Theoretical Economics 13 (2018), 1319–1367 1555-7561/20181319

The optimal degree of monetary discretion in a new Keynesian model with private information

Yuichiro Waki
School of Economics, University of Queensland

Richard Dennis
Department of Economics, Adam Smith Business School, University of Glasgow

Ippei Fujiwara
Faculty of Economics, Keio University and Crawford School of Public Policy, ANU

This paper considers the optimal degree of monetary discretion 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 should 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, it is a transient phenomenon and some discretion is granted eventually.

Keywords. Rules versus discretion, monetary policy, new Keynesian model, private information, delegation, mechanism design, inflation targeting.

JEL classification. D82, E52, E58.

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

Yuichiro Waki: y.waki@uq.edu.au
Richard Dennis: richard.dennis@glasgow.ac.uk
Ippei Fujiwara: ippei.fujiwara@keio.jp

This study was conducted as a part of the project "Social Security, Taxation, and Public Finance" undertaken at the Research Institute of Economy, Trade and Industry (RIETI). We thank Ctirad Slavik, Ryo Jinnai, Giuseppe Moscarini, Begona Dominguez, Marco Bassetto, Andres Bellofatto, 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. We also thank two anonymous referees for their detailed feedback. Fujiwara and Waki are also grateful for financial support from JSPS KAKENHI Grant-in-Aid for Scientific Research (A) Grant 15H01939.

© 2018 The Authors. Licensed under the Creative Commons Attribution-NonCommercial License 4.0. Available at http://econtheory.org. https://doi.org/10.3982/TE2369

private information (Canzoneri 1985) or "nonruleable" information that is difficult to encode in a rule (Kocherlakota 2015), 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 noncontractible 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 et al. (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.

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 mechanism can be expressed as a function of the central bank's private information (its type) and of the previous period's promised inflation. In other words, unlike AAK, we find the optimal mechanism to be history-dependent.

Despite its history dependence, at each point in time, the optimal mechanism takes a rather simple form of "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 the previous period's promised inflation so that the central bank is incentivized to deliver inflation that is, on average, consistent with the previous 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 a subset of types.

How does the optimal degree of discretion vary with the previous 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 the previous period's promised inflation departs from this value, the smaller is the degree of discretion that 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 the previous period's promised inflation depends crucially on the value of promised inflation. At the welfare-maximizing value for 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 so as 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 central bank there is an interval for promised inflation within which that type has discretion. In such an interval, inflation, the output gap, and the continuation mechanism depend on history only through the current value of private information, and the history dependence is disposed of. This characteristic resembles the "amnesia" property that Kocherlakota (1996) finds in a full-information limited-commitment model of risk-sharing.

Using a quadratic social welfare function, we characterize how the optimal degree of discretion changes over time. We find that some discretion is always granted in the ergodic set of promised inflation. This implies that, regardless of the initial condition for promised inflation, no discretion is at most a short-run, transient phenomenon, and some discretion is eventually granted. Interestingly, no discretion is given only when the initial promised inflation 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 begins. Therefore, some discretion is always granted in the fully optimal mechanism and also in the optimal timeless perspective mechanism, which takes the stationary distribution of promised inflation as its initial distribution.

To implement the second-best with a nondirect mechanism, we propose a history-dependent inflation targeting regime that stipulates a band for permissible inflation that varies with the central bank's promised inflation. This regime allows the central bank to constrain its future self through its choice of promised inflation, mitigating the time-inconsistency problem. We show that an appropriately designed regime implements the outcome of the optimal direct mechanism. Importantly, unlike AAK, there are situations where a lower limit on inflation imposes a binding constraint on the central bank's choice.

It is crucial to constrain the central bank's policy actions in a history-dependent manner. If the band for permissible inflation is required to be history-independent, then it is optimal to set the band as wide as possible so that it never binds. Hence, both strict inflation targeting and fixed-range targeting are inferior to full flexibility. Although it is a common practice to focus on a Markov perfect equilibrium when the policy-maker is unable to commit, limiting attention to Markov perfect mechanisms in our setting

results in the misleading conclusion that legislation that constrains the policy-maker's actions is not useful for improving welfare.

While highlighting the inefficiency of inflation targeting regimes that do not have state-contingent inflation bands, our results also have implications for the efficacy of forward guidance. Our optimal mechanism is a dynamic one in which the upper and lower edges of the permissible inflation range are state-contingent, depending upon promised inflation. Central bank types for which the inflation range does not bind dispense with history and conduct policy with discretion. For these central bank types, providing forward guidance is not central to the policy process, because the need to deliver on the promised inflation is not constraining. As a consequence, our model implies a connection between forward guidance and the degree of discretion. Forward guidance has greater value and takes on greater importance when the degree of discretion is low, which occurs when the previous period's promised inflation is further from its ergodic mean. An implication of our model, therefore, is that central banks should provide forward guidance in situations where the economy is far from its ergodic mean, which might occur after the economy has been hit by large shocks.

Finally, we use a numerical example to illustrate visually how the optimal mechanism in our private information framework differs from two benchmarks that are analyzed in the new Keynesian policy literature. One benchmark is the full-information optimal mechanism, which is identical to the optimal policy with commitment in the literature; the other benchmark is the optimal discretionary policy. Compared to the optimal discretionary policy, the optimal mechanism uses promised inflation to stabilize the fluctuations in the output gap and to allow inflation to vary with the private information. Still, compared to the full-information solution, promised inflation in the optimal mechanism changes less with the private information, because letting promised inflation move away from the peak of the social welfare function requires a tighter band for permissible inflation in the next period, which is costly. As a result, inflation in the optimal mechanism moves less with the private information than in the full-information solution.

The remainder of this paper is organized as follows. Section 2 reviews related literature. Section 3 describes the setup and illustrates how private information enters the model. Section 4 formulates an optimal (direct) mechanism design problem. In Section 5, we discuss theoretical results. Section 6 presents the numerical results that emerge from the benchmark policies and from the optimal mechanism. Section 7 offers concluding comments. Appendices A–D contain technical material, including proofs of theoretical results and complete descriptions of how the various solutions were computed. Additional Appendices E and F are available in a supplementary file on the journal website, http://econtheory.org/supp/2369/supplement.pdf.

2. Related literature

We build on the literature on 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] McCallum (2003), in discussing AAK's paper, stresses the importance of using a forward-looking Phillips curve instead of a static one.[2] 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). This literature has generally focused on settings in which society cannot directly constrain the central bank's action set and in which granting the central bank some discretion is suboptimal.[3] 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, e.g., Bernanke et al. (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.[4] Using an approach akin to that of Athey et al. (2004) and extending the results of AAK, we show that our problem can be formulated recursively as a function-valued dynamic programming problem in which the previous period's promised inflation 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 et al. 1990), which is common in the dynamic contracting literature.[5]

[1] Canzoneri (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.

[2] He writes, "there is a significant reason to depart from the AAK specification of the example—and also, I believe, their formal model. That is that it does not include a feature that has been at the center of much recent work in monetary economics, namely, forward-looking behavior by the economy's individual agents."

[3] In 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 the central bank and that the central bank is required to minimize it, e.g., Svensson (1997) and Jensen (2002). Neither approach allows society to remove certain actions from the central bank's choice set.

[4] For static delegation problems, Holmstrom (1984), Alonso and Matouschek (2008), and Amador and Bagwell (2013) give some sufficient conditions for interval delegation to be optimal. Sleet (2004) and Amador et al. (2006) consider dynamic delegation problems, and the latter study finds the optimality of interval delegation. AAK consider a repeated delegation problem, but as the optimal mechanism is shown to be static, their problem reduces to a static delegation problem.

[5] A similar result is obtained in Sleet (2004) in a two-type hidden information model of optimal taxation. We consider a monetary-policy model with a continuum of types. Atkeson (1991) shows a similar result in a hidden action model of optimal international lending, but his model allows for monetary transfer between the lender and the borrower. 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.

3. The setup

Our framework is similar to the canonical framework that is used in the new Keynesian policy literature to analyze optimal policy without commitment (see, e.g., Woodford 2003 and Galì 2008). 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.

We depart from the canonical new Keynesian framework in two aspects. First, in our framework the central bank possesses private information. Specifically, at the beginning of each period, the central bank observes privately the state of the economy, θ, which is drawn from a compact interval $\Theta := [\underline{\theta}, \overline{\theta}] \subset \mathbb{R}$, according to an independent and identically distributed (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 observe the state, θ.

Second, society can design and commit to a mechanism that limits the central bank's action each period by specifying the set of acceptable actions from which the central bank must choose. After the central bank observes θ, public communication takes place. 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. Society imposes an interest rate mechanism, which specifies a particular value for the nominal interest rate, i, as a function of the history of the central bank's messages. Although the central bank can send any messages it wants, it is not allowed to set the nominal rate at a level that is inconsistent with its message history. The central bank has some discretion when the interest rate specified by the mechanism varies with its current message, whereas it has no discretion when the mechanism specifies a single value irrespective of its current message.

Inflation, π, and the output gap, x, are determined as a rational expectation equilibrium (REE) outcome. The interest rate mechanism, the central bank's reporting strategy, and the stochastic process for private information determine a stochastic process for messages and nominal interest rates, $\{(m_t, i_t)\}_{t=0}^{\infty}$. Private agents recognize correctly this joint process and optimize, taking other agents' actions and aggregate variables as given, and a REE realizes. In Appendix E, we detail the model's microfoundation and show that a REE given $\{(m_t, i_t)\}_{t=0}^{\infty}$ is a sequence of measurable functions of message history for inflation and the output gap, $\{(\pi_t, x_t)\}_{t=0}^{\infty}$, that satisfy the new Keynesian Phillips curve and the dynamic investment-saving (IS) equation in any period and for any message history that can occur under the central bank's reporting strategy. The new Keynesian Phillips curve (NKPC) is given by

$$\pi_t = \kappa x_t + \beta \mathbb{E}_t^P[\pi_{t+1}], \tag{1}$$

where κ is a strictly positive parameter and $\beta \in (0, 1)$ is the household's preference discount factor. This is a standard log-linear NKPC without a cost-push shock, and is forward-looking in that it involves expected future inflation.[6] The dynamic IS equation

[6]The 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).

is given by

$$x_t = \mathbb{E}_t^P[x_{t+1}] - \zeta\big(i_t - \mathbb{E}_t^P[\pi_{t+1}] - r^n\big),$$

where $\zeta > 0$ is the elasticity of intertemporal substitution and $r^n > 0$ is the constant natural rate of interest. The conditional expectation operator, $\mathbb{E}_t^P$, in the two equations is based on the information available to the private sector in period t after the communication stage and on the central bank's reporting strategy. We focus on bounded equilibria and, therefore, there are compact intervals $\Pi := [\underline{\pi}, \overline{\pi}]$ and $X = [\underline{x}, \overline{x}]$, such that $\pi_t \in \Pi$ and $x_t \in X$ always hold. We assume that X contains both $(\underline{\pi} - \beta\overline{\pi})/\kappa$ and $(\overline{\pi} - \beta\underline{\pi})/\kappa$.

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. The central bank is benevolent, so $R(\pi, x, \theta)$ also equals its momentary payoff.[7] The 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, a dependence that can arise, for example, if the redistributional effects of inflation are time-varying.[8] Although we do not explicitly model such effects and our welfare specification is reduced form, it allows us to model the way the state of the economy affects social welfare in a flexible manner. An example of R is the quadratic specification

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

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

We assume that society is able to commit while the central bank is not. Due to its inability to commit, the central bank is unable to manage the private sector's expectations by committing to a particular sequence of interest rates in the future. Society can improve welfare because the mechanism can limit the set of interest rates from which the central bank can choose in the future, thereby influencing the private sector's inflation expectations. However, an overly restrictive mechanism 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 the interest rate mechanism so as to maximize social welfare?"

3.1 *Discussion*

Although our framework shares much in common with AAK, a distinct feature of our setup is the forward-looking NKPC. In contrast, AAK's benchmark example assumes a static Phillips curve,

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

[7] Appendix F.2 extends our setup by incorporating an inflation bias.

[8] AAK 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."

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 (1) as more relevant, because it is a centerpiece of many new Keynesian models that are widely used in central banks. Moreover, it captures an important channel for policy, allowing central banks to use forward guidance to manage inflation expectations.

This difference in the Phillips curve results in a difference in the nature of a time-inconsistency problem. In AAK and Sleet (2001), the central bank is subject to an "inflation bias" (Barro and Gordon 1983): once the private sector forms inflation expectations, the central bank is tempted to generate higher than expected inflation to stimulate real economic activity. In contrast, the forward-looking NKPC gives rise to a "stabilization bias" (see, e.g., Svensson 1997). When higher inflation is desirable today, the central bank may promise to make future policy more accommodative to raise inflation expectations, thereby incentivizing price setters to raise prices today. In the future, however, because the gains from the promise are already realized, the central bank faces the temptation to stabilize inflation and the output gap by not delivering on its promise.

4. Optimal mechanism design problem

Once society has specified an interest rate mechanism, the problem becomes a dynamic communication game with incomplete information, with the central bank being the only strategic player. Society's problem is to choose an interest rate mechanism so that the best equilibrium given the mechanism yields the highest social welfare.[9] We focus on a public perfect Bayesian equilibrium, where the central bank's reporting strategy is a public one that depends only on its message history and the current private information, i.e., it does not depend on past private information.

Appendix E shows that the revelation principle applies and, therefore, that we can focus on direct mechanisms that satisfy two constraints that we describe below. A direct mechanism is a sequence of measurable functions of reports, $\{(\pi_t, x_t, i_t)\}_{t=0}^{\infty}$, such that $(\pi_t, x_t, i_t) : \Theta^{t+1} \to \Pi \times X \times \mathbb{R}$ for any t.[10]

The first constraint that society must respect is, clearly, that it must choose an incentive-compatible direct mechanism. A reporting strategy is a sequence of measurable functions $\sigma := \{\sigma_t\}_{t=0}^{\infty}$ with $\sigma_t : \Theta^{t+1} \to \Theta$ 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 direct mechanism is said to be *incentive-compatible* if and only if, for any report history θ^{t-1}, for any current type θ_t,

[9] In general, there are multiple rational expectation equilibria for a given stochastic process of the nominal interest rates. Here we are assuming either that society, as the mechanism designer, can choose the best equilibrium or that all agents in the economy can coordinate on it.

[10] We abstract from the zero lower bound on the nominal interest rate.

and for any reporting strategy σ,

$$\begin{aligned}
&R\big(\pi_t\big(\theta^{t-1},\theta_t\big),x_t\big(\theta^{t-1},\theta_t\big),\theta_t\big)\\
&\quad+\beta\sum_{s=t+1}^{\infty}\int_{\Theta^{s-t}}\beta^{s-t-1}R\big(\pi_s\big(\theta^{t-1},\theta_t^s\big),x_s\big(\theta^{t-1},\theta_t^s\big),\theta_s\big)\mu^{s-t}\big(d\theta_{t+1}^s\big)\\
&\geq R\big(\pi_t\big(\theta^{t-1},\sigma_0(\theta_t)\big),x_t\big(\theta^{t-1},\sigma_0(\theta_t)\big),\theta_t\big)\\
&\quad+\beta\sum_{s=t+1}^{\infty}\int_{\Theta^{s-t}}\beta^{s-t-1}R\big(\pi_s\big(\theta^{t-1},\sigma^{s-t}\big(\theta_t^s\big)\big),x_s\big(\theta^{t-1},\sigma^{s-t}\big(\theta_t^s\big)\big),\theta_s\big)\mu^{s-t}\big(d\theta_{t+1}^s\big),
\end{aligned}\tag{4}$$

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.[11] The set of these inequalities is referred to as the *incentive-compatibility constraint*.

The second constraint that a direct mechanism must satisfy is the rational expectation equilibrium condition. Society must choose a direct mechanism such that for any t and any report history $\theta^t:=(\theta_0,\theta_1,\dots,\theta_t)\in\Theta^{t+1}$, the NKPC

$$\pi_t\big(\theta^t\big)=\kappa x_t\big(\theta^t\big)+\beta\int_{\theta_{t+1}}\pi_{t+1}\big(\theta^t,\theta_{t+1}\big)p(\theta_{t+1})\,d\theta_{t+1}\tag{5}$$

and the dynamic IS equation

$$\begin{aligned}
x_t\big(\theta^t\big)=&\int_{\theta_{t+1}}x_{t+1}\big(\theta^t,\theta_{t+1}\big)p(\theta_{t+1})\,d\theta_{t+1}\\
&-\zeta\left\{i_t\big(\theta^t\big)-\int_{\theta_{t+1}}\pi_{t+1}\big(\theta^t,\theta_{t+1}\big)p(\theta_{t+1})\,d\theta_{t+1}-r^n\right\}
\end{aligned}\tag{6}$$

hold.

As in the canonical new Keynesian model, the dynamic IS equation (6) is redundant because for any $\{(\pi_t,x_t)\}_{t=0}^{\infty}$ that satisfies the new Keynesian Phillips curve (5), we can find $\{i_t\}_{t=0}^{\infty}$ so that the dynamic IS equation is satisfied. Hence, we drop the the dynamic IS equation from the constraint and drop the nominal interest rate from the choice.

The mechanism design problem we consider is one that maximizes social welfare, which is given by

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

by choosing $\{(\pi_t,x_t)\}_{t=0}^{\infty}$ subject to the incentive-compatibility constraint (4) and the NKPC (5). For simplicity, we refer to $\{(\pi_t,x_t)\}_{t=0}^{\infty}$ as a mechanism and refer to the NKPC as the *feasibility constraint*.

[11] This history is recursively defined as $\sigma^0(\theta_t^t):=\sigma_0(\theta_t)$ 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 $s\geq t+1$.

4.1 *Recursive formulation*

Both the incentive-compatibility and the feasibility constraints are forward-looking constraints. We rewrite these constraints to obtain recursive formulations.

