'(x_S(\pi)) = \kappa F'((\pi - \kappa x_S(\pi))/\beta)$. Because both B' and F' are continuous and strictly decreasing, $x_S(\cdot)$ is a strictly increasing, continuous function. The same FONC $B'(x_S(\pi)) = \kappa F'(\pi_S^e(\pi))$ then implies that $\pi_S^e(\cdot)$ is also a strictly increasing, continuous, function.

We now show that $x_S(\cdot)$ and $\pi_S^e(\cdot)$ are piecewise C^1 . Combining the FONC and the NKPC, we obtain $\pi = (B')^{-1}(\kappa F'(\pi_S^e(\pi))) + \beta \pi_S^e(\pi)$. Let $G(\pi^e) := (B')^{-1}(\kappa F'(\pi^e)) + \beta \pi^e$. Then G is a strictly increasing continuous function, and it is C^1 except on a finite set of points and its right- and the left-derivatives exist, allowing for $+\infty$. Both the right- and the left-derivatives of G , D_+G and D_-G must be strictly positive.²⁵ Because $\pi = G(\pi_S^e(\pi))$, the composite function $G \circ \pi_S^e$ must be C^1 . Taking the right-derivative of $\pi = G(\pi_S^e(\pi))$ at an arbitrary π , we obtain

$$1 = D_+(G \circ \pi_S^e)(\pi) = \lim_{\Delta \downarrow 0} \frac{G(\pi_S^e(\pi + \Delta)) - G(\pi_S^e(\pi))}{\pi_S^e(\pi + \Delta) - \pi_S^e(\pi)} \times \frac{\pi_S^e(\pi + \Delta) - \pi_S^e(\pi)}{\Delta}.$$

If $D_+G = +\infty$ at $\pi^e = \pi_S^e(\pi)$, then we must have

$$\lim_{\Delta \downarrow 0} \frac{\pi_S^e(\pi + \Delta) - \pi_S^e(\pi)}{\Delta} = 0,$$

and $D_+\pi_S^e(\pi) = 0$. Otherwise, $D_+G > 0$ is finite at $\pi^e = \pi_S^e(\pi)$ and thus $D_+\pi_S^e(\pi) = 1/D_+G(\pi_S^e(\pi)) < +\infty$. Therefore, the right-derivative of π_S^e always exists. Analogously, the left-derivative always exists, and it is equal to $1/D_-G(\pi^e(\pi)) \in \mathbb{R}_{++}$ when $D_-G > 0$ is finite at $\pi^e = \pi_S^e(\pi)$, and to 0 when $D_-G = \infty$. Because G is C^1 except on a finite set of points and

²⁴They depend on the function F , but we suppress this dependence to simplify the notation.

²⁵ D_+ and D_- denote the right- and left-derivative operators, respectively.

$\pi^e(\cdot)$ is strictly increasing, $D_+G(\pi^e(\pi)) = D_-G(\pi^e(\pi))$ except on a finite set, and thus $\pi_S^e(\cdot)$ is piecewise C^1 . Because $x_S(\pi) = (\pi - \beta\pi_S^e(\pi))/\kappa$, $x_S(\cdot)$ is also piecewise C^1 .

Because $S' = B'(x(\pi))/\kappa$ and $x_S(\pi)$ is a continuous function that is piecewise C^1 , S is a C^1 function with a piecewise C^1 derivative. ■

It follows that \tilde{R} has the same properties as S .

Corollary 1 \tilde{R} is a strictly concave C^1 function with a piecewise C^1 derivative.

The single-crossing condition and the monotone-hazard condition in AAK follow from Assumption 2 and the definition of \tilde{R} , which together imply $\tilde{R}_{\pi\theta}(\pi, \theta) = \tilde{R}_{\theta\pi}(\pi, \theta) = 1$.

Lemma 3 Under Assumption 2, \tilde{R} satisfies the single-crossing condition: $\tilde{R}_{\pi\theta}(\pi, \theta) > 0$.

Lemma 4 Under Assumptions 2 and 3, the pair (\tilde{R}, p) satisfies the monotone hazard condition in AAK: For any $\pi(\cdot)$ that is non-decreasing,

$$\frac{1 - P(\theta)}{p(\theta)} \tilde{R}_{\theta\pi}(\pi(\theta), \theta) \text{ is strictly decreasing in } \theta,$$

and

$$\frac{P(\theta)}{p(\theta)} \tilde{R}_{\theta\pi}(\pi(\theta), \theta) \text{ is strictly increasing in } \theta.$$

Now we are ready to show that a solution to (P3) satisfies $\delta(\theta) = 0$ for all θ .

Proposition 7 For a given π_-^e , a solution to (P3), $(\pi(\cdot), \delta(\cdot))$, satisfies (i) $\delta(\theta) = 0$ for all θ , and (ii) $\pi(\cdot)$ is continuous.

Proof. In the proofs of Lemmas 1, 2, and 3 in AAK, they consider variations that keep the value of $\int_{\underline{\theta}}^{\bar{\theta}} \pi(\theta)p(\theta)d\theta$ (in their notation, ‘ x ’) unchanged. Therefore, these variations satisfy the constraint set in (P3). The single-crossing condition and the monotone hazard condition imply their Lemmas 1, 2, and 3. This means that their Proposition 1 holds for (P3), and that $\delta(\theta)$ is at its upper-bound for all θ . Thus (i) holds. Their Lemma 3 implies (ii). ■

The next corollary follows from Proposition 7.

Corollary 2 For a given π_-^e , there is a solution to (P2) such that $k(\theta) = q(\theta) = 0$ for all θ , and that $\pi(\cdot)$ is continuous.

Let $\pi_*(.)$ denote a solution to (P2) together with $(k(.), q(.)) = 0$. Defining the composite functions $x_* = x_S \circ \pi_*$, $\pi_*^e = \pi_S^e \circ \pi_*$, and $W_* = F \circ \pi_S^e \circ \pi_*$, then $(\pi_*, x_*, \pi_*^e, W_*)$ satisfies all the constraint in (P1), and the value of the objective function it achieves is the same as the maximized value of the relaxed problem (P2). This implies the following proposition.

Proposition 8 $(\pi_*(.), x_*(.), \pi_*^e(.), W_*(.))$ defined above is a solution to (P1).

This proves the third and fourth parts of Proposition 2 and Proposition 5.

7.2 Remaining proofs for Propositions 2(i), 3, and 4

7.2.1 π_* takes a simple form

First we show that π_* must be either constant or of the form

$$\pi_*(\theta) = \begin{cases} \pi_D(\theta_1), & \forall \theta \in [\underline{\theta}, \theta_1], \\ \pi_D(\theta), & \forall \theta \in (\theta_1, \theta_2), \\ \pi_D(\theta_2), & \forall \theta \in [\theta_2, \bar{\theta}]. \end{cases} \quad (21)$$

for a well-behaved function π_D .

