

Federal Reserve Bank of Minneapolis
Research Department Staff Report 326

Revised August 2004

The Optimal Degree of Discretion in Monetary Policy*

Susan Athey

Stanford University
and National Bureau of Economic Research

Andrew Atkeson

University of California, Los Angeles,
Federal Reserve Bank of Minneapolis,
and National Bureau of Economic Research

Patrick J. Kehoe

Federal Reserve Bank of Minneapolis,
University of Minnesota,
and National Bureau of Economic Research

ABSTRACT

How much discretion should the monetary authority have in setting its policy? This question is analyzed in an economy with an agreed-upon social welfare function that depends on the randomly fluctuating state of the economy. The monetary authority has private information about that state. In the model, well-designed rules trade off society's desire to give the monetary authority discretion to react to its private information against society's need to guard against the time inconsistency problem arising from the temptation to stimulate the economy with unexpected inflation. Although this dynamic mechanism design problem seems complex, society can implement the optimal policy simply by legislating an inflation cap that specifies the highest allowable inflation rate. The more severe the time inconsistency problem and the less important is private information, the smaller is the optimal degree of discretion. As either the time inconsistency problem becomes sufficiently severe or private information becomes sufficiently unimportant, the optimal degree of discretion is none.

*The authors thank the editor and the referees for very useful comments, Kathy Rolfe for excellent editorial assistance, and the NSF for generous assistance. The views expressed are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

Suppose that society can credibly impose on the monetary authority rules governing the conduct of monetary policy. How much discretion should be left to the monetary authority in setting its policy? The conventional wisdom from policymakers is that optimal outcomes can be achieved only if some discretion is left in the hands of the monetary authority. Starting with Kydland and Prescott (1977), most of the academic literature has contradicted that view. In summarizing this literature, Taylor (1983) and Canzoneri (1985) argue that when the monetary authority does not have private information about the state of the economy, the debate is settled: there should be no discretion; that is, the best outcomes can be achieved by rules that specify the action of the monetary authority as a function of observables. The unsettled question in this debate is Canzoneri's: What about when the monetary authority does have private information? What, then, is the optimal degree of monetary policy discretion?

To answer this question, we use a model of monetary policy similar to that of Kydland and Prescott (1977) and Barro and Gordon (1983). In our *legislative approach to monetary policy* we suppose that society designs the optimal rules governing the conduct of monetary policy by the monetary authority. The model includes an agreed-upon social welfare function that depends on the random state of the economy. We begin with the assumption that the monetary authority observes the state and individual agents do not. In the context of our model, we say that the monetary authority has *discretion* if its policy is allowed to vary with its private information.¹

The assumption of private information creates a tension between discretion and time inconsistency.² Tight constraints on discretion mitigate the time inconsistency problem in which the monetary authority is tempted to claim repeatedly that the current state of the economy justifies a monetary stimulus to output. However, tight constraints leave little room for the monetary authority to fine tune its policy to its private information. Loose constraints allow the monetary authority to do that fine tuning, but they also allow more room for the monetary authority to stimulate the economy with a surprise inflation. These constraints may vary with publicly observed variables, but the relevant question is, how tight should they be? How much discretion should be allowed?

Our purpose here is to answer this question by finding the constraints on monetary policy

¹Our approach here is different from that in the early literature on rules vs. discretion, as is our notion of discretion. In the early literature it is assumed that society has no mechanism for committing to rules governing monetary policy. As does Taylor (1983), we find the legislative approach appealing for advanced economies.

²For some potential empirical support for the idea that the Federal Reserve possesses some nontrivial private information, see the work of Romer and Romer (2000). As we discuss below, we interpret this private information in our economy along the lines of Sleet and Yeltekin (2003) and Sleet (forthcoming).

that, in the presence of private information, optimally resolve the tension between discretion and time inconsistency. Formally, we cast this problem as a dynamic mechanism design problem. Canzoneri (1985) conjectures that because of the dynamic nature of the problem, the resulting optimal social contract with regard to monetary policy is likely to be quite complex. We find that, in fact, it is quite simple. For a broad class of economies, the optimal social contract is static and can be implemented by setting an *inflation cap*, an upper limit on the permitted inflation rate.

More formally, our model can be described as follows. Each period, the monetary authority observes one of a continuum of possible privately observed states of the economy. These states are i.i.d. over time. In terms of current payoffs, the monetary authority prefers to choose higher inflation when higher values of this state are realized and lower inflation when lower values are realized. Here a *mechanism* specifies what monetary policy is chosen each period as a function of the history of the monetary authority's reports of its private information. We say that a mechanism is *static* if policies depend only on the current report by the monetary authority and *dynamic* if policies depend also on the history of past reports.

Our main technical result is that, as long as a monotone hazard condition is satisfied, the optimal mechanism is static. We also give examples in which this monotone hazard condition fails and the optimal mechanism is dynamic.

We then show that our result on the optimality of a static mechanism implies that the optimal policy has one of two forms: either it has bounded discretion, or it has no discretion. Under *bounded discretion*, there is a cutoff state: for any state less than this, the monetary authority chooses its *static best response*, which is an inflation rate that increases with the state, and for any state greater than this cutoff state, the monetary authority chooses a constant inflation rate. Under *no discretion*, the monetary authority chooses some constant inflation rate regardless of its information.

We then show that we can implement the optimal policy as a repeated static equilibrium of a game in which the monetary authority chooses its policy subject to an inflation cap and in which individual agents' expectations of future inflation do not vary with the monetary authority's policy choice. In general, the inflation cap would vary with observable states, but to keep the model simple, we abstract from observable states, and the inflation cap is a single number. Depending on the realization of the private information, sometimes the cap will bind, and sometimes it will not.

These results imply that the optimal constraints on discretion take the form of an inflation cap. The monetary authority is allowed to choose any inflation rate below this cap but is constrained

from choosing an inflation rate above it. We say that a given inflation cap implies less discretion than another cap if it is more likely to bind. We show that the more severe the time inconsistency problem, the smaller is the optimal degree of discretion. Likewise, we show that the less important is private information, the smaller is the optimal degree of discretion. It is immediate that we can equivalently implement the optimal policy by choosing ranges on acceptable inflation rates. These ranges will decrease as the time inconsistency problem becomes more severe relative to the importance of private information.

Here the rationale for discretion depends in a critical way on the monetary authority having some private information that the other agents in the economy do not have. Of course, if the amount of such private information is thought to be very small in actual economies, relative to time inconsistency problems, then our work argues that in such economies the logical case for a sizable amount of discretion is weak, and the monetary authority should follow a rather tightly specified rule.

One interpretation of our work is that we solve for the optimal inflation targets. As such, our work is related to the burgeoning literature on inflation targeting. (See the work of Cukierman and Meltzer (1986), Bernanke and Woodford (1997), and Faust and Svensson (2001), among many others.) In terms of the practical application of inflation targets, Bernanke and Mishkin (1997) discuss how inflation targets often take the form of ranges or limits on acceptable inflation rates similar to the ranges that we derive. Indeed, our work here provides one theoretical rationale for the type of constrained discretion advocated by Bernanke and Mishkin.

Here we have assumed that the monetary authority maximizes the welfare of society. As such, the monetary authority is viewed as the conduit through which society exercises its will. An alternative approach is to view the monetary authority as an individual or an organization motivated by concerns other than that of society's well-being. If, for example, the monetary authority is motivated in part by its own wages, then, as Walsh (1995) has shown, it is possible to implement the full-information, full-commitment solution. Hence, with such a setup, there are no binding incentive problems in monetary policy to begin with. As Persson and Tabellini (1993) note, there are a host of reasons such contracts are either difficult or impossible to implement, and the main issue for research following this approach is why such contracts are, at best, rarely used.

Our work is related to several literatures. It is related to some work on private information in monetary policy games. (See, for example, that of Backus and Driffill (1985); Ireland (2000); Da

Costa and Werning (2001); Sleet (2001); Angeletos, Hellwig, and Pavan (2003); Sleet and Yeltekin (2003); and Stokey (2003).) The most closely related of these is the work of Sleet (2001), who considers a dynamic general equilibrium model in which the monetary authority sees a noisy signal about future productivity before it sets the money growth rate. Sleet finds that, depending on parameters, the optimal mechanism may be static, as we find here, or it may be dynamic.

Our work is also related to a large literature on dynamic contracting. Our result on the optimality of a static mechanism is quite different from what is typically found in this literature, namely, that static mechanisms are not optimal. (See, for example, Green (1987), Atkeson and Lucas (1992), and Kocherlakota (1996).) We discuss the relation between our work and both of these literatures in more detail after we present our results.

At a technical level, we draw heavily on the literature on recursive approaches to dynamic games. We use the technique of Abreu, Pearce, and Stacchetti (1990), which has been applied to monetary policy games by Chang (1998) and is related to the policy games studied by Phelan and Stacchetti (2001), Albanesi and Sleet (2002), and Albanesi, Chari, and Christiano (2003).

The mechanism design problem that we study, at an abstract level, is related to some work on supporting collusive outcomes in cartels by Athey, Bagwell, and Sanchirico (2004), work on risk-sharing with nonpecuniary penalties for default by Rampini (2003), and work on the tradeoff between flexibility and commitment in savings plans for consumers with hyperbolic discounting by Amador, Angeletos, and Werning (2003). Our paper is both substantively and technically quite different from these papers. We discuss the details of the relation after we present our results.

1. The Economy

A. The Model

Here we describe our simple model of monetary policy. The economy has a monetary authority and a continuum of individual agents. The time horizon is infinite, with periods indexed $t = 0, 1, \dots$

At the beginning of each period, agents choose individual action z_t from some compact set. We interpret z as (the growth rate of) an individual's nominal wage and let x_t denote the (growth of the) average nominal wage. Next, the monetary authority observes the current realization of its private information about the state of the economy. This private information θ_t is an i.i.d., mean 0 random variable with support $\theta \in [\underline{\theta}, \bar{\theta}]$, with a strictly positive density $p(\theta)$ and a distribution function $P(\theta)$. Given this private information, the monetary authority chooses money growth μ_t in

some large compact set $[\underline{\mu}, \bar{\mu}]$.

The monetary authority maximizes a social welfare function $R(x_t, \mu_t, \theta_t)$ that depends on the average nominal wage x_t , the monetary growth rate μ_t , and a privately observed shock θ_t . We interpret θ_t to be private information of the monetary authority regarding the impact of a monetary stimulus on social welfare in the current period. Throughout, we assume that R is strictly concave in μ and twice continuously differentiable.

A leading interpretation of the private information in our economy follows that of Sleet and Yeltekin (2003) and Sleet (forthcoming). Individual consumers 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 consumers do not observe. (Informally, we imagine that this private information takes resources to acquire, so that while agents in the economy feasibly can acquire this information, the costs involved in doing so outweigh the benefits.) When we pose our optimal policy problem as a mechanism design problem, we are presuming that the mechanism designer is a separate agent that has no independent information of its own. We interpret the society's objective as a weighted average of the preferences of the heterogeneous consumers.

As a *benchmark example*, we use this function:

$$(1) \quad R(x_t, \mu_t, \theta_t) = -\frac{1}{2} \left[(U + x_t - \mu_t)^2 + (\mu_t - \alpha\theta_t)^2 \right].$$

We interpret (1) as the reduced form that results from a monetary authority which maximizes a social welfare function that depends on unemployment, inflation, and the monetary authority's private information θ . Each period, inflation π_t is equal to the money growth rate μ_t chosen by the monetary authority. Unemployment is determined by a Phillips curve. The unemployment rate is given by

$$(2) \quad u_t = U + x_t - \mu_t$$

where U is a positive constant, which we interpret as the natural rate of unemployment. Social welfare in period t is a function of u_t and π_t and the shock θ_t . Our benchmark example is derived from a quadratic objective function which has the form

$$(3) \quad -\frac{u_t^2}{2} - \frac{(\pi_t - \alpha\theta_t)^2}{2}$$

similar to that used by Kydland and Prescott (1977) and Barro and Gordon (1983). Using (2) and $\pi_t = \mu_t$ in (3), we obtain (1). Here the monetary authority's private information is about the social

cost of inflation, but we develop our model for general specifications of the social welfare function $R(x_t, \mu_t, \theta_t)$ which subsume (1) as a special case. Notice that in our general formulation we allow for the current payoff to vary with expected inflation, through x_t ; actual inflation, through μ_t ; and the state θ_t . This formulation thus subsumes many other versions of the Kydland-Prescott and Barro-Gordon models in the literature.³

Throughout, a *policy* for the monetary authority in any given period, denoted $\mu(\cdot)$, specifies the money growth rate $\mu(\theta)$ for each level of the shock θ . For any x , we define the *static best response* to be the policy $\mu^*(\theta; x)$ that solves $R_\mu(x, \mu(\theta), \theta) = 0$. We assume that if $x = \int \mu(\theta)p(\theta) d\theta$, then

$$(4) \quad \int R_x(x, \mu(\theta), \theta)p(\theta) d\theta < 0.$$

B. Two Ramsey Benchmarks

Before we analyze the economy in which the monetary authority has private information, it is useful to consider two alternative economies. We think of the optimal policies in these economies as benchmarks for that in the private information economy.

One benchmark, the *Ramsey policy*, denoted $\mu^R(\cdot)$, yields the highest payoff that can be achieved in an economy with full information. The gap between that Ramsey payoff and the payoff in the economy with private information measures the welfare loss due to private information.

The other benchmark, the *expected Ramsey policy*, denoted μ^{ER} , is the optimal policy when the policy is restricted to not depend on private information. In our environment, there is no publicly observed shock to the economy; hence, this policy is a constant. The expected Ramsey policy is a useful benchmark because it is the best policy that can be achieved by a rule which specifies policies as a function only of observables. This policy is analogous to the strict targeting rule discussed by Canzoneri (1985).

For the Ramsey policy benchmark, consider an economy with full information with the following timing scheme. Before the shock θ is realized, the monetary authority commits to a schedule for money growth rates $\mu(\cdot)$. Next, individual agents choose their nominal wages z with associated average nominal wages x . Then the state θ is realized and the money growth rate $\mu(\theta)$ is

³Note that the inflation rate that enters the period t social welfare function is the current inflation rate, that from period $t - 1$ to period t . As noted by a number of authors, this formulation captures the distortions in a sticky price model with multiple sectors. As the current inflation rate rises or falls, the prices of goods in sectors that can currently change prices rise or fall relative to the prices in sectors that cannot. Movements in the current inflation rate thus create resource allocation distortions.

Also, for simplicity, our formulation abstracts from direct costs due to future inflation. One interpretation of this feature is that it captures what happens in the cashless limit of a sticky price model.

implemented. The optimal allocations and policies in this economy solve the *Ramsey problem*:

$$\max_{x, \mu(\cdot)} \int R(x, \mu(\theta), \theta) p(\theta) d\theta$$

subject to $x = \int \mu(\theta) p(\theta) d\theta$. For our example (1), the Ramsey policy is $\mu^R(\theta) = \alpha\theta/2$. Note that the Ramsey policy has the monetary authority choosing a money growth rate which is increasing in its private information. Thus, with full information, it is optimal to have the monetary authority fine tune its policy to the state. This feature of the environment leads to a tension in the economy with private information between allowing the monetary authority discretion for fine tuning and experiencing the resulting time inconsistency problems.