4.1.1 *Incentive compatibility* First, the incentive-compatibility constraint (4) can be written recursively by adding the agent's continuation, or promised, utility as a choice variable. Let $\underline{U} := \int_\Theta \{\min_{x,\pi} R(\pi, x, \theta)\} p(\theta)\, d\theta/(1-\beta)$, $\overline{U} := \int_\Theta \{\max_{x,\pi} R(\pi, x, \theta)\} p(\theta)\, d\theta/(1-\beta)$ and $U := [\underline{U}, \overline{U}]$, so that 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), it can be shown that a mechanism $\{(\pi_t, x_t)\}_{t=0}^\infty$ is incentive compatible if and only if the following statements hold:

(i) There exists a sequence of measurable functions $\{w_t\}_{t=-1}^\infty$ with $w_t : \Theta^t \to U$ for all $t \geq -1$, such that for all $t \geq 0$ and θ^t,

$$w_{t-1}(\theta^{t-1}) = \int_\Theta \left[R\left(\pi_t(\theta^t), x_t(\theta^t), \theta_t\right) + \beta w_t(\theta^t)\right] p(\theta_t)\, d\theta_t. \tag{8}$$

(ii) For all t, θ^{t-1}, θ_t, and $\theta' \neq \theta_t$,

$$\begin{aligned} &R\left(\pi_t(\theta^t), x_t(\theta^t), \theta_t\right) + \beta w_t(\theta^t) \\ &\quad \geq R\left(\pi_t(\theta^{t-1}, \theta'), x_t(\theta^{t-1}, \theta'), \theta_t\right) + \beta w_t(\theta^{t-1}, \theta'). \end{aligned} \tag{9}$$

Here $\{w_t\}_{t=0}^\infty$ has an interpretation as promised utility.

4.1.2 *Feasibility* Second, we show that the feasibility constraint (5) can also be written recursively by adding a new variable to the mechanism. Note that a direct mechanism $\{(\pi_t, x_t)\}_{t=0}^\infty$ satisfies the feasibility constraint if and only if there exists a sequence of measurable functions of reports $\{\pi_t^e\}_{t=0}^\infty$ such that

$$\pi_t(\theta^t) = \kappa x_t(\theta^t) + \beta \pi_t^e(\theta^t) \tag{10}$$

and

$$\pi_t^e(\theta^t) = \int_\Theta \pi_{t+1}(\theta^t, \theta_{t+1}) p(\theta_{t+1})\, d\theta_{t+1}. \tag{11}$$

By treating (x_t, π_t, π_t^e) as choice variables in period t, (10) amounts to a static constraint in period t. In period $t+1$, the previously chosen π_t^e imposes a constraint, reflected in (11), 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 promised inflation in the previous period. We therefore refer to π^e as *promised inflation*.

4.1.3 *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 (7) subject to constraints (8), (9), (10), and (11). However, because the period-0 inflation choice is not subject to a constraint like (11), there is an

asymmetry between period 0 and all other periods, and the problem is not fully recursive.

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

$$\pi^e_{-1} = \int_\Theta \pi_0(\theta)\, p(\theta), \tag{12}$$

where π^e_{-1} is a given number in Π and represents the promised inflation 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 it is 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^e_{-1}. For any $\pi^e_{-1} \in \Pi$, we say that $\{(x_t, \pi_t, \pi^e_t, w_t)\}_{t=0}^\infty$ is *feasible from* π^e_{-1} if and only if it satisfies (10), (11), and (12), and that it is *incentive-feasible from* π^e_{-1} if and only if it is feasible from π^e_{-1} and satisfies (8) and (9).[12]

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. As a result, the optimal mechanism is obtained in the form of time-invariant Markovian policy functions. Then we characterize its properties. We find that depending on the previous 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 disposed of for types that have discretion. Finally, we propose an inflation targeting rule that achieves the same outcome as the optimal mechanism. It is crucial to design a history-dependent rule, and history-independent rules are shown to achieve lower welfare than imposing no rule.

In what follows, we make the following assumptions on the return function, R, and the density function, p.

Assumption 1. *The return function R is separable in (π, θ) and x:*

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

Furthermore, A and B are twice continuously differentiable and satisfy $A_{\pi\pi} < 0$, $A_{\pi\theta} > 0$, and $B''(x) < 0$. Moreover, for any θ, $A_{\pi\theta}(\cdot; \theta)$ is nonincreasing.[13]

[12] For any $\pi^e_{-1} \in \Pi$, it is straightforward to prove that the constraint set is nonempty. For a given $\pi^e_{-1} \in \Pi$, consider a mechanism such that $\pi_t(\theta^t) = \pi^e_{-1}$ and $x_t(\theta^t) = (1-\beta)\pi^e_{-1}/\kappa$ 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 $(\underline{\pi} - \beta\overline{\pi})/\kappa$ and $(\overline{\pi} - \beta\underline{\pi})/\kappa$, this mechanism satisfies the NKPC after any history. Because this mechanism is independent of history, it is incentive-compatible.

[13] For univariate functions we use prime ($'$) and double prime ($''$) to denote its derivative and second derivative, respectively. For a bivariate function $f(x, y)$, we denote its partial derivatives by f_i ($i \in \{x, y\}$) and second derivatives by f_{ij} ($i, j \in \{x, y\}$).

ASSUMPTION 2. *For any nondecreasing function π of θ, $A_{\pi\theta}(\pi(\theta);\theta)(1-P(\theta))/p(\theta)$ is strictly decreasing in θ and $A_{\pi\theta}(\pi(\theta);\theta)P(\theta)/p(\theta)$ is strictly increasing in θ.*

ASSUMPTION 3. *The function A satisfies the following properties:*

(i) For any π, $A(\pi;\cdot)$ is absolutely continuous.

(ii) There is an integrable function $a:\Theta\to\mathbb{R}_+$ such that $|A_\theta(\pi;\theta)|<a(\theta)$ for all π and almost all θ.

The condition $A_{\pi\theta}>0$ in Assumption 1 is akin to the single-crossing condition in the mechanism design literature. Assumption 2 is akin to the monotone-hazard condition assumed in AAK and reduces to the standard condition when the cross derivative of A is constant (e.g., A is a sum of $\pi\times\theta$ and a function of π). The quadratic specification in (2) satisfies Assumption 1 and, together with, e.g., a uniform distribution, Assumption 2. We use the quadratic specification and a uniform distribution for θ later in our numerical experiments. Assumption 3 allows us to obtain an integral representation of utility (Milgrom and Segal 2002).

Proofs of lemmas and propositions are presented in Appendix A.

5.1 *Dynamic programming*

Let Ω be the set of pairs (π^e_{-1},w_{-1}), such that there exists a sequence of measurable functions $\{(x_t,\pi_t,\pi^e_t,w_t)\}_{t=0}^\infty$ that is incentive-feasible from π^e_{-1} and that attains social welfare of w_{-1}. For any $\pi^e_-\in\Pi$, the maximized social welfare given π^e_- is given by

$$\overline{W}(\pi^e_-)=\sup_{w_-\text{ s.t. }(\pi^e_-,w_-)\in\Omega} w_-, \tag{13}$$

which implies that if we obtain Ω, we also obtain the maximized social welfare. However, Ω is 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 et al. (1990). In our setup, there are a continuum of types and continuous action spaces, and the measurability restriction is difficult to impose on the Abreu–Pearce–Stacchetti-type set operator.

Instead, we directly characterize $\overline{W}$ as follows. We first define a Bellman operator, $\mathbb{T}$, that is a monotone contraction mapping on a Banach space of functions with π^e_- being the state variable. Its fixed point, $\tilde{W}$, is shown to satisfy $\overline{W}\le\tilde{W}$. We then show that the optimal policy function associated with $\mathbb{T}$ and $\tilde{W}$ is continuous, and that, for any π^e_-, we can generate from the policy function an incentive-feasible mechanism that attains $\tilde{W}(\pi^e_-)$. Therefore, $\tilde{W}\le\overline{W}$. Taken together, $\overline{W}=\tilde{W}$ holds.

We consider a Bellman operator on the Banach space of bounded functions on Π, $\mathcal{B}(\Pi)$, endowed with the sup norm $\|\cdot\|$. Define a Bellman operator $\mathbb{T}$, for all $F\in\mathcal{B}(\Pi)$ and for all $\pi^e_-\in\Pi$, as

$$\mathbb{T}F(\pi^e_-)=\sup_{\pi(\cdot),x(\cdot),\pi^e(\cdot),w(\cdot)}\int_{\underline{\theta}}^{\overline{\theta}}\big\{R\big(\pi(\theta),x(\theta),\theta\big)+\beta w(\theta)\big\}p(\theta)\,d\theta \tag{14}$$

subject to

$$\pi_-^e = \int_{\underline{\theta}}^{\overline{\theta}} \pi(\theta)\, p(\theta)\, d\theta \tag{15}$$

$$\pi(\theta) = \kappa x(\theta) + \beta \pi^e(\theta) \quad \forall \theta \tag{16}$$

$$R\big(\pi(\theta), x(\theta), \theta\big) + \beta w(\theta) \geq R\big(\pi\big(\theta'\big), x\big(\theta'\big), \theta\big) + \beta w\big(\theta'\big) \quad \forall \theta, \theta' \neq \theta \tag{17}$$

$$w(\theta) \leq F\big(\pi^e(\theta)\big) \quad \forall \theta, \tag{18}$$

and (π, x, π^e, w) are measurable functions. For any $F \in \mathcal{B}(\Pi)$ and $\pi_-^e \in \Pi$, we refer to the maximization problem in (14) as the $\mathbb{T}F(\pi_-^e)$ problem. It is straightforward that the operator $\mathbb{T}$ satisfies Blackwell's sufficient condition and, therefore, the following lemma obtains (see Stokey et al. 1989).

Lemma 1. *The Bellman operator $\mathbb{T}$ is a monotone β-contraction mapping on $(\mathcal{B}(\Pi), \|\cdot\|)$.*

The fixed point of $\mathbb{T}$, which we denote by $\tilde{W}$, is shown to be weakly larger than $\overline{W}$.

Lemma 2. *For any positive integer n, $\overline{W} \leq (\mathbb{T})^n \overline{W} \leq \lim_{n\uparrow\infty} (\mathbb{T})^n \overline{W} = \tilde{W}$.*

An intuition is that the recursive problem in (14) relaxes the problem in (13), because the constraint (18) with $F = \overline{W}$ is weaker than the constraint that requires promised inflation and promised utility to satisfy $(\pi^e(\theta), w(\theta)) \in \Omega$. The relaxation of a constraint, which is akin to that used in Athey et al. (2004), weakly increases the value, and $\tilde{W} \geq \overline{W}$ obtains.

Now we prove that for each π_-^e, there is an incentive-feasible mechanism that attains $\tilde{W}(\pi_-^e)$, and therefore, that $\overline{W} \geq \tilde{W}$. We begin by showing that the Bellman operator $\mathbb{T}$ preserves concavity, that the optimal policy function is continuous if F is concave, and that if F is concave, then (18) is always satisfied with equality. We need another assumption for this purpose.

Assumption 4. *In the problem (P1), which relaxes some constraints in the problem (14) and is defined in Appendix A.2, (i) a maximum is attained and (ii) the bound constraints do not bind, i.e., a solution satisfies $(\pi(\theta), x(\theta), \pi^e(\theta)) \in \operatorname{int}(\Pi) \times \operatorname{int}(X) \times \operatorname{int}(\Pi)$ for all θ.*

Lemma 3. *Let $\mathcal{V}(\Pi)$ be the set of weakly concave functions on Π. Suppose that $F \in \mathcal{V}(\Pi)$. Under Assumptions 1, 2, 3, and 4, (i) $\mathbb{T}F$ is a strictly concave C^1 function, (ii) the optimal policy correspondence is a quadruple of continuous functions $(\pi, x, \pi^e, w) : \Theta \times \Pi \to \Pi \times X \times \Pi \times U$, and (iii) $w(\theta, \pi_-^e) = F(\pi^e(\theta, \pi_-^e))$ for all (θ, π_-^e).*

Here we only briefly illustrate our approach; proofs can be found in Appendix A. Under the stated assumptions, the relaxed problem (P1) has the same structure as the mechanism design problem in AAK except that, for given π_-^e, $\pi(\cdot, \pi_-^e)$ is not required to be a piecewise C^1 function of θ in (P1). We generalize AAK's theoretical results in Appendix B so that we can apply them to the problem (P1). The problem (P1) is then shown to be equivalent to the original $\mathbb{T}F(\pi_-^e)$ problem, and we proceed to characterize its solution.

We do not use the approach in Amador and Bagwell (2013) because it requires a stronger assumption on the function $\overline{W}$, which is an endogenous object in our setting (see Appendix F.1 for details).

Because $\mathcal{V}(\Pi)$ is closed and the set of strictly concave C^1 functions is a subset of $\mathcal{V}(\Pi)$, the following corollary obtains.

COROLLARY 1. *Under Assumptions 1, 2, 3, and 4, the fixed point $\tilde{W}$ is a strictly concave C^1 function and Lemma 3(ii) and (iii) hold for $F = \tilde{W}$.*

Finally, we use the optimal policy function to construct, for each π_{-}^{e}, an incentive-feasible mechanism that attains $\tilde{W}(\pi_{-}^{e})$. Let $(\pi_*, x_*, \pi_*^e, w_*)$ be the optimal policy function associated with $\tilde{W}$. For each $\pi_{-1}^{e} \in \Pi$, construct a mechanism recursively, for $t = 0$, as

$$\big(\pi_0(\theta_0), x_0(\theta_0), \pi_0^e(\theta_0), w_0(\theta_0)\big) = \big(\pi_*\big(\theta_0, \pi_{-1}^e\big), x_*\big(\theta_0, \pi_{-1}^e\big), \pi_*^e\big(\theta_0, \pi_{-1}^e\big), w_*\big(\theta_0, \pi_{-1}^e\big)\big),$$

$$w_{-1} = \int_\Theta \big[R\big(\pi_0(\theta_0), x_0(\theta_0), \theta_0\big) + \beta w_0(\theta_0)\big] p(\theta_0)\, d\theta_0,$$

and for any $t \geq 1$ and θ^t, as

$$\begin{aligned}&\big(\pi_t\big(\theta^t\big), x_t\big(\theta^t\big), \pi_t^e\big(\theta^t\big), w_t\big(\theta^t\big)\big)\\ &\quad = \big(\pi_*\big(\theta_t, \pi_{t-1}^e\big(\theta^{t-1}\big)\big), x_*\big(\theta_t, \pi_{t-1}^e\big(\theta^{t-1}\big)\big), \pi_*^e\big(\theta_t, \pi_{t-1}^e\big(\theta^{t-1}\big)\big), w_*\big(\theta_t, \pi_{t-1}^e\big(\theta^{t-1}\big)\big)\big).\end{aligned}$$

This mechanism is clearly incentive-feasible from π_{-1}^e, and attains $\tilde{W}(\pi_{-1}^e)$ because $w_{-1} = \tilde{W}(\pi_{-1}^e)$. Moreover, because $(\pi_*, x_*, \pi_*^e, w_*)$ are continuous, the mechanism constructed above is θ^t measurable. Therefore, $\overline{W} \geq \tilde{W}$. Thus, we have established the following proposition.

PROPOSITION 1. *Under Assumptions 1, 2, 3, and 4, the fixed point of $\mathbb{T}$ is $\overline{W}$.*

The mechanism that we have constructed using the optimal policy function is optimal because it attains $\overline{W}(\pi_{-}^{e})$ for any π_{-}^{e}.

Under the same set of assumptions, the optimal policy function for inflation, π_*, has a simple, cutoff representation. We turn to this point next.

5.2 *Optimal degree of discretion*

We characterize the optimal degree of discretion using the optimal policy function for inflation, π_*. As a benchmark, we consider the inflation rate that the central bank would choose if it were given a certain form of policy flexibility, and compare it to the inflation rate prescribed by the optimal mechanism. This approach is analogous to AAK: they define the central bank's "static best response," i.e., 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 does not 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 $\pi^e_- = \int \pi(\theta)p(\theta)\,d\theta$, but faces the optimal mechanism in all subsequent periods. The one-shot discretionary best response is the inflation rate that would be chosen by the central bank in this hypothetical situation.

DEFINITION 1. The one-shot discretionary best response is a function $\pi_D : \Theta \to \Pi$ such that, for all θ, $\pi_D(\theta)$ solves

$$\max_{\pi}\Big\{A(\pi;\theta) + \max_{(x,\pi^e):\pi=\kappa x+\beta\pi^e}\big(B(x)+\beta\overline{W}(\pi^e)\big)\Big\}. \tag{19}$$

The one-shot discretionary best response is well behaved.

LEMMA 4. *The function $\pi_D(\cdot)$ is a strictly increasing, continuous function.*

It is natural to interpret that, for a given π^e_-, a type θ has discretion when $\pi_*(\theta;\pi^e_-) = \pi_D(\theta)$. We say that the central bank has *full discretion* at π^e_- if $\pi_*(\cdot,\pi^e_-)$ coincides with $\pi_D(\cdot)$. 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 2. *Let $\pi^{e*}_- := \int \pi_D(\theta)p(\theta)\,d\theta$. The point $\overline{W}$ is maximized uniquely at $\pi^e_- = \pi^{e*}_-$, and the policy function satisfies $\pi_*(\theta;\pi^{e*}_-) = \pi_D(\theta)$ for all θ.*

From Proposition 2, we can think of π^{e*}_- as the most desirable initial condition: if $\pi^e_{-1} = \pi^{e*}_-$, then social welfare from time zero onward is maximized. Furthermore, if the central bank were allowed to renege on previously promised inflation, then it would behave as if promised inflation were π^{e*}_-. Therefore, for each π^e_-, we can interpret the gain from reneging, $\overline{W}(\pi^{e*}_-) - \overline{W}(\pi^e_-)$, as the severity of the time-inconsistency problem at π^e_-. Because $\overline{W}$ 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 alternatives: *no discretion* and *bounded discretion*.