To this end we replace the incentive compatibility constraint in (P3) with the local incentive compatibility constraint: (i) $\pi(.)$ is non-decreasing in θ , (ii)

$$\tilde{R}_\pi(\pi(\theta), \theta) \frac{\partial \pi(\theta)}{\partial \theta} + \frac{\partial \delta(\theta)}{\partial \theta} = 0,$$

whenever $\partial \pi(\theta)/\partial \theta$ and $\partial \delta(\theta)/\partial \theta$ exist, and (iii)

$$\lim_{\theta \uparrow \theta'} \tilde{R}(\pi(\theta), \theta') + \delta(\theta) = \lim_{\theta \downarrow \theta'} \tilde{R}(\pi(\theta), \theta') + \delta(\theta),$$

for all θ' at which these derivatives don't exist. (This definition is taken from AAK).

Because Proposition 7 implies $\partial \delta(\theta)/\partial \theta = 0$ for all θ , the incentive-compatibility constraint implies

$$\tilde{R}_\pi(\pi_*(\theta), \theta) \frac{\partial \pi_*(\theta)}{\partial \theta} = 0,$$

whenever $\partial \pi_*(\theta)/\partial \theta$ exists. Therefore, when the partial derivative exists, it is either $\partial \pi_*(\theta)/\partial \theta = 0$ or $\tilde{R}_\pi(\pi_*(\theta), \theta) = 0$.

Let $\pi_D(\cdot; F) : \Pi \rightarrow \Pi$ be the *one-shot discretionary best response given F*: for each θ , $\pi_D(\theta; F)$ solves

$$\max_{\pi} \left\{ A(\pi) + \theta\pi + \max_{(x, \pi^e): \pi = \kappa x + \beta \pi^e} \{ B(x) + \beta F(\pi^e) \} \right\} = \max_{\pi} \tilde{R}(\pi, \theta).$$

This is an unconstrained optimization of a strictly concave function with the first-order condition $\tilde{R}_{\pi}(\pi_D(\theta; F), \theta) = 0$. Therefore $\pi_*(\theta) = \pi_D(\theta; F)$ if and only if $\tilde{R}_{\pi}(\pi_*(\theta), \theta) = 0$. It follows that $\pi_*(\theta) = \pi_D(\theta)$ whenever $\partial \pi_*(\theta) / \partial \theta$ exists and is non-zero.

As π_* has to satisfy Assumption 1, we need to show that π_D is a piecewise C^1 function. The next lemma establishes this. For simplicity, we drop the dependence of π_D on F hereafter.

Lemma 5 $\pi_D(\cdot)$ is a strictly increasing, continuous, piecewise C^1 function.

Proof. Note that $\tilde{R}_{\pi}(\pi, \theta) = A'(\pi) + \theta + S'(\pi)$. Then, for each θ , $\pi_D(\theta)$ is a solution to

$$\theta = -(A'(\pi) + S'(\pi)).$$

The RHS is a strictly increasing, continuous, piecewise C^1 function, and goes to ∞ ($-\infty$) as $\pi \uparrow \infty$ ($\pi \downarrow -\infty$). Therefore we can invert this relationship to obtain a strictly increasing, continuous function, $\pi_D(\cdot)$, that is piecewise C^1 . The right- and left-derivatives of π_D are

$$(D_+ \pi_D(\theta), D_- \pi_D(\theta)) = \left(\frac{-1}{A''(\pi_D(\theta)) + D_+ S'(\pi_D(\theta))}, \frac{-1}{A''(\pi_D(\theta)) + D_- S'(\pi_D(\theta))} \right).$$

Note that, since $D_+ S' \leq 0$, $D_- S' \leq 0$, $A''(\pi) < \bar{A}'' < 0$ for some constant \bar{A}'' , the RHS is finite. Thus π_D is differentiable at θ if and only if S' is differentiable at $\pi_D(\theta)$. As S' is C^1 except on a finite set of points and π_D is strictly increasing, π_D is piecewise C^1 . ■

Because π_* is continuous (Proposition 7 (ii)), it follows that it must be either constant or of the form in equation (21). Later we show that either $\theta_1 = \underline{\theta}$ or $\theta_2 = \bar{\theta}$ must hold.

7.2.2 $\mathbb{T}F$ is single-peaked

Now we show that $\mathbb{T}F$ is single-peaked, that it is strictly increasing on the left of its peak, and that it is strictly decreasing on the right of its peak.

Recall that the maximized value of (P3) is the same as that of (P1):

$$\mathbb{T}F(\pi_-^e) = \max_{\pi(\cdot)} \int_{\underline{\theta}}^{\bar{\theta}} \tilde{R}(\pi(\theta); \theta) p(\theta) d\theta \quad (22)$$

subject to equation (13) and

$$\tilde{R}(\pi(\theta); \theta) \geq \tilde{R}(\pi(\theta'); \theta), \quad \forall \theta, \theta' \neq \theta.$$

We call the problem on the RHS of equation (22) the problem (P4).

Proposition 9 *The function $\mathbb{T}F : \bar{\Pi} \rightarrow \mathbb{R}$ is uniquely maximized at $\pi_-^e = \pi_-^{e*} := \int_{\underline{\theta}}^{\bar{\theta}} \pi_D(\theta) p(\theta) d\theta$.*

Proof. The objective function in (P4) is maximized if and only if $\pi = \pi_D$ a.e., because \tilde{R} is strictly concave in π for each θ . $\pi = \pi_D$ satisfies the constraint (13) if and only if $\pi_-^e = \pi_-^{e*}$. Thus the function $\mathbb{T}F$ is maximized at $\pi_-^e = \pi_-^{e*}$ and the maximum is unique. ■

Proposition 9 together with Assumption 6 implies Proposition 3.

Corollary 3 *$\mathbb{T}F$ is strictly increasing for $\pi_-^e < \pi_-^{e*}$ and is strictly decreasing for $\pi_-^e > \pi_-^{e*}$.*

Proof. Let $\pi_1^e < \pi_2^e < \pi_-^{e*}$. We show that $\mathbb{T}F(\pi_1^e) < \mathbb{T}F(\pi_2^e)$. Let $\pi_*(\cdot; \pi_1^e)$ be a solution to (P4) at $\pi_-^e = \pi_1^e$, then it is of the form of equation (21) for some θ_1 and θ_2 . Notice that $\theta_2 < \bar{\theta}$, because otherwise $\pi_*(\cdot; \pi_1^e) \geq \pi_D(\cdot)$ with strict equality for $\theta > \theta_1$, and the expected value of $\pi_*(\cdot; \pi_1^e)$ satisfies

$$\pi_1^e = \int_{\underline{\theta}}^{\bar{\theta}} \pi_*(\theta; \pi_1^e) p(\theta) d\theta \geq \int_{\underline{\theta}}^{\bar{\theta}} \pi_D(\theta) p(\theta) d\theta = \pi_-^{e*},$$

which is a contradiction.