For the other benchmark, consider an economy in which the monetary authority is restricted to choosing money growth μ that does not vary with its private information. The equilibrium allocations and policies in the economy with these constraints solve the *expected Ramsey problem*:

$$(5) \quad \max_{x, \mu} \int R(x, \mu, \theta) p(\theta) d\theta$$

subject to $x = \mu$. For our example (1), the expected Ramsey policy is $\mu^{ER} = 0$.

For our example (1), the Ramsey policy obviously yields strictly higher welfare than does the expected Ramsey policy. More generally, when $R_{\mu\theta}(x, \mu, \theta) > 0$, the Ramsey policy $\mu^R(\cdot)$ is strictly increasing in θ and yields strictly higher welfare than does the expected Ramsey policy.

C. The Dynamic Mechanism Design Problem

To analyze the problem of finding the optimal degree of discretion, we use the tools of dynamic mechanism design. Without loss of generality, we formulate the problem as a direct revelation game. In this problem, society specifies a *monetary policy*, the money growth rate as a function of the history of the monetary authority's reports of its private information. Given the specified monetary policy, the monetary authority chooses a strategy for reporting its private information. Individual agents choose their wages as functions of the history of reports of the monetary authority.

A monetary policy in this environment is a sequence of functions $\{\mu_t(h_t, \hat{\theta}_t) \mid \text{all } h_t, \hat{\theta}_t\}_{t=0}^{\infty}$, where $\mu_t(h_t, \hat{\theta}_t)$ specifies the money growth rate that will be chosen in period t following the history $h_t = (\hat{\theta}_0, \hat{\theta}_1, \dots, \hat{\theta}_{t-1})$ of past reports together with the current report $\hat{\theta}_t$. The monetary authority chooses a reporting strategy $\{m_t(h_t, \theta_t) \mid \text{all } h_t, \theta_t\}_{t=0}^{\infty}$ in period 0, where θ_t is the current realization of private information and $m_t(h_t, \theta_t) \in [\underline{\theta}, \bar{\theta}]$ is the reported private information in t . As is standard, we restrict attention to public strategies, those that depend only

on public histories and the current private information, not on the history of private information.⁴ Also, from the Revelation Principle, we need only restrict attention to truth-telling equilibria in which $m_t(h_t, \theta_t) = \theta_t$ for all h_t and θ_t .

In each period, each agent chooses the action z_t as a function of the history of reports h_t . Since agents are competitive, the history need not include either agents' individual past actions or the aggregate of their past actions.⁵

Each agent chooses nominal wage growth equal to expected inflation. For each history h_t , with monetary policy $\mu_t(h_t, \cdot)$ given, agents set $z_t(h_t)$ equal to expected inflation:

$$(6) \quad z_t(h_t) = \int \mu_t(h_t, \theta) p(\theta) d\theta$$

where we have used the fact that agents expect the monetary authority to report truthfully, so that $m_t(h_t, \theta_t) = \theta_t$. Aggregate wages are defined by $x_t(h_t) = z_t(h_t)$.

The optimal monetary policy maximizes the discounted sum of social welfare:

$$(7) \quad (1 - \beta) \sum_{t=0}^{\infty} \int \beta^t R(x_t(h_t), \mu_t(h_t, \theta_t), \theta_t) p(\theta_t) d\theta_t$$

where the future histories h_t are recursively generated from the choice of monetary policy $\mu_t(\cdot, \cdot)$ in the natural way, starting from the null history. The term $1 - \beta$ normalizes the discounted payoffs to be in the same units as the per-period payoffs.

A *perfect Bayesian equilibrium* of this revelation game is a monetary policy, a reporting strategy, a strategy for wage-setting by agents $\{z_t(\cdot)\}_{t=0}^{\infty}$, and average wages $\{x_t(\cdot)\}_{t=0}^{\infty}$ such that (6) is satisfied in every period following every history h_t , average wages equal individual wages in that $x_t(h_t) = z_t(h_t)$, and the monetary policy is incentive-compatible in the standard sense that, in every period, following every history h_t and realization of the private information θ_t , the monetary authority prefers to report $m_t(h_t, \theta_t) = \theta_t$ rather than any other value $\hat{\theta} \in [\underline{\theta}, \bar{\theta}]$. Note that since average wages $x_t(h_t)$ always equal wages of individual agents $z_t(h_t)$, we need only record average wages from now on.

Note that this definition of a perfect Bayesian equilibrium includes no notion of optimality for society. Instead, it simply requires that in response to a given monetary policy, private agents respond optimally and truth-telling for the monetary authority is incentive-compatible. The

⁴For a discussion of the large class of environments for which this restriction does not alter the set of equilibrium payoffs, see Fudenberg and Tirole (1991).

⁵For details of why this is true, see the work of Chari and Kehoe (1990).

set of perfect Bayesian equilibria outcomes are the set of incentive-compatible outcomes that are implementable by some monetary policy.

The mechanism design problem is to choose a monetary policy, a reporting strategy, and a strategy for average wages, the outcomes of which maximize social welfare (7) subject to the constraint that these strategies are incentive-compatible.

D. A Recursive Formulation

Here we formulate the problem of characterizing the solution to this mechanism design problem recursively. The repeated nature of the model implies that the set of incentive-compatible payoffs that can be obtained from any period t on is the same that can be obtained from period 0. Thus, the payoff to any incentive-compatible outcome for the repeated game can be broken down into payoffs from current actions for the players and continuation payoffs that are themselves drawn from the set of incentive-compatible payoffs. Following this logic, Abreu, Pearce, and Stacchetti (1990) show that the set of incentive-compatible payoffs can be found using a recursive method that we exploit here.

In our environment, this recursive method is as follows. Consider an operator on sets of the following form. Let W be some compact subset of the real line, and let \bar{w} be the largest element of W . The set W may be interpreted as a candidate set of incentive-compatible levels of social welfare. In our recursive formulation, the current actions are average wages x and a report $\hat{\theta} = m(\theta)$ for every realized value of the state θ . For each possible report $\hat{\theta}$, there is a corresponding *continuation payoff* $w(\hat{\theta})$ that represents the discounted utility for the monetary authority from next period on. Clearly, these continuation payoffs cannot vary directly with the privately observed state θ .

We say that the actions x and $\mu(\cdot)$ and the continuation payoff $w(\cdot)$ are *enforceable by* W if

$$(8) \quad w(\hat{\theta}) \in W \text{ for all } \hat{\theta} \in [\underline{\theta}, \bar{\theta}]$$

$$(9) \quad x = \int \mu(\theta)p(\theta) d\theta$$

and the incentive constraints

$$(10) \quad (1 - \beta)R(x, \mu(\theta), \theta) + \beta w(\theta) \geq (1 - \beta)R(x, \mu(\hat{\theta}), \theta) + \beta w(\hat{\theta})$$

are satisfied for all θ and all $\hat{\theta}$, where $\mu(\theta) \in [\underline{\mu}, \bar{\mu}]$. Constraint (8) requires that each continuation payoff $w(\hat{\theta})$ be drawn from the candidate set of incentive-compatible payoffs W , while constraint (9)

requires that average wages equal expected inflation. Constraint (10) requires that for each privately observed state θ , the monetary authority prefer to report the truth θ rather than any other message $\hat{\theta}$. That is, the monetary authority prefers the money growth rate $\mu(\theta)$ and the continuation value $w(\theta)$ rather than a money growth rate $\mu(\hat{\theta})$ and its corresponding continuation value $w(\hat{\theta})$.

The payoff corresponding to $x, \mu(\cdot)$, and $w(\cdot)$ is

$$(11) \quad V(x, \mu(\cdot), w(\cdot)) = \int [(1 - \beta)R(x, \mu(\theta), \theta) + \beta w(\theta)] p(\theta) d\theta.$$

Define the operator T that maps a set of payoffs W into a new set of payoffs

$$(12) \quad T(W) = \left\{ v \mid \text{there exist } x_v, \mu_v(\cdot), w_v(\cdot) \text{ enforceable by } W \right. \\ \left. \text{s.t. } v = V(x_v, \mu_v(\cdot), w_v(\cdot)) \right\}.$$

As demonstrated by Abreu, Pearce, and Stacchetti (1990), the set of incentive-compatible payoffs is the largest set W that is a fixed point of this operator:

$$(13) \quad W^* = T(W^*).$$

For any given candidate set of incentive-compatible payoffs W , we are interested in finding the largest payoff that is enforceable by W , or the largest element $\bar{v} \in T(W)$. We find this payoff by solving the following problem, termed the *best payoff problem*:

$$(14) \quad \bar{v} = \max_{x, \mu(\cdot), w(\cdot)} \int [(1 - \beta)R(x, \mu(\theta), \theta) + \beta w(\theta)] p(\theta) d\theta$$

subject to the constraint that $x, \mu(\cdot)$, and $w(\cdot)$, are enforceable by W , in that they satisfy (8)-(10). Throughout, we assume that $\mu(\cdot)$ is a piecewise, continuously differentiable function.

The best payoff problem is a mechanism design problem of choosing an incentive-compatible allocation $x, \mu(\cdot), w(\cdot)$ which maximizes utility. Following the language of mechanism design, we refer to θ as the *type* of the monetary authority, which changes every period. When we solve this problem with $W = W^*$, (13) implies that the resulting payoff is the highest incentive-compatible payoff. We will prove our main result in Proposition 1 for any W . Hence, we will not have to explicitly solve the fixed-point problem of finding W^* .

Moreover, to prove our main result we also need only focus on the best payoff problem, which gives the highest payoff that can be obtained from period 0 onward. For completeness, however,

notice that given some $w_0(\theta)$ from the best payoff problem, a period 1 policy and continuation value, $\mu_{w_0(\theta)}(\cdot)$ and $w_{w_0(\theta)}(\cdot)$, that satisfy

$$(15) \quad w_0(\theta) = \int \left[(1 - \beta)R(x_{w_0(\theta)}, \mu_{w_0(\theta)}(z), z) + \beta w_{w_0(\theta)}(z) \right] p(z) dz$$

exist by the definition of T . Equation (15) and its analog for other periods is sometimes referred to as a *promise-keeping constraint*. In our approach we never explicitly need to mention this constraint since it is built into the definition of the operator T .

2. Characterizing the Optimal Mechanism

Now we solve the best payoff problem and use the solution to characterize the optimal mechanism. Our main result here is that under two simple conditions, a single-crossing condition and a monotone hazard condition, the optimal mechanism is static. To highlight the importance of the monotone hazard condition for this result, in Appendix B we give three examples which show that if the monotone hazard condition is violated, the optimal mechanism is dynamic.

A. Preliminaries

We begin with some definitions. In our recursive formulation, we say that a mechanism is *static* if the continuation value $w(\theta) = \bar{w}$ for (almost) all θ . We say that a mechanism is *dynamic* if $w(\theta) < \bar{w}$ for some set of θ which is realized with strictly positive probability.

Our characterization of the solution to the best payoff problem does not depend on the exact value of β . Hence, to simplify the notation, we suppress explicit dependence on β and think of the term β as being subsumed in the w function and $1 - \beta$ as being subsumed in the R function.

We assume that the preferences are differentiable and satisfy a standard *single-crossing assumption*, that

$$(A1) \quad R_{\mu\theta}(x, \mu, \theta) > 0.$$

This implies that higher types have a stronger preference for current inflation. Standard arguments can be used to show that the static best response $\mu^*(\theta; x)$ is strictly increasing in θ .

Under the single-crossing assumption (A1), a standard lemma lets us replace the global incentive constraints (10) with some local versions of them. We say that an allocation is *locally incentive-compatible* if it satisfies three conditions: $\mu(\cdot)$ is nondecreasing in θ ;

$$(16) \quad R_\mu(x, \mu(\theta), \theta) \frac{d\mu(\theta)}{d\theta} + \frac{dw(\theta)}{d\theta} = 0$$

wherever $d\mu(\theta)/d\theta$ and $dw(\theta)/d\theta$ exist; and for any point θ_i at which these derivatives do not exist,

$$(17) \quad \lim_{\theta \nearrow \theta_i} R(x, \mu(\theta), \theta_i) + w(\theta) = \lim_{\theta \searrow \theta_i} R(x, \mu(\theta), \theta_i) + w(\theta).$$

Standard arguments give the following result: under the single-crossing assumption (A1), the allocation $(x, \mu(\cdot), w(\cdot))$ satisfies the incentive constraints (10) if and only if the allocation is locally incentive-compatible. (See, for example, Fudenberg and Tirole's (1991) text.)

Given any incentive-compatible allocation, we define the *utility of the allocation* at θ to be

$$U(\theta) = R(x, \mu(\theta), \theta) + w(\theta).$$

Local incentive-compatibility implies that $U(\cdot)$ is continuous and differentiable almost everywhere, with derivative $U'(\theta) = R_\theta(x, \mu(\theta), \theta)$. Integrating $U'(\cdot)$ from $\underline{\theta}$ up to θ gives that

$$(18) \quad U(\theta) = U(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} R_\theta(x, \mu(z), z) dz$$

while integrating $U'(\cdot)$ from $\bar{\theta}$ down to θ gives that

$$(19) \quad U(\theta) = U(\bar{\theta}) - \int_{\theta}^{\bar{\theta}} R_\theta(x, \mu(z), z) dz.$$

With integration by parts, it is easy to show that for interval endpoints $\theta_1 < \theta_2$,

$$(20) \quad \int_{\theta_1}^{\theta_2} U(\theta)p(\theta) d\theta = P(\theta_2)U(\theta_2) - P(\theta_1)U(\theta_1) - \int_{\theta_1}^{\theta_2} R_\theta(x, \mu(\theta), \theta)P(\theta) d\theta.$$

Using (18) and (20), we can write the value of the objective function $\int_{\underline{\theta}}^{\bar{\theta}} U(\theta)p(\theta) d\theta$ as

$$(21) \quad U(\underline{\theta}) + \int_{\underline{\theta}}^{\bar{\theta}} \frac{1 - P(\theta)}{p(\theta)} R_\theta(x, \mu(\theta), \theta)p(\theta) d\theta \text{ or}$$

$$U(\bar{\theta}) - \int_{\underline{\theta}}^{\bar{\theta}} \frac{P(\theta)}{p(\theta)} R_\theta(x, \mu(\theta), \theta)p(\theta) d\theta.$$

Next we make some joint assumptions on the probability distribution and the return function.