DEFINITION 2. 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 $\pi_*(\cdot;\pi^e_-) = \max\{c, \pi_D(\cdot)\}$ or $\pi_*(\cdot;\pi^e_-) = \min\{c, \pi_D(\cdot)\}$ for some constant c.

Note that bounded discretion is equivalent to a cutoff 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 in the next proposition.

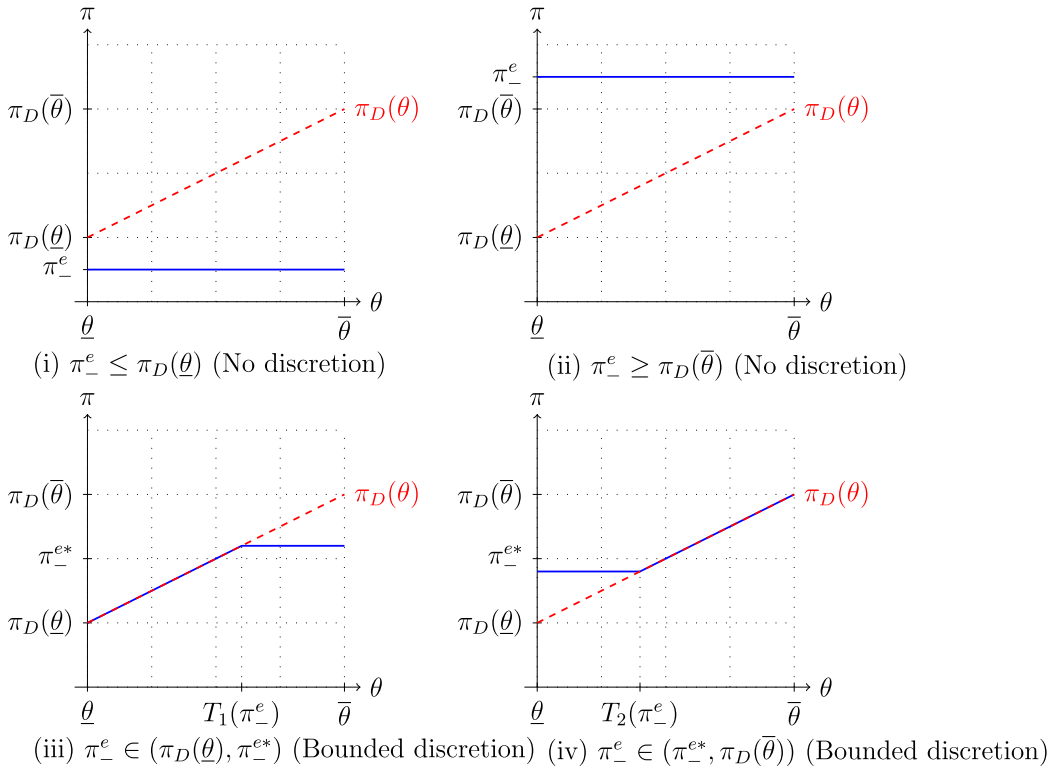

FIGURE 1. Policy function $\pi_*(\theta, \pi_-^e)$ as a function of θ for different values of π_-^e.

PROPOSITION 3. *There exist two strictly increasing, continuous threshold functions* $T_1 : (\pi_D(\underline{\theta}), \pi_-^{e*}) \to \Theta$ *and* $T_2 : (\pi_-^{e*}, \pi_D(\overline{\theta})) \to \Theta$ *such that the policy function for inflation,* π_**, has one of the following features.*

(i) Full discretion: $\pi_*(\theta; \pi_-^e) = \pi_D(\theta)$ *for all* θ *if* $\pi_-^e = \pi_-^{e*}$.

(ii) No discretion: $\pi_*(\theta; \pi_-^e) = \pi_-^e$ *for all* θ *if* $\pi_-^e \le \pi_D(\underline{\theta})$ *or* $\pi_-^e \ge \pi_D(\overline{\theta})$.

(iii) Bounded discretion: if $\pi_-^e \in (\pi_D(\underline{\theta}), \pi_-^{e*})$,

$$\pi_*(\theta; \pi_-^e) = \min\{\pi_D(T_1(\pi_-^e)), \pi_D(\theta)\},$$

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

$$\pi_*(\theta; \pi_-^e) = \max\{\pi_D(T_2(\pi_-^e)), \pi_D(\theta)\}.$$

The threshold functions are defined so that $\pi_-^e = \int_{\underline{\theta}}^{\overline{\theta}} \pi_*(\theta; \pi_-^e)\, p(\theta)\, d\theta$ *for all* π_-^e.[14]

Proposition 3 reveals three important properties. First, there is no discretion when π_-^e is sufficiently far from π_-^{e*} (either above or below). This is depicted in panels (i) and (ii) in Figure 1. Dashed lines represent π_D and solid lines represent $\pi_*(\cdot; \pi_-^e)$ for a

[14]The exact specifications of T_1 and T_2 are given in Appendix A.3.3.

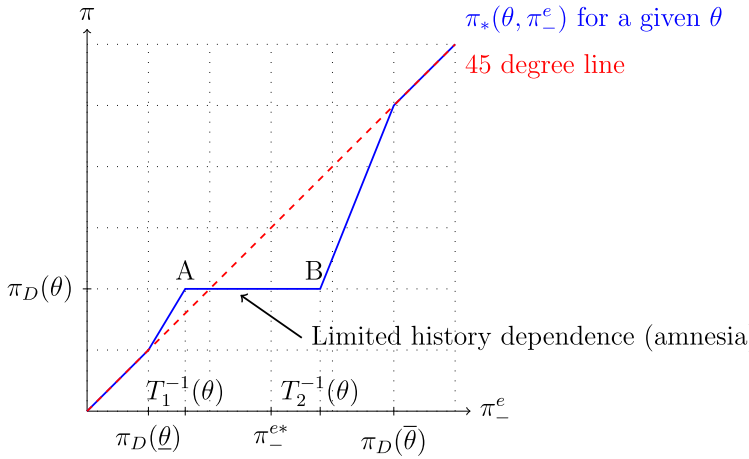

FIGURE 2. Policy function $\pi_*(\theta; \pi_-^e)$ as a function of π_-^e, given θ.

given π_-^e. Second, when π_-^e is not too far from π_-^{e*}, the policy function exhibits bounded discretion, or a cutoff property. When π_-^e is lower 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.

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 resides 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 all types that have discretion, the optimal mechanism therefore features *amnesia* in the sense of Kocherlakota (1996). Intuitively, society imposes an interval restriction on the central bank's choice of inflation; the upper and lower limits of this interval are history-dependent. The history-dependent upper and lower limits do not bind on types whose preferred inflation choice is contained in the interval. These types have discretion and choose their most preferred inflation.

Finally, Proposition 4 characterizes the policy functions for the output gap and promised inflation, x_* and π_*^e, respectively. The output gap and promised inflation in the optimal mechanism depend on the central bank's type only through the inflation rate. Therefore, when the optimal mechanism implies amnesia for the inflation rate, it does so for the output gap and promised inflation too.

PROPOSITION 4. *There exists a pair of strictly increasing, continuous functions,* $(g_x, g_{\pi^e}) : \Pi \to X \times \Pi$, *such that the policy functions for the output gap and promised inflation satisfy, for all* (θ, π_-^e),

$$x_*(\theta; \pi_-^e) = g_x(\pi_*(\theta; \pi_-^e)) \quad \text{and} \quad \pi_*^e(\theta; \pi_-^e) = g_{\pi^e}(\pi_*(\theta; \pi_-^e)).$$

These functions, g_x *and* g_{π^e}, *satisfy, for any* $\pi \in \Pi$,

$$(g_x(\pi), g_{\pi^e}(\pi)) = \arg\max_{x, \pi^e} B(x) + \beta \overline{W}(\pi^e) \quad \text{subject to } \pi = \kappa x + \beta \pi^e.$$

5.3 *Dynamics of the optimal mechanism*

Because the optimal mechanism is inherently dynamic, it is interesting to see how it changes over time in the short and long run.

Recall that some discretion is given when π_-^e lies between $\pi_D(\underline{\theta})$ and $\pi_D(\overline{\theta})$, and that inflation, $\pi_*(\theta; \pi_-^e)$, takes a value between $\pi_D(\underline{\theta})$ and $\pi_D(\overline{\theta})$ whenever some discretion is granted. According to Proposition 4, current period's promised inflation is an increasing function of current inflation, $\pi_*^e(\theta; \pi_-^e) = g_{\pi^e}(\pi_*(\theta; \pi_-^e))$, where g_{π^e} is an increasing function, and, therefore, it lies in an interval between $g_{\pi^e}(\pi_D(\underline{\theta}))$ and $g_{\pi^e}(\pi_D(\overline{\theta}))$ when some discretion is granted. If this interval is contained in $(\pi_D(\underline{\theta}), \pi_D(\overline{\theta}))$, once some discretion is given, the continuation mechanism always prescribes some discretion afterward, and the probability of visiting the no-discretion region is zero. Moreover, because the fully optimal mechanism starts from the optimal initial condition π_-^{e*}, which is contained in $(\pi_D(\underline{\theta}), \pi_D(\overline{\theta}))$, it never visits the no-discretion region, and therefore some discretion is always given. The following proposition shows these properties for the quadratic social welfare specification in (2).

PROPOSITION 5 (Dynamics of the optimal mechanism). *Assume that the social welfare function takes a quadratic form specified in (2), that the type distribution satisfies* $E[\theta] = 0$ *and is symmetric around the mean, and that Assumptions 2 and 4 hold. Then* $\pi_-^{e*} = 0$ *and the following statements hold.*

(i) If $\pi_-^e \le \pi_D(\underline{\theta})$, *promised inflation strictly increases:* $\pi^e(\theta; \pi_-^e) > \pi_-^e$ *for all* θ.

(ii) If $\pi_-^e \ge \pi_D(\overline{\theta})$, *promised inflation strictly decreases:* $\pi^e(\theta; \pi_-^e) < \pi_-^e$ *for all* θ.

(iii) If π_-^e *lies in the interval* $(\pi_D(\underline{\theta}), \pi_D(\overline{\theta}))$, *promised inflation stays within the same interval.*

Proposition 5 has several implications for the short-run and long-run dynamics of the optimal mechanism. In the short run, when the previous period's promised inflation is in the no-discretion region, promised inflation drifts toward the region in which some discretion is granted, while when the previous period's promised inflation is in the nonzero-discretion region, promised inflation varies with the current type of the central bank, but stays in the same nonzero-discretion region. These short-run dynamics imply that no matter what π_{-1}^e is, the optimal mechanism eventually reaches the region where

some discretion is granted and that no discretion is only a short-run phenomenon. The support of the long-run stationary distribution of the state variable (promised inflation) is contained in the interval $(\pi_D(\underline{\theta}), \pi_D(\overline{\theta}))$ for which some discretion is given. In the fully optimal mechanism, which starts with $\pi_-^e = \pi_-^{e*}$, some discretion is always given.

The new Keynesian monetary-policy literature often considers an optimal *timeless perspective* policy (Woodford 1999). For this quadratic specification, we can characterize the optimal timeless mechanism using Proposition 5. According to the timeless perspective, the optimal mechanism is assumed to have been in effect for a sufficiently long period of time before time 0, which requires treating the long-run stationary distribution of the state variables as the distribution of these variables at time 0. Proposition 5 implies that the support of this stationary distribution is contained in the nonzero-discretion region and that for the optimal timeless perspective mechanism, promised inflation never escapes from that region.

5.4 *Implementation by inflation targeting*

The new Keynesian policy literature has analyzed the optimal policy without commitment using a Markov perfect equilibrium (MPE) in the setting where the central bank chooses inflation and the output gap subject to the NKPC. Here we show that the optimal mechanism can be implemented by a history-dependent, yet simple, *inflation targeting rule* in that setting, and that imposing a history-independent rule is actually inferior to imposing no rule.

The inflation targeting rule is designed as follows. Society specifies a time-invariant correspondence, $\Gamma : M \rightrightarrows \Pi$, of the previous period's message that constrains the central bank's choice of inflation in the current period. In other words, with m_- denoting the previous period's message, the current period's inflation must satisfy $\pi \in \Gamma(m_-)$. Then the payoff-relevant state for the central bank is (θ, m_-). A *Markov perfect equilibrium under an inflation targeting rule* Γ consists of (i) the central bank's policy function, $(\pi_{\mathrm{IT}}, x_{\mathrm{IT}}, m_{\mathrm{IT}}) : \Theta \times M \to \Pi \times X \times M$, and (ii) the central bank's (ex post) value function $V_{\mathrm{IT}} : \Theta \times M \to \mathbb{R}$ such that the following statements hold.

(i) For all $(\theta, m_-) \in \Theta \times \Pi$, taking the function π_{IT} as given,

$$V_{\mathrm{IT}}(\theta, m_-) = \max_{\pi, x, m} R(\pi, x, \theta) + \beta \mathbb{E}\big[V_{\mathrm{IT}}\big(\theta', m\big)\big]$$

subject to $\pi = \kappa x + \beta \mathbb{E}[\pi_{\mathrm{IT}}(\theta', m)]$ and $\pi \in \Gamma(m_-)$.

(ii) The policy function for the above problem is $(\pi_{\mathrm{IT}}, x_{\mathrm{IT}}, m_{\mathrm{IT}})$.

Ex ante social welfare in this equilibrium is simply given by $W_{\mathrm{IT}}(m_-) = \mathbb{E}[V_{\mathrm{IT}}(\theta', m_-)]$. It is worth noting that the expected inflation in the NKPC is correctly specified as the private sector's expectation, because given the Markov strategy for inflation, π_{IT}, and a message in the current period, m, the best forecast for next period's inflation by the private sector is given by $\mathbb{E}[\pi_{\mathrm{IT}}(\theta', m)]$ and is independent of the private sector's belief about the current private information θ.

Now we describe the inflation targeting rule that implements the optimal mechanism. For this purpose, we confine our attention to the rule in which the central bank is asked to announce its promised inflation, i.e., $M = \Pi$. We say that Γ implements the optimal policy if there exists a Markov perfect equilibrium given Γ such that $(\pi_{\mathrm{IT}}, x_{\mathrm{IT}}, m_{\mathrm{IT}}) = (\pi_*, x_*, \pi_*^e)$, and $W_{\mathrm{IT}} = \overline{W}$. Note that the condition $m_{\mathrm{IT}} = \pi_*^e$ means that promised inflation announced by the central bank is always delivered. Proposition 6 shows that an appropriately chosen Γ can implement the optimal policy.

PROPOSITION 6. *Define* $\overline{\Gamma} : \Pi \rightrightarrows \Pi$ *as follows:*

(i) If $m_- \leq \pi_D(\underline{\theta})$, *set* m_- *as the upper limit.*

(ii) If $m_- \in (\pi_D(\underline{\theta}), \pi_-^{e*})$, *set* $\pi_D(T_1(m_-))$ *as the upper limit.*

(iii) If $m_- = \pi_-^{e*}$, *set neither an upper nor a lower limit.*

(iv) If $m_- \in (\pi_-^{e*}, \pi_D(\overline{\theta}))$, *set* $\pi_D(T_2(m_-))$ *as the lower limit.*

(v) If $m_- \geq \pi_D(\overline{\theta})$, *set* m_- *as the lower limit.*

Then $\overline{\Gamma}$ *implements the optimal policy.*

There are many inflation-range targeting rules other than $\overline{\Gamma}$ that implement the optimal policy, and $\overline{\Gamma}$ is the largest among them. The smallest correspondence is $\underline{\Gamma}(m_-) = \{\pi \in \Pi | \pi = \pi_*(\theta; m_-) \text{ for some } \theta\}$. It is straightforward to see that a necessary and sufficient condition for Γ to implement the optimal policy is $\underline{\Gamma}(m_-) \subset \Gamma(m_-) \subset \overline{\Gamma}(m_-)$ for all $m_- \in \Pi$. Note that, for a given m_-, the 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 (Holmstrom 1977, Holmstrom 1984, Alonso and Matouschek 2008, and Amador and Bagwell 2013).

Proposition 6 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 an upper bound on inflation 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 the 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 a stabilization bias and not an inflation bias. When the previous period's promised inflation 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.

Throughout the paper, we have assumed that the central bank is benevolent and that its payoff is identical to social welfare. In Appendix F.2, we generalize the setup to incorporate an inflation bias. Society finds it optimal to tighten the upper limit on inflation when the central bank has an inflation bias compared to when the central bank is benevolent. As a result, social welfare, as a function of previous period's promised inflation, is maximized at a value that is lower than π_-^{e*}.

5.4.1 *Importance of history dependence* To illustrate the importance of history dependence of an inflation targeting rule, let us consider a fixed-range inflation targeting rule such that $\Gamma(m_-)$ is a constant interval that is independent of m_-.

For simplicity, assume that the social welfare function takes the quadratic form specified in (2), that the type distribution satisfies $E[\theta] = 0$ and is symmetric around the mean, and that Assumptions 2 and 4 hold. Then, because the model features symmetry around 0, we can without loss of generality focus on fixed-range inflation targeting rules that take the form $\Gamma(m_-) = [-c, c]$ with $c \geq 0$. This form encapsulates the strict zero inflation targeting of always setting $\pi_t = \pi_-^{e*} = 0$ when c is set to zero, and the "optimal discretionary policy" in the new Keynesian policy literature when c is set to so large a number that the target inflation range never binds.

Because the targeting rule is independent of the the previous period's message, the payoff-relevant state in a MPE is θ only, and we drop the message m from the set of choice variables. In an MPE given Γ, the value and the policy functions must satisfy that (i) for all $\theta \in \Theta$, $V_{\mathrm{IT}}(\theta) = \max_{\pi,x} R(\pi, x, \theta) + \beta \mathbb{E}[V_{\mathrm{IT}}(\theta')]$ subject to $\pi = \kappa x + \beta \mathbb{E}[\pi_{\mathrm{IT}}(\theta')]$ and $\pi \in [-c, c]$, and (ii) $(\pi_{\mathrm{IT}}, x_{\mathrm{IT}})$ is the policy function for the above problem.