Because $\pi_2^e \in (\pi_1^e, \pi_-^{e*})$ and π_D is strictly increasing, there exists $\theta_3 \in (\theta_2, \bar{\theta})$ such that

$$\pi_2^e = \int_{\underline{\theta}}^{\theta_1} \pi_D(\theta_1) p(\theta) d\theta + \int_{\theta_1}^{\theta_3} \pi_D(\theta) p(\theta) d\theta + \int_{\theta_3}^{\bar{\theta}} \pi_D(\theta_3) p(\theta) d\theta.$$

For such θ_3 , define $\pi'(\cdot)$ as follows: $\pi'(\theta) = \pi_*(\theta; \pi_1^e)$ for all $\theta < \theta_2$, $\pi(\theta) = \pi_D(\theta)$ for all $\theta \in [\theta_2, \theta_3)$, and $\pi'(\theta) = \pi_D(\theta_3)$, for all $\theta \in [\theta_3, \bar{\theta}]$. Then π' satisfies the constraints in (P4) at

$\pi_-^e = \pi_2^e$. Hence $\mathbb{T}F(\pi_2^e) \geq \int_{\underline{\theta}}^{\bar{\theta}} \tilde{R}(\pi'(\theta); \theta)p(\theta)d\theta$. Because π' and $\pi_*(\cdot; \pi_-^e)$ are identical up to θ_2 ,

$$\begin{aligned} & \int_{\underline{\theta}}^{\bar{\theta}} \tilde{R}(\pi'(\theta); \theta)p(\theta)d\theta - \mathbb{T}F(\pi_1^e) \\ &= \int_{\theta_2}^{\theta_3} \{\tilde{R}(\pi_D(\theta); \theta) - \tilde{R}(\pi_D(\theta_2); \theta)\}p(\theta)d\theta + \int_{\theta_3}^{\bar{\theta}} \{\tilde{R}(\pi_D(\theta_3); \theta) - \tilde{R}(\pi_D(\theta_2); \theta)\}p(\theta)d\theta. \end{aligned}$$

The first integral on the RHS is strictly positive. The second integral on the RHS is also strictly positive, because for all $\theta > \theta_3$, $\pi_D(\theta) > \pi_D(\theta_3) > \pi_D(\theta_2)$, and the concavity of \tilde{R} implies

$$\tilde{R}(\pi_D(\theta); \theta) > \tilde{R}(\pi_D(\theta_3); \theta) > \tilde{R}(\pi_D(\theta_2); \theta)$$

for all $\theta > \theta_3$. Therefore $\int_{\underline{\theta}}^{\bar{\theta}} \tilde{R}(\pi'(\theta); \theta)p(\theta)d\theta > \mathbb{T}F(\pi_1^e)$, establishing $\mathbb{T}F(\pi_2^e) > \mathbb{T}F(\pi_1^e)$.

An analogous argument shows that $\mathbb{T}F$ is strictly decreasing for $\pi_-^e > \pi_-^{e*}$. ■

7.2.3 Proof of Proposition 4

Lemma 6 (i) For $\pi_-^e < \pi_-^{e*}$, then $\pi_*(\cdot)$ is either constant or has the form in equation (21) with $\theta_1 = \underline{\theta}$. (ii) For $\pi_-^e > \pi_-^{e*}$, $\pi_*(\cdot)$ is either constant or has the form in equation (21) with $\theta_2 = \bar{\theta}$.

Proof. Suppose to the contrary that, for some $\pi_-^e < \pi_-^{e*}$, $\pi_*(\cdot)$ has the form in equation (21) with $\theta_1 > \underline{\theta}$. Fix such $\pi_-^e < \pi_-^{e*}$. Because $\mathbb{T}F$ is strictly increasing by Corollary 3, replacing the first constraint in P4 with

$$\pi_-^e \geq \int_{\underline{\theta}}^{\bar{\theta}} \pi(\theta)p(\theta)d\theta,$$

must not increase the maximized value. Let $\pi_{**}(\cdot)$ be such that $\pi_{**}(\theta) = \pi_*(\theta)$ for all $\theta > \theta_1$ and $\pi_{**}(\theta) = \pi_D(\theta)$ for all $\theta \leq \theta_1$. Then π_{**} is locally incentive compatible and $\pi_-^e > \int_{\underline{\theta}}^{\bar{\theta}} \pi_{**}(\theta)p(\theta)d\theta$.

Moreover, the objective function increases by

$$\int_{\underline{\theta}}^{\theta_1} \{\tilde{R}(\pi_{**}(\theta); \theta) - \tilde{R}(\pi_*(\theta); \theta)\}p(\theta)d\theta = \int_{\underline{\theta}}^{\theta_1} \{\tilde{R}(\pi_D(\theta); \theta) - \tilde{R}(\pi_D(\theta_1); \theta)\}p(\theta)d\theta > 0,$$

because the integrand is strictly positive for all $\theta < \theta_1$. This is a contradiction, and thus $\theta_1 = \underline{\theta}$ must hold. Part (ii) can be shown in the same way, and we omit the proof. ■

We introduce two kinds of threshold functions, $T_1 : (\pi_D(\underline{\theta}), \pi_-^{e*}) \rightarrow [\underline{\theta}, \bar{\theta}]$ and $T_2 : (\pi_-^{e*}, \pi_D(\bar{\theta})) \rightarrow$

$[\underline{\theta}, \bar{\theta}]$, which are implicitly defined by

$$\begin{aligned}\pi_-^e &= \int_{\underline{\theta}}^{T_1(\pi_-^e)} \pi_D(\theta)p(\theta)d\theta + [1 - P(T_1(\pi_-^e))]\pi_D(T_1(\pi_-^e)) \quad \text{for } \pi_-^e < \pi_-^{e*}, \\ \pi_-^e &= P(T_2(\pi_-^e))\pi_D(T_2(\pi_-^e)) + \int_{T_2(\pi_-^e)}^{\bar{\theta}} \pi_D(\theta)p(\theta)d\theta \quad \text{for } \pi_-^e > \pi_-^{e*}.\end{aligned}$$

Lemma 7 *Both T_1 and T_2 are strictly increasing, continuous, piecewise C^1 functions. Composite functions $\pi_D \circ T_1$ and $\pi_D \circ T_2$ are C^1 and their derivatives are $[1 - P(T_1(\pi_-^e))]^{-1}$ and $P(T_2(\pi_-^e))^{-1}$, respectively. When T_1 and T_2 are differentiable, so is π_D , and their derivatives are given by $\frac{\partial T_1(\pi_-^e)}{\partial \pi_-^e} = [1 - P(T_1(\pi_-^e))]^{-1} / \frac{\partial \pi_D(T_1(\pi_-^e))}{\partial \theta}$ and $\frac{\partial T_2(\pi_-^e)}{\partial \pi_-^e} = P(T_2(\pi_-^e))^{-1} / \frac{\partial \pi_D(T_2(\pi_-^e))}{\partial \theta}$.*

Proof. We show this for T_1 only. Let

$$H(b) := \int_{\underline{\theta}}^b \pi_D(\theta)p(\theta)d\theta + [1 - P(b)]\pi_D(b).$$

Since π_D is a strictly increasing, continuous, piecewise C^1 function, so is H . Thus H has an inverse that is strictly increasing and continuous. It follows that $T_1(\pi_-^e) = H^{-1}(\pi_-^e)$.