Assume that, for any action profile $x, \mu(\cdot)$ with $\mu(\cdot)$ nondecreasing,

$$(A2a) \quad \frac{1 - P(\theta)}{p(\theta)} R_{\theta\mu}(x, \mu(\theta), \theta) \text{ is strictly decreasing in } \theta$$

$$(A2b) \quad \frac{P(\theta)}{p(\theta)} R_{\theta\mu}(x, \mu(\theta), \theta) \text{ is strictly increasing in } \theta.$$

We refer to assumptions (A2a) and (A2b) together as (A2) and, in a slight abuse of terminology, refer to them as the *monotone hazard condition*. In our benchmark example (1), $R_{\theta\mu}(x, \mu(\theta), \theta) = 1$, so that (A2) reduces to the standard monotone hazard condition familiar from the mechanism design literature, that $[1 - P(\theta)]/p(\theta)$ be strictly decreasing and $P(\theta)/p(\theta)$ be strictly increasing.

B. Showing That the Optimal Mechanism Is Static

Here we show that the optimal mechanism is static. In the next section, we characterize the optimal static mechanism.

Proposition 1. Under assumptions (A1) and (A2), the optimal mechanism is static.

The approach we take in proving Proposition 1 is different from the standard approach used by Fudenberg and Tirole (1991, Chapter 7.3) for solving a mathematically-related principal-agent problem. To motivate our approach, we first show why the standard approach does not work for our problem.

The best payoff problem can be written as follows: Choose $\mu(\theta)$ to maximize social welfare

$$U(\underline{\theta}) + \int_{\underline{\theta}}^{\bar{\theta}} \frac{1 - P(\theta)}{p(\theta)} R_{\theta}(x, \mu(\theta), \theta) p(\theta) d\theta$$

subject to the constraints that (i) $x = \int \mu(\theta) p(\theta) d\theta$, (ii) $\mu(\theta)$ is nondecreasing, and (iii) the continuation values defined by

$$w(\theta) \equiv U(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} R_{\theta}(x, \mu(z), z) dz - R(x, \mu(\theta), \theta)$$

satisfy $w(\theta) \leq \bar{w}$ for all θ . Alternatively, we can write the best payoff problem as choosing $\mu(\theta)$ to maximize

$$U(\bar{\theta}) - \int_{\underline{\theta}}^{\bar{\theta}} \frac{P(\theta)}{p(\theta)} R_{\theta}(x, \mu(\theta), \theta) p(\theta) d\theta$$

subject to the constraints (i), (ii), and (iii), the continuation values defined by

$$w(\theta) \equiv U(\bar{\theta}) - \int_{\theta}^{\bar{\theta}} R_{\theta}(x, \mu(z), z) dz - R(x, \mu(\theta), \theta)$$

satisfy $w(\theta) \leq \bar{w}$ for all θ .

The standard approach to solving either version of this problem is to guess that the analog of constraints (ii) and (iii) do not bind, take the corresponding first-order conditions of either of these problems to find the implied $\mu(\cdot)$, and then verify that constraints (ii) and (iii) are in fact satisfied at that choice of $\mu(\cdot)$. If we take that approach here, we see that it fails. The first-order conditions with respect to $\mu(\theta)$ are

$$(22) \quad \frac{1 - P(\theta)}{p(\theta)} R_{\theta\mu}(x, \mu(\theta), \theta) = \lambda$$

for the first version of the best payoff problem and

$$(23) \quad -\frac{P(\theta)}{p(\theta)} R_{\theta\mu}(x, \mu(\theta), \theta) = \lambda$$

for the second version of the best payoff problem, where λ is the Lagrange multiplier on constraint (i). The solution to these first-order conditions (22) and (23), from the relaxed problem in which we have dropped constraints (ii) and (iii), implies a decreasing $\mu(\cdot)$ schedule. To see why, note, for example, that the left side of equation (22) is the increment to social welfare from marginally increasing $\mu(\cdot)$ at some particular θ and adjusting the continuation values $w(\cdot)$ for $\theta' \geq \theta$ to preserve incentive-compatibility, while the right side is the cost in terms of welfare from raising expected inflation x . Under assumption (A2a), the benefits of raising $\mu(\cdot)$ are higher for low values of θ than for high values of θ . Thus, in the relaxed problem, it is optimal to have a downward-sloping $\mu(\cdot)$ schedule. Similar logic applies to (23). Clearly, then, the solution to the relaxed problem violates at least one of the dropped constraints (ii) or (iii), and hence, we cannot use this standard approach.

We also cannot use the “ironing” approach designed to deal with cases in which the monotonicity constraint (ii) binds, because in our problem, the constraint that binds is constraint (iii), which is not dealt with in that approach. Instead, in the proof of Proposition 1 that follows, we use a variational argument to show that constraint (iii) binds for all θ at the solution to the best payoff problem.

The key feature of our problem that leads to the failure of the standard approach is that the continuation value enters positively into the payoff of both society and the monetary authority. Mathematically, these continuation values are analogous to the transfers between the principal and the agent in the standard principal-agent problem presented by Fudenberg and Tirole (1991, Chapter 7.3). In that problem, the transfers enter positively into the agent’s problem but negatively into the principal’s problem. This difference between our problem and the principal-agent problem is the key reason the standard approach does not work.

Before proving Proposition 1, we sketch our basic argument. Our discussion of the first-order conditions of the relaxed problem (22) and (23) suggests that given any strictly increasing $\mu(\cdot)$ schedule, a variation that flattens this schedule will improve welfare if it is feasible in the sense that the associated continuation value satisfies constraint (iii). Our proof of Lemma 1 formalizes this logic.

Our objective is to show that the optimal continuation value $w(\cdot)$ is constant at \bar{w} . We prove this by contradiction. We start with the observation that $w(\cdot)$ is piecewise-differentiable since $\mu(\cdot)$ is piecewise-differentiable and (16) holds. We first show that $w(\cdot)$ must be a step function. If not, there is some interval over which $w'(\theta)$ is nonzero, and hence, from local incentive-compatibility,

$\mu(\cdot)$ is strictly increasing. In Lemma 2, we show that a variation that flattens $\mu(\cdot)$ over that interval is feasible. From Lemma 1, we know it is welfare-improving.

We next show that $w(\cdot)$ must be continuous, and since it is a step function, it must be constant. We prove this by showing that if $w(\cdot)$ is discontinuous at some point θ , then (17) implies that $\mu(\cdot)$ must be increasing in the sense that it jumps up at that point. In Lemma 3, we show that a variation that flattens $\mu(\cdot)$ in a neighborhood of that point is feasible, and again from Lemma 1, we know that it is welfare-improving.

It is convenient in the proof to use a definition of *increasing on an interval* which covers the cases we deal with in Lemmas 2 and 3. This definition subsumes the case of Lemma 2 in which $d\mu(\theta)/d\theta > 0$ for some interval and the case of Lemma 3 in which $\mu(\cdot)$ jumps up at $\tilde{\theta}$. We say that $\mu(\cdot)$ is *increasing on* (θ_1, θ_2) if $\mu(\cdot)$ is weakly increasing on this interval and there is some $\tilde{\theta}$ in this interval such that $\mu(\theta) < \tilde{\mu}$ for $\theta < \tilde{\theta}$ and $\mu(\theta) > \tilde{\mu}$ for $\theta > \tilde{\theta}$, where $\tilde{\mu}$ is the conditional mean of $\mu(\cdot)$ on this interval, namely,

$$(24) \quad \tilde{\mu} = \frac{\int_{\theta_1}^{\theta_2} \mu(\theta) p(\theta) d\theta}{P(\theta_2) - P(\theta_1)}.$$

In words, on this interval, the function $\mu(\cdot)$ is weakly increasing and is strictly below its conditional mean $\tilde{\mu}$ up to $\tilde{\theta}$ and strictly above its conditional mean after $\tilde{\theta}$.⁶ Throughout, we will also say that the policy $\mu(\cdot)$ is *flat* at some particular point θ if the derivative $\mu'(\theta)$ exists and equals zero at that point.

Consider now some dynamic mechanism $(x, \mu(\cdot), w(\cdot))$ in which the policy $\mu(\cdot)$ is increasing on some interval, say, (θ_1, θ_2) . In our variation, we marginally move the function $\mu(\cdot)$ toward its conditional mean on this interval and adjust the continuation values to preserve incentive-compatibility. In particular, our variation moves our original policy $\mu(\cdot)$ marginally toward a policy $\tilde{\mu}(\cdot)$ defined by

$$(25) \quad \tilde{\mu}(\theta) = \left[\begin{array}{l} \tilde{\mu} \text{ if } \theta \in (\theta_1, \theta_2) \\ \mu(\theta) \text{ otherwise} \end{array} \right].$$

This policy $\tilde{\mu}(\cdot)$ differs from the original policy $\mu(\cdot)$ only on the interval (θ_1, θ_2) , and there the original policy $\mu(\cdot)$ is replaced by the conditional mean $\tilde{\mu}$ of the original policy over the interval.

⁶Observe that this definition of *increasing* is stronger than the definition of a function being *weakly increasing* on an interval because our definition rules out a function that is constant over the interval. But our definition is weaker than the definition of a function being *strictly increasing* over an interval because ours allows for subintervals over which $\mu(\cdot)$ is constant.

Clearly, the expected inflation under $\tilde{\mu}(\cdot)$ is the same as the expected inflation under the original policy.

We let $(x(a), \mu(\cdot; a), w(\cdot; a))$ and $U(\cdot; a)$ denote our variation and the associated utility. The policy $\mu(\cdot; a)$ in our variation is a convex combination of the policy $\tilde{\mu}(\cdot)$ and the original policy $\mu(\cdot)$ and is defined by

$$(26) \quad \mu(\theta; a) = a\tilde{\mu}(\theta) + (1 - a)\mu(\theta)$$

for $a \in [0, 1]$. (For a graph of $\mu(\cdot; a)$, see Figure 1.) Clearly, the expected inflation in our variation $\tilde{x}(a)$ equals that of the original allocation x for all $a \in [0, 1]$.

The delicate part of the variation is to construct the continuation value $w(\cdot; a)$ so as to satisfy the feasibility constraint $w(\theta; a) \leq \bar{w}$ for all θ , in addition to incentive-compatibility. It turns out that we can ensure feasibility if we use one of two ways to adjust continuation values. In the *up variation*, we leave the continuation values unchanged below θ_1 and pass up any changes induced by our variation in the policy to higher types by suitably adjusting the continuation values to maintain incentive-compatibility. In the *down variation*, we leave the continuation values unchanged above θ_2 and pass down any changes induced by our variation in the policy to lower types by suitably adjusting the continuation values to maintain incentive-compatibility.

In the up variation, we determine the continuation values by substituting $U(\theta; a) = R(x, \mu(\theta; a), \theta) + w(\theta; a)$ into (18) to get that $w(\theta; a)$ is defined by

$$(27) \quad w(\theta; a) = U(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} R_{\theta}(x, \mu(z; a), z) dz - R(x, \mu(\theta; a), \theta).$$

In the down variation, we use (19) in a similar way to get that $w(\theta; a)$ is defined by

$$(28) \quad w(\theta; a) = U(\bar{\theta}) - \int_{\theta}^{\bar{\theta}} R_{\theta}(x, \mu(z; a), z) dz - R(x, \mu(\theta; a), \theta).$$

By construction, these variations are incentive-compatible. In the following lemma, we show that, if either variation is feasible, it improves welfare.

Lemma 1. Assume (A1) and (A2), and let $(x, \mu(\cdot), w(\cdot))$ be an allocation in which $\mu(\cdot)$ is increasing on some interval (θ_1, θ_2) . Then the up variation and the down variation both increase the objective function (21).

Proof. To see that the up variation improves welfare, use (21) to write the value of the objective function under this variation as

$$(29) \quad V(a) = U(\underline{\theta}) + \int_{\underline{\theta}}^{\bar{\theta}} \frac{1 - P(\theta)}{p(\theta)} R_{\theta}(x, a\tilde{\mu}(\theta) + (1 - a)\mu(\theta), \theta) p(\theta) d\theta.$$

To evaluate the effect on welfare of a marginal change of this type, take the derivative of $\tilde{V}(a)$ and evaluate it at $a = 0$ to get

$$(30) \quad \frac{dV(0)}{da} = \int_{\underline{\theta}}^{\bar{\theta}} \frac{1 - P(\theta)}{p(\theta)} R_{\theta\mu}(x, \mu(\theta), \theta) [\tilde{\mu}(\theta) - \mu(\theta)] p(\theta) d\theta$$

which, with the form of $\tilde{\mu}(\cdot)$, reduces to

$$(31) \quad \frac{dV(0)}{da} = \int_{\theta_1}^{\theta_2} \frac{1 - P(\theta)}{p(\theta)} R_{\theta\mu}(x, \mu(\theta), \theta) [\tilde{\mu} - \mu(\theta)] p(\theta) d\theta.$$

If we divide (31) by the positive constant $P(\theta_2) - P(\theta_1)$, then we can interpret (31) to be the expectation of the product of two functions, namely $f(\theta)$ defined as $[1 - P(\theta)]R_{\theta\mu}(x, \mu(\theta), \theta)/p(\theta)$ and $g(\theta)$ defined as $\tilde{\mu} - \mu(\theta)$, where the $p(\theta)/[P(\theta_2) - P(\theta_1)]$ is the density of θ over the interval (θ_1, θ_2) . The function f is strictly decreasing by assumption (A2a). Because the function $\mu(\theta)$ is increasing on the interval (θ_1, θ_2) , the function g is decreasing on this interval in the sense that $g(\theta)$ is weakly decreasing and lies strictly below its conditional mean for $\theta < \tilde{\theta}$ and strictly above its conditional mean for $\theta > \tilde{\theta}$. By the definition of a covariance, we know that $Efg = \text{cov}(f, g) + (Ef)(Eg)$, where the expectation is taken with respect to the density $p(\theta)/[P(\theta_2) - P(\theta_1)]$. By the construction of $\tilde{\mu}$ in (24), we know that $Eg = 0$, so that $Efg = \text{cov}(f, g)$, which is clearly positive because f is strictly decreasing and g is decreasing on the interval (θ_1, θ_2) . Thus, (31) is strictly positive, and the variation improves welfare.

The down variation also improves welfare. The value of the objective function under this variation is

$$V(a) = U(\bar{\theta}) - \int_{\underline{\theta}}^{\bar{\theta}} \frac{P(\theta)}{p(\theta)} R_{\theta}(x, a\tilde{\mu}(\theta) + (1 - a)\mu(\theta), \theta) p(\theta) d\theta.$$

Hence,

$$(32) \quad \frac{dV(0)}{da} = \int_{\theta_1}^{\theta_2} \frac{P(\theta)}{p(\theta)} R_{\theta\mu}(x, \mu(\theta), \theta) [\mu(\theta) - \tilde{\mu}] p(\theta) d\theta > 0$$

by arguments similar to those given before. *Q.E.D.*

To gain some intuition for how these variations improve welfare, consider the up variation and the expression for the change in welfare (31). We show how the total effect on welfare resulting from this flattening of the inflation schedule can be thought of as arising from two effects: a *positive* effect that comes from raising inflation for low types and a *negative* effect that comes from lowering inflation for high types. Our assumption (A2a) ensures that the positive effect outweighs the negative effect.