Interestingly, given history independence, imposing bounds on the central bank's choice of inflation is harmful for welfare, because ex ante welfare in the MPE, $W_{\mathrm{IT}} = \mathbb{E}[V_{\mathrm{IT}}(\theta)]$, weakly increases with c. The logic is as follows. First, it is straightforward to establish that $\mathbb{E}[\pi_{\mathrm{IT}}(\theta')] = 0$ for any c, due to symmetry around 0. Hence, for each θ, the pair $(\pi_{\mathrm{IT}}(\theta), x_{\mathrm{IT}}(\theta))$ solves $\max_{\pi,x} R(\pi, x, \theta)$ subject to $\pi = \kappa x$ and $\pi \in [-c, c]$. Because lower c implies a tighter constraint, when c is reduced, the expected momentary welfare in the MPE, $\mathbb{E}[R(\pi_{\mathrm{IT}}(\theta), x_{\mathrm{IT}}(\theta), \theta)]$, never increases, and it strictly decreases if the bound constraint $\pi \in [-c, c]$ binds for some type θ. Finally, ex ante welfare satisfies $W_{\mathrm{IT}} = \mathbb{E}[V_{\mathrm{IT}}(\theta)] = \mathbb{E}[R(\pi_{\mathrm{IT}}(\theta), x_{\mathrm{IT}}(\theta), \theta)]/(1 - \beta)$ and, therefore, it is weakly increasing in c. Hence, if society can only use history-independent inflation targeting rules, it is optimal to give full discretion to the central bank, i.e., to set c so large that the bound constraint never binds.

The reason why imposing a history-independent inflation targeting rule only hurts welfare is that it does not help society to mitigate the time-inconsistency problem of the central bank. When the inflation targeting rule is given by $\overline{\Gamma}$, the next period's constraint set for the central bank depends on its current-period promised inflation, and the central bank can use promised inflation to restrict the action chosen by its future self, thereby mitigating the time-inconsistency problem. Fixed-range inflation targeting rules do not provide such a mechanism and the time-inconsistency problem is left unsolved. They only restrict the contemporaneous choice of the central bank, which is unambiguously harmful for welfare.

6. Numerical example

In this section, we use a numerical example to compare the optimal mechanism with two benchmarks. 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 (i.e., a 25 percent markup), 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, h_t denotes hours worked, and ν, the inverse of the Frisch labor-supply elasticity, is set to 0.50.

We use the quadratic specification in (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., Woodford 2003).

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 correspond approximately to $[-2\%, 2\%]$ per annum. For computational purposes, this continuous density is approximated using a uniform grid containing 41 points.

6.2 *Benchmarks*

The first alternative is the aforementioned optimal discretionary policy, and the second alternative is the full-information solution. These two benchmarks are commonly analyzed in the new Keynesian literature (see Clarida et al. (1999) and Woodford (2003)) and are detailed in Appendix D.

6.2.1 *Optimal discretionary policy* The optimal discretionary policy exhibits no history dependence because there is no state variable. Policy functions for inflation and the output gap are strictly increasing in θ, but that for promised inflation is independent of θ and is constant at zero. As we discussed in Section 5.4.1, welfare achieved by this policy serves as an upper bound for welfare achievable by any fixed-range inflation targeting rule.

6.2.2 *Full-information benchmark* The full-information problem is to choose a mechanism to maximize social welfare (7) subject to the feasibility constraint (10) and (11). The solution corresponds to the *optimal commitment policy* in the new Keynesian policy literature.

For each of the three variables—inflation, the output gap, and promised inflation—the policy function has a constant coefficient on the previous period's promised inflation, π_-^e, that is independent of θ and an intercept that varies with θ. This property follows from the fact that the full-information problem is an example of an optimal linear-quadratic regulator problem.

All policy functions are strictly increasing in θ and π_-^e. Optimal inflation naturally increases with θ because higher θ implies that the marginal social welfare of inflation is higher. To incentivize the price setters to raise their prices, the central bank raises

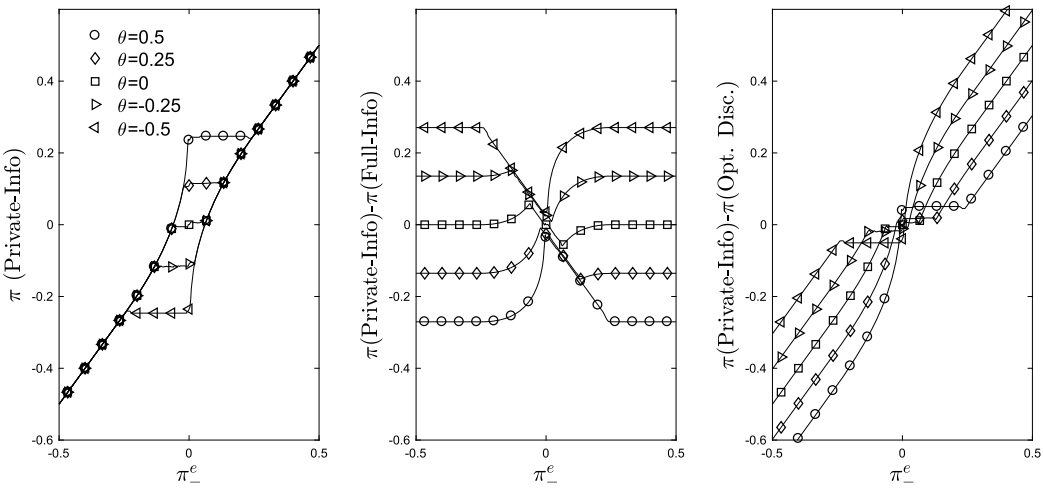

FIGURE 3. Policy function comparison: π.

the current marginal cost by stimulating the current output gap and the private sector's inflation expectations by promising higher future inflation. Hence, both the output gap and the promised inflation increase with θ. Inflation increases with the previous period's promised inflation because inflation has to be on average equal to promised inflation, and both the output gap and the expected inflation move positively with inflation for the same reason. Finally, promised inflation moves less than one for one with the previous period's promised inflation, implying that promised inflation settles in a compact interval around zero in the long run.

6.3 *Comparison*

Considering inflation first, the left panel in Figure 3 displays the policy function for inflation, π_*, as a function of the state for different values of θ. The middle and right panels then plot the differences between the optimal mechanism and the full-information solution, and between the optimal mechanism and the optimal discretionary policy, respectively. To make these plots visible, we report these policy functions for just five 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% per quarter.

Looking at the policy functions shown in the left panel it is clear that the optimal mechanism exhibits properties that are largely consistent with the theoretical results for the continuous-type model despite the fact that we use a discrete-type model for computation.[15]

[15]Figure 3 reveals some behaviors that differ from the theoretical predictions of the continuous-type model. First, we observe full discretion not only at one point ($\pi_-^e = \pi_-^{e*} = 0$), but on an interval around 0 (though it is small). Second, on this interval, the policy function is not flat, but increasing slightly. 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

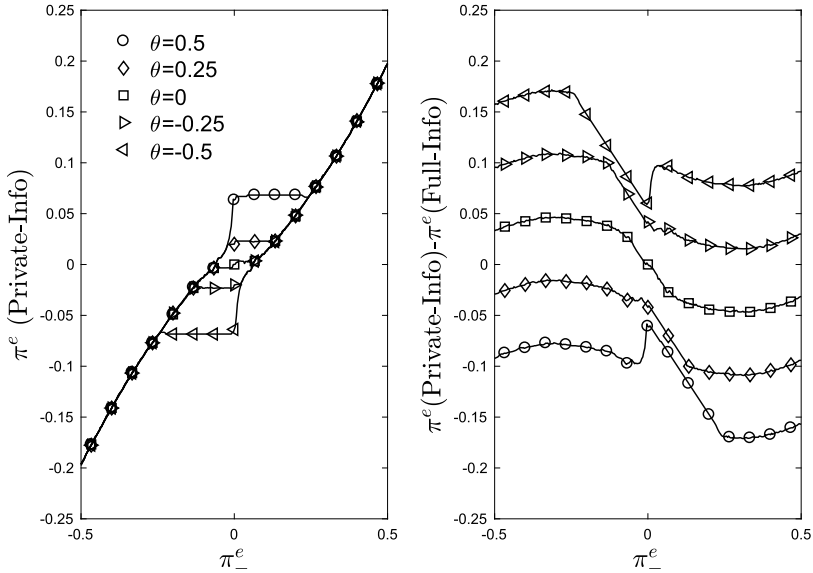

FIGURE 4. Policy function comparison: π^e.

Inflation in the optimal mechanism moves *less* with θ than it does in the full-information solution. Recall that the full-information solution increases monotonically with θ. The middle panel of Figure 3 shows that, at each π^e_-, the optimal mechanism tends to be higher (lower) than the full-information solution for low (high, respectively) values of θ. Loosely speaking, for each value of the previous period's promised inflation, the optimal mechanism is more compressed vertically than is the full-information solution. In contrast, when a type θ has discretion in the optimal mechanism, inflation moves *more* with θ than it does in the optimal discretionary policy. In the right panel of Figure 3, one can see that, for each value of θ, the flat part has the same sign as θ. Therefore, the optimal history-dependent mechanism allows inflation to vary more with the desirable inflation θ than in the optimal discretionary policy, and it is welfare-enhancing.

Next we turn to promised inflation. In Figure 4, the left panel depicts the optimal mechanism and the right panel depicts the difference from the full-information solution.[16] As shown in Proposition 4, the policy function for promised inflation is a monotone transformation of that for inflation. Therefore, they are qualitatively very similar. The right panel implies that promised inflation in the optimal mechanism move *less* with θ than in the full-information solution. Thus, one feature of the optimal mechanism is that it exhibits less volatility in expected inflation than the full-information solution. 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

computational method. The interval of full discretion indeed becomes small as we increase the number of types.

[16]Promised inflation is always zero under the optimal discretionary policy and is omitted.

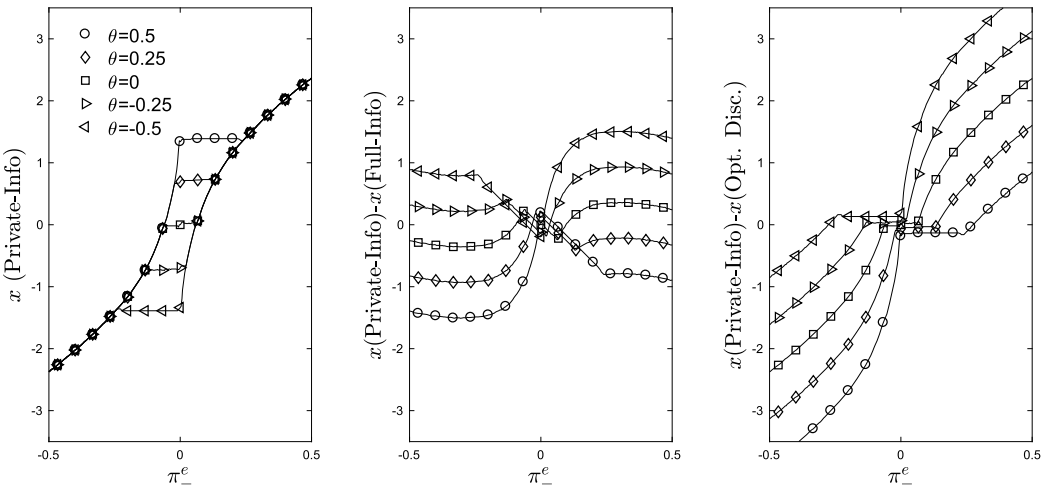

FIGURE 5. Policy function comparison: x.

welfare. Accordingly, the private information solution makes less use of promised inflation.

Figure 5 depicts the behavior of the output gap. First observe in the right panel that for each value of θ, the flat part has the *opposite* sign to θ. This implies that the output gap moves *less* with θ in the optimal mechanism when a type has discretion than it does in the optimal discretionary policy. This difference is due to the fact that the optimal mechanism uses θ-dependent promised inflation while the optimal discretionary policy does not. By generating a positive correlation between promised inflation and θ, the optimal mechanism can reduce the output gap fluctuations relative to the optimal discretionary policy. The middle panel of Figure 5 shows that the difference from the full-information solution is non-monotonic in θ. The reason is simple. Recall the NKPC, $\pi = \kappa x + \beta \pi^e$. We have seen that both inflation, π, and promised inflation, π^e, move less with θ in the optimal mechanism than they do in the full-information solution. Hence, the output gap, x, may move more or less with θ.

As a final point, we use the dynamic IS equation to calculate the nominal interest rate, assuming that the natural rate of interest is constant at $1/\beta - 1$ (approximately 1% in this example), and plot it as a function of π^e_- in Figure 6. In the optimal discretionary policy (the right panel), the nominal rate moves negatively with θ and is independent of promised inflation. In the full-information solution (the middle panel), the nominal rate depends negatively on promised inflation and θ, and the dependence is linear. In the optimal mechanism (the left panel), types that do not have discretion choose the same value for the nominal rate, while those that have discretion choose a type-specific but history-independent nominal rate, just as they do for inflation and the output gap. This suggests that one might also implement the optimal mechanism by imposing history-dependent upper and lower limits on the nominal interest rate. When promised inflation is lower than π^{e*}_-, a lower limit is imposed so that high-θ types are constrained and

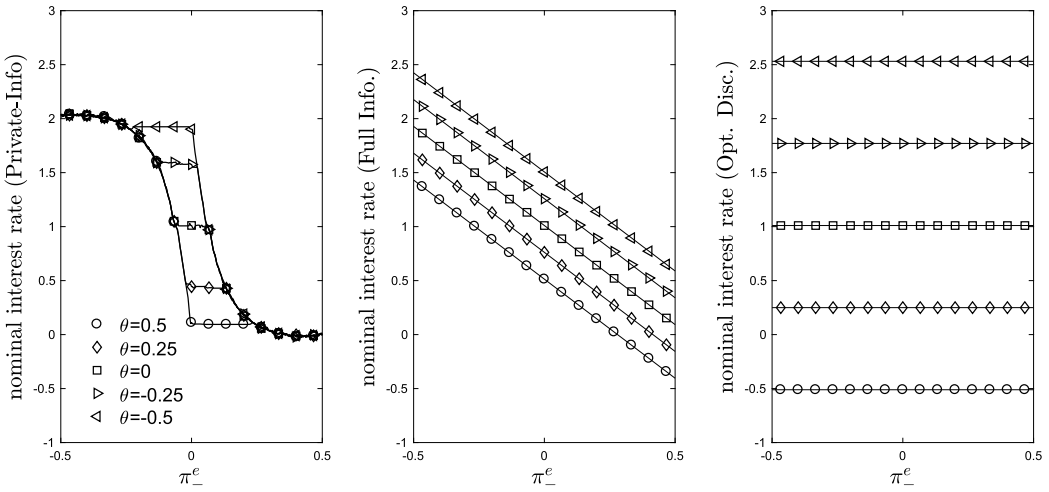

FIGURE 6. The nominal interest rate.

are unable to choose high inflation. When promised inflation is higher than π^{e*}_-, an upper bound is imposed and low-θ types must inflate more than they desire by choosing a low nominal interest rate.

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 the previous period's promised inflation 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 the previous period's promised inflation happens to maximize social welfare, no discretion should be granted when the previous period's promised inflation is sufficiently far from the peak of social welfare, and some discretion should be granted for all intermediate values of the previous period's promised inflation. 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 regime is suboptimal. One way to encode history dependence is to make the upper and lower bounds contingent on promised inflation announced by the central bank in the previous period, a form of inflation target. 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.[17] 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 i.i.d. 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.[18] 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 target ranges for the year-ahead inflation in an attempt to bring about disinflation (Bernanke et al. 2001). Other central banks, such as the Reserve Bank of New Zealand, have changed their target ranges for inflation based on the history of outcomes generated. Our analysis suggests that fixed inflation targets, while practical, lack the sophistication needed to optimally trade off the gains and losses from discretion.

Appendix

Appendix A contains the proofs of the theoretical results presented in the main paper. Appendix B extends the results in Athey et al. (2005). Appendix C describes our numerical procedure. Appendix D presents the two benchmarks.

Appendix A: Proofs

A.1 *Proof of Lemma 2*

Fix any $\pi_-^e \in \Pi$. For any W_- such that $(\pi_-^e, W_-) \in \Omega$, we can find measurable functions (π, x, π^e, w) of θ such that (15), (16), (17), and (18) are satisfied with $F = \overline{W}$. Hence,

$$\mathbb{T}\overline{W}(\pi_-^e) \geq \sup_{W_- \text{ s.t. } (\pi_-^e, W_-) \in \Omega} W_- = \overline{W}(\pi_-^e).$$

This establishes $\overline{W} \leq \mathbb{T}\overline{W}$. Because $\mathbb{T}$ is a monotone operator, it follows that $(\mathbb{T})^n\overline{W} \leq (\mathbb{T})^{n+1}\overline{W}$ for any integer $n \geq 0$. Because $\mathbb{T}$ is a contraction mapping with the fixed point $\tilde{W}$, we have $\overline{W} \leq (\mathbb{T})^n\overline{W} \leq \lim_{n\uparrow\infty}(\mathbb{T})^n\overline{W} = \tilde{W}$.

[17] We 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 $\theta_t = \rho\theta_{t-1} + e_t$.

[18] A 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.

A.2 *Proofs of Lemma 3(iii) and Proposition 4*

We assume Assumptions 1, 2, 3, and 4.