T_1 is piecewise C^1 : Since $\pi_-^e = H(T_1(\pi_-^e))$ and H is piecewise C^1 , we have

$$\begin{aligned}1 &= D_+(H \circ T_1)(\pi_-^e) = D_+H(T_1(\pi_-^e)) \times D_+T_1(\pi_-^e), \\ 1 &= D_-(H \circ T_1)(\pi_-^e) = D_-H(T_1(\pi_-^e)) \times D_-T_1(\pi_-^e).\end{aligned}$$

Therefore $D_+T_1 = D_-T_1$ except on a finite set, and T_1 is piecewise C^1 .

$h := \pi_D \circ T_1$ is C^1 : Note that

$$\begin{aligned}1 &= D_+(H \circ T_1)(\pi_-^e) = [1 - P(T_1(\pi_-^e))] \times D_+h(\pi_-^e), \\ 1 &= D_-(H \circ T_1)(\pi_-^e) = [1 - P(T_1(\pi_-^e))] \times D_-h(\pi_-^e).\end{aligned}$$

It follows that $D_-h(\pi_-^e) = D_+h(\pi_-^e) = [1 - P(T_1(\pi_-^e))]^{-1}$. Since T_1 is continuous, the rightmost term is continuous in π_-^e . This proves $h = \pi_D \circ T_1$ is C^1 .

Since $h := \pi_D \circ T_1$ is C^1 , it follows that T_1 is differentiable whenever π_D is, and that the product of $\partial \pi_D(T_1(\pi_-^e)) / \partial \theta$ and $\partial T_1(\pi_-^e) / \partial \pi_-^e$ equals $[1 - P(T_1(\pi_-^e))]^{-1}$. ■

Proposition 10 *For each π_-^e , there is a unique solution to (P4) and it has the form described*

in Proposition 4.

Proof. We have already seen that a solution is unique for $\pi_-^e = \pi_-^{e*}$. Consider $\pi_-^e < \pi_-^{e*}$. Then π_* at π_-^e is either constant or satisfies:

$$\pi_*(\theta) = \begin{cases} \pi_D(\theta), & \forall \theta \in [\underline{\theta}, \theta^\#], \\ \pi_D(\theta^\#), & \forall \theta \in [\theta^\#, \bar{\theta}], \end{cases} \quad (23)$$

for some $\theta^\#$, and $\pi_-^e = \int_{\underline{\theta}}^{\bar{\theta}} \pi_*(\theta)p(\theta)d\theta$. Note that when π_* is not constant,

$$\pi_-^e = \int_{\underline{\theta}}^{\theta_1^\#} \pi_D(\theta)p(\theta)d\theta + [1 - P(\theta_1^\#)]\pi_D(\theta^\#). \quad (24)$$

The RHS of equation (24) is strictly increasing in $\theta_1^\#$, and takes values from $\pi_D(\underline{\theta})$ to π_-^{e*} . This implies that for any $\pi_-^e \leq \pi_D(\underline{\theta})$, π_* has to be constant and satisfies $\pi_*(\theta) = \pi_-^e$ for all θ .

For $\pi_-^e \in (\pi_D(\underline{\theta}), \pi_-^{e*})$, either π_* is constant or it has the form in equation (23) with $\theta_1^\# = T_1(\pi_-^e)$. We show that a constant π_* is not a solution. Let π be the rule in equation (23) with $\theta_1^\# = T_1(\pi_-^e)$. Then for $\pi_-^e = \int_{\underline{\theta}}^{\bar{\theta}} \pi(\theta)p(\theta)d\theta$,

$$\begin{aligned} & \int_{\underline{\theta}}^{\bar{\theta}} \tilde{R}(\pi(\theta); \theta)p(\theta)d\theta - \int_{\underline{\theta}}^{\bar{\theta}} \tilde{R}(\pi_-^e; \theta)p(\theta)d\theta \\ &= \int_{\underline{\theta}}^{\theta_1^\#} [\tilde{R}(\pi_D(\theta); \theta) - \tilde{R}(\pi_-^e; \theta)]p(\theta)d\theta + \int_{\theta_1^\#}^{\bar{\theta}} [\tilde{R}(\pi_D(\theta_1^\#); \theta) - \tilde{R}(\pi_-^e; \theta)]p(\theta)d\theta. \end{aligned}$$

The first term is strictly positive. The second term is strictly positive, because $\pi_D(\theta) > \pi_D(\theta_1^\#) > \pi_-^e$ for all $\theta \geq \theta_1^\#$ and \tilde{R} is strictly concave. This proves that there is unique solution to (P4) for each $\pi_-^e < \pi_-^{e*}$. The proof for $\pi_-^e > \pi_-^{e*}$ is analogous. ■

This proposition together with Assumption 6 implies Proposition 4.

7.2.4 Proof of Proposition 2(i)

Proof. Note that, denoting $U(\theta) = \tilde{R}(\pi_*(\theta), \theta)$,

$$\mathbb{T}F(\pi_-^e) = U(\underline{\theta}) + \int_{\underline{\theta}}^{\bar{\theta}} \frac{1 - P(\theta)}{p(\theta)} \tilde{R}_\theta(\pi_*(\theta), \theta)p(\theta)d\theta.$$

We begin by showing that the first derivative of $\mathbb{T}F$ is continuous.

For $\pi_-^e < \pi_D(\underline{\theta})$, we have $\pi_*(\theta) = \pi_-^e$ for all θ and hence

$$\begin{aligned}\frac{\partial \mathbb{T}F(\pi_-^e)}{\partial \pi_-^e} &= \frac{\partial \tilde{R}(\pi_-^e, \underline{\theta})}{\partial \pi_-^e} + \int_{\underline{\theta}}^{\bar{\theta}} \frac{1 - P(\theta)}{p(\theta)} \tilde{R}_{\theta\pi}(\pi_-^e, \theta) p(\theta) d\theta \\ &= \frac{\partial \tilde{R}(\pi_-^e, \underline{\theta})}{\partial \pi_-^e} + (E[\theta] - \underline{\theta}).\end{aligned}$$

The last expression on the RHS is continuous for $\pi_-^e < \pi_D(\underline{\theta})$, because \tilde{R} is C^1 , and is also strictly decreasing in π_-^e . The same result obtains for the left-derivative of $\mathbb{T}F$ at $\pi_-^e = \pi_D(\underline{\theta})$ with the first term on the RHS replaced by the left-derivative of \tilde{R} at $\pi_D(\underline{\theta})$, which is zero. Therefore, the left-derivative of $\mathbb{T}F$ is continuous for $\pi_-^e \leq \pi_D(\underline{\theta})$, and equals $E[\theta] - \underline{\theta}$ at $\pi_-^e = \pi_D(\underline{\theta})$.