For any type, the flattening affects both the current payoff R and the continuation value w . The impact of increasing a on the current payoff for type θ is

$$R_\mu(x, \mu(\theta), \theta) [\tilde{\mu}(\theta) - \mu(\theta)].$$

In the up variation, the impact of increasing a on the continuation value for this type is

$$(33) \quad \frac{d\tilde{w}(\theta; 0)}{da} = \int_{\underline{\theta}}^{\theta} R_{\theta\mu}(x, \mu(z), z) [\tilde{\mu}(z) - \mu(z)] dz - R_\mu(x, \mu(\theta), \theta) [\tilde{\mu}(\theta) - \mu(\theta)].$$

Hence, the impact on the utility of type θ is simply the sum of these pieces and is given by

$$(34) \quad \frac{d\tilde{U}(\theta; 0)}{da} = \int_{\underline{\theta}}^{\theta} R_{\theta\mu}(x, \mu(z), z) [\tilde{\mu}(z) - \mu(z)] dz.$$

Notice from (34) that any change in the policy for some particular type z affects the utility of all types θ above z . Thus, each term

$$(35) \quad [1 - P(z)] R_{\theta\mu}(x, \mu(z), z) [\tilde{\mu}(z) - \mu(z)]$$

in the integral (30) can be thought of as the sum of the change in utility for all types z and above resulting from the change in the inflation schedule for the type z . Under our single-crossing assumption, $R_{\theta\mu}(x, \mu(\theta), \theta) > 0$, so the impact of changing the policy at θ depends on the sign of $\tilde{\mu}(\theta) - \mu(\theta)$. Recall that outside the interval (θ_1, θ_2) , $\tilde{\mu}(\theta) = \mu(\theta)$, so that the value of (35) is zero. Inside the interval (θ_1, θ_2) , $\tilde{\mu}(\theta) = \tilde{\mu}$, where $\tilde{\mu}$ is the conditional mean on this interval. By definition of the type $\tilde{\theta}$, on the interval $(\theta_1, \tilde{\theta})$, $\tilde{\mu} - \mu(\theta) > 0$, and on the interval $(\tilde{\theta}, \theta_2)$, $\tilde{\mu} - \mu(\theta) < 0$. Therefore, our variation has both a positive effect and a negative effect on welfare.

The positive effect of flattening the inflation schedule comes from increasing the policy of those types θ below $\tilde{\theta}$ and then passing this change up to higher types. The negative effect of the flattening comes from decreasing the policy for those types θ above $\tilde{\theta}$. Under assumption (A2a), the positive effect outweighs the negative effect.

In the down variation, the intuition for the derivative (32) is the same as that for (31), except that, in this variation, a change in the inflation rate chosen by type θ affects the continuation value of all types below θ .

The following lemma proves that if $w(\cdot)$ is not a step function, then $\mu(\cdot)$ is increasing on some interval, and there is a feasible variation that flattens $\mu(\cdot)$ and improves welfare.

Lemma 2. Under (A1) and (A2), in the optimal mechanism, the continuation value function $w(\cdot)$ is a step function.

Proof. Since by assumption $\mu(\cdot)$ is piecewise-differentiable, we know from (16) that $w(\cdot)$ is too. By way of contradiction, assume that $w(\cdot)$ is not a step function. Hence, there is an interval over which $w'(\theta)$ exists and does not equal zero. Clearly, then, there is a subinterval (θ_1, θ_2) over which $w'(\theta)$ is either strictly positive or strictly negative, and $w(\theta) \leq \bar{w} - \varepsilon$ for some $\varepsilon > 0$. From local incentive-compatibility, we know that

$$R_\mu(x, \mu(\theta), \theta) \frac{d\mu(\theta)}{d\theta} + \frac{dw(\theta)}{d\theta} = 0$$

so that regardless of the sign of $w'(\theta)$, we have that $\mu'(\theta) > 0$ on this interval. Hence, $\mu(\cdot)$ is increasing on (θ_1, θ_2) in the sense defined above. From Lemma 1, we know that, if the up and down variations are feasible, then they both improve welfare.

To complete the proof, we show that either the up variation or the down variation is always feasible. Under the up variation, (26) and (27) imply that $w(\theta; a)$ equals $w(\theta)$ for $\theta \leq \theta_1$ and

$$w(\theta) + \Delta(a)$$

for $\theta \geq \theta_2$, where

$$(36) \quad \Delta(a) \equiv \int_{\theta_1}^{\theta_2} [R_\theta(x, \mu(z; a), z) - R_\theta(x, \mu(z), z)] dz.$$

See Figure 2 for a graph of $w(\theta; a)$ in the up variation. This graph illustrates several features of $w(\theta; a)$: it coincides with $w(\theta)$ for $\theta \leq \theta_1$, it differs from $w(\theta)$ by the constant $\Delta(a)$ for $\theta \geq \theta_2$, and it jumps at both θ_1 and θ_2 . This last feature follows from (17) and the fact that $\mu(\theta; a)$ jumps at these points. Notice in the graph that $w(\theta) \leq \bar{w} - \varepsilon$ for $\theta \in (\theta_1, \theta_2)$.

Under the down variation, (26) and (28) imply that $w(\theta; a)$ equals

$$(37) \quad w(\theta) - \Delta(a)$$

for $\theta \leq \theta_1$ and $w(\theta)$ for $\theta \geq \theta_2$. See Figure 3 for a graph of $w(\theta; a)$ in the down variation.

To ensure that the continuation value satisfies feasibility, we do the following. We use the up variation when term $\Delta(a) \leq 0$ and the down variation when that term is positive. By doing so, we ensure that outside the interval (θ_1, θ_2) the continuation value under this variation is no larger than the original continuation value $w(\theta)$, which, by assumption, is feasible. We know that inside the interval (θ_1, θ_2) , $w(\theta) \leq \bar{w} - \varepsilon$. Since R is continuous in μ , we can choose a small enough to ensure that $w(\theta; a) \leq \bar{w}$. *Q.E.D.*

In the next lemma, we show that the optimal policy $w(\theta)$ is continuous. Since we know from Lemma 2 that $w(\cdot)$ is a step function, we conclude that $w(\cdot)$ is a constant. Optimality implies that this constant is \bar{w} .

Lemma 3. Under (A1) and (A2), $w(\theta)$ is continuous for all θ .

Proof. In Appendix A, we prove that $w(\cdot)$ is continuous by contradiction. We show that if $w(\cdot)$ jumps at some point $\tilde{\theta}$, then the same up variation and down variation we used in Lemma 1 will improve welfare. The only difficult part of the proof is showing that when the appropriate interval (θ_1, θ_2) is selected that contains the jump point $\tilde{\theta}$, the associated continuation values are feasible. Here it may turn out that the feasibility constraint binds inside the interval (θ_1, θ_2) , in that the original allocation has $w(\theta) = \bar{w}$ for some θ in (θ_1, θ_2) . Thus, we cannot simply shrink the size of the weight a in the variation to ensure feasibility on (θ_1, θ_2) , as we did in the proof of Lemma 2. Instead we show that the variation is feasible inside the interval (θ_1, θ_2) by direct calculations that we relegate to Appendix A. *Q.E.D.*

Together Lemmas 2 and 3 establish Proposition 1, that under our assumptions, the optimal mechanism is static. Our characterization of optimal policy relied on the monotone hazard assumption (A2). Under this assumption, we showed that the dynamic mechanism design problem has a static solution. In Appendix B, we give three simple examples in which the monotone hazard condition (A2) is violated yet the dynamic mechanism design problem does not have a static solution. In the first two of these examples (A2) fails because $(1 - P(\theta))/p(\theta)$ is not monotone, in the third (A2a) fails because $R_{\mu\theta}$ is increasing at a sufficiently rapid rate.

3. The Optimal Degree of Discretion

So far we have demonstrated that the optimal mechanism is static. Now we characterize the optimal static mechanism. We show three results: The optimal policy has either bounded discretion or no discretion. A policy with either bounded discretion or no discretion can be implemented by society setting an upper limit, or *cap*, on the inflation rate which the monetary authority is allowed to choose. And the optimal degree of discretion is decreasing the more severe is the time inconsistency problem and the less important is private information.

A. Characterizing the Optimal Policy

In the optimal static mechanism, the monetary policy $\mu(\cdot)$ maximizes

$$(38) \quad \int R(x, \mu(\theta), \theta) p(\theta) d\theta$$

subject to the constraints that $x = \int \mu(\theta) p(\theta) d\theta$ and $R(x, \mu(\theta), \theta) \geq R(x, \mu(\hat{\theta}), \theta)$ for all $\theta, \hat{\theta}$.

We say that a monetary policy $\mu(\cdot)$ has *bounded discretion* if it takes the form

$$(39) \quad \mu(\theta) = \begin{cases} \mu^*(\theta; x) & \text{if } \theta \in [\underline{\theta}, \theta^*) \\ \mu^* = \mu^*(\theta^*, x) & \text{if } \theta \in [\theta^*, \bar{\theta}] \end{cases}$$

where $\mu^*(\theta; x)$ is the static best response given wages $x = \int \mu(\theta) p(\theta) d\theta$. Thus, for $\theta < \theta^*$, the monetary authority chooses the static best response, and for $\theta \geq \theta^*$, the monetary authority chooses the upper limit μ^* . A policy has *no discretion* if $\mu(\theta) = \mu$ for some constant μ , so that regardless of θ , the monetary authority chooses the same growth rate. Clearly, the best policy with no discretion is the expected Ramsey policy.⁷

We show that the optimal policy has either bounded discretion or no discretion. Here, as before, we can replace the global incentive constraint in (38) with the local incentive constraints, with the restriction that $w(\theta) = \bar{w}$. In particular, (17) implies that $\mu(\cdot)$ is continuous, while (16), namely, $R_\mu d\mu/d\theta = 0$, implies that for all θ , $\mu(\theta)$ is either flat or equal to the static best response. Clearly, if $\mu(\cdot)$ is flat everywhere, it is a constant; hence, it equals the expected Ramsey policy, which by definition is the best constant policy. If $\mu(\cdot)$ is not flat everywhere, then it must be of the following form for some θ_1 and θ_2 :

$$(40) \quad \mu(\theta) = \begin{cases} \mu_1 = \mu^*(\theta_1; x) & \text{if } \theta \in [\underline{\theta}, \theta_1) \\ \mu^*(\theta; x) & \text{if } \theta \in [\theta_1, \theta_2] \\ \mu_2 = \mu^*(\theta_2; x) & \text{if } \theta \in (\theta_2, \bar{\theta}] \end{cases}$$

where $x = \int \mu(\theta) p(\theta) d\theta$. In words, the policy must be constant up to some point $\theta_1 \geq \underline{\theta}$ and equal to the static best response of type θ_1 ; it must be equal to the static best response of each type $\theta \in [\theta_1, \theta_2]$ with $\theta_2 \leq \bar{\theta}$; and then it must be constant and equal to the static best response of type θ_2 .

⁷Note that the best policy with no discretion, the *Ramsey policy*, will not typically be a special case of a policy with bounded discretion. Specifically, when $\theta^* = \underline{\theta}$, the form (39) yields one particular policy with no discretion, namely, $\mu(\theta) = \mu^*(\underline{\theta}; x)$ for all θ . But this policy does not typically coincide with the expected Ramsey policy μ^{ER} since the best response of the lowest type is not typically the expected Ramsey policy.

In the following proposition, we show that if the optimal policy is not the expected Ramsey policy, then it must be of the form (40) with θ_1 equal to $\underline{\theta}$, so that the policy's form reduces to the bounded discretion form in (39).

Proposition 2. Under assumptions (A1) and (A2), the optimal policy $\mu(\cdot)$ has either bounded discretion or no discretion.

Proof. We have argued that if the optimal policy is constant, then it must be an expected Ramsey policy, which has no discretion. If the optimal policy is not constant, then it must be of the form (40). But $\mu(\theta)$ having the form (40) with $\theta_1 > \underline{\theta}$ cannot be optimal. To see this, observe that an alternative policy $\tilde{\mu}(\theta)$ of the same form would exist with $\tilde{\theta}_1 < \theta_1$ and $\tilde{\theta}_2 = \theta_2$. We illustrate this alternative policy in Figure 4. This alternative policy $\tilde{\mu}(\theta)$ would be closer to $\mu^*(\theta, x)$ wherever it differs from $\mu(\theta)$ and would satisfy $\int \tilde{\mu}(\theta)p(\theta) d\theta < \int \mu(\theta)p(\theta) d\theta = x$. Hence, this alternative policy $\tilde{\mu}(\theta)$ would be strictly preferred to $\mu(\theta)$; the change from $\mu(\theta)$ to $\tilde{\mu}(\theta)$ directly improves welfare for all types $\theta < \theta_1$, with x held fixed. The change also reduces x , which by (4) contributes to improving total welfare. More formally, observe that the marginal impact on welfare of a marginal reduction in θ_1 is given by $d\tilde{V}$ equal to

$$\int_{\underline{\theta}}^{\theta_1} \left[R_{\mu}(x, \mu^*(\theta_1; x), \theta) \frac{\partial \mu^*(\theta_1; x)}{\partial \theta} \Delta \theta_1 \right] p(\theta) d\theta + \int_{\underline{\theta}}^{\bar{\theta}} [R_x(x, \mu(\theta), \theta) \Delta x] p(\theta) d\theta$$

which is positive since $R_{\mu}(x, \mu^*(\theta_1; x), \theta) < 0$, $\partial \mu^*(\theta_1; x) / \partial \theta > 0$, $\Delta \theta_1 < 0$, $\Delta x < 0$, and (4). *Q.E.D.*

B. Implementing Optimal Policy with an Inflation Cap or a Range of Inflation Rates

We have characterized the solution to a dynamic mechanism design problem. We now imagine implementing the resulting outcome with an *inflation cap*, a highest level of allowable inflation $\bar{\pi}$. We imagine that society legislates this highest allowable level and that doing so restricts the monetary authority's choices to be $\mu_t \leq \bar{\pi}$. If this cap is appropriately set and agents simply play the repeated one-shot equilibrium of the resulting game with this inflation cap, then the monetary authority will optimally choose the outcome of the mechanism design problem. In this sense, the repeated one-shot game with an inflation cap implements the policy that solves the best payoff problem.

The intuition for this result—that a policy with either bounded discretion or no discretion can be implemented by setting an upper limit on permissible inflation rates—is simple. In our environment, the only potentially beneficial deviations from either type of policy are ones that raise

inflation. Under bounded discretion, the types in $[\underline{\theta}, \theta^*)$ are choosing their static best response to wages and, hence, have no incentive to deviate, whereas the types in $(\theta^*, \bar{\theta}]$ have an incentive to deviate to a higher rate than $\bar{\pi}$. Similarly, from Proposition 3 (stated and proved below), we know that if the expected Ramsey policy is optimal, then at this policy all types have an incentive to deviate to higher rates of inflation. Hence, an inflation cap of $\bar{\pi} = \mu^{ER}$ implements such a policy. (For completeness, we formalize this argument in Appendix C.)