We take the following steps. First, we fix a $F \in \mathcal{V}(\Pi)$ and an arbitrary $\pi_-^e \in \Pi$. We then relax the problem in (14) and characterize its solution. The relaxed problem is a static problem in which society chooses inflation, π, as a function of θ, and a "money burning" variable, $\delta \leq 0$, also as a function of θ. By generalizing AAK's theorems and applying them to the relaxed problem, we show that it has a unique solution, such that π is continuous and that $\delta(\theta) = 0$ for all θ. When the optimal money burning is zero, the original problem (14) has a solution such that (a) inflation is identical to the optimal inflation in the relaxed problem, (b) the output gap and promised inflation satisfy the properties in Proposition 4 with $\overline{W}$ in the statement being replaced with F, and (c) the continuation utility hits the upper bound, $w(\theta) = F(\pi^e(\theta))$. Hence, Lemma 3(iii) follows, and Proposition 4 is implied once we show Lemma 3(i).

Fix an arbitrary weakly concave function $F \in \mathcal{V}(\Pi)$ and arbitrary $\pi_-^e \in \Pi$.

To relax the problem in (14), we define the function S of π: for any $\pi \in \Pi$,

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

subject to $\pi = \kappa x + \beta\pi^e$. (When obvious, we suppress the dependence of S on F.) Then, for any quadruple (π, x, π^e, w) that satisfies the NKPC and the constraint (18), there exists $\delta \leq 0$ such that

$$A(\pi;\theta) + B(x) + \beta w = A(\pi;\theta) + S(\pi; F) + \delta.$$

Here δ represents money burning: $S(\pi; F)$ is the maximal welfare from feasible pairs of the output gap and promised inflation for given inflation, π, and $\delta < 0$ implies that, conditional on π, the output gap, promised inflation, and the continuation utility are at suboptimal levels, possibly to incentivize the central bank.

In the relaxed problem, we allow society to choose δ arbitrarily subject only to the nonpositivity constraint. Let

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

Then the following problem (P1) relaxes the problem in (14):

$$\sup_{\pi(\cdot),\delta(\cdot)} \int_{\underline{\theta}}^{\overline{\theta}} \big[\tilde{R}\big(\pi(\theta);\theta\big) + \delta(\theta)\big] p(\theta)\, d\theta, \tag{P1}$$

subject to the constraints

$$\pi_-^e = \int \pi(\theta)\, p(\theta)\, d\theta,$$

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

$$\delta(\theta) \leq 0 \quad \forall\theta.$$

As we state in Assumption 4(i), we assume that a maximum is attained in problem (P1) and replace sup with max in what follows.

Problem (P1) has a unique solution in which $\pi(\cdot)$ is a continuous function of θ, and $\delta(\theta) = 0$ for all θ. To show these results, we exploit the fact that, interpreting $\tilde{R}$ as the return function in the social welfare function and interpreting δ as the continuation value, problem (P1) has the same structure as the best payoff problem in AAK for which the optimal $\delta(\cdot)$ is shown to be constant at its upper bound.[19] AAK, however, assume that $\pi(\cdot)$ is a piecewise C^1 function, while we do not. We generalize their results to dispense with that assumption.

LEMMA 5 (AAK's result extended). *Suppose $\tilde{R}$ and p satisfy the following conditions.*

(i) The function $\tilde{R}$ is strictly concave and $\tilde{R}_\theta$ exists.

(ii) The single-crossing condition: $\tilde{R}_{\theta\pi}(\pi; \theta) > 0$ for any (π, θ).

(iii) The monotone-hazard condition: For any nondecreasing $\pi(\cdot)$,

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

and

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

(iv) The Milgrom–Segal condition:

(a) For any π, $\tilde{R}(\pi; \cdot)$ is absolutely continuous.

(b) There is an integrable function $a : \Theta \to \mathbb{R}_+$ such that $|\tilde{R}_\theta(\pi; \theta)| < a(\theta)$ for all π and almost all θ.

Assume further that a maximum is attained and the bound constraints are slack in problem (P1). Then a solution to problem (P1) is unique, and in the solution, (i) $\pi(\cdot)$ is a continuous function and (ii) $\delta(\cdot)$ is constant at its upper bound of zero.

A proof can be found in Appendix B.

The assumptions in Lemma 5 are implied by Assumptions 1, 2, 3, and 4, and $F \in \mathcal{V}(\Pi)$.

LEMMA 6.

(i) Assumption 1 implies $\tilde{R}_\theta = A_\theta$, and if we further assume $F \in \mathcal{V}$, $\tilde{R}$ is strictly concave.

(ii) Assumption 1 implies the single-crossing condition for $\tilde{R}$.

[19] One difference is 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. However, the constancy of optimal $\delta(\cdot)$ in AAK is unaffected by this difference.

(iii) Assumptions 1 and 2 imply the monotone-hazard condition for $(\tilde{R}, p)$.

(iv) Assumption 3 implies the Milgrom–Segal condition for $\tilde{R}$.

PROOF. Because $\tilde{R}(\pi;\theta) = A(\pi;\theta) + S(\pi)$, we have $\tilde{R}_\theta = A_\theta$ and $\tilde{R}_{\theta\pi} = A_{\theta\pi}$ for any function $S(\cdot)$. The first part of (i) and parts (ii)–(iv) follow immediately. To prove the remaining part of (i), consider the problem that defines S for $F \in \mathcal{V}(\Pi)$. Because the graph of the constraint correspondence is convex and the objective function $B(x) + \beta F(\pi^e)$ is weakly concave, S is weakly concave. Because $A(\pi;\theta)$ is strictly concave in π for any θ, it follows that $\tilde{R}(\pi;\theta) = A(\pi;\theta) + S(\pi)$ is strictly concave in π for any θ. □

Therefore, under the set of assumptions made in Lemma 3, the solution to problem (P1), $(\pi_*(\cdot), \delta_*(\cdot))$, satisfies that $\delta_*(\theta) = 0$ for all θ and that $\pi_*(\cdot)$ is continuous (Lemma 5).

To construct a solution to the original problem from the solution to problem (P1), we use a solution to the problem that defines the function S. We turn to its characterization next.

LEMMA 7. *Suppose that* $F \in \mathcal{V}(\Pi)$. *Then* S *is weakly concave and the optimal policy correspondence*

$$\left\{\left(x, \pi^e\right) \in X \times \Pi \middle| \pi = \kappa x + \beta\pi^e \text{ and } S(\pi) = B(x) + \beta F\left(\pi^e\right)\right\}$$

is single-valued and can be written as a pair of functions of π, $(g_x(\cdot), g_{\pi^e}(\cdot))$, *where both* $g_x(\cdot)$ *and* $g_{\pi^e}(\cdot)$ *are nondecreasing, continuous functions.*[20] *If* F *is further strictly concave and differentiable,* S *is strictly concave and* C^1, *both* g_x *and* g_{π^e} *are strictly increasing, and the envelope condition holds:* $S'(\pi) = B'(g_x(\pi))/\kappa = F'(g_{\pi^e}(\pi))$.

PROOF. Suppose first that $F \in \mathcal{V}(\Pi)$. Weak concavity of S is already established in Lemma 6. Note that $S(\pi) = \max_{\pi^e} B((\pi - \beta\pi^e)/\kappa) + \beta F(\pi^e)$. Because F is weakly concave, $f(\pi, \pi^e) := B((\pi - \beta\pi^e)/\kappa) + \beta F(\pi^e)$ is strictly concave in π^e for given π. Therefore, the solution to the right-hand side problem is unique for each π and we denote the policy function by $g_{\pi^e}(\cdot)$. The maximum theorem implies the continuity of $g_{\pi^e}(\cdot)$. It is straightforward to check that f has increasing differences in (π, π^e) and it follows that $g_{\pi^e}(\cdot)$ is weakly increasing. The same properties of g_x can be shown by considering $\max_x B(x) + \beta F((\pi - \kappa x)/\beta)$. Clearly, for each π, (x, π^e) is in the policy correspondence if and only if $(x, \pi^e) = (g_x(\pi), g_{\pi^e}(\pi))$.

If F is further strictly concave and differentiable, first-order necessary conditions (FONC) imply $B'(g_x(\pi))/\kappa = F'(g_{\pi^e}(\pi))$. Because both B' and F' are strictly decreasing, and because $g_x(\cdot)$ and $g_{\pi^e}(\cdot)$ must satisfy $\pi = \kappa g_x(\pi) + \beta g_{\pi^e}(\pi)$, both $g_x(\cdot)$ and $g_{\pi^e}(\cdot)$ must be strictly increasing. The Benveniste–Scheinkman theorem implies $S'(\pi) = B'(g_x(\pi))/\kappa = F'(g_{\pi^e}(\pi))$. Because $B'(g_x(\cdot))$ is strictly decreasing and continuous, so is S'. Therefore, S is a strictly concave C^1 function. □

[20]The functions S, g_x, and g_{π^e} depend on the function F, but we suppress this dependence to simplify the notation.

Define the composite functions $x_* = g_x \circ \pi_*$, $\pi_*^e = g_{\pi^e} \circ \pi_*$, and $w_* = F \circ g_{\pi^e} \circ \pi_*$. Then $(\pi_*, x_*, \pi_*^e, w_*)$ satisfies all the constraints in the original problem, and the value of the objective function it achieves is the same as the maximized value of the relaxed problem (P1). By construction, the quadruple of functions $(\pi_*(\cdot), x_*(\cdot), \pi_*^e(\cdot), w_*(\cdot))$ defined above is continuous in θ. Because problem (P1) has a unique solution, it is a unique solution to the original problem.

This proves the second and the third parts of Lemma 3. Proposition 4 follows once we show $\overline{W} \in \mathcal{V}(\Pi)$.

A.3 *Proofs of Lemma 3(i) and Propositions 2 and 3*

We again take an arbitrary weakly concave function $F \in \mathcal{V}(\Pi)$. Now we characterize the optimal policy function for inflation given F and the value function $\mathbb{T}F$. For each π_-^e, the former, together with a constant function $\delta(\theta) = 0$, constitutes the solution to problem (P1), and the latter equals the maximized value of problem (P1). Hence, we use problem (P1) with $\delta(\theta) = 0$ being imposed:

$$\mathbb{T}F\big(\pi_-^e\big) = \max_{\pi(\cdot)} \int_{\underline{\theta}}^{\overline{\theta}} \tilde{R}\big(\pi(\theta); \theta\big) p(\theta)\, d\theta \tag{P2}$$

subject to (15) and $\tilde{R}(\pi(\theta); \theta) \geq \tilde{R}(\pi(\theta'); \theta)$ for all θ and $\theta' \neq \theta$. We call this problem (P2).

We first show that at each π_-^e, the solution to (P2) has at most two thresholds, θ_1 and θ_2: it is constant below θ_1, coincides with the one-shot discretionary best response between θ_1 and θ_2, and is constant above θ_2. Using this property, we show that $\mathbb{T}F$ has a peak when π_-^e is at the expected value of the one-shot discretionary best response, that it is strictly increasing on the left of the peak, and that it is strictly decreasing on the right of the peak. The shape of $\mathbb{T}F$ in turn implies that at each π_-^e, the solution to problem (P2) has at most one cutoff point and has the form described in Proposition 3. The cutoff point is shown to vary with the state variable π_-^e, implying that the optimal policy function $\pi_*(\theta; \pi_-^e)$ given F is a continuous function of (θ, π_-^e) (Lemma 3(ii)). Finally, using the properties of the optimal policy function, the value function $\mathbb{T}F$ is shown to be a strictly concave C^1 function, establishing Lemma 3(i).

A.3.1 π_* *takes a simple form*

Define the *one-shot discretionary best response given* F as a policy correspondence $\pi_D(\cdot; F)$ for the maximization problem

$$\max_{\pi}\Big[A(\pi; \theta) + \max_{(x, \pi^e): \pi = \kappa x + \beta \pi^e} \big\{B(x) + \beta F\big(\pi^e\big)\big\}\Big] = \max_{\pi} \tilde{R}(\pi, \theta).$$

Clearly $\pi_D(\cdot; F)$ maps from Θ to Π.

The next lemma establishes that $\pi_D(\cdot; F)$ is a nondecreasing, continuous function. For simplicity, we drop the dependence of π_D on F hereafter.

Lemma 8. *If $F \in \mathcal{V}(\Pi)$ and Assumption 1 holds, then $\pi_D(\cdot)$ is a weakly increasing, continuous function. If F is a strictly concave, differentiable function, then $\pi_D(\cdot)$ is strictly increasing.*

PROOF. Suppose that $F \in \mathcal{V}(\Pi)$. Because $\tilde{R}(\cdot;\theta)$ is strictly concave, its maximizer is unique and, thus, π_D is a function. The maximum theorem implies the continuity of π_D. Because $A_{\pi\theta} > 0$ (Assumption 1), $\tilde{R}$ has increasing differences in (π,θ) and, thus, $\pi_D(\cdot)$ is weakly increasing. Suppose further that F is strictly concave and differentiable. Then $\tilde{R}$ is strictly concave and C^1 in π. Thus, for each π, the FONC must be satisfied: $0 = \tilde{R}_\pi(\pi_D(\theta);\theta) = A_\pi(\pi_D(\theta);\theta) + S'(\pi_D(\theta))$. Because $A_{\pi\theta} > 0$, $\pi_D(\cdot)$ must be strictly increasing. □

Now we show that, for any π^e_-, $\pi_*(\cdot)$ must be either constant or of the form

$$\pi_*(\theta) = \begin{cases} \pi_D(\theta_1) & \forall\theta \in [\underline{\theta}, \theta_1] \\ \pi_D(\theta) & \forall\theta \in (\theta_1, \theta_2) \\ \pi_D(\theta_2) & \forall\theta \in [\theta_2, \overline{\theta}]. \end{cases} \tag{20}$$

Suppose that $\pi_*(\cdot)$ is not constant. Then $\pi_*(\cdot)$ is strictly increasing on some interval (θ_1,θ_2). Then for any triple $(\theta,\theta',\theta'')$ such that $\theta_1 < \theta'' < \theta < \theta' < \theta_2$, the incentive-compatibility constraint implies $\tilde{R}(\pi_*(\theta'');\theta) < \tilde{R}(\pi_*(\theta);\theta)$ and $\tilde{R}(\pi_*(\theta');\theta) < \tilde{R}(\pi_*(\theta);\theta)$. Because π_* is continuous, $\tilde{R}(\cdot;\theta)$ must have a peak at $\pi = \pi_*(\theta)$. Therefore, $\pi_*(\cdot) = \pi_D(\cdot)$ in an interval in which $\pi_*(\cdot)$ is strictly increasing, and $\pi_*(\cdot)$ is constant when $\pi_*(\cdot) \neq \pi_D(\cdot)$. Because $\pi_*(\cdot)$ is continuous (Lemma 5(ii)), it follows that it must be either constant or of the form in (20). Later we show that either $\theta_1 = \underline{\theta}$ or $\theta_2 = \overline{\theta}$ must hold.

A.3.2 *$\mathbb{T}F$ is single-peaked* It follows that $\mathbb{T}F$ is single-peaked, that it is strictly increasing to the left of its peak, and that it is strictly decreasing to the right of its peak.

LEMMA 9. *The function $\mathbb{T}F : \Pi \to \mathbb{R}$ is uniquely maximized at $\pi^{e*}_- := \int_{\underline{\theta}}^{\overline{\theta}} \pi_D(\theta)\,p(\theta)\,d\theta$.*

PROOF. The objective function in problem (P2) is maximized if and only if $\pi = \pi_D$ a.e., because $\tilde{R}$ is strictly concave in π for each θ. However, $\pi = \pi_D$ satisfies the constraint (15) 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. □

COROLLARY 2. *The function $\mathbb{T}F$ is strictly increasing for $\pi^e_- < \pi^{e*}_-$ and is strictly decreasing for $\pi^e_- > \pi^{e*}_-$.*

PROOF. Let $\pi^e_1 < \pi^e_2 < \pi^{e*}_-$. We show that $\mathbb{T}F(\pi^e_1) < \mathbb{T}F(\pi^e_2)$. Let $\pi_*(\cdot;\pi^e_1)$ be a solution to problem (P2) at $\pi^e_- = \pi^e_1$. Then it is of the form of (20) for some θ_1 and θ_2. Notice that $\theta_2 < \overline{\theta}$, because otherwise $\pi_*(\cdot;\pi^e_1) \geq \pi_D(\cdot)$ with strict equality for $\theta > \theta_1$, and the expected value of $\pi_*(\cdot;\pi^e_1)$ satisfies

$$\pi^e_1 = \int_{\underline{\theta}}^{\overline{\theta}} \pi_*\big(\theta;\pi^e_1\big)\,p(\theta)\,d\theta \geq \int_{\underline{\theta}}^{\overline{\theta}} \pi_D(\theta)\,p(\theta)\,d\theta = \pi^{e*}_-,$$

which is a contradiction.

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

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

This is because if $\pi_D(\theta) = \pi_D(\theta_2)$ for all $\theta > \theta_2$, then

$$\pi_1^e = \int_{\underline{\theta}}^{\theta_1} \pi_D(\theta_1) p(\theta)\, d\theta + \int_{\theta_1}^{\overline{\theta}} \pi_D(\theta) p(\theta)\, d\theta > \pi_-^{e*},$$

which is a contradiction.