For $\pi_-^e \in [\pi_D(\underline{\theta}), \pi_-^{e*}]$, $U(\underline{\theta}) = \tilde{R}(\pi_D(\underline{\theta}), \underline{\theta})$ is independent of π_-^e , and

$$\mathbb{T}F(\pi_-^e) = U(\underline{\theta}) + \int_{\underline{\theta}}^{T_1(\pi_-^e)} \frac{1 - P(\theta)}{p(\theta)} \pi_D(\theta) p(\theta) d\theta + \pi_D(T_1(\pi_-^e)) \int_{T_1(\pi_-^e)}^{\bar{\theta}} \frac{1 - P(\theta)}{p(\theta)} p(\theta) d\theta.$$

Recall that $h := \pi_D \circ T_1$ is C^1 . Therefore, for $\pi_-^e \in (\pi_D(\underline{\theta}), \pi_-^{e*})$, the derivative of the RHS exists and equal to

$$\begin{aligned}\frac{\partial h(\pi_-^e)}{\partial \pi_-^e} \int_{T_1(\pi_-^e)}^{\bar{\theta}} \{1 - P(\theta)\} d\theta &= \frac{1}{1 - P(T_1(\pi_-^e))} \int_{T_1(\pi_-^e)}^{\bar{\theta}} [1 - P(\theta)] d\theta \\ &= \int_{T_1(\pi_-^e)}^{\bar{\theta}} \frac{1 - P(\theta)}{p(\theta)} \frac{p(\theta)}{1 - P(T_1(\pi_-^e))} d\theta.\end{aligned}$$

The last expression on the RHS is continuous for $\pi_-^e \in (\pi_D(\underline{\theta}), \pi_-^{e*})$, because T_1 is a continuous function, and is also strictly decreasing in π_-^e because $(1 - P(\theta))/p(\theta)$ is strictly decreasing. Again the analogous equations obtain for $D_+\mathbb{T}F$ at $\pi_-^e = \pi_D(\underline{\theta})$ and for $D_-\mathbb{T}F$ at $\pi_-^e = \pi_-^{e*}$, and their values are respectively $E[\theta] - \underline{\theta}$ and 0.

Taken together, the first derivative of $\mathbb{T}F$ is continuous for $\pi_-^e < \pi_-^{e*}$, and the left-derivative of $\mathbb{T}F$ is zero at $\pi_-^e = \pi_-^{e*}$. Using the same argument for $\pi_-^e \geq \pi_-^{e*}$, one can show that the first derivative of $\mathbb{T}F$ is continuous for $\pi_-^e > \pi_-^{e*}$, and that the right-derivative of $\mathbb{T}F$ is zero at $\pi_-^e = \pi_-^{e*}$. This proves that $\mathbb{T}F$ is a C^1 function. We have also shown that $\partial \mathbb{T}F / \partial \pi_-^e$ is strictly decreasing, implying that $\mathbb{T}F$ is strictly concave.

Now we turn to the second derivative of $\mathbb{T}F$. For $\pi_-^e < \pi_D(\underline{\theta})$, the left- and the right-

derivatives of $\partial\mathbb{T}F/\partial\pi_-^e$ are the left- and the right-derivatives of $\tilde{R}(\pi_-^e, \underline{\theta})$ with respect to the first argument. Because \tilde{R}_π is piecewise C^1 , so is $\partial\mathbb{T}F/\partial\pi_-^e$. The left-derivative of $\partial\mathbb{T}F/\partial\pi_-^e$ at $\pi_-^e = \pi_D(\underline{\theta})$ is $D_- \tilde{R}_\pi(\pi_D(\underline{\theta}), \underline{\theta})$, which is finite.

Consider $\pi_-^e \in (\pi_D(\underline{\theta}), \pi_-^{e*})$. Let $J(x) := \int_x^{\bar{\theta}} [1 - P(\theta)] d\theta / [1 - P(x)]$ for $x \in [\underline{\theta}, \bar{\theta}]$, then it is C^1 on $(\underline{\theta}, \bar{\theta})$, and $\partial\mathbb{T}F/\partial\pi_-^e = J \circ T_1$. Because $T_1 : (\pi_D(\underline{\theta}), \pi_-^{e*}) \rightarrow (\underline{\theta}, \bar{\theta})$ is piecewise C^1 , the derivative of $\partial\mathbb{T}F/\partial\pi_-^e$ is piecewise C^1 . The right-derivative of $\partial\mathbb{T}F/\partial\pi_-^e$ is

$$\begin{aligned} D_+ \frac{\partial\mathbb{T}F(\pi_-^e)}{\partial\pi_-^e} &= D_+ \left[\frac{1}{1 - P(T_1(\pi_-^e))} \int_{T_1(\pi_-^e)}^{\bar{\theta}} [1 - P(\theta)] d\theta \right] \\ &= D_+ T_1 \times \frac{p(T_1)}{1 - P(T_1)} \left[\int_{T_1(\pi_-^e)}^{\bar{\theta}} \frac{1 - P(\theta)}{p(\theta)} \frac{p(\theta)}{1 - P(T_1)} d\theta - \frac{1 - P(T_1)}{p(T_1)} \right], \end{aligned}$$

and the analogous equation obtains for the left-derivative.

What remains to show is whether the right-derivative of $\partial\mathbb{T}F/\partial\pi_-^e$ exists at $\pi_-^e = \pi_D(\underline{\theta})$ and whether the left-derivative exists at $\pi_-^e = \pi_-^{e*}$, allowing for $-\infty$.