Clearly, we can also implement the optimal policy with a range of inflation rates denoted $[\underline{\pi}, \bar{\pi}]$. The top end of such a range is the inflation cap, $\bar{\pi}$, discussed above. The bottom end of the range, $\underline{\pi}$, is simply the optimal policy chosen by the lowest type $\underline{\theta}$ in the optimal static mechanism. Under a policy of bounded discretion, $\underline{\pi} < \bar{\pi}$, while under a policy of no discretion, $\underline{\pi} = \bar{\pi}$.

C. Linking Discretion with Time Inconsistency and Private Information

So far we have shown that the optimal policy has either bounded discretion or no discretion. Here we link the optimal degree of discretion to the severity of the time inconsistency problem relative to the importance of private information. We show that the more severe the time inconsistency problem, the smaller is the optimal degree of discretion. Likewise, we show that the less important the private information, the smaller is the optimal degree of discretion.

The literature using general equilibrium models to study optimal monetary policies suggests a qualitative way to measure the severity of the time inconsistency problem. In most of this literature, the time inconsistency problem is severe in that the static Nash equilibrium is always at the highest feasible inflation rate $\bar{\mu}$. This result follows because the static best response of the monetary authority to any given level of expected inflation is always above that level and thus, the monetary authority is always tempted to generate a monetary surprise. Examples of the models with the more severe problems include those of Ireland (1997); Chari, Christiano, and Eichenbaum (1998); and Sleet (2001). In the rest of the literature the problem is less severe, in that the static Nash equilibrium is interior. Examples of the models with the less severe problems include those of Chang (1998), Nicolini (1998), and Albanesi, Chari, and Christiano (2003).

In our reduced-form model, we can mimic the general equilibrium models with the more severe problems by choosing a payoff function R for which $R_{\mu}(x, x, \theta) > 0$ for all θ . That is, in response to any choice of wages x , the monetary authority wants to choose inflation higher than x , regardless of its type. Under (A1) this condition is equivalent to requiring that the static best response function satisfies $\mu^*(\underline{\theta}, x) \geq x$ for all $x \in [\underline{\mu}, \bar{\mu}]$. We show in the next proposition that this

condition implies that the optimal policy has no discretion.

We can mimic the general equilibrium models with the less severe problems by choosing a payoff function R for which the static Nash equilibrium best response is interior. For such a payoff function the optimal policy will typically depend on parameters. When the time inconsistency problem is sufficiently mild, however, we can show a general result, namely that optimal policy must have bounded discretion. Here by mild we mean that when wages are set at the expected Ramsey level, the lowest type wants to set inflation at some level lower than the expected Ramsey level. Technically, we can state this condition as the static best response satisfies $\mu^*(\underline{\theta}, \mu^{ER}) < \mu^{ER}$ or equivalently the payoff function satisfies $R_\mu(\mu^{ER}, \mu^{ER}, \underline{\theta}) < 0$.

We summarize this discussion as follows.

Proposition 3. Assume (A1) and (A2) then *i*) if the static best response satisfies $\mu^*(\underline{\theta}, x) \geq x$ for all $x \in [\underline{\mu}, \bar{\mu}]$ then the optimal policy has no discretion and *ii*) if the static best response satisfies $\mu^*(\underline{\theta}, \mu^{ER}) < \mu^{ER}$ then the optimal policy has bounded discretion.

Proof. Under (A1) and (A2), the optimal mechanism is static. To prove *i*) note that in any equilibrium with bounded discretion

$$(41) \quad x = \int_{\underline{\theta}}^{\theta^*} \mu^*(\theta, x) p(\theta) d\theta + (1 - P(\theta^*)) \mu^*(\theta^*, x)$$

Under (A1), $\mu^*(\theta, x)$ is strictly increasing in θ whenever $\mu^*(\theta, x) < \bar{\mu}$. Thus, $\mu^*(\underline{\theta}, x) \geq x$ for all $x \in [\underline{\mu}, \bar{\mu}]$ implies that whenever $\theta^* > \underline{\theta}$, the right side of (41) is greater than the left side of (41) for any $x < \bar{\mu}$. The only feasible policies of the bounded discretion form must have $\theta^* = \underline{\theta}$ or $x = \bar{\mu}$ and hence reduce to policies with no discretion. The optimal policy with no discretion, the expected Ramsey policy, by definition yields higher welfare.

We prove *ii*) by contradiction. Assume that $\mu^*(\underline{\theta}, \mu^{ER}) < \mu^{ER}$, but that the optimal policy has no discretion. The variation used in Proposition 2 immediately implies that such a policy cannot be optimal. Thus, the optimal policy must have bounded discretion. *Q.E.D.*

In Proposition 3 we have characterized the form of the optimal policy for two cases for which this can be done independently of parameters. To characterize the optimal policy in the remaining case *iii*) in which $\mu^*(\underline{\theta}, \mu^{ER}) > \mu^{ER}$ but there exists an x such that $\mu^*(\underline{\theta}, x) < x$, we turn to our benchmark example (1).

In general, the choice of the optimal inflation cap depends on the importance of private information relative to the severity of the time inconsistency problem. In our benchmark example the

parameter α indexes the importance of private information and the parameter U indexes the severity of the time inconsistency problem. To see why α indexes the importance of private information note that the Ramsey policy is $\mu^R(\theta) = \alpha\theta/2$, so that the slope increases with α . Hence, as α increases, the Ramsey policy responds more to the private information θ and the gap in welfare between the Ramsey policy and the expected Ramsey policy grows. To see why U indexes the severity of the time inconsistency problem, note that the Nash inflation rate is $x^N = U$, and the Nash policies are $\mu^*(\theta; U) = U + \alpha\theta/2$. The Ramsey inflation rate is $x^R = 0$, and the Ramsey policies are $\mu^R(\theta) = \alpha\theta/2$. Thus, for each type θ , the Nash policies are simply the Ramsey policies shifted up by U . As U gets smaller, the Nash policies converge to the Ramsey policies. When U is zero, the Nash and Ramsey policies coincide.

When the objective function satisfies (1), the condition $\mu^*(\underline{\theta}; \mu^{ER}) < \mu^{ER}$ in Proposition 3 reduces to $U/\alpha < -\underline{\theta}$, where $\underline{\theta}$ is a negative number. Proposition 3 thus implies that bounded discretion is optimal when private information is important relative to the severity of the time inconsistency problem. We characterize the optimal mechanism more fully in the benchmark case in the next proposition to get a more precise link between the severity of the time inconsistency problem and the optimal degree of discretion.

For policies of the bounded discretion form (39), we think of θ^* as indexing the degree of discretion. If $\theta^* = \bar{\theta}$, then all types θ are on their static best responses, and, hence, we say there is complete discretion. As θ^* decreases, fewer types are on their static best responses, and, hence, we say there is less discretion. We then have the following proposition.

Proposition 4. Assume (1), (A1), and (A2a). If $U/\alpha = 0$, then the optimal policy has complete discretion. If $U/\alpha \in (0, -\underline{\theta})$, then that policy has bounded discretion with $\theta^* < \bar{\theta}$. The optimal degree of discretion θ^* is decreasing in U/α . As U/α approaches $-\underline{\theta}$, the cutoff θ^* approaches $\underline{\theta}$. If $U/\alpha \geq -\underline{\theta}$, then the optimal policy is the expected Ramsey policy with no discretion.

We prove this proposition by direct calculations that we provide in Appendix D. Figure 5 illustrates the proposition for two economies with different degrees relative importance of private information and severity of time inconsistency problems, $(U/\alpha)_H > (U/\alpha)_L$. In these two economies, we denote the optimal policies by $\mu_H(\cdot)$ indexed by θ_H^* and $\mu_L(\cdot)$ indexed by θ_L^* , along with the inflation caps $\bar{\pi}_H$ and $\bar{\pi}_L$.

4. Comparison to the Literature

Our result on the optimality of a static mechanism is quite different from what is typically found in dynamic contracting problems, namely, that static mechanisms are not optimal. Using a recursive approach, we have shown how our dynamic mechanism design problem reduces to a simple quasi-linear mechanism design problem. Our results are thus also directly comparable to the large literature on mechanism design with broad applications, including those in industrial organization, public finance, and auctions. (See Fudenberg and Tirole's 1991 book for an introduction to mechanism design and its applications.) In this comparison, the continuation values in our framework correspond to the contractual compensation to the agent in the mechanism design literature. Our result that the optimal mechanism is static, so that the continuation values do not vary with type, stands in contrast to the standard result in the mechanism design literature that under the optimal contract, the compensation to the agent varies with the agent's type. In this sense, our result is also quite different from what is found in the mechanism design literature.

We believe that the key feature of our model that distinguishes it from much of the dynamic incentive literature is the constraint

$$(42) \quad w(\theta) \leq \bar{w}.$$

The implication of this constraint is that it is not possible to trade off the continuation values of one type against another as it is in many other models. To highlight the importance of this constraint we consider a highly stylized example in Appendix E that replaces the constraint (42) with

$$(43) \quad \int w(\theta)p(\theta)d\theta \leq \bar{w}$$

and shows that the answer differs radically in that the optimal value of w varies with θ . Now what is the difference between constraints (42) and (43)? In providing incentives under (43) it is possible to trade off a low continuation value for one type against a high continuation value for another. This feature is common in a wide variety of incentive problems, and in such problems the optimal incentive scheme has $w(\theta)$ varying with the type θ . In contrast, when providing incentives under (42) it is not possible to make this tradeoff: a low value of $w(\theta)$ for one type does not allow us to raise the value of $w(\theta)$ for some other type. Hence, under (42) using $w(\theta)$ to provide incentives is akin to burning money.

There is a large class of dynamic incentive models that might usefully be thought of as *debt models*. Early versions of these include the private debt models of Green (1987), Thomas and

Worall (1990), Atkeson (1991) and Atkeson and Lucas (1992, 1995) while later versions include the government debt models of Sleet (forthcoming) and Sleet and Yeltekin (2003). All of these models share the feature that optimal contracts are dynamic because in each of these settings it is possible to trade a low continuation for one type against a high continuation value for another type. In this sense the debt models share many of the features of models with constraints of the form (43) rather than models with constraints of the form (42).

Having a constraint like (42) rather than (43) is important for our result that the optimal mechanism is static, but it is not sufficient, for at least two reasons. First, even in our model, we have given examples in which the optimal mechanism is dynamic when our monotone hazard condition is violated. Second, the information structure is also important. In our model, private agents receive no direct information about the state. If private agents receive a noisy signal about the state before the monetary authority takes its action, then our results go through pretty much unchanged; the noisy signal is just a publicly observed variable upon which the inflation cap is conditioned. If, however, private agents receive a noisy signal about the information the monetary authority received after the monetary authority takes its action, then dynamic mechanisms in which continuation values vary with this signal may be optimal.

Sleet (2001) considers such an information structure and shows that the optimality of the dynamic mechanism depends on the parameters governing the noise. He finds that when the public signal about the monetary authority's information is sufficiently noisy, it is not optimal to have the monetary authority's action depend on its private information; hence, the optimal mechanism is static. In contrast, when this public signal is sufficiently precise, the optimal mechanism is dynamic. The logic of why a dynamic mechanism is optimal is roughly similar to that in the literature in industrial organization following Green and Porter (1984) on optimal collusive agreements that are supported by periodic reversion to price wars, even though these price wars lower all firms' profits.

Our work here is also related to some of the repeated game literature in industrial organization about supporting collusion in oligopolies. Athey and Bagwell (2001) and Athey, Bagwell, and Sanchirico (2004) solve for the best trigger strategy-type equilibria in games with hidden information about cost types. Athey and Bagwell (2001) show that, in general, the best equilibrium is dynamic (nonstationary). In this equilibrium, a firm which sets low prices gets a lower discounted value of profits from then on. Athey, Bagwell, and Sanchirico (2004) show that when strategies are restricted to be strongly symmetric, so that all firms receive the same continuation values even though the

took observably different actions, a different result emerges. In particular, under some conditions, the best equilibrium is stationary and entails pooling of all cost types. When those conditions fail, and when firms are sufficiently patient, there may be a set of stationary and nonstationary equilibria that yield the same payoffs. (The latter result relies heavily on the Revenue Equivalence Theorem from auction theory.)

5. Conclusion

What is the optimal degree of discretion in monetary policy? For economies in which private information is not important or time inconsistency problems are severe the optimal degree of discretion is zero. For economies in which private information is important and time inconsistency problems are less severe, it is not zero, but bounded. More generally, the optimal degree of discretion is decreasing the more severe is the time inconsistency problem and the less important is private information. For all of these economies, the optimal policy can be implemented by enforcing a simple inflation cap.

In our simple model, the optimal inflation cap is a single number because there is no publicly observed state. If the model were extended to have a publicly observed state, then the optimal policy would respond to this state, but not to the private information. To implement optimal policy, society would need to specify a rule for setting the inflation cap, where the cap would vary with public information. Equivalently, society could specify a rule for setting ranges for acceptable inflation, where these ranges would vary with public information. We interpret these rules as a type of inflation targeting that is broadly similar to the types actually practiced by a fair number of countries. (For a discussion of inflation targeting in practice, see Bernanke and Mishkin (1997).)

To keep our theoretical model simple, we have abstracted from exotic events which are both unforeseen and unquantifiable. Anyone interpreting the implications of our results for an actual society, therefore, should keep in mind that to handle such exotic events, the optimal policy rule would need to be adapted to deal with such events, perhaps by the addition of some type of escape clauses.

References

- Abreu, Dilip; Pearce, David; and Stacchetti, Ennio. 1990. Toward a Theory of Discounted Repeated Games with Imperfect Monitoring. *Econometrica* 58(5, September), 1041–63.
- Albanesi, Stefania; Chari, V. V.; and Christiano, Lawrence. 2003. Expectation Traps and Monetary Policy. *Review of Economic Studies* 70(4, October), 715–41.
- Albanesi, Stefania, and Sleet, Christopher. 2002. Optimal Policy with Endogenous Fiscal Constitutions. Manuscript, Fuqua School of Business, Duke University.
- Amador, Manuel; Angeletos, George-Marios; and Werning, Iván. 2003. Optimal Commitment. Manuscript, Massachusetts Institute of Technology.
- Angeletos, George-Marios; Hellwig, Christian; and Pavan, Alessandro. 2003. Coordination and Policy Traps. Manuscript, Massachusetts Institute of Technology.
- Athey, Susan, and Bagwell, Kyle. 2001. Optimal Collusion with Private Information. *RAND Journal of Economics* 32(3, Autumn), 428–65.
- Athey, Susan; Bagwell, Kyle; and Sanchirico, Chris. 2004. Collusion and Price Rigidity. *Review of Economic Studies* 71(2), 317–49.
- Atkeson, Andrew. 1991. International Lending with Moral Hazard and Risk of Repudiation, *Econometrica*, 59(4), 1069–89.
- Atkeson, Andrew, and Lucas, Robert E., Jr. 1992. On Efficient Distribution with Private Information. *Review of Economic Studies* 59(3, July), 427–53.
- Backus, David, and Driffill, John. 1985. Inflation and Reputation. *American Economic Review* 75(3, June), 530–38.
- Barro, Robert J., and Gordon, David B. 1983. Rules, Discretion and Reputation in a Model of Monetary Policy. *Journal of Monetary Economics* 12(1, July), 101–21.
- Bernanke, Ben S., and Mishkin, Frederic S. 1997. Inflation Targeting: A New Framework for Monetary Policy? *Journal of Economic Perspectives* 11(2, Spring), 97–116.
- Bernanke, Ben S., and Woodford, Michael. 1997. Inflation Forecasts and Monetary Policy. *Journal of Money, Credit, and Banking*, Part 2, 29(4, November), 653–84.
- Canzoneri, Matthew B. 1985. Monetary Policy Games and the Role of Private Information. *American Economic Review* 75(5, December), 1056–70.
- Chang, Roberto. 1998. Credible Monetary Policy in an Infinite Horizon Model: Recursive Approaches. *Journal of Economic Theory* 81(2, August), 431–61.