For such θ_3, define an alternative choice $\pi_{\mathrm{ALT}}(\cdot)$ as $\pi_{\mathrm{ALT}}(\theta) = \pi_*(\theta; \pi_1^e)$ for all $\theta < \theta_2$, $\pi_{\mathrm{ALT}}(\theta) = \pi_D(\theta)$ for all $\theta \in [\theta_2, \theta_3)$, and $\pi_{\mathrm{ALT}}(\theta) = \pi_D(\theta_3)$, for all $\theta \in [\theta_3, \overline{\theta}]$. Then π_{ALT} satisfies the constraints in problem (P2) at $\pi_-^e = \pi_2^e$. Hence, $\mathbb{T}F(\pi_2^e) \geq \int_{\underline{\theta}}^{\overline{\theta}} \tilde{R}(\pi_{\mathrm{ALT}}(\theta); \theta) p(\theta)\, d\theta$. Because π_{ALT} and $\pi_*(\cdot; \pi_-^e)$ are identical up to θ_2,

$$\begin{aligned}
&\int_{\underline{\theta}}^{\overline{\theta}} \tilde{R}\big(\pi_{\mathrm{ALT}}(\theta); \theta\big) p(\theta)\, d\theta - \mathbb{T}F\big(\pi_1^e\big) \\
&\quad = \int_{\theta_2}^{\theta_3} \big\{\tilde{R}\big(\pi_D(\theta); \theta\big) - \tilde{R}\big(\pi_D(\theta_2); \theta\big)\big\} p(\theta)\, d\theta \\
&\qquad + \int_{\theta_3}^{\overline{\theta}} \big\{\tilde{R}\big(\pi_D(\theta_3); \theta\big) - \tilde{R}\big(\pi_D(\theta_2); \theta\big)\big\} p(\theta)\, d\theta.
\end{aligned}$$

The first integral on the right-hand side is strictly positive. The second integral on the right-hand side is also strictly positive, because for all $\theta > \theta_3$, $\pi_D(\theta) \geq \pi_D(\theta_3) > \pi_D(\theta_2)$, and the concavity of $\tilde{R}$ implies $\tilde{R}(\pi_D(\theta); \theta) \geq \tilde{R}(\pi_D(\theta_3); \theta) > \tilde{R}(\pi_D(\theta_2); \theta)$ for all $\theta > \theta_3$. Therefore, $\int_{\underline{\theta}}^{\overline{\theta}} \tilde{R}(\pi_{\mathrm{ALT}}(\theta); \theta) p(\theta)\, d\theta > \mathbb{T}F(\pi_1^e)$, 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 $\pi_-^e > \pi_-^{e*}$. □

A.3.3 $\pi_*(\cdot)$ *satisfies the cutoff property in Proposition 3*

LEMMA 10. *(i) For $\pi_-^e < \pi_-^{e*}$, $\pi_*(\cdot)$ is either constant or has the form in (20) with $\theta_1 = \underline{\theta}$. (ii) For $\pi_-^e > \pi_-^{e*}$, $\pi_*(\cdot)$ is either constant or has the form in (20) with $\theta_2 = \overline{\theta}$.*

PROOF. Suppose to the contrary that, for some $\pi_-^e < \pi_-^{e*}$, $\pi_*(\cdot)$ has the form in (20) with $\theta_1 > \underline{\theta}$. Fix such $\pi_-^e < \pi_-^{e*}$.

First consider the case where $\pi_D(\theta) = \pi_D(\theta_1)$ for all $\theta < \theta_1$. Then we can replace θ_1 with $\underline{\theta}$ without loss of generality.

Suppose next that for some $\tilde{\theta} < \theta_1$, $\pi_D(\theta) < \pi_D(\theta_1)$ for all $\theta \leq \tilde{\theta}$. Because $\mathbb{T}F$ is strictly increasing by Corollary 2, replacing the first constraint in problem (P2) with $\pi_-^e \geq \int_{\underline{\theta}}^{\overline{\theta}} \pi(\theta) p(\theta)\, d\theta$ must not increase the maximized value. Let an alternative choice $\pi_{\mathrm{ALT}}(\cdot)$ be such that $\pi_{\mathrm{ALT}}(\theta) = \pi_*(\theta)$ for all $\theta > \theta_1$ and $\pi_{\mathrm{ALT}}(\theta) = \pi_D(\theta)$ for all $\theta \leq \theta_1$.

Then π_{ALT} is incentive compatible and $\pi_-^e > \int_{\underline{\theta}}^{\overline{\theta}} \pi_{\mathrm{ALT}}(\theta)\, p(\theta)\, d\theta$. Moreover, the objective function increases by

$$\int_{\underline{\theta}}^{\theta_1} \big\{\tilde{R}\big(\pi_{\mathrm{ALT}}(\theta);\theta\big) - \tilde{R}\big(\pi_*(\theta);\theta\big)\big\}\, p(\theta)\, d\theta$$

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

because the integrand is nonnegative for all $\theta \le \theta_1$ and is strictly positive for all $\theta < \tilde{\theta}$. 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 threshold functions, $T_1 : (\pi_D(\underline{\theta}), \pi_-^{e*}) \to [\underline{\theta}, \overline{\theta}]$ and $T_2 : (\pi_-^{e*}, \pi_D(\overline{\theta})) \to [\underline{\theta}, \overline{\theta}]$. Let

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

$$J(b) := P(b)\pi_D(b) + \int_{b}^{\overline{\theta}} \pi_D(\theta)\, p(\theta)\, d\theta,$$

and define, for $\pi_-^e \in (\pi_D(\underline{\theta}), \pi_-^{e*})$,

$$T_1\big(\pi_-^e\big) := \min\big\{b \in \Theta | H(b) = \pi_-^e\big\},$$

and, for $\pi_-^e \in (\pi_-^{e*}, \pi_D(\overline{\theta}))$,

$$T_2\big(\pi_-^e\big) := \min\big\{b \in \Theta | J(b) = \pi_-^e\big\}.$$

They are well defined because both H and J are continuous and onto functions.

LEMMA 11. *Suppose that $F \in \mathcal{V}(\Pi)$. Then both T_1 and T_2 are strictly increasing functions, and the composite functions $\pi_D \circ T_1$ and $\pi_D \circ T_2$ are differentiable and their derivatives are $[1 - P(T_1(\pi_-^e))]^{-1}$ and $P(T_2(\pi_-^e))^{-1}$, respectively. If F is strictly concave and differentiable, T_1 and T_2 are also continuous, and the functions $\pi_D \circ T_1$ and $\pi_D \circ T_2$ are C^1.*

PROOF. We show this for T_1 only. Note that $H(T_1(\pi_-^e)) = \pi_-^e$ for all $\pi_-^e \in (\pi_D(\underline{\theta}), \pi_-^{e*})$. Therefore, the composite function $H \circ T_1$ must be strictly increasing and continuous. Because H is a continuous, weakly increasing function, it follows that T_1 is strictly increasing. If F is strictly concave and differentiable, H is a strictly increasing and continuous function, and, therefore, T_1 must also be strictly increasing and continuous.

Next we show that the composite function $h := \pi_D \circ T_1$ is differentiable. Because $H(T_1(\pi_-^e)) = \pi_-^e$, the right and left derivatives of $H \circ T_1$ are equal to 1 and satisfy

$$1 = (H \circ T_1)'\big(\pi_-^e+\big) = \big[1 - P\big(T_1\big(\pi_-^e\big)\big)\big] \times h'\big(\pi_-^e+\big),$$
$$1 = (H \circ T_1)'\big(\pi_-^e-\big) = \big[1 - P\big(T_1\big(\pi_-^e\big)\big)\big] \times h'\big(\pi_-^e-\big).$$

It follows that $h'(\pi_-^e+) = h'(\pi_-^e-) = [1 - P(T_1(\pi_-^e))]^{-1}$. If F is strictly concave and differentiable, T_1 is a continuous function and, therefore, h is a C^1 function. □

LEMMA 12. *For each π_-^e, there is a unique solution to problem (P2) and it has the form described in Proposition 3.*

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

$$\pi_*(\theta) = \begin{cases} \pi_D(\theta) & \forall \theta \in [\underline{\theta}, \theta^\#), \\ \pi_D(\theta^\#) & \forall \theta \in [\theta^\#, \overline{\theta}] \end{cases} \tag{21}$$

for some $\theta^\#$ and $\pi_-^e = \int_{\underline{\theta}}^{\overline{\theta}} \pi_*(\theta) p(\theta)\, d\theta$. Clearly, for any $\pi_-^e \le \pi_D(\underline{\theta})$, π_* has to be constant and it satisfies $\pi_*(\overline{\theta}) = \pi_-^e$ for all θ.

Consider $\pi_-^e \in (\pi_D(\underline{\theta}), \pi_-^{e*})$. Then either π_* is constant and equal to π_-^e or it has the form in (21) with $\theta^\# = T_1(\pi_-^e)$. We show that a constant π_* is not a solution. When π has the form in (21) with $\theta^\# = T_1(\pi_-^e)$,

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

The first term is strictly positive. The second term is strictly positive, because $\pi_D(\theta) \ge \pi_D(\theta^\#) > \pi_-^e$ for all $\theta \ge \theta^\#$ and $\tilde{R}$ is strictly concave. This proves that there is a unique solution to problem (P2) for each $\pi_-^e < \pi_-^{e*}$, and it has the form described in Proposition 3. The proof for $\pi_-^e > \pi_-^{e*}$ is analogous. □

Clearly $\pi_*(\theta; \pi_-^e)$ is continuous in (θ, π_-^e). Therefore, the composite functions $x_* = g_x \circ \pi_*$, $\pi_*^e = g_{\pi^e} \circ \pi_*$, and $w_* = F \circ g_{\pi^e} \circ \pi_*$ are also continuous. This implies Lemma 3(ii).

A.3.4 *Proof of Lemma 3(i)* Note that, denoting $U(\theta) = \tilde{R}(\pi_*(\theta), \theta)$,

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

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, hencem

$$(\mathbb{T}F)'\big(\pi_-^e\big) = \tilde{R}_\pi\big(\pi_-^e, \underline{\theta}\big) + \int_{\underline{\theta}}^{\overline{\theta}} \frac{1 - P(\theta)}{p(\theta)} A_{\theta\pi}\big(\pi_-^e, \theta\big) p(\theta)\, d\theta.$$

The right-hand side is continuous because $\tilde{R}$ is C^1 and A is C^2. It is also strictly decreasing because $\tilde{R}_\pi$ is strictly decreasing and $A_{\theta\pi}$ is nonincreasing in π. 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 right-hand

side 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 $\pi_-^e \le \pi_D(\underline{\theta})$ and equals $\int_{\underline{\theta}}^{\overline{\theta}}((1-P(\theta))/p(\theta))A_{\theta\pi}(\pi_D(\underline{\theta}), \theta)p(\theta)\,d\theta$ 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

$$\mathbb{T}F(\pi_-^e) = U(\underline{\theta}) + \int_{\underline{\theta}}^{T_1(\pi_-^e)} \frac{1-P(\theta)}{p(\theta)} \tilde{R}_\theta\big(\pi_D(\theta); \theta\big)p(\theta)\,d\theta + \int_{T_1(\pi_-^e)}^{\overline{\theta}} \frac{1-P(\theta)}{p(\theta)} \tilde{R}_\theta\big(\pi_D\big(T_1(\pi_-^e)\big); \theta\big)p(\theta)\,d\theta.$$

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

$$\begin{aligned} &h'(\pi_-^e)\int_{T_1(\pi_-^e)}^{\overline{\theta}} \frac{1-P(\theta)}{p(\theta)} A_{\theta\pi}\big(\pi_D\big(T_1(\pi_-^e)\big); \theta\big)p(\theta)\,d\theta \\ &= \frac{1}{1-P\big(T_1(\pi_-^e)\big)}\int_{T_1(\pi_-^e)}^{\overline{\theta}} \frac{1-P(\theta)}{p(\theta)} A_{\theta\pi}\big(\pi_D\big(T_1(\pi_-^e)\big); \theta\big)p(\theta)\,d\theta \\ &= \int_{T_1(\pi_-^e)}^{\overline{\theta}} \frac{1-P(\theta)}{p(\theta)} A_{\theta\pi}\big(\pi_D\big(T_1(\pi_-^e)\big); \theta\big)\frac{p(\theta)}{1-P\big(T_1(\pi_-^e)\big)}\,d\theta. \end{aligned}$$

The last expression on the right-hand side is strictly decreasing in π_-^e because $A_{\theta\pi}$ is nonincreasing in its first argument and $A_{\theta\pi}(\pi_D(T_1(\pi_-^e)); \theta) \times (1-P(\theta))/p(\theta)$ is strictly decreasing in θ for a given π_-^e. Again the analogous equations obtain for the right derivative of $\mathbb{T}F$ at $\pi_-^e = \pi_D(\underline{\theta})$ and for the left derivative of $\mathbb{T}F$ at $\pi_-^e = \pi_-^{e*}$, and their values are, respectively, $\int_{\underline{\theta}}^{\overline{\theta}}((1-P(\theta))/p(\theta))A_{\theta\pi}(\pi_D(\underline{\theta}), \theta)p(\theta)\,d\theta$ and 0.

Taken together, the first derivative of $\mathbb{T}F$ exists 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 $\pi_-^e \ge \pi_-^{e*}$, one can show that the first derivative of $\mathbb{T}F$ exists 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.

Proposition 1 follows from Lemma 3. Together with Proposition 1, Lemma 9 implies Proposition 2, and Lemmas 11 and 12 imply Proposition 3.

A.4 *Proof of Proposition 5*

It is clear that $\overline{W}$ is maximized at 0 and, therefore, that $\pi_-^{e*} = 0$.

Consider $\pi_-^e \le \pi_D(\underline{\theta})$. The derivative of $\overline{W}'$ at π_-^e is given by

$$\overline{W}'(\pi_-^e) = \int_{\underline{\theta}}^{\overline{\theta}} \big\{A_\pi(\pi_-^e; \theta) + S'(\pi_-^e)\big\}p(\theta)\,d\theta.$$

The envelope condition implies $S'(\pi) = \overline{W}'(g_{\pi^e}(\pi))$. When $A(\pi;\theta) = -(\pi-\theta)^2/2$ and $E[\theta] = 0$, this implies

$$\overline{W}'(\pi_-^e) = \int_{\underline{\theta}}^{\overline{\theta}} \{-\pi_-^e + \theta + \overline{W}'(g_{\pi^e}(\pi_-^e))\} p(\theta)\, d\theta = -\pi_-^e + \overline{W}'(g_{\pi^e}(\pi_-^e)).$$

Because $\pi_-^e \le \pi_D(\underline{\theta}) < \pi_-^{e*} = 0$, we obtain $\overline{W}'(\pi_-^e) > \overline{W}'(g_{\pi^e}(\pi_-^e))$. From the strict concavity of $\overline{W}$, it follows that $\pi_-^e < g_{\pi^e}(\pi_-^e)$. Using the same line of arguments, one can show $\pi_-^e > g_{\pi^e}(\pi_-^e)$ if $\pi_-^e \ge \pi_D(\overline{\theta})$.

It follows that $\pi_D(\underline{\theta}) < g_{\pi^e}(\pi_D(\underline{\theta}))$ and $\pi_D(\overline{\theta}) > g_{\pi^e}(\pi_D(\overline{\theta}))$. Therefore, for all θ and $\pi_-^e \in (\pi_D(\underline{\theta}), \pi_D(\overline{\theta}))$, we have $\pi^e(\theta;\pi_-^e) \in [g_{\pi^e}(\pi_D(\underline{\theta})), g_{\pi^e}(\pi_D(\overline{\theta}))] \subset (\pi_D(\underline{\theta}), \pi_D(\overline{\theta}))$.

A.5 *Proof of Proposition 6*

Let

$$V(\theta, m_-) = \max_{(\pi,x,m)\in\Pi\times X\times\Pi} R(\pi, x, \theta) + \beta\overline{W}(m) \tag{22}$$

subject to $\pi = \kappa x + \beta m$ and $\pi \in \overline{\Gamma}(m_-)$. Proposition 4 then implies that the maximization problem on the right-hand side is equivalent to the problem in (19) with the additional constraint $\pi \in \overline{\Gamma}(m_-)$. Because the objective function in (19) is strictly concave in inflation, the optimal inflation choice, given (θ, m_-), is (i) $\pi_D(\theta)$ if $\pi_D(\theta) \in \overline{\Gamma}(m_-)$, (ii) the smallest element of $\overline{\Gamma}(m_-)$ if $\pi_D(\theta) \le \pi$ for all $\pi \in \overline{\Gamma}(m_-)$, and (iii) the largest element of $\overline{\Gamma}(m_-)$ if $\pi_D(\theta) \ge \pi$ for all $\pi \in \overline{\Gamma}(m_-)$. This implies that π_* solves the problem in (19) with the additional constraint $\pi \in \overline{\Gamma}(m_-)$. Proposition 4 then implies that, for each (θ, m_-), $(\pi_*(\theta, m_-), x_*(\theta, m_-), \pi_*^e(\theta, m_-))$ solves the problem in (22).

Recall that $\overline{W}$ is the fixed point of $\mathbb{T}$. Hence,

$$\overline{W}(m_-) = \mathbb{E}\big[R\big(\pi_*(\theta, m_-), x_*(\theta, m_-), \theta\big) + \beta\overline{W}\big(\pi_*^e(\theta, m_-)\big)\big]$$

and, for the function V we defined above, $\overline{W}(m_-) = \mathbb{E}[V(\theta, m_-)]$ holds. Because for any m, $\pi_*^e(\theta, m) = \mathbb{E}[\pi_*(\theta', m)]$, the first constraint in the problem in (22) can be replaced by $\pi = \kappa x + \beta\mathbb{E}[\pi_*(\theta', m)]$. Therefore, the problem in (22) is identical to

$$V(\theta, m_-) = \max_{(\pi,x,m)\in\Pi\times X\times\Pi} R(\pi, x, \theta) + \beta\mathbb{E}\big[V\big(\theta', m\big)\big]$$

subject to $\pi = \kappa x + \beta\mathbb{E}[\pi_*(\theta', m)]$ and $\pi \in \overline{\Gamma}$. This establishes that the policy functions (π_*, x_*, π_*^e) and the value function V in (22) constitute an MPE given $\overline{\Gamma}$.