First we show that the right-derivative of $\partial\mathbb{T}F/\partial\pi_-^e$ exists at $\pi_-^e = \pi_D(\underline{\theta})$. To this end, we prove that as $\Delta \downarrow 0$,

$$\left(\frac{\partial\mathbb{T}F(\pi_D(\underline{\theta}) + \Delta)}{\partial\pi_-^e} - \frac{\partial\mathbb{T}F(\pi_D(\underline{\theta}))}{\partial\pi_-^e} \right) / \Delta$$

converges to a constant. Denote $T_1(\pi_D(\underline{\theta}) + \Delta)$ by $T_1(\Delta)$ for simplicity, this term equals

$$\begin{aligned} & \frac{-1}{\Delta} \left[\int_{\underline{\theta}}^{T_1(\Delta)} [1 - P(\theta)] d\theta + \left(1 - \frac{1}{1 - P(T_1(\Delta))}\right) \int_{T_1(\Delta)}^{\bar{\theta}} [1 - P(\theta)] d\theta \right] \\ &= - \left[\frac{\int_{\underline{\theta}}^{T_1(\Delta)} [1 - P(\theta)] d\theta}{T_1(\Delta) - \underline{\theta}} \frac{T_1(\Delta) - \underline{\theta}}{\Delta} + \frac{1 - \frac{1}{1 - P(T_1(\Delta))}}{T_1(\Delta) - \underline{\theta}} \frac{T_1(\Delta) - \underline{\theta}}{\Delta} \int_{T_1(\Delta)}^{\bar{\theta}} [1 - P(\theta)] d\theta \right] \\ &= \frac{T_1(\Delta) - \underline{\theta}}{\Delta} \left[\frac{\frac{1}{1 - P(T_1(\Delta))} - 1}{T_1(\Delta) - \underline{\theta}} \int_{T_1(\Delta)}^{\bar{\theta}} [1 - P(\theta)] d\theta - \frac{\int_{\underline{\theta}}^{T_1(\Delta)} [1 - P(\theta)] d\theta}{T_1(\Delta) - \underline{\theta}} \right]. \end{aligned}$$

It is easy to show that $D_+ T_1(\pi_D(\underline{\theta}))$ exists and equals $1/D_+ \pi_D(\underline{\theta})$. Thus the first term $(T_1(\Delta) - \underline{\theta})/\Delta$ converges to $D_+ T_1(\pi_D(\underline{\theta}))$ as $\Delta \downarrow 0$, which equals $1/D_+ \pi_D(\underline{\theta}) = -D_+ \tilde{R}_\pi(\pi_D(\underline{\theta}), \underline{\theta})$ because $\tilde{R}_\pi(\pi_D(\underline{\theta}), \underline{\theta}) = 0$. Terms in the brackets converges to

$$\frac{p(\underline{\theta})}{(1 - P(\underline{\theta}))^2} \int_{\underline{\theta}}^{\bar{\theta}} [1 - P(\theta)] d\theta - (1 - P(\underline{\theta})) = p(\underline{\theta}) \int_{\underline{\theta}}^{\bar{\theta}} [1 - P(\theta)] d\theta - (1 - P(\underline{\theta}))$$

as $\Delta \downarrow 0$. Thus the right-derivative of $\partial\mathbb{T}F/\partial\pi_-^e$ exists and is finite at $\pi_-^e = \pi_D(\underline{\theta})$.

Now we show that $D_- \partial \mathbb{T}F / \partial \pi_-^e$ exists at π_-^{e*} . Let $T_1(\Delta) = T_1(\pi_-^{e*} - \Delta)$, then

$$\begin{aligned} \frac{1}{\Delta} \left\{ \frac{\partial \mathbb{T}F(\pi_-^{e*})}{\partial \pi_-^e} - \frac{\partial \mathbb{T}F(\pi_-^{e*} - \Delta)}{\partial \pi_-^e} \right\} &= \frac{-1}{\Delta} \frac{1}{1 - P(T_1(\Delta))} \int_{T_1(\Delta)}^{\bar{\theta}} [1 - P(\theta)] d\theta \\ &= \left[\frac{-1}{p(T_1(\Delta))} \frac{1}{1 - P(T_1(\Delta))} \frac{\int_{T_1(\Delta)}^{\bar{\theta}} \frac{(1-P(\theta))/p(\theta)}{(1-P(T_1(\Delta)))/p(T_1(\Delta))} p(\theta) d\theta}{\bar{\theta} - T_1(\Delta)} \right] \\ &\quad \times \left\{ (1 - P(T_1(\Delta))) \frac{\bar{\theta} - T_1(\Delta)}{\Delta} \right\}. \end{aligned}$$

The terms in the square brackets diverges to $-\infty$ as $\Delta \downarrow 0$. The remaining terms converge to a finite, strictly positive number, and thus the right-hand side diverges to $-\infty$. To see this, observe that, by definition of T_1 ,

$$\Delta = \pi_-^{e*} - (\pi_-^{e*} - \Delta) = \int_{T_1(\Delta)}^{\bar{\theta}} \{\pi_D(\theta) - \pi_D(T_1(\Delta))\} p(\theta) d\theta,$$

and that

$$\begin{aligned} 1 &= \lim_{\Delta \downarrow 0} \frac{\int_{T_1(\Delta)}^{\bar{\theta}} \{\pi_D(\theta) - \pi_D(T_1(\Delta))\} p(\theta) d\theta}{\Delta} \\ &= \lim_{\Delta \downarrow 0} \left[\frac{\int_{T_1(\Delta)}^{\bar{\theta}} \{\pi_D(\theta) - \pi_D(T_1(\Delta))\} \frac{p(\theta)}{1 - P(T_1(\Delta))} d\theta}{\bar{\theta} - T_1(\Delta)} \right] \times (1 - P(T_1(\Delta))) \frac{\bar{\theta} - T_1(\Delta)}{\Delta}. \end{aligned}$$

As $\Delta \downarrow 0$, terms in the square brackets converges to $D_- \pi_D(\bar{\theta}) = -1/D_- \tilde{R}_\pi(\pi_D(\bar{\theta}), \bar{\theta})$, which is a finite, strictly positive number. This implies that $\lim_{\Delta \downarrow 0} (1 - P(T_1(\Delta))) \frac{\bar{\theta} - T_1(\Delta)}{\Delta}$ exists and is strictly positive.

Using a symmetric argument, one can show the same properties hold for $\pi_-^e > \pi_-^{e*}$. ■

7.3 Proof of Proposition 6

Proof. Conditions 1 and 3 are clearly satisfied. It follows from condition 1 that $m = \pi_{IT}^e(m; m_-)$ for any $m \in \bar{\Pi}$. To prove condition 2, observe that condition 1 implies $m = \pi_{IT}^e(m; m_-)$ for any $m \in \bar{\Pi}$, and replace the function $\pi_{IT}^e(m; m_-)$ in the constraint with m . Then Proposition 5 implies that the maximization problem in condition 2 is the problem in equation (20), with the additional constraint $\pi \in \bar{\Gamma}(m_-)$. Because $\bar{\Gamma}(m_-)$ is an interval for each π_-^e , for a given θ , $\pi_D(\theta)$ is either in $\bar{\Gamma}(m_-)$, or smaller than any element in $\bar{\Gamma}(m_-)$, or larger

than any elements in $\bar{\Gamma}(m_-)$. Because the objective function in equation (20) is strictly concave in inflation, the optimal inflation choice for given (θ, m_-) is (i) $\pi_D(\theta)$ if $\pi_D(\theta) \in \bar{\Gamma}(m_-)$, (ii) the smallest element of $\bar{\Gamma}(m_-)$ if $\pi_D(\theta) \leq \pi$ for all $\pi \in \bar{\Gamma}(m_-)$, and (iii) the largest element of $\bar{\Gamma}(m_-)$ if $\pi_D(\theta) \geq \pi$ for all $\pi \in \bar{\Gamma}(m_-)$. This implies that π_* solves this problem, and condition 2 is met for $(\pi_{IT}, x_{IT}, m_{IT}) = (\pi_*, x_*, \pi_*^e)$. ■

Appendix B: Benchmark Problems

In this appendix, we focus on the quadratic specification in equation (2).