- Chari, V. V.; Christiano, Lawrence J.; and Eichenbaum, Martin. 1998. Expectation Traps and Discretion. *Journal of Economic Theory* 81(2, August), 462–92.
- Chari, V. V., and Kehoe, Patrick J. 1990. Sustainable Plans. *Journal of Political Economy* 98(4, August), 783–802.
- Cukierman, Alex, and Meltzer, Allan H. 1986. A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information. *Econometrica* 54(5, September), 1099–1128.
- Da Costa, Carlos, and Werning, Iván. 2001. On the Optimality of the Friedman Rule with Heterogeneous Agents and Non-Linear Income Taxation. Manuscript, Massachusetts Institute of Technology.
- Faust, Jon, and Svensson, Lars E. O. 2001. Transparency and Credibility: Monetary Policy with Unobservable Goals. *International Economic Review* 42(2, May), 369–97.
- Fudenberg, Drew, and Tirole, Jean. 1991. *Game Theory*. Cambridge, Mass.: MIT Press.
- Green, Edward J. 1987. Lending and the Smoothing of Uninsurable Income. *Contractual Arrangements for Intertemporal Trade*, 3–25. Minnesota Studies in Macroeconomics series, vol. 1. Minneapolis: University of Minnesota Press.
- Green, Edward J., and Porter, Robert H. 1984. Noncooperative Collusion under Imperfect Price Information. *Econometrica* 52(1, January), 87–100.
- Ireland, Peter N. 1997. Sustainable Monetary Policies. *Journal of Economic Dynamics and Control* 22(1, November), 87–108.
- Ireland, Peter N. 2000. Expectations, Credibility, and Time-Consistent Monetary Policy. *Macroeconomic Dynamics* 4(4, December), 448–66.
- Kocherlakota, Narayana R. 1996. Implications of Efficient Risk Sharing without Commitment. *Review of Economic Studies* 63(4, October), 595–609.
- Kydland, Finn E., and Prescott, Edward C. 1977. Rules Rather Than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy* 85(3, June), 473–91.
- Nicolini, Juan Pablo. 1998. More on the Time Consistency of Monetary Policy. *Journal of Monetary Economics* 41(2, April), 333–50.
- Persson, Torsten, and Tabellini, Guido. 1993. Designing Institutions for Monetary Stability. *Carnegie-Rochester Conference Series on Public Policy* 39(December), 53–84.
- Phelan, Christopher, and Stacchetti, Ennio. 2001. Sequential Equilibria in a Ramsey Tax Model. *Econometrica* 69(6, November), 1491–1518.

- Rampini, Adriano A. 2003. Default and Aggregate Income. Manuscript, Kellogg School of Management, Northwestern University.
- Romer, Christina D., and Romer, David H. 2000. Federal Reserve Information and the Behavior of Interest Rates. *American Economic Review* 90(3, June), 429–57.
- Sleet, Christopher. 2001. On Credible Monetary Policy and Private Government Information. *Journal of Economic Theory* 99(1–2, July–August), 338–76.
- Sleet, Christopher. Optimal Taxation with Private Government Information. *Review of Economic Studies* forthcoming.
- Sleet, Christopher, and Yeltekin, Sevin. 2003. Credible Monetary Policy with Private Government Preferences. Manuscript, Kellogg School of Management, Northwestern University.
- Stokey, Nancy L. 2003. “Rules versus Discretion” After Twenty-Five Years. In *NBER Macroeconomics Annual 2002*, vol. 17, Mark Gertler and Kenneth Rogoff eds., 9–45. Cambridge, Mass.: MIT Press.
- Taylor, John B. 1983. Rules, Discretion and Reputation in a Model of Monetary Policy: Comments. *Journal of Monetary Economics* 12(1, July), 123–25.
- Thomas, Jonathan and Worrall, Tim. 1990. Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-Agent Problem. *Journal of Economic Theory*, 51(2), 367–90.
- Walsh, Carl E. 1995. Optimal Contracts for Central Bankers. *American Economic Review* 85(1, March), 150–67.

Appendix A: Proof of Lemma 3

Proof. By way of contradiction, suppose that $w(\cdot)$ jumps at some point $\tilde{\theta}$. Since $w(\cdot)$ is a step function, $w'(\theta) = 0$ in some interval (θ_1, θ_2) containing $\tilde{\theta}$. Clearly, this implies that either $w(\theta) < \bar{w}$ for all θ in $(\theta_1, \tilde{\theta})$ or $w(\theta) < \bar{w}$ for all θ in $(\tilde{\theta}, \theta_2)$. We know from (16) that at any point θ in the intervals $(\theta_1, \tilde{\theta})$ and $(\tilde{\theta}, \theta_2)$, either $\mu'(\theta) = 0$, so that $\mu(\cdot)$ is flat, or $R_\mu(x, \mu(\theta), \theta) = 0$, so that $\mu(\theta)$ equals the static best response. By continuity of the static best response, we can choose the points θ_1 and θ_2 to be close enough to $\tilde{\theta}$ so that $\mu(\cdot)$ either is constant on the interval $(\theta_1, \tilde{\theta})$ or equals the static best response on this interval, and similarly for the interval $(\tilde{\theta}, \theta_2)$.

Consider first the hard case, namely, when $\mu(\cdot)$ is constant on both $(\theta_1, \tilde{\theta})$ and $(\tilde{\theta}, \theta_2)$. Let (μ_1, w_1) denote the allocation on $(\theta_1, \tilde{\theta})$ and (μ_2, w_2) denote the allocation on $(\tilde{\theta}, \theta_2)$. By the continuity of R_μ , we can choose this interval (θ_1, θ_2) small enough so that if $R_\mu(x, \mu_1, \tilde{\theta})$ is strictly positive, then so is $R_\mu(x, \mu_1, \theta_1)$, and if $R_\mu(x, \mu_2, \tilde{\theta})$ is strictly negative, then so is $R_\mu(x, \mu_2, \theta_2)$.

Suppose that for the chosen interval (θ_1, θ_2) , the term $\Delta(a)$, defined in (36), is negative for small a . If the up variation is feasible, then we know it improves welfare, based on the same logic as in the proof of Lemma 2. By construction, the up variation is incentive-compatible. This variation is feasible outside the interval (θ_1, θ_2) , based on the logic of the proof of Lemma 2. We complete the proof for this case by showing that the variation is also feasible inside the interval (θ_1, θ_2) .

Suppose, initially, that $R(x, \mu_1, \tilde{\theta}) > R(x, \mu_2, \tilde{\theta})$. From (17) we have that $w_1 < w_2$, and from the feasibility of the original allocation that $w_2 \leq \bar{w}$. This case is illustrated in Figure 6a. For $\theta \in (\theta_1, \tilde{\theta})$, we thus know that $w(\theta; a) \leq \bar{w}$ for sufficiently small a .

For $\theta \in (\tilde{\theta}, \theta_2)$, if $w_2 < \bar{w}$, then since a is sufficiently small, $w(\theta; a) \leq \bar{w}$. If $w_2 = \bar{w}$, we show that $\partial \tilde{w}(\theta; 0)/\partial a$ is negative for $\theta \in (\tilde{\theta}, \theta_2)$ as follows. Differentiate (27) to obtain that $\partial \tilde{w}(\theta; 0)/\partial a$ equals

$$(44) \quad (\tilde{\mu} - \mu_1) \int_{\theta_1}^{\tilde{\theta}} R_{\theta\mu}(x, \mu_1, z) dz + (\tilde{\mu} - \mu_2) \int_{\tilde{\theta}}^{\theta_2} R_{\theta\mu}(x, \mu_2, z) dz - R_\mu(x, \mu_2, \tilde{\theta})(\tilde{\mu} - \mu_2).$$

Using $\int_{\tilde{\theta}}^{\theta_2} R_{\theta\mu}(x, \mu_2, z) dz = R_\mu(x, \mu_2, \theta_2) - R_\mu(x, \mu_2, \tilde{\theta})$ and an analogous expression for $\int_{\theta_1}^{\tilde{\theta}} R_{\theta\mu}(x, \mu_1, z) dz$, we can write (44) as

$$(45) \quad [R_\mu(x, \mu_1, \tilde{\theta}) - R_\mu(x, \mu_1, \theta_1)](\tilde{\mu} - \mu_2) - R_\mu(x, \mu_2, \tilde{\theta})(\tilde{\mu} - \mu_2).$$

We will show that (45) is negative. To do so, we begin by noting that $\Delta'(0) < 0$. This is true because $\Delta(0) = 0$, and we have assumed that $\Delta(a)$ is negative for small a . Using the form of $\mu(\theta)$ on the interval (θ_1, θ_2) , we have that

$$\Delta'(0) = (\tilde{\mu} - \mu_1) \int_{\theta_1}^{\tilde{\theta}} R_{\theta\mu}(x, \mu_1, \theta) d\theta + (\tilde{\mu} - \mu_2) \int_{\tilde{\theta}}^{\theta_2} R_{\theta\mu}(x, \mu_2, \theta) d\theta < 0.$$

Substituting for the integrals, we can write this inequality as

$$(46) \quad [R_\mu(x, \mu_1, \tilde{\theta}) - R_\mu(x, \mu_1, \theta_1)](\tilde{\mu} - \mu_1) + [R_\mu(x, \mu_2, \theta_2) - R_\mu(x, \mu_2, \tilde{\theta})](\tilde{\mu} - \mu_2) < 0.$$

Comparing the inequality in (46) with the expression in (45), we can see that a sufficient condition for (45) to be negative is that

$$(47) \quad R_\mu(x, \mu_2, \theta_2)(\tilde{\mu} - \mu_2) > 0.$$

We now show that (47) holds. Note that since $\mu(\cdot)$ is increasing on the interval (θ_1, θ_2) , it follows by definition that $\tilde{\mu} < \mu_2$, since, by construction, $\tilde{\mu}$ is the conditional mean of $\mu(\theta)$ on this interval. Thus, (47) is positive if $R_\mu(x, \mu_2, \theta_2)$ is negative. To see that $R_\mu(x, \mu_2, \theta_2)$ is negative, note that since $w_1 \leq \bar{w}$ and $w_2 = \bar{w}$, the incentive constraint $R(x, \mu_1, \tilde{\theta}) + w_1 = R(x, \mu_2, \tilde{\theta}) + \bar{w}$ implies that $R(x, \mu_1, \tilde{\theta}) \geq R(x, \mu_2, \tilde{\theta})$. Since $\mu_2 > \mu_1$ and R is strictly concave, we know that $R_\mu(x, \mu_2, \tilde{\theta}) < 0$. By our construction of the interval, since $R_\mu(x, \mu_2, \tilde{\theta})$ is strictly negative, so is $R_\mu(x, \mu_2, \theta_2)$. Thus, for this case, the up variation is feasible, incentive-compatible, and welfare-improving. An analogous argument holds when $R(x, \mu_1, \tilde{\theta}) < R(x, \mu_2, \tilde{\theta})$ and $\bar{w} \geq w_1 > w_2$, as in Figure 6b.

So far we have considered the case when $\mu(\cdot)$ is constant on both sides of $\tilde{\theta}$ and the term $\Delta(a)$ is negative for small a . In the case when $\mu(\cdot)$ is constant on both sides of $\tilde{\theta}$ but the term $\Delta(a)$ is positive for small a , we use the down variation and an analogous argument.

The case when $\mu(\cdot)$ is constant on one side of $\tilde{\theta}$ and equal to the static best response on the other side of $\tilde{\theta}$ is the easy case. Suppose, for example, that $\mu(\cdot)$ equals the static best response for θ on some interval $(\theta_1, \tilde{\theta})$. Here we simply take the relevant interval to be $(\theta_1, \tilde{\theta})$, from some point θ_1 just below the jump point $\tilde{\theta}$ up to the jump point $\tilde{\theta}$. Clearly, $\mu(\cdot)$ is increasing on the interval $(\theta_1, \tilde{\theta})$. We claim that $w(\cdot)$ is uniformly bounded below \bar{w} , and so Lemma 2 immediately applies.

We prove that $w(\cdot)$ is uniformly bounded below \bar{w} on $(\theta_1, \tilde{\theta})$ as follows. Since $\mu(\cdot)$ jumps up at $\tilde{\theta}$, it lies strictly above the static best response for some interval $(\tilde{\theta}, \theta_2)$, so that $\lim_{\theta \nearrow \tilde{\theta}} R(x, \mu(\theta), \tilde{\theta}) > \lim_{\theta \searrow \tilde{\theta}} R(x, \mu(\theta), \tilde{\theta})$. Hence, from condition (17) in local incentive-compatibility, we know that $\lim_{\theta \nearrow \tilde{\theta}} w(\theta) < \lim_{\theta \searrow \tilde{\theta}} w(\theta)$. Thus, for $\theta \in (\theta_1, \tilde{\theta})$, $w(\theta)$ is uniformly bounded below \bar{w} .

With an analogous argument, we can rule out the case in which $\mu(\theta)$ equals the static best response for θ on the other side of the jump point, on some interval $(\tilde{\theta}, \theta_2)$. *Q.E.D.*

Appendix B: Optimal Policy without Monotone Hazards

Here we give three examples in which our monotone hazard condition is violated and in which the optimal mechanism is dynamic. In the first two examples, we assume that the hazard $[1 - P(\theta)]/p(\theta)$ is decreasing in θ at all points except the point θ_1 , where the hazard jumps up. We also assume that $P(\theta)/p(\theta)$ is increasing throughout. In the third example, we shed light on the role of $R_{\mu\theta}$ in (A2) by assuming that the hazard $[1 - P(\theta)]/p(\theta)$ is decreasing throughout but that $[1 - P(\theta)]R_{\mu\theta}/p(\theta)$ is not.

For the first two examples assume that at the point θ_1 , we assume that

$$(48) \quad \int_{\underline{\theta}}^{\theta_1} \frac{1 - P(\theta)}{P(\theta_1)} d\theta < \int_{\theta_1}^{\tilde{\theta}} \frac{1 - P(\theta)}{1 - P(\theta_1)} d\theta.$$

To interpret this inequality, note that the left side is the conditional mean of the function $[1 - P(\theta)]/p(\theta)$ over the interval $[\underline{\theta}, \theta_1]$ while the right side is the conditional mean of this function over the interval $(\theta_1, \bar{\theta}]$. Clearly, for any distribution for which $[1 - P(\theta)]/p(\theta)$ is decreasing throughout $[\underline{\theta}, \bar{\theta}]$, this inequality is reversed.