Appendix B: Extending the results in Athey et al. (2005)

Problem (P1) is given by

$$\max_{\pi(\cdot),\delta(\cdot)} \int_{\underline{\theta}}^{\overline{\theta}} \big[\tilde{R}\big(\pi(\theta);\theta\big) + \delta(\theta)\big] p(\theta)\, d\theta$$

subject to

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

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

$$\delta(\theta) \leq 0 \quad \forall \theta.$$

Athey et al. (2005) show that a solution to this problem satisfies (i) $\delta(\theta) = 0$ for all θ and (ii) that $\pi(\cdot)$ is continuous, assuming that $\pi(\cdot)$ is a piecewise C^1 function.[21] They also show that $\pi(\cdot)$ is either constant or a truncated version of the "static best response," $\pi_D(\cdot)$ ($\pi_D(\theta) = \arg\max_\pi \tilde{R}(\pi;\theta)$ for all θ). In their setting, $\tilde{R}$ is a function that is exogenously specified and, therefore, by making certain assumptions on $\tilde{R}$, the static best response is shown to be C^1. Hence, a truncated version of the static best response is piecewise C^1. In our setting, $\tilde{R}$ is an endogenous object the properties of which depend on those of a value function. If we were to assume that $\pi(\cdot)$ is a piecewise C^1 function, we generally need a value function to be C^2 so as to guarantee that $\pi_D(\cdot)$ is C^1, and it requires much stronger assumptions.

Instead, we extend the results of AAK so that we do not have to assume that $\pi(\cdot)$ is a piecewise C^1 function.

B.1 *Preliminaries*

To generalize AAK's proofs of their Lemma 1–3, it is useful to establish key equations they use in the proofs under the set of assumptions in Lemma 5. We restate these assumptions below.

Assumption 5. *For any θ, $\tilde{R}(\pi;\theta)$ is a strictly concave function of π and $\tilde{R}_\theta$ exists.*

Assumption 6 (Single-crossing condition). *The function $\tilde{R}_{\theta\pi}$ exists and $\tilde{R}_{\theta\pi}(\pi;\theta) > 0$ for any (π,θ).*

Assumption 7 (Monotone hazard condition). *For any nondecreasing $\pi(\cdot)$,*

$$\frac{1-P(\theta)}{p(\theta)}\tilde{R}_{\theta\pi}\big(\pi(\theta);\theta\big) \quad \textit{is strictly decreasing in } \theta$$

and

$$\frac{P(\theta)}{p(\theta)}\tilde{R}_{\theta\pi}\big(\pi(\theta);\theta\big) \quad \textit{is strictly increasing in } \theta.$$

Assumption 8 (Milgrom–Segal condition). *The return function satisfies the following properties:*

[21] In their setting, the return function $\tilde{R}$ depends on average inflation, π_{-}^{e}, and it is a choice variable. However, their results hold irrespective of this difference.

(i) For any π, $\tilde{R}(\pi;\cdot)$ is absolutely continuous.

(ii) There is an integrable function $b:\Theta\to\mathbb{R}_+$ such that $|\tilde{R}_\theta(\pi;\theta)|<b(\theta)$ for all π and almost all θ.

It follows that the static best response is a weakly increasing, continuous function.

LEMMA 13. *Under Assumptions 5, 6, and 8, $\max_\pi \tilde{R}(\pi;\theta)$ has a unique solution for each θ, and the policy function $\pi_D(\cdot)$ is weakly increasing and continuous.*

The proof is essentially identical to that of Lemma 8 and is omitted.

Let $U(\theta):=\tilde{R}(\pi(\theta);\theta)+\delta(\theta)$ be the utility of type θ. Then it has an integral representation.

LEMMA 14 (Integral representation (Theorem 2 in Milgrom and Segal (2002))).

Let (π,δ) be an incentive-compatible mechanism. Then under Assumptions 6 and 8, U is absolutely continuous and, for all θ,

$$U(\theta)=U(\underline{\theta})+\int_{\underline{\theta}}^{\theta}\tilde{R}_\theta\big(\pi(z);z\big)\,dz \tag{24}$$

and

$$U(\theta)=U(\overline{\theta})-\int_{\theta}^{\overline{\theta}}\tilde{R}_\theta\big(\pi(z);z\big)\,dz.$$

COROLLARY 3. *Let (π,δ) be an incentive-compatible mechanism. Then under Assumptions 6 and 8,*

$$\lim_{\theta'\uparrow\theta}\tilde{R}\big(\pi(\theta');\theta\big)+\delta\big(\theta'\big)=\lim_{\theta'\downarrow\theta}\tilde{R}\big(\pi(\theta');\theta\big)+\delta\big(\theta'\big)=U(\theta) \tag{25}$$

for all θ.

PROOF. First suppose that π is continuous at θ. Because U is absolutely continuous and $\delta(\theta)=U(\theta)-\tilde{R}(\pi(\theta);\theta)$, it follows that δ is also continuous at θ. Then (25) holds at θ.

Next suppose that π is discontinuous at θ. Let $\theta'<\theta$. Then

$$\begin{aligned}\tilde{R}\big(\pi(\theta');\theta\big)+\delta\big(\theta'\big)&=\tilde{R}\big(\pi(\theta');\theta'\big)+\int_{\theta'}^{\theta}\tilde{R}_\theta\big(\pi(\theta');z\big)\,dz+\delta\big(\theta'\big)\\&=U\big(\theta'\big)+\int_{\theta'}^{\theta}\tilde{R}_\theta\big(\pi(\theta');z\big)\,dz,\end{aligned}$$

because $\tilde{R}(\pi,\cdot)$ is absolutely continuous and differentiable. Therefore,

$$\begin{aligned}\big|\tilde{R}\big(\pi(\theta'),\theta\big)+\delta\big(\theta'\big)-U(\theta)\big|&\le\big|U\big(\theta'\big)-U(\theta)\big|+\left|\int_{\theta'}^{\theta}\tilde{R}_\theta\big(\pi(\theta'),z\big)\,dz\right|\\&\le\big|U\big(\theta'\big)-U(\theta)\big|+\int_{\theta'}^{\theta}b(z)\,dz.\end{aligned}$$

Because U is continuous and b is an integrable function, the right-hand side converges to zero as $\theta' \uparrow \theta$. Using the same argument for $\theta' > \theta$, it follows that (25) holds at θ. □

Finally, the incentive-compatibility constraints can be replaced with a monotonicity condition for $\pi(\cdot)$ and an integral representation.

LEMMA 15. *The set (π, δ) is incentive compatible if and only if (i) π is nondecreasing in θ and (ii) (24) holds for all θ.*

PROOF. The only-if part is standard. To show the if part, we assume that (i) and (ii) hold. Pick any θ and $\theta' \neq \theta$. Without loss of generality, assume $\theta' > \theta$. Then

$$\begin{aligned}\tilde{R}\big(\pi(\theta'), \theta\big) + \delta(\theta') &= U(\theta') - \int_{\theta}^{\theta'} \tilde{R}_\theta\big(\pi(\theta'); z\big)\, dz \\ &= U(\theta) + \int_{\theta}^{\theta'} \tilde{R}_\theta\big(\pi(z); z\big)\, dz - \int_{\theta}^{\theta'} \tilde{R}_\theta\big(\pi(\theta'); z\big)\, dz \\ &= U(\theta) + \int_{\theta}^{\theta'} \big(\tilde{R}_\theta\big(\pi(z); z\big) - \tilde{R}_\theta\big(\pi(\theta'); z\big)\big)\, dz.\end{aligned}$$

Condition (i) implies $\pi(\theta') \geq \pi(z)$ for all $z \leq \theta'$. Together with the single-crossing condition, $\tilde{R}_\theta(\pi(z); z) - \tilde{R}_\theta(\pi(\theta'); z) \leq 0$ for all $z \in [\theta, \theta']$. It follows that $\tilde{R}(\pi(\theta'); \theta) + \delta(\theta') \leq U(\theta)$. Because θ and θ' are arbitrarily chosen, this implies that (π, δ) is incentive compatible. □

B.2 *Proof of Lemma 5*

Now we generalize AAK's Lemmas 1–3. Throughout, we assume that Assumptions 5, 6, 7, and 8 hold.

LEMMA 16 (AAK's Lemma 1). *Let $(\pi(\cdot), \delta(\cdot))$ be a mechanism such that all the constraints in problem (P1) hold and that $\pi(\cdot)$ is increasing on some interval (θ_1, θ_2).[22] Then the up variation and the down variation both increase the objective function in problem (P1).*

AAK's proof of Lemma 1 remains valid because it only uses the integral representation.

Because AAK assume that $\pi(\cdot)$ is piecewise C^1, it automatically follows that each discontinuity point of $\pi(\cdot)$ is isolated. Therefore, for any jump discontinuity point of $\pi(\cdot)$, they can choose two intervals on either side of it so that both $\pi(\cdot)$ and $\delta(\cdot)$ are continuous in each of them. AAK's proofs of their Lemmas 2 and 3 use this implication. Here, although the incentive-compatibility condition implies that $\pi(\cdot)$ is weakly increasing and, therefore, that it has at most countably many jump discontinuity points, it is still possible that $\pi(\cdot)$ jumps on a dense subset of some interval in Θ. However, in the optimal mechanism, each discontinuity point of $\pi(\cdot)$ is isolated, as the next lemma shows.

[22]The function π is said to be *increasing on an interval* (θ_1, θ_2) if and only if (i) π is nondecreasing on (θ_1, θ_2) and (ii) there is some $\tilde{\theta}$ in this interval such that $\pi(\theta) < \tilde{\pi}$ for all $\theta \in (\theta_1, \tilde{\theta})$ and $\pi(\theta) > \tilde{\pi}$ for all $\theta \in (\tilde{\theta}, \theta_2)$, where $\tilde{\pi} := \int_{\theta_1}^{\theta_2} \pi(\theta)\, p(\theta)\, d\theta / (P(\theta_2) - P(\theta_1))$ is the conditional mean of π on this interval.

LEMMA 17. *In an optimal mechanism, each discontinuity point of $\pi(\cdot)$ is isolated.*

PROOF. Suppose to the contrary that there is an interval $(a, b) \subset [\underline{\theta}, \overline{\theta}]$ of which discontinuity points of π form a dense subset. Then π is strictly increasing on (a, b) and, therefore, it is impossible to have $\pi(\theta) = \pi_D(\theta)$ for all $\theta \in (a, b)$ because the static best response π_D is a continuous function. It follows that there is $\tilde{\theta} \in (a, b)$ such that either (i) $\lim_{\theta \downarrow \tilde{\theta}} \pi(\theta) > \pi_D(\tilde{\theta})$ or (ii) $\lim_{\theta \uparrow \tilde{\theta}} \pi(\theta) < \pi_D(\tilde{\theta})$.

Let $\tilde{\theta} \in (a, b)$ be such that condition (i) holds. Then we can find $c > \tilde{\theta}$ and $d > c$ in a neighborhood of $\tilde{\theta}$ such that $\pi(c) > \pi_D(\theta)$ for all $\theta \in (c, d)$. Because π is strictly increasing on (c, d), we have, for all $\theta, \theta' \in (c, d)$,

$$\theta > \theta' \quad \Rightarrow \quad \pi(\theta) > \pi\big(\theta'\big) > \pi(c) > \pi_D(\theta)$$

and, thus,

$$\tilde{R}\big(\pi(\theta); \theta\big) < \tilde{R}\big(\pi\big(\theta'\big); \theta\big).$$

The incentive-compatibility constraint then requires $\delta(\theta) > \delta(\theta')$ for any $\theta, \theta' > \theta$ in (c, d), i.e., $\delta(\cdot)$ is strictly increasing on (c, d). Thus we can find a subinterval of (c, d) such that π is strictly increasing and that $\delta(\cdot)$ is strictly increasing and is bounded away from the upper bound of 0 by some $\epsilon > 0$. Using the line of arguments in AAK's Lemma 2, either the up or down variation applied to such a subinterval is feasible, which contradicts the optimality of (π, w).

Next consider $\tilde{\theta} \in (a, b)$ for which condition (ii) holds. Then we can find $d < \tilde{\theta}$ and $c < d$ in a neighborhood of $\tilde{\theta}$ such that $\pi(d) < \pi_D(\theta)$ for all $\theta \in (c, d)$. Because π is strictly increasing on (c, d), we have, for all $\theta, \theta' \in (c, d)$,

$$\theta < \theta' \quad \Rightarrow \quad \pi(\theta) < \pi\big(\theta'\big) < \pi(d) < \pi_D(\theta)$$

and, thus,

$$\tilde{R}\big(\pi(\theta); \theta\big) < \tilde{R}\big(\pi\big(\theta'\big); \theta\big).$$

The incentive-compatibility constraint then requires $\delta(\theta) > \delta(\theta')$ for any $\theta, \theta' < \theta$ in (c, d), i.e., $\delta(\cdot)$ is strictly decreasing on (c, d). There exists a subinterval of (c, d) such that π is strictly increasing and that $\delta(\cdot)$ is strictly decreasing and is bounded away from 0 by some $\epsilon > 0$. Again, either the up or down variation applied to such a subinterval is feasible, which contradicts the optimality of (π, δ). □

LEMMA 18 (AAK's Lemma 2 extended). *In an optimal mechanism, $\delta(\cdot)$ is continuous except on at most countable, isolated jump discontinuities, and for any subinterval (θ_1, θ_2) in which $\delta(\cdot)$ is continuous, $\delta(\cdot)$ is constant.*

PROOF. AAK's proof of Lemma 2 remains mostly valid except that it uses the assumption that $\pi(\cdot)$ is piecewise C^1 when (and only when) showing that $\pi(\cdot)$ is increasing on some interval if $\delta(\cdot)$ is continuous but not constant over it. Therefore, we show that $\pi(\cdot)$ is increasing on some interval if $\delta(\cdot)$ is continuous but not constant over it.

Note that the integral representation implies

$$\delta(\theta) = U(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} \tilde{R}_\theta\big(\pi(z); z\big)\, dz - \tilde{R}\big(\pi(\theta); \theta\big).$$

From the previous lemma, all discontinuity points of $\pi(\cdot)$ are isolated and $\pi(\cdot)$ is continuous except for the discontinuity points. It follows that $\delta(\cdot)$ is continuous except on jump discontinuity points that are isolated.

Suppose that there is an interval in which $\delta(\cdot)$ is continuous but not constant. Then there exists an interval over which $\delta(\cdot)$ is continuous and either strictly increasing or strictly decreasing. This further implies that there is a nonempty subinterval (θ_1, θ_2) of it such that $\delta(\theta) \leq -\epsilon$ for some $\epsilon > 0$ for all $\theta \in (\theta_1, \theta_2)$. Therefore, from the integral representation,

$$0 \neq \delta\big(\theta'\big) - \delta(\theta) = \int_{\theta}^{\theta'} \tilde{R}_\theta\big(\pi(z), z\big)\, dz - \big(\tilde{R}\big(\pi\big(\theta'\big); \theta'\big) - \tilde{R}\big(\pi(\theta); \theta\big)\big)$$

for any θ and $\theta' > \theta$ in this interval. It follows that δ cannot be constant over any subinterval of (θ_1, θ_2) and that π is continuous on (θ_1, θ_2).[23] Therefore, π is strictly increasing on (θ_1, θ_2) and, therefore, is increasing on an interval (θ_1, θ_2) as defined in AAK. □

Lemma 19 (AAK's Lemma 3 extended). *In the optimal mechanism, both $\pi(\cdot)$ and $\delta(\cdot)$ are continuous.*

Proof. AAK's proof of Lemma 3 first defines two types of possible discontinuities and then rules out each type. In their first type of discontinuity, $\pi(\cdot)$ and potentially $\delta(\cdot)$ jump at a point $\tilde{\theta}$, and both take constant values in some intervals $(\theta_1, \tilde{\theta})$ and $(\tilde{\theta}, \theta_2)$ on either side of $\tilde{\theta}$. In their second type of discontinuity, $\pi(\cdot)$ and $\delta(\cdot)$ both jump at $\tilde{\theta}$, and $\pi(\cdot)$ coincides with the static best response in some interval $(\theta_1, \tilde{\theta})$ or $(\tilde{\theta}, \theta_2)$ on either side of the jump point $\tilde{\theta}$.

Because each discontinuity point of $\pi(\cdot)$ is isolated and the first type of discontinuity is independent of the nature of the static best response, the first type of discontinuity is ruled out as in AAK's proof. However, when ruling out the second type of discontinuity, AAK use the strict monotonicity of $\pi(\cdot)$. Their argument is as follows. (A) One can pick up either interval $(\theta_1, \tilde{\theta})$ or $(\tilde{\theta}, \theta_2)$ on which π equals the static best response. (B) Because the static best response is strictly increasing, π is increasing and $\delta(\cdot) < 0$ on that interval. (C) Then either the up variation or the down variation is feasible. It is only part (B) that uses the strict monotonicity of the static best response.

However, the strict monotonicity of the static best response is stronger than is necessary. Let us define the second type of discontinuity slightly differently as follows: $\pi(\cdot)$

[23] Continuity is shown as follows. Observe that

$$\int_{\theta}^{\theta'} \tilde{R}_\theta\big(\pi(z); z\big)\, dz - \big(\delta\big(\theta'\big) - \delta(\theta)\big) = \tilde{R}\big(\pi\big(\theta'\big); \theta'\big) - \tilde{R}\big(\pi(\theta); \theta\big).$$

Because the left-hand side is continuous in θ and θ', so is the right-hand side. Because R is continuous, neither θ nor θ' can be a discontinuity point of π.

and $\delta(\cdot)$ both jump at the point $\tilde{\theta}$, and $\pi(\cdot)$ is strictly increasing in some interval $(\theta_1, \tilde{\theta})$ or $(\tilde{\theta}, \theta_2)$ on either side of the jump point $\tilde{\theta}$. Clearly, if $\tilde{\theta}$ is a discontinuity point of π but does not satisfy this condition, there are $\theta_1 < \tilde{\theta}$ and $\theta_2 > \tilde{\theta}$ such that π is constant on $(\theta_1, \tilde{\theta})$ and $(\tilde{\theta}, \theta_2)$. Therefore, $\tilde{\theta}$ is a discontinuity point of the first type. It follows that any discontinuity point of π is either of the first or the second type.