B.1. Full-information problem

The full-information problem has the following recursive formulation:

$$W^{FI}(\pi_-^e) = \max_{\pi(\cdot), x(\cdot), \pi^e(\cdot)} \int_{\theta} \left\{ -\frac{1}{2}(\pi(\theta) - \theta)^2 - \frac{b}{2}x(\theta)^2 + \beta W^{FI}(\pi^e(\theta)) \right\} p(\theta) d\theta,$$

with the constraints given by

$$\pi^e(\theta) = \frac{1}{\beta}\pi(\theta) - \frac{\kappa}{\beta}x(\theta),$$

for all θ , and

$$\pi_-^e = \int_{\Theta} \pi(\theta) p(\theta) d\theta.$$

Because the return function is quadratic and the constraints are linear, the value function is quadratic and the policy function is linear. For simplicity, we have disposed of the compactness of Π and X and assumed that π and x can be chosen from the real line.

B.2. Optimal discretionary policy

We can solve analytically for the optimal discretionary policy, because the problem is linear-quadratic. (A Markov perfect equilibrium is unique when the return function has the form in equation (2).) This policy depends only on the current shock θ , and is given by

$$(\pi^{MP}(\theta), x^{MP}(\theta)) = \left(\frac{\kappa^2/b}{1 + \kappa^2/b} \theta, \frac{\kappa/b}{1 + \kappa^2/b} \theta \right).$$

The welfare delivered by this policy is given by

$$W^{MP} = \frac{1}{1 - \beta} E \left[-\frac{1}{2} \frac{1}{1 + \kappa^2/b} \theta^2 \right] = -\frac{1}{2(1 - \beta)(1 + \kappa^2/b)} E[\theta^2].$$

Appendix C: Computing the private information solution

This appendix details the algorithm we use to compute the private information solution. We numerically implement the Bellman operator \mathbb{T} with discrete types. We also discretize the choice sets for inflation, the output gap, and expected inflation, and introduce lotteries/randomization over these sets to convexify the problem. As a result the problem becomes a concave dynamic programming problem, and we apply the method proposed by Fukushima and Waki (2013).

We begin with convexifying the Bellman operator \mathbb{T} in equation (18). Let \hat{X} , $\hat{\Pi}$, and $\hat{\Theta}$ be grids over X , Π , and Θ , respectively. We assume that grids are such that $co(\hat{X}) = X$ and $co(\hat{\Pi}) = \Pi$. Let \hat{p} denote a discrete approximation of density p . In each state π_-^e , for each θ , the mechanism designer chooses a lottery γ^x over \hat{X} , and a lottery γ^π over $\hat{\Pi}$, in addition to $(\pi(\cdot), x(\cdot), \pi^e(\cdot), W(\cdot))$.

We define the Bellman operator \mathbb{T}^l as follows: for all $\pi_-^e \in \bar{\Pi}$,

$$\mathbb{T}^l F(\pi_-^e) = \max_{\theta} \sum_{\theta} \hat{p}(\theta) \left[\sum_{\pi_i \in \hat{\Pi}} \gamma^\pi(\pi_i|\theta) \{A(\pi_i) + \theta\pi_i\} + \sum_{x_i \in \hat{X}} \gamma^x(x_i|\theta) B(x_i) + \beta W(\theta) \right]$$

subject to the feasibility constraints,

$$\pi_-^e = \sum_{\theta} \hat{p}(\theta) \pi(\theta) \quad (25)$$

$$\pi(\theta) = \beta \pi^e(\theta) + \kappa x(\theta), \quad \forall \theta \in \hat{\Theta}, \quad (26)$$

the lottery constraints,

$$\gamma^\pi(\pi_i|\theta) \in [0, 1], \quad \forall \theta \in \hat{\Theta}, \pi_i \in \hat{\Pi}, \quad (27)$$

$$\gamma^x(x_i|\theta) \in [0, 1], \quad \forall \theta \in \hat{\Theta}, x_i \in \hat{X}, \quad (28)$$

$$\sum_i \gamma^\pi(\pi_i|\theta) = \sum_i \gamma^x(x_i|\theta) = 1, \quad \forall \theta \in \hat{\Theta}, \quad (29)$$

the consistency constraints,

$$\pi(\theta) = \sum_{\pi_i \in \hat{\Pi}} \gamma^\pi(\pi_i|\theta)\pi_i, \quad \forall \theta \in \hat{\Theta}, \quad (30)$$

$$x(\theta) = \sum_{x_i \in \hat{X}} \gamma^x(x_i|\theta)x_i, \quad \forall \theta \in \hat{\Theta}, \quad (31)$$

the incentive-compatibility constraint,

$$\begin{aligned} & \sum_{\pi_i \in \hat{\Pi}} \gamma^\pi(\pi_i|\theta)\{A(\pi_i) + \theta\pi_i\} + \sum_{x_i \in \hat{X}} \gamma^x(x_i|\theta)B(x_i) + \beta W(\theta) \\ \geq & \sum_{\pi_i \in \hat{\Pi}} \gamma^\pi(\pi_i|\theta')\{A(\pi_i) + \theta\pi_i\} + \sum_{x_i \in \hat{X}} \gamma^x(x_i|\theta')B(x_i) + \beta W(\theta'), \quad \forall (\theta, \theta') \in \hat{\Theta}^2, \end{aligned} \quad (32)$$

and the upper-bound constraint,

$$W(\theta) \leq F(\pi^e(\theta)), \quad \forall \theta \in \hat{\Theta}. \quad (33)$$

Equations (27) to (29) require that, for each θ , both $\gamma^\pi(\cdot|\theta)$ and $\gamma^x(\cdot|\theta)$ are lotteries over $\hat{\Pi}$ and \hat{X} , respectively. Equations (30) and (31) requires that π and x are achieved on average by γ^π and γ^x .