It is easy to show that a two-piece uniform distribution with $p(\theta) = \rho_1$ if $\theta \leq \theta_1$ and $p(\theta) = \rho_2$ if $\theta > \theta_1$ will satisfy (48) if ρ_2 is chosen to be sufficiently small relative to ρ_1 . In this case, illustrated in Figure 7, the function $[1 - P(\theta)]/p(\theta)$ will jump up sufficiently at θ_1 so that the conditional mean of this function over the higher interval $[\theta_1, \bar{\theta}]$ is larger than the conditional mean over the lower interval $[\underline{\theta}, \theta_1]$.

In the first example, the *linear example*, we make the calculations trivial by assuming that $R(x, \mu, \theta) = (\theta - \underline{\theta})\mu + r(x)$ with $r(x) = -x^2/2$. In the second example, which is the *benchmark example* of (1), we assume that

$$(49) \quad R(x, \mu, \theta) = -\frac{1}{2} \left[(U + x - \mu)^2 + (\mu - \theta)^2 \right].$$

In the third example, the *discrete example*, $R(x, \mu, \theta) = g(\theta)\mu + r(x)$ with g an increasing nonlinear function.

All three of these examples satisfy the single-crossing property (A1). In the first two of them, $R_{\theta\mu} = 1$, so that the conditions (A2) reduce to the standard monotone hazard conditions. Note that for the first two examples, any distribution that satisfies (48) is inconsistent with the monotone hazard condition (A2a).

The Linear Example

Notice that any solution to the mechanism design problem must have the two-piece form

$$(50) \quad (\mu(\theta), w(\theta)) = \begin{cases} (\mu_1, w_1) & \text{for } \theta \in [\underline{\theta}, \theta_1) \\ (\mu_2, w_2) & \text{for } \theta \in [\theta_1, \bar{\theta}] \end{cases}.$$

This follows because the arguments used in Lemmas 1 and 2 can be applied separately to the intervals $[\underline{\theta}, \theta_1)$ and $(\theta_1, \bar{\theta}]$ and because for any $\theta > \underline{\theta}$, the static best response to any x in the interval $[\underline{\mu}, \bar{\mu}]$ is a constant, namely, the upper limit $\bar{\mu}$. Since this policy must satisfy the incentive constraint $(\theta_1 - \underline{\theta})\mu_1 + w_1 = (\theta_1 - \underline{\theta})\mu_2 + w_2$, the monotonicity condition $\mu_1 \leq \mu_2$ implies that $w_1 \geq w_2$. Thus, we know that $w_1 = \bar{w}$ and that the constraint $w_2 \leq \bar{w}$ will be automatically satisfied by any monotonic policy.

The mechanism design problem then reduces to the linear problem of choosing μ_1 , μ_2 , and x to maximize

$$r(x) + \bar{w} + \mu_1 \int_{\underline{\theta}}^{\theta_1} \frac{1 - P(\theta)}{p(\theta)} p(\theta) d\theta + \mu_2 \int_{\theta_1}^{\bar{\theta}} \frac{1 - P(\theta)}{p(\theta)} p(\theta) d\theta$$

subject to the constraints that $\underline{\mu} \leq \mu_1 \leq \mu_2 \leq \bar{\mu}$ and that $x = P(\theta_1)\mu_1 + [1 - P(\theta_1)]\mu_2$. If (48) holds and if the lower and upper limits $\underline{\mu}, \bar{\mu}$ include the expected Ramsey policy, then the optimal

policy will have either $\underline{\mu} = \mu_1 < \mu_2$ or $\mu_1 < \mu_2 = \bar{\mu}$. To see this, consider spreading out the policy by decreasing μ_1 by Δ_1 and increasing μ_2 by Δ_2 so that the change in expected inflation $[1 - P(\theta_1)]\Delta_2 - P(\theta_1)\Delta_1$ is zero. The associated welfare change can be written as

$$(51) \quad \left[- \int_{\underline{\theta}}^{\theta_1} \frac{1 - P(\theta)}{P(\theta_1)} d\theta + \int_{\theta_2}^{\bar{\theta}} \frac{1 - P(\theta)}{1 - P(\theta_1)} d\theta \right] P(\theta_1)\Delta_1 > 0$$

where the inequality follows from (48). Hence, the solution must have $\mu_1 < \mu_2$, and from the incentive constraint, we then know that $w_2 < w_1 = \bar{w}$. Thus, the solution to the mechanism design problem is necessarily dynamic.

The Benchmark Example

Assume that the policy $\mu(\cdot)$, which solves the static mechanism design problem, has bounded discretion and that $\theta_1 > \theta^*$, so that the jump point in the hazard occurs on the flat portion of that policy. (We can construct a numerical example in which this assumption holds.) We will show that there is a dynamic mechanism that improves on the optimal static mechanism. The basic idea is to use a variation that spreads out the policy as a function of type instead of flattens it as we did in Lemmas 1 and 2.

This variation is similar to the one in the linear example. Consider an alternative policy that lowers inflation for types at or below θ_1 , raises it for types above θ_1 , and keeps expected inflation constant:

$$\tilde{\mu}(\theta) = \begin{cases} \mu(\theta) - \Delta_0 & \text{if } \theta \leq \theta_1 \\ \mu(\theta) + \Delta_1 & \text{if } \theta > \theta_1 \end{cases}$$

with $\Delta_0, \Delta_1 > 0$ and $[1 - P(\theta_1)]\Delta_1 - P(\theta_1)\Delta_0 = 0$, so that expected inflation is constant. Note that this alternative policy $\tilde{\mu}(\cdot)$ is monotonically increasing since $\mu(\cdot)$ must be monotonically increasing. Our variation is a marginal shift from $\mu(\cdot)$ toward $\tilde{\mu}(\cdot)$ defined as $\mu(\theta; a) = a\tilde{\mu}(\theta) + (1 - a)\mu(\theta)$ for each θ . Welfare is given by

$$V(a) = R(x, \mu(\underline{\theta}; a), \underline{\theta}) + \bar{w} + \int_{\underline{\theta}}^{\bar{\theta}} \frac{1 - P(\theta)}{p(\theta)} R_{\theta\mu}(x, \mu(\theta; a), \theta) p(\theta) d\theta.$$

The impact of this variation on welfare is given by

$$(52) \quad \frac{\partial V(0)}{\partial a} = -\Delta_0 R_{\mu}(x, \mu(\underline{\theta}), \underline{\theta}) - \Delta_0 \int_{\underline{\theta}}^{\theta_1} \frac{1 - P(z)}{p(z)} R_{\theta\mu}(x, \mu(z), z) p(z) dz \\ + \Delta_1 \int_{\theta_1}^{\bar{\theta}} \frac{1 - P(z)}{p(z)} R_{\theta\mu}(x, \mu(z), z) p(z) dz.$$

Since $\mu(\theta)$ has bounded discretion, $R_{\mu}(x, \mu(\underline{\theta}), \underline{\theta}) = 0$. In our quadratic example, $R_{\theta\mu}(x, \mu(z), z) = 1$; hence, (52) reduces to (51), which we know from (48) is positive.

It is straightforward, but somewhat tedious, to show that the associated continuation values $w(\theta; a)$ defined by

$$R(x, \mu(\underline{\theta}; a), \underline{\theta}) + \bar{w} + \int_{\underline{\theta}}^{\theta} R_{\theta}(x, \mu(z; a)) dz - R(x, \mu(\theta; a), \theta)$$

have $\partial w(\theta; 0) / \partial a \leq 0$ for all θ and $\partial w(\theta; 0) / \partial a < 0$ for $\theta > \theta_1$. To do so, we use the facts that $R_{\mu}(x, \mu(\underline{\theta}), \underline{\theta}) = 0$ and that $\theta_1 > \theta^*$, so that $\mu(\theta) = \mu(\theta_1)$ for $\theta \geq \theta_1$. These results imply that this variation both improves welfare and is feasible. Thus, the optimal mechanism must be dynamic.

Note that if $\mu(\cdot)$ has no discretion, then we need a different condition on the distribution to show that the static mechanism is not optimal. This is because when $\mu(\cdot)$ has no discretion, we can have $R_{\mu}(x, \mu(\underline{\theta}), \underline{\theta}) > 0$, and the above argument that $\partial w(\theta; 0) / \partial a \leq 0$ for all θ does not go through. When $\mu(\cdot)$ has no discretion, the analog of the condition (48) is that at $x = \mu = \mu^{ER}$, there exists a θ_1 such that

$$R_{\mu}(\mu^{ER}, \mu^{ER}, \underline{\theta}) + \int_{\underline{\theta}}^{\theta_1} \frac{1 - P(z)}{P(\theta_1)} dz < \int_{\theta_1}^{\bar{\theta}} \frac{1 - P(z)}{1 - P(\theta_1)} dz.$$

With this condition, the optimal mechanism is dynamic rather than static. Note that, in our linear example, this distinction did not come up because in the linear example, our utility function is such that $R_{\mu}(x, \mu(\underline{\theta}), \underline{\theta}) = 0$ with no discretion.

The Discrete Example

Let the types be θ_i for $i = 1, 2, 3$ with associated probabilities p_i and let $P_i = \sum_{j=0}^i p_j$. It is easy to show that under the discrete analog of (A1) the only relevant incentive constraints are

$$(53) \quad R(x, \mu_i, \theta_i) + w_i \geq (1 - \beta)R(x, \mu_{i+1}, \theta_i) + \beta w_{i+1}$$

for $i = 1$ and 2 . The discrete analog of (A2) for types θ_2 and θ_3 is

$$\frac{1 - P_1}{p_2} [R_{\mu}(x, \mu_2, \theta_2) - R_{\mu}(x, \mu_2, \theta_1)] > \frac{1 - P_2}{p_3} [R_{\mu}(x, \mu_3, \theta_3) - R_{\mu}(x, \mu_3, \theta_2)]$$

which here reduces to

$$(54) \quad \frac{1 - P_1}{p_2} [g(\theta_2) - g(\theta_1)] > \frac{1 - P_2}{p_3} [g(\theta_3) - g(\theta_2)]$$

We now give an example in which the hazard $(1 - P_i) / p_{i+1}$ is monotone but g is increasing sufficiently convex so that (54) is violated and the optimal policy is dynamic. Suppose that $\mu_2 = \mu_3$ is part of a candidate optimal policy. Consider the variation of decreasing μ_1 and μ_2 by Δ and increasing μ_3 by $(p_1 + p_2)\Delta / p_3$ so that expected inflation x is constant. We can maintain incentives by keeping w_1 and w_2 unchanged and lowering w_3 by $\theta_3\Delta / p_3$. This variation leads to a change in welfare of

$$(p_1 + p_2)g(\theta_3) - (1 + p_2)g(\theta_2) - p_1g(\theta_1)$$

With a uniform distribution, $p_i = 1/3$, and $g(\theta_1) = 1$, $g(\theta_2) = 2$ this variation is welfare improving as long as $g(\theta_3) > 9/2$.

We can summarize these three examples as follows. In each example we showed that welfare could be improved relative to a static policy by raising inflation for high types and lowering inflation for low types so as to keep expected inflation constant. In the first two examples this improved welfare because there were sufficiently few high types relative to low types so that we could raise the inflation a lot for the types who valued it relatively more and lower only a little for the types who valued it relatively less. In the third example, even though the distribution of types is uniform, the high types valued inflation so much more than the lower types that by raising their inflation and lowering it for the lower types we could raise welfare.

Appendix C: Implementation with an Inflation Cap

Here we prove that the equilibrium outcome in an economy with an inflation cap is the optimal outcome of the mechanism design problem. We show this result formally using the following one-shot game in which we drop time subscripts.

With an inflation cap of $\bar{\pi}$ in the current period, the problem of the monetary authority at a given θ is as follows: Given aggregate wages x , choose money growth $\mu(\theta)$ for this state θ to maximize $R(x, \mu, \theta)$ subject to $\mu(\theta) \leq \bar{\pi}$. The private agents' decisions on wages are summarized by $x = \int \mu(\theta)p(\theta)$.

An equilibrium of this one-shot game consists of aggregate wages x and a money growth policy $\mu(\cdot)$ such that (i) with x given, $\mu(\cdot)$ satisfies $\mu(\theta) \leq \bar{\pi}$, and (ii) $x = \int \mu(\theta)p(\theta)$. We denote the optimal choice of the monetary authority as $\mu^*(\cdot; x, \bar{\pi})$. This notation reflects the fact that the monetary authority is choosing a static best response to x given that its choice set is restricted by $\bar{\pi}$, which we call the *inflation cap*.

To implement the best equilibrium in the dynamic game, we choose $\bar{\pi}$ as follows. Whenever the expected Ramsey policy is optimal, we choose the inflation cap to be

$$(55) \quad \bar{\pi} = \mu^{ER}.$$

Whenever bounded discretion is optimal, we choose the cap $\bar{\pi}$ to be the money growth rate chosen by the cutoff type θ^* :

$$(56) \quad \bar{\pi} = \mu^*(\theta^*, x^*)$$

where x^* is the equilibrium inflation rate with this level of bounded discretion.

Proposition 5. Assume (A1), (A2), and that the inflation cap $\bar{\pi}$ is set according to (55) and (56). Then the equilibrium outcome of the one-shot game with the inflation cap for each period coincides with the optimal equilibrium outcome of the dynamic game.

Proof. To establish this result, we first show that the monetary authority will choose the upper bound $\bar{\pi} = \mu^{ER}$ when the expected Ramsey policy is optimal in the dynamic game. Note that Proposition 3 implies that whenever the expected Ramsey policy is optimal, $\mu^{ER} \leq \mu^*(\underline{\theta}; \mu^{ER})$. Also, recall that the single-crossing assumption (A1) implies that the best response is strictly increasing in θ , so that (??) holds. Thus, $\mu^*(\underline{\theta}; \mu^{ER}) \leq \mu^*(\theta; \mu^{ER})$ for all θ . Hence, at the expected Ramsey policies and the associated inflation rate, all types want to deviate by increasing their inflation above μ^{ER} ; hence, the constraint $\bar{\pi} = \mu^{ER}$ binds, and all types choose the expected Ramsey levels.

We next show that if bounded discretion is optimal in the dynamic game, then in the associated static game with the inflation cap, all types choose the bounded discretion policies. For all types $\theta \leq \theta^*$, the policies under bounded discretion are simply the static best responses, and these clearly coincide with those in the static game. For all types θ above θ^* , the policies under bounded discretion are the static best responses of the θ^* type, namely, $\mu^*(\theta; x^*)$, where x^* is the equilibrium expected inflation rate under bounded discretion. Under assumption (A1), the static best responses are increasing in the type, so that the best response of any type $\theta \geq \theta^*$ is above $\mu^*(\theta; x^*)$. Thus, in the one-shot game with the inflation cap, the constraint (56) binds for such types. Thus, the equilibrium outcomes of the two games coincide. *Q.E.D.*

Appendix D: Proof of Proposition 4

The optimal policy with bounded discretion is found as the solution to the problem of choosing θ^* and x to maximize

$$(57) \quad \int_{\underline{\theta}}^{\theta^*} R(x, \mu^*(\theta; x), \theta) p(\theta) d\theta + \int_{\theta^*}^{\bar{\theta}} R(x, \mu^*(\theta^*; x), \theta) p(\theta) d\theta$$

$$x = \int_{\underline{\theta}}^{\theta^*} \mu^*(\theta, x) p(\theta) d\theta + \int_{\theta^*}^{\bar{\theta}} \mu^*(\theta^*; x) p(\theta) d\theta.$$

Letting λ be the Lagrange multiplier on (57), the first order conditions for θ^* and x imply that the derivative of the objective function with respect to θ^* is

$$\int_{\underline{\theta}}^{\theta^*} R_x(x, \mu^*(\theta; x), \theta) p(\theta) d\theta + \int_{\theta^*}^{\bar{\theta}} R_x(x, \mu^*(\theta^*; x), \theta) p(\theta) d\theta$$

$$+ \left[\int_{\theta^*}^{\bar{\theta}} R_\mu(x, \mu^*(\theta^*; x), \theta) \frac{p(\theta)}{1 - P(\theta^*)} d\theta \right] \left[1 - \int_{\underline{\theta}}^{\theta^*} \frac{\partial \mu^*(\theta^*, x)}{\partial x} p(\theta) d\theta \right].$$

Using our functional forms and $x = \int \mu(\theta) p(\theta) d\theta$ we can simplify this derivative to

$$(58) \quad \left[\int_{\theta^*}^{\bar{\theta}} (\theta - \theta^*) \frac{p(\theta)}{1 - P(\theta^*)} d\theta \right] \left[1 - \frac{P(\theta^*)}{2} \right] - \frac{U}{\alpha}$$

We can show that, under (A2), this derivative is strictly decreasing in θ^* as follows. Integration by parts gives

$$\int_{\theta^*}^{\bar{\theta}} (\theta - \theta^*) p(\theta) d\theta = \int_{\theta^*}^{\bar{\theta}} (1 - P(\theta)) d\theta,$$

so that (58) is equivalent to

$$\left[\int_{\theta^*}^{\bar{\theta}} \frac{1 - P(\theta)}{p(\theta)} \frac{p(\theta)}{1 - P(\theta^*)} d\theta \right] \left[1 - \frac{P(\theta^*)}{2} \right] - \frac{U}{\alpha},$$

and this expression is clearly strictly decreasing in θ^* under (A2).

The fact that (58) is strictly decreasing in θ^* implies that there are three possible cases characterizing the optimal policy with bounded discretion, depending on the value of U/α . In the first case, the derivative (58) is positive for all θ^* and the solution is $\theta^* = \bar{\theta}$. Since the left side of (58) equals zero when $\theta^* = \bar{\theta}$, this case occurs only when $U/\alpha = 0$. As is clear, in this case, there is no time inconsistency problem and the Ramsey policy is incentive compatible. In the second case, the derivative (58) is negative for all θ^* and the solution is $\theta^* = \underline{\theta}$. Since the derivative (58) evaluated at $\theta^* = \underline{\theta}$ reduces to $-\underline{\theta} - U/\alpha$, this case occurs when $U/\alpha \geq -\underline{\theta} > 0$. Note that in this case, the optimal policy with bounded discretion specifies a constant inflation rate and hence is dominated, at least weakly, by the expected Ramsey policy with no discretion. Hence, we say that in this case, the optimal policy has no discretion. In the third case, there is an interior θ^* that sets the derivative (58) to zero. This case occurs when $0 < U/\alpha < -\underline{\theta}$. Clearly, in this case, the value of θ^* characterizing the optimal degree of discretion is decreasing in U/α .

Finally, to complete the proof, we must show that when $0 < U/\alpha < -\underline{\theta}$, the optimal policy with bounded discretion dominates the expected Ramsey policy. To do so, we use Proposition 3 part *ii*). Note that when $U/\alpha < -\underline{\theta}$ we have

$$\mu^*(\underline{\theta}, \mu^{ER}) = \frac{U + \alpha \underline{\theta}}{2} < \mu^{ER} = 0.$$

The result then follows directly from Proposition 3.

Appendix E: The Role of the Constraint $w(\theta) \leq \bar{w}$

Here we develop a highly stylized example (about traffic congestion) that illustrates the importance of the constraint

$$(59) \quad w(\theta) \leq \bar{w}$$

in generating our result that the optimal policy is static. In the example we replace the constraint (53) with the constraint

$$(60) \quad \int w(\theta) p(\theta) d\theta \leq \bar{w}.$$

and shows that the optimal mechanism differs radically. To be concrete consider a mechanism design problem of choosing $\mu(\cdot)$ and $w(\cdot)$ to solve

$$\max \int_{\underline{\theta}}^{\bar{\theta}} [R(x, \mu(\theta), \theta) + w(\theta)]p(\theta)d\theta$$

$$(61) \quad R(x, \mu(\theta), \theta) + w(\theta) \geq R(x, \mu(\hat{\theta}), \theta) + w(\hat{\theta})$$

$$(62) \quad x \geq \int \mu(\theta)p(\theta)d\theta$$

and (60). One interpretation of this problem is as follows. A large number of people want to share a road. Each person differs in their desire to use the road, as indexed by the privately observed θ . Let $\mu(\theta)$ denote the time that type θ is allowed to drive. Let x denote the average traffic on the road, as denoted by (62). Because of congestion people dislike higher average traffic (R is decreasing in x). Let $w(\theta)$ denote the toll to drive $\mu(\theta)$. Constraint (60) is a budget constraint on tolls — where \bar{w} is the money needed to operate the road, possibly zero.

It is easy to see that here the optimal $w(\theta)$ varies with θ . Specifically, it is possible to choose $w(\theta)$ in such a way as to support the first best. (Here we are assuming A1, so that the first best schedule for $\mu(\theta)$ is upward sloping. To see this result drop the incentive constraint (61) and solve for the first best $\mu^*(\theta)$, then use the local incentive compatibility condition to construct the $w^*(\theta)$ function (up to the constant $w^*(\underline{\theta})$) that makes $\mu^*(\theta)$ incentive compatible. Finally, choose the constant $w^*(\underline{\theta})$ to satisfy (60).) Clearly, the answer to this problem is very different from the answer to our problem because here the optimal $w(\theta)$ varies with θ , while in ours it does not and $w(\theta) = \bar{w}$.

Note that the result that the first best is incentive compatible is special to this functional form in which payoffs are linear in w . If instead we had

$$\max \int_{\underline{\theta}}^{\bar{\theta}} [R(x, \mu(\theta), \theta) + U(w(\theta))]p(\theta)d\theta$$

with U concave then we would have the standard tradeoff between insurance (or redistribution) and incentives.

How could we interpret our model and results in this road congestion context? Suppose that it is infeasible to use tolls and the only way to ration road use is to make individuals wait to get on the road. Let $t(\theta) \geq 0$ be the amount of time someone has to wait to drive $\mu(\theta)$ and let $w(\theta) = \bar{w} - t(\theta)$ be the associated utility from waiting $t(\theta)$. Then $t(\theta) \geq 0$ is, of course, equivalent to $w(\theta) \leq \bar{w}$. In this context we get a very different answer than when you can use tolls. Under A1 and A2, the optimal scheme is to have no one wait ($t(\theta) = 0$) and let everyone drive as much as they like subject to a cap, $\mu(\theta) \leq \mu^*$.

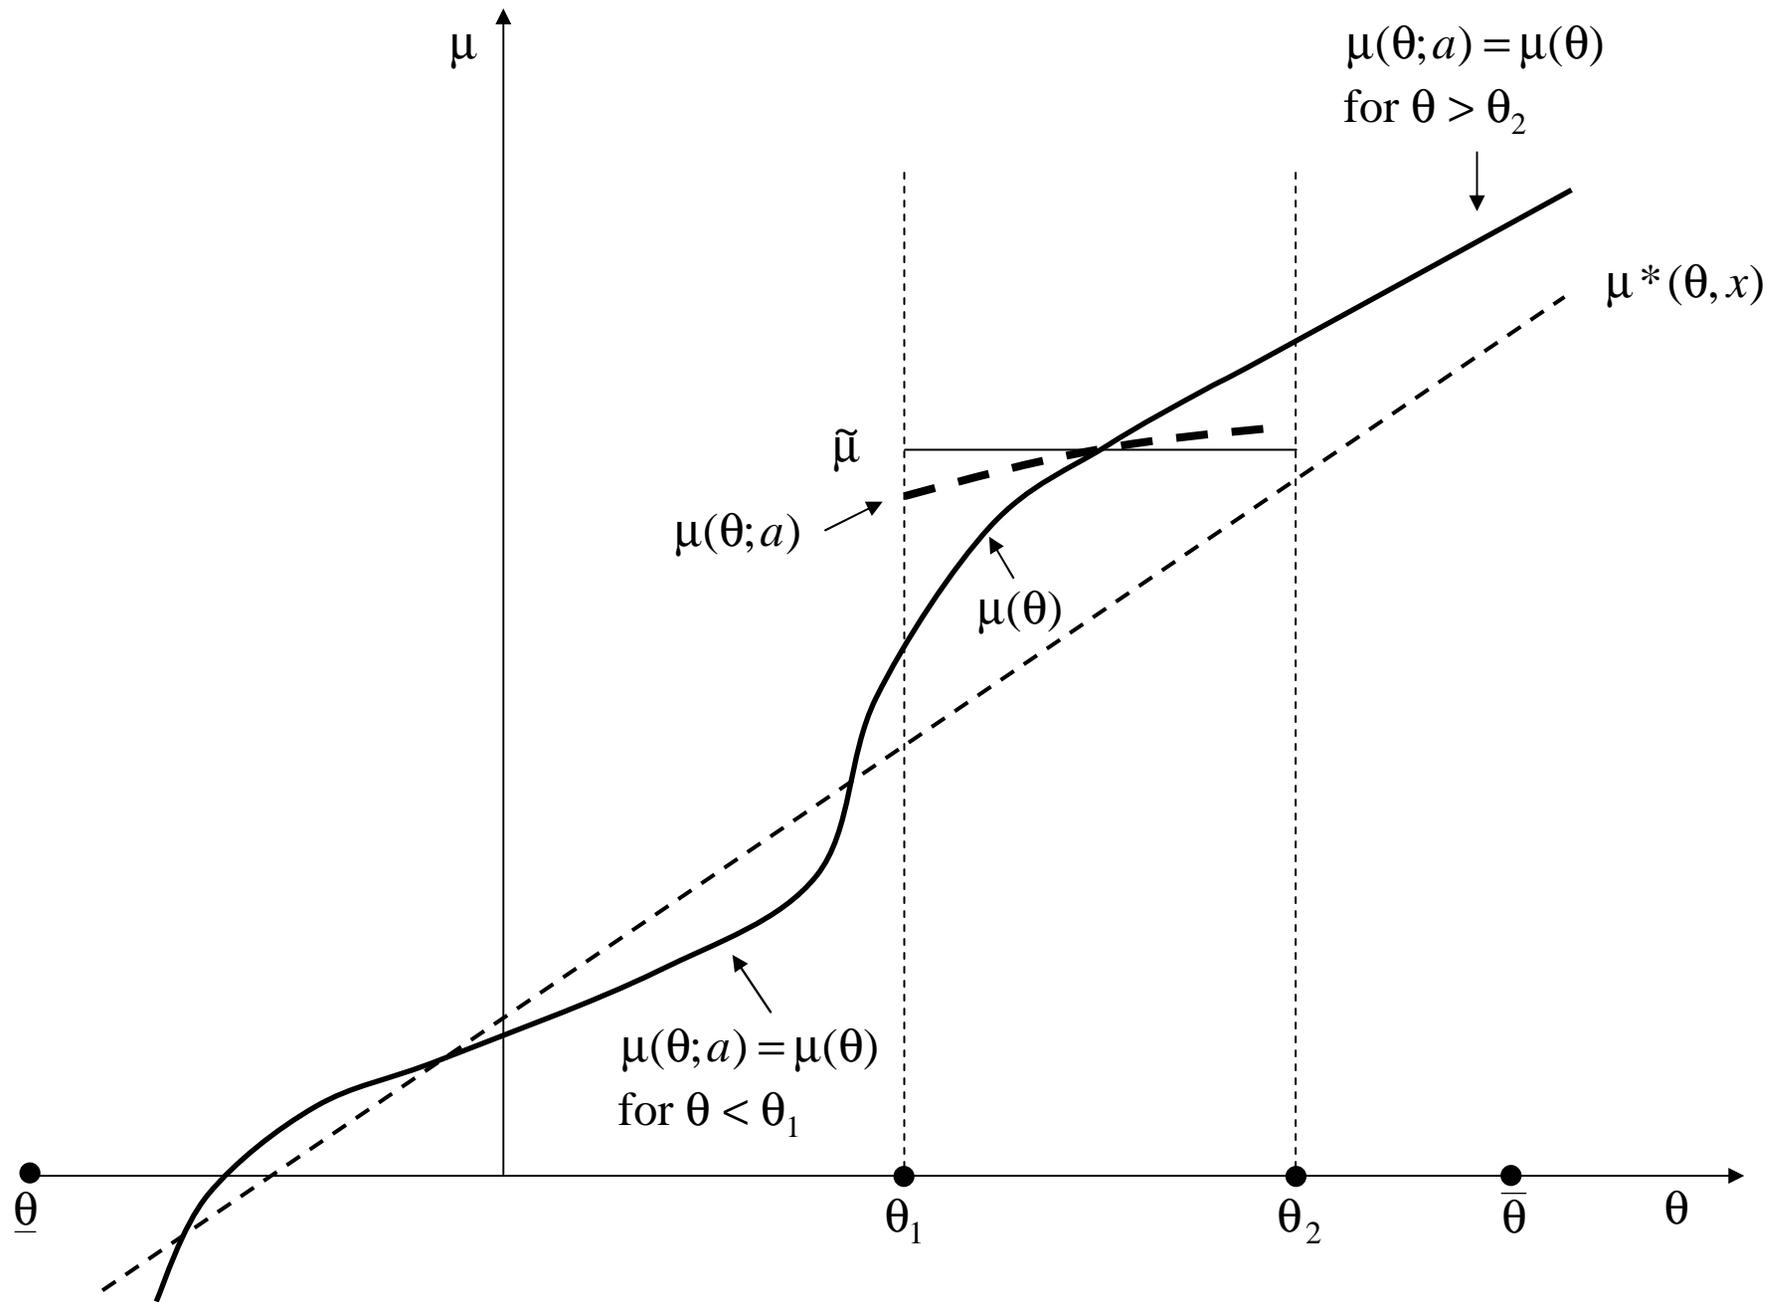


Figure 1: A Welfare-Improving Variation in μ