Suppose that $\tilde{\theta}$ is a discontinuity point of the second type, and that $\pi(\cdot)$ is strictly increasing in $(\theta_1, \tilde{\theta})$ and is flat in $(\tilde{\theta}, \theta_2)$. Because $\delta(\cdot)$ is constant on an interval in which it is continuous and discontinuity points of $\delta(\cdot)$ are isolated, we can choose θ_1 and θ_2 so that $\delta(\cdot)$ is constant at δ_1 over $(\theta_1, \tilde{\theta})$ and at δ_2 over $(\tilde{\theta}, \theta_2)$. Because $\delta(\cdot)$ is flat in $(\theta_1, \tilde{\theta})$, the incentive-compatibility condition implies that for any θ and θ' in this interval,

$$\tilde{R}\big(\pi(\theta); \theta\big) \geq \tilde{R}\big(\pi\big(\theta'\big); \theta\big).$$

Because $\pi(\cdot)$ is continuous and strictly increasing in $(\theta_1, \tilde{\theta})$, it follows that $\pi(\cdot)$ coincides with the static best response on the same interval. Recall that the incentive-compatibility condition implies

$$\lim_{\theta \uparrow \tilde{\theta}} \tilde{R}\big(\pi(\theta); \tilde{\theta}\big) + \delta(\theta) = \lim_{\theta \downarrow \tilde{\theta}} \tilde{R}\big(\pi(\theta); \tilde{\theta}\big) + \delta(\theta).$$

As in AAK, it implies that

$$\tilde{R}\big(\pi(\tilde{\theta}), \tilde{\theta}\big) + \delta_1 = \tilde{R}\Big(x, \lim_{\theta \downarrow \tilde{\theta}} \pi(\theta), \tilde{\theta}\Big) + \delta_2,$$

and that $\delta_1 < \delta_2 \leq 0$. Therefore, either the up variation or the down variation applied to $(\theta_1, \tilde{\theta})$ is feasible. □

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 set for inflation and introduce lotteries/randomization over the set to convexify the problem.

Our numerical implementation is based on problem (P1). Consider a concave function F on Π. For each $\pi_i \in \hat{\Pi}$, let

$$S(\pi_i; F) := \max_{(x, \pi^e): \pi_i = \kappa x + \beta \pi^e} \big\{B(x) + \beta F\big(\pi^e\big)\big\}.$$

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

$$\mathbb{T}^l F\big(\pi_-^e\big) = \max_{\gamma^\pi, \delta} \sum_\theta \hat{p}(\theta) \Big[\sum_{\pi_i \in \hat{\Pi}} \gamma^\pi(\pi_i|\theta) \big\{A(\pi_i) + \theta \pi_i + S(\pi_i; F)\big\} + \delta(\theta) \Big]$$

subject to the feasibility constraint

$$\pi_-^e = \sum_\theta \hat{p}(\theta) \Big(\sum_{\pi_i \in \hat{\Pi}} \gamma^\pi(\pi_i|\theta) \pi_i \Big),$$

the lottery constraints,

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

the incentive-compatibility constraint,

$$\sum_{\pi_i \in \hat{\Pi}} \gamma^{\pi}(\pi_i|\theta)\big\{A(\pi_i) + \theta\pi_i + S(\pi_i; F)\big\} + \delta(\theta)$$
$$\geq \sum_{\pi_i \in \hat{\Pi}} \gamma^{\pi}\big(\pi_i|\theta'\big)\big\{A(\pi_i) + \theta\pi_i + S(\pi_i; F)\big\} + \delta\big(\theta'\big)$$

for all $(\theta, \theta') \in \hat{\Theta}^2$, and the upper-bound constraint

$$\delta(\theta) \leq 0 \quad \forall \theta \in \hat{\Theta}.$$

The operator $\mathbb{T}^l$ satisfies Blackwell's sufficient condition and, thus, is a contraction mapping. The dynamic programming problem is a concave problem, allowing us to apply the method in Fukushima and Waki (2013) to compute its solution. 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:

Step 1. Fix a compact interval $\Pi^e \subset \Pi$ and a finite grid $\hat{\Pi}^e$ on Π^e.

Step 2. Set the initial condition $F_0 = \underline{U}$ for value function iteration.

Step 3. For each $n \geq 1$, take the following actions.

(a) First compute $S(\pi_i; F_{n-1})$ for all $\pi_i \in \hat{\Pi}$.

(b) Then compute $\mathbb{T}^l F_{n-1}(\pi_-^e)$ for all $\pi_-^e \in \hat{\Pi}^e$, and define F_n as the piecewise linear interpolant of $\mathbb{T}F_{n-1}(\pi_-^e)$ on $\hat{\Pi}^e$.

Step 4. Stop if a prespecified stopping criterion is satisfied: $||F_n - F_{n-1}|| < \epsilon$ for some constant ϵ.

Step 5. Use the computed value function F_n to check whether the solution is interior and whether $\delta(\theta) = 0$ for all θ. If these conditions are satisfied, use F_n as an estimate for the true value function.

It follows that $\{F_n\}_{n=0}^{\infty}$ is a nondecreasing sequence of concave, continuous, piecewise linear functions that are bounded by the fixed point of $\mathbb{T}^l$ from above. Therefore, the convergence is guaranteed.

Let $\{\gamma_*^{\pi}(\cdot|\theta, \pi_-^e), \delta_*(\theta, \pi_-^e)\}_{(\theta, \pi_-^e) \in \hat{\Theta} \times \hat{\Pi}^e}$ be the solution at a fixed point $F = F_*$. In our numerical example, we find $\delta_*(\theta, \pi_-^e) = 0$ for all (θ, π_-^e). When plotting the policy functions, we use

$$\pi\big(\theta; \pi_-^e\big) = \sum_{\pi_i \in \hat{\Pi}} \gamma_*^{\pi}\big(\pi_i|\theta, \pi_-^e\big)\pi_i,$$

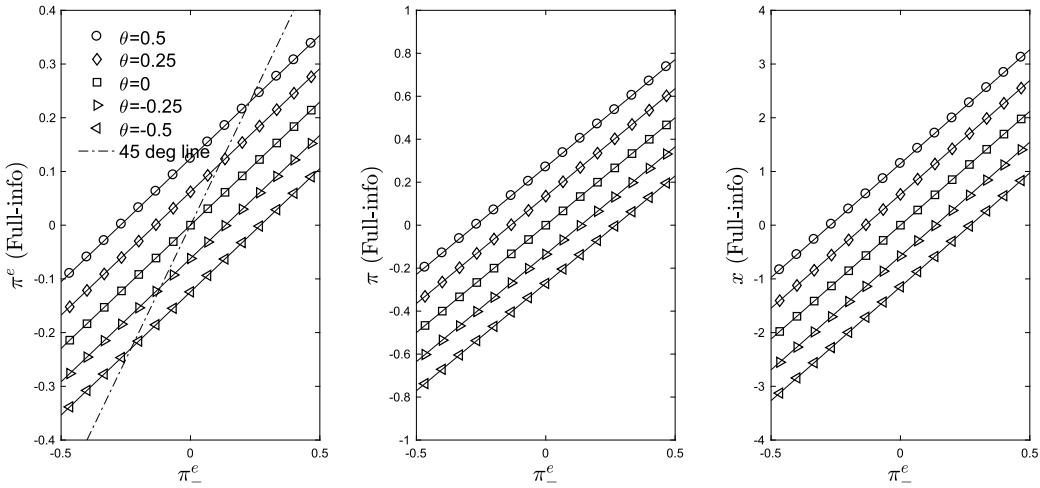

(a) Full-information solution

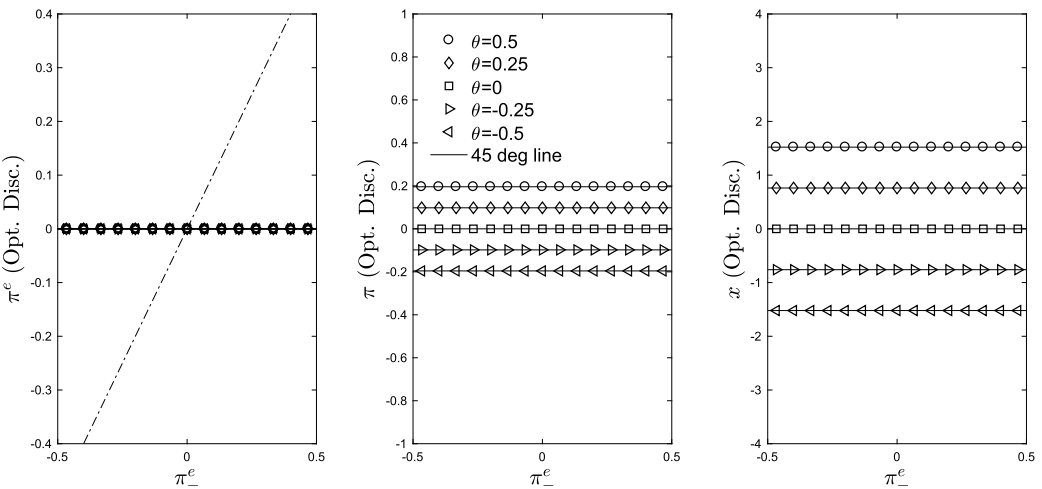

(b) Optimal discretionary policy

FIGURE 7. Benchmarks.

$$x(\theta; \pi^e_-) = \sum_{\pi_i \in \hat{\Pi}} \gamma^{\pi}_*(\pi_i|\theta, \pi^e_-) g_x(\pi_i; F_*),$$

$$\pi^e(\theta; \pi^e_-) = \sum_{\pi_i \in \hat{\Pi}} \gamma^{\pi}_*(\pi_i|\theta, \pi^e_-) g_{\pi^e}(\pi_i; F_*).$$

APPENDIX D: BENCHMARK PROBLEMS

In this appendix, we define two benchmark problems using the quadratic specification in (2). Their solutions are depicted in Figure 7.

D.1 *Full-information problem*

The full-information problem has the recursive formulation

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

subject to $\pi(\theta) = \kappa x(\theta) + \beta \pi^e(\theta)$ 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.

D.2 *Optimal discretionary policy*

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

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

The welfare delivered by this policy is given by

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

References

Abreu, Dilip, David Pearce, and Ennio Stacchetti (1990), "Toward a theory of discounted repeated games with imperfect monitoring." *Econometrica*, 58, 1041–1063. Available at https://ideas.repec.org/a/ecm/emetrp/v58y1990i5p1041-63.html. [1323, 1330]

Alonso, Ricardo and Niko Matouschek (2008), "Optimal delegation." *Review of Economic Studies*, 75, 259–293. Available at https://ideas.repec.org/a/oup/restud/v75y2008i1p259-293.html. [1323, 1338]

Amador, Manuel and Kyle Bagwell (2013), "The theory of optimal delegation with an application to tariff caps." *Econometrica*, 81, 1541–1599. Available at https://ideas.repec.org/a/ecm/emetrp/v81y2013i4p1541-1599.html. [1323, 1332, 1338]

Amador, Manuel, Iván Werning, and George-Marios Angeletos (2006), "Commitment vs. flexibility." *Econometrica*, 74, 365–396. Available at https://ideas.repec.org/a/ecm/emetrp/v74y2006i2p365-396.html. [1323]

Athey, Susan, Andrew Atkeson, and Patrick J. Kehoe (2005), "The optimal degree of discretion in monetary policy." *Econometrica*, 73, 1431–1475. Available at https://ideas.repec.org/a/ecm/emetrp/v73y2005i5p1431-1475.html. [1320, 1345, 1355, 1356]

Athey, Susan, Kyle Bagwell, and Chris Sanchirico (2004), "Collusion and price rigidity." *Review of Economic Studies*, 71, 317–349. Available at https://ideas.repec.org/a/oup/restud/v71y2004i2p317-349.html. [1323, 1331]

Atkeson, Andrew (1991), "International lending with moral hazard and risk of repudiation." *Econometrica*, 59, 1069–1089. Available at https://ideas.repec.org/a/ecm/emetrp/v59y1991i4p1069-89.html. [1323]

Barro, Robert J. and David B. Gordon (1983), "Rules, discretion and reputation in a model of monetary policy." *Journal of Monetary Economics*, 12, 101–121. Available at https://ideas.repec.org/a/eee/moneco/v12y1983i1p101-121.html. [1326]

Bernanke, Ben S., Thomas Laubach, Frederic S. Mishkin, and Adam S. Posen (2001), *Inflation Targeting: Lessons From the International Experience*. Princeton University Press, Princeton, New Jersey. [1323, 1345]

Bernanke, Ben S. and Frederic S. Mishkin (1997), "Inflation targeting: A new framework for monetary policy?" *Journal of Economic Perspectives*, 11, 97–116. Available at https://ideas.repec.org/a/aea/jecper/v11y1997i2p97-116.html. [1320]

Calvo, Guillermo A. (1983), "Staggered prices in a utility-maximizing framework." *Journal of Monetary Economics*, 12, 383–398. Available at https://ideas.repec.org/a/eee/moneco/v12y1983i3p383-398.html. [1324]

Canzoneri, Matthew B. (1985), "Monetary policy games and the role of private information." *American Economic Review*, 75, 1056–1070. Available at https://ideas.repec.org/a/aea/aecrev/v75y1985i5p1056-70.html. [1320, 1322, 1323]

Clarida, Richard, Jordi Gali, and Mark Gertler (1999), "The science of monetary policy: A new Keynesian perspective." *Journal of Economic Literature*, 37, 1661–1707. Available at https://ideas.repec.org/a/aea/jeclit/v37y1999i4p1661-1707.html. [1340]

Fukushima, Kenichi and Yuichiro Waki (2013), "A polyhedral approximation approach to concave numerical dynamic programming." *Journal of Economic Dynamics and Control*, 37, 2322–2335. Available at https://ideas.repec.org/a/eee/dyncon/v37y2013i11p2322-2335.html. [1362]

Galì, Jordi (2008), *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton University Press, Princeton, New Jersey. [1324]

Gertler, Mark and John Leahy (2008), "A Phillips curve with an Ss foundation." *Journal of Political Economy*, 116, 533–572. Available at https://ideas.repec.org/a/ucp/jpolec/v116y2008i3p533-572.html. [1324]

Green, Edward J. (1987), "Lending and the smoothing of uninsurable income." In *Contractual Arrangement in Intertemporal Trade* (E. C. Prescott and N. Wallace, eds.), 3–25, University of Minnesota Press, Minneapolis, MN. [1328]

Halac, Marina and Pierre Yared (2014), "Fiscal rules and discretion under persistent shocks." *Econometrica*, 82, 1557–1614. Available at https://ideas.repec.org/a/wly/emetrp/v82y2014ip1557-1614.html. [1345]

Holmstrom, Bengt (1977), *On Incentives and Control in Organizations*. Ph.D. Dissertation. Stanford University. [1338]

Holmstrom, Bengt (1984), "On the theory of delegation." In *Bayesian Models in Economic Theory* (Marcel Boyer and Richard Kihlstrom, eds.). Studies in Bayesian Econometrics. North-Holland. [1323, 1338]

Jensen, Henrik (2002), "Targeting nominal income growth or inflation?" *American Economic Review*, 92, 928–956. Available at https://ideas.repec.org/a/aea/aecrev/v92y2002i4p928-956.html. [1323]

Kocherlakota, Narayana R. (1996), "Implications of efficient risk sharing without commitment." *Review of Economic Studies*, 63, 595–609. Available at https://ideas.repec.org/a/oup/restud/v63y1996i4p595-609..html. [1321, 1335]

Kocherlakota, Narayana R. (2015), "Rules versus discretion: A reconsideration." Speech delivered at Korea-America Economic Association Boston, Massachusetts. [1320]

Kurozumi, Takushi (2008), "Optimal sustainable monetary policy." *Journal of Monetary Economics*, 55, 1277–1289. Available at https://ideas.repec.org/a/eee/moneco/v55y2008i7p1277-1289.html. [1323]

Kydland, Finn E. and Edward C. Prescott (1977), "Rules rather than discretion: The inconsistency of optimal plans." *Journal of Political Economy*, 85, 473–491. Available at https://ideas.repec.org/a/ucp/jpolec/v85y1977i3p473-91.html. [1320]

McCallum, Bennett T. (2003), "Comments on paper by athey, atkeson, and kehoe, 'The optimal degree of monetary policy discretion'." Discussion at International Research Forum on Monetary Policy Second Conference, Washington, D.C. [1323]

Milgrom, Paul and Ilya Segal (2002), "Envelope theorems for arbitrary choice sets." *Econometrica*, 70, 583–601. [1330, 1357]

Rotemberg, Julio J. (1982), "Sticky prices in the United States." *Journal of Political Economy*, 90, 1187–1211. Available at https://ideas.repec.org/a/ucp/jpolec/v90y1982i6p1187-1211.html. [1324]

Sleet, Christopher (2001), "On credible monetary policy and private government information." *Journal of Economic Theory*, 99, 338–376. Available at https://ideas.repec.org/a/eee/jetheo/v99y2001i1-2p338-376.html. [1322, 1323, 1326]

Sleet, Christopher (2004), "Optimal taxation with private government information." *Review of Economic Studies*, 71, 1217–1239. Available at https://ideas.repec.org/a/oup/restud/v71y2004i4p1217-1239.html. [1320, 1323]

Stokey, Nancy L., Robert E. Jr. Lucas, and Edward C. Prescott (1989), *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge, Massachusetts, and London, England. [1331]

Svensson, Lars E. O. (1997), "Optimal inflation targets, "Conservative" central banks, and linear inflation contracts." *American Economic Review*, 87, 98–114. Available at https://ideas.repec.org/a/aea/aecrev/v87y1997i1p98-114.html. [1323, 1326]

Woodford, Michael (1999), "Commentary: How should monetary policy be conducted in an era of price stability?" Proceedings—Economic Policy Symposium—Jackson Hole, 277–316. Available at https://ideas.repec.org/a/fip/fedkpr/y1999p277-316.html. [1337]

Woodford, Michael (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, Princeton, New Jersey. [1323, 1324, 1340]

Co-editor Giuseppe Moscarini handled this manuscript.

Manuscript received 12 January, 2016; final version accepted 16 February, 2018; available online 21 February, 2018.