This operator \mathbb{T}^l satisfies Blackwell's sufficient condition, and thus is a contraction mapping. This dynamic programming problem is a concave problem, allowing us to apply the method in Fukushima and Waki (2013) to compute its solution. Because $\bar{\Pi}$ is unknown, we use a sufficiently large interval $\Pi^e \subset \Pi$ as the state space, and then check that the computed solution is interior. This leads us to the following algorithm based on Fukushima and Waki (2013):

1. Fix a compact interval $\Pi^e \subset \Pi$ and a finite grid $\hat{\Pi}^e$ on Π^e .
2. Set the initial condition $\bar{v}_0 = \min U$ for value function iteration. Let \mathbb{T}_L^l be the numerical Bellman operator that approximates \mathbb{T}^l from below (see Fukushima and Waki, 2013). Then $\bar{v}_n := (\mathbb{T}_L^l)^n \bar{v}_0$ is increasing in n and converges uniformly to the fixed point of \mathbb{T}_L^l .
3. Stop if a pre-specified stopping criterion is satisfied: $\|\bar{v}_n - \bar{v}_{n-1}\| < \epsilon$ for some constant ϵ .
4. Use the computed value function \bar{v}_n to check whether the solution is interior. If it is, use \bar{v}_n as an estimate for the true value function.

References

- [1] Abreu, D., Pearce, D., and E. Stacchetti, (1990), “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, 58, 5, pp. 1041–1063.
- [2] Alonso, R., and N. Matouschek, (2008), “Optimal Delegation,” *Review of Economic Studies*, 75, 1, pp. 259–293.
- [3] Amador, M., and K. Bagwell, (2013), “The Theory of Optimal Delegation with An Application to Tariff Caps,” *Econometrica*, 81, 4, pp. 1541–1599.
- [4] Amador, M., Werning, I., and G.-M. Angeletos, (2006), “Commitment vs. Flexibility,” *Econometrica*, 74, 2, pp. 365–396.
- [5] Athey, S., Atkeson, A., and P. Kehoe, (2005), “The Optimal Degree of Discretion in Monetary Policy,” *Econometrica*, 73, 5, pp. 1431–1475.
- [6] Athey, S., K. Bagwell, and C. Sanchirico, (2004), “Collusion and Price Rigidity,” *Review of Economic Studies*, 71, 2, pp. 317–349.
- [7] Atkeson, A. (1991), “International Lending with Moral Hazard and Risk of Repudiation,” *Econometrica*, 59, 4, pp. 1069–1089.
- [8] Barro, R. J., and D. B. Gordon, (1983), “Rules, discretion and reputation in a model of monetary policy,” *Journal of Monetary Economics*, 12, 1, pp. 101–121.
- [9] Bernanke, B., and F. Mishkin, (1997), “Inflation Targeting: A New Framework for Monetary Policy?” *Journal of Economic Perspectives*, 11, 2, pp. 97–116.
- [10] Bernanke, B., Laubach, T., Mishkin, F., and A. Posen, (2001), *Inflation Targeting: Lessons from the International Experience*, Princeton University Press, Princeton, New Jersey.
- [11] Canzoneri, M., (1985), “Monetary Policy Games and the Role of Private Information,” *American Economic Review*, 75, 5, pp. 1056–1070.
- [12] Calvo, G., (1983), “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, 12, pp. 383–398.

- [13] Clarida, R., Galí, J., and M. Gertler, (1999), “The Science of Monetary Policy: A New Keynesian Perspective,” *Journal of Economic Literature*, 37, pp. 1661–1707.
- [14] Dennis, R., and Söderström, (2006), “How Important Is Precommitment for Monetary Policy?” *Journal of Money, Credit, and Banking*, 38, 4, pp.847–872.
- [15] Fukushima, K., and Y. Waki, (2013), “A Polyhedral Approximation Approach to Concave Numerical Dynamic Programming,” *Journal of Economic Dynamics and Control*, 37, 11, pp. 2322–2335.
- [16] Galí, J., (2008), *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*, Princeton University Press, Princeton, New Jersey.
- [17] Gertler, M., and J. Leahy (2008), “A Phillips Curve with an Ss Foundation,” *Journal of Political Economy*, 116, 3, pp. 533–572.
- [18] Green, E. J. (1987), “Lending and the Smoothing of Uninsurable Income,” in *Contractual Arrangement in Intertemporal Trade*, eds. E. C. Prescott and N. Wallace, pp. 3–25. University of Minnesota Press, Minneapolis, MN.
- [19] Halac, M., and P. Yared, (2015), “Fiscal Rules and Discretion under Persistent Shocks,” *Econometrica*, 82, 5, pp. 1557–1614.
- [20] Holmström, B., (1977), “On incentives and control in organizations,” Ph.D. Dissertation, Stanford University.
- [21] Holmström, B., (1984), “On the Theory of Delegation, ” in *Bayesian Models in Economic Theory*, eds. M. Boyer and R. E. Kihlstrom.
- [22] Jensen, H., (2002), “Targeting Nominal Income Growth of Inflation?” *American Economic Review*, 92, 4, pp. 928–956.
- [23] Kocherlakota, N., (1996), “Implications of Efficient Risk Sharing Without Commitment,” *Review of Economic Studies*, 63, pp. 595–610.
- [24] Kydland, F., and E. Prescott, (1977), “Rules Rather than Discretion: The Inconsistency of Optimal Plans,” *Journal of Political Economy*, 85, 3, pp. 473–491.

- [25] Kurozumi, T., (2008), “Optimal Sustainable Monetary Policy,” *Journal of Monetary Economics*, 55, 7, pp. 1277-1289.
- [26] Melumad, N. D., T. Shibano, (1991), “Communication in settings with no transfers,” *RAND Journal of Economics*, 22, 2, pp. 173-198.
- [27] Rotemberg, J., (1982), “Sticky Prices in the United States,” *Journal of Political Economy*, 90, 6, pp.1187–1211.
- [28] Rotemberg, J., and M. Woodford, (1997), “An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy,” in Bernanke, B., and J. Rotemberg, eds, *NBER Macroeconomics Annual*, 12, MIT Press, Cambridge, Massachusetts.
- [29] Sleet, C., (2001), “On Credible Monetary Policy and Private Government Information,” *Journal of Economic Theory*, 99, pp. 338–376.
- [30] Sleet, C., (2004), “Optimal Taxation with Private Government Information,” *Review of Economic Studies*, 71, pp. 1217–1239.
- [31] Stokey, N. L., and R. E. Lucas, Jr., with E. C. Prescott, (1989), *Recursive Methods in Economic Dynamics*, Harvard University Press, Cambridge, Massachusetts, and London, England.
- [32] Svensson, L. E. O., (1997), “Optimal Inflation Targets, “Conservative” Central Banks, and Linear Inflation Contracts,” *American Economic Review*, 87, 1, pp. 98–114.
- [33] Thomas, J., and T. Worrall, (1990), “Income Fluctuation and Asymmetric Information: An Example of A Repeated Principal-Agent Problem,” *Journal of Economic Theory*, 51, 2, pp. 367-390.
- [34] Woodford, M., (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press, Princeton, New Jersey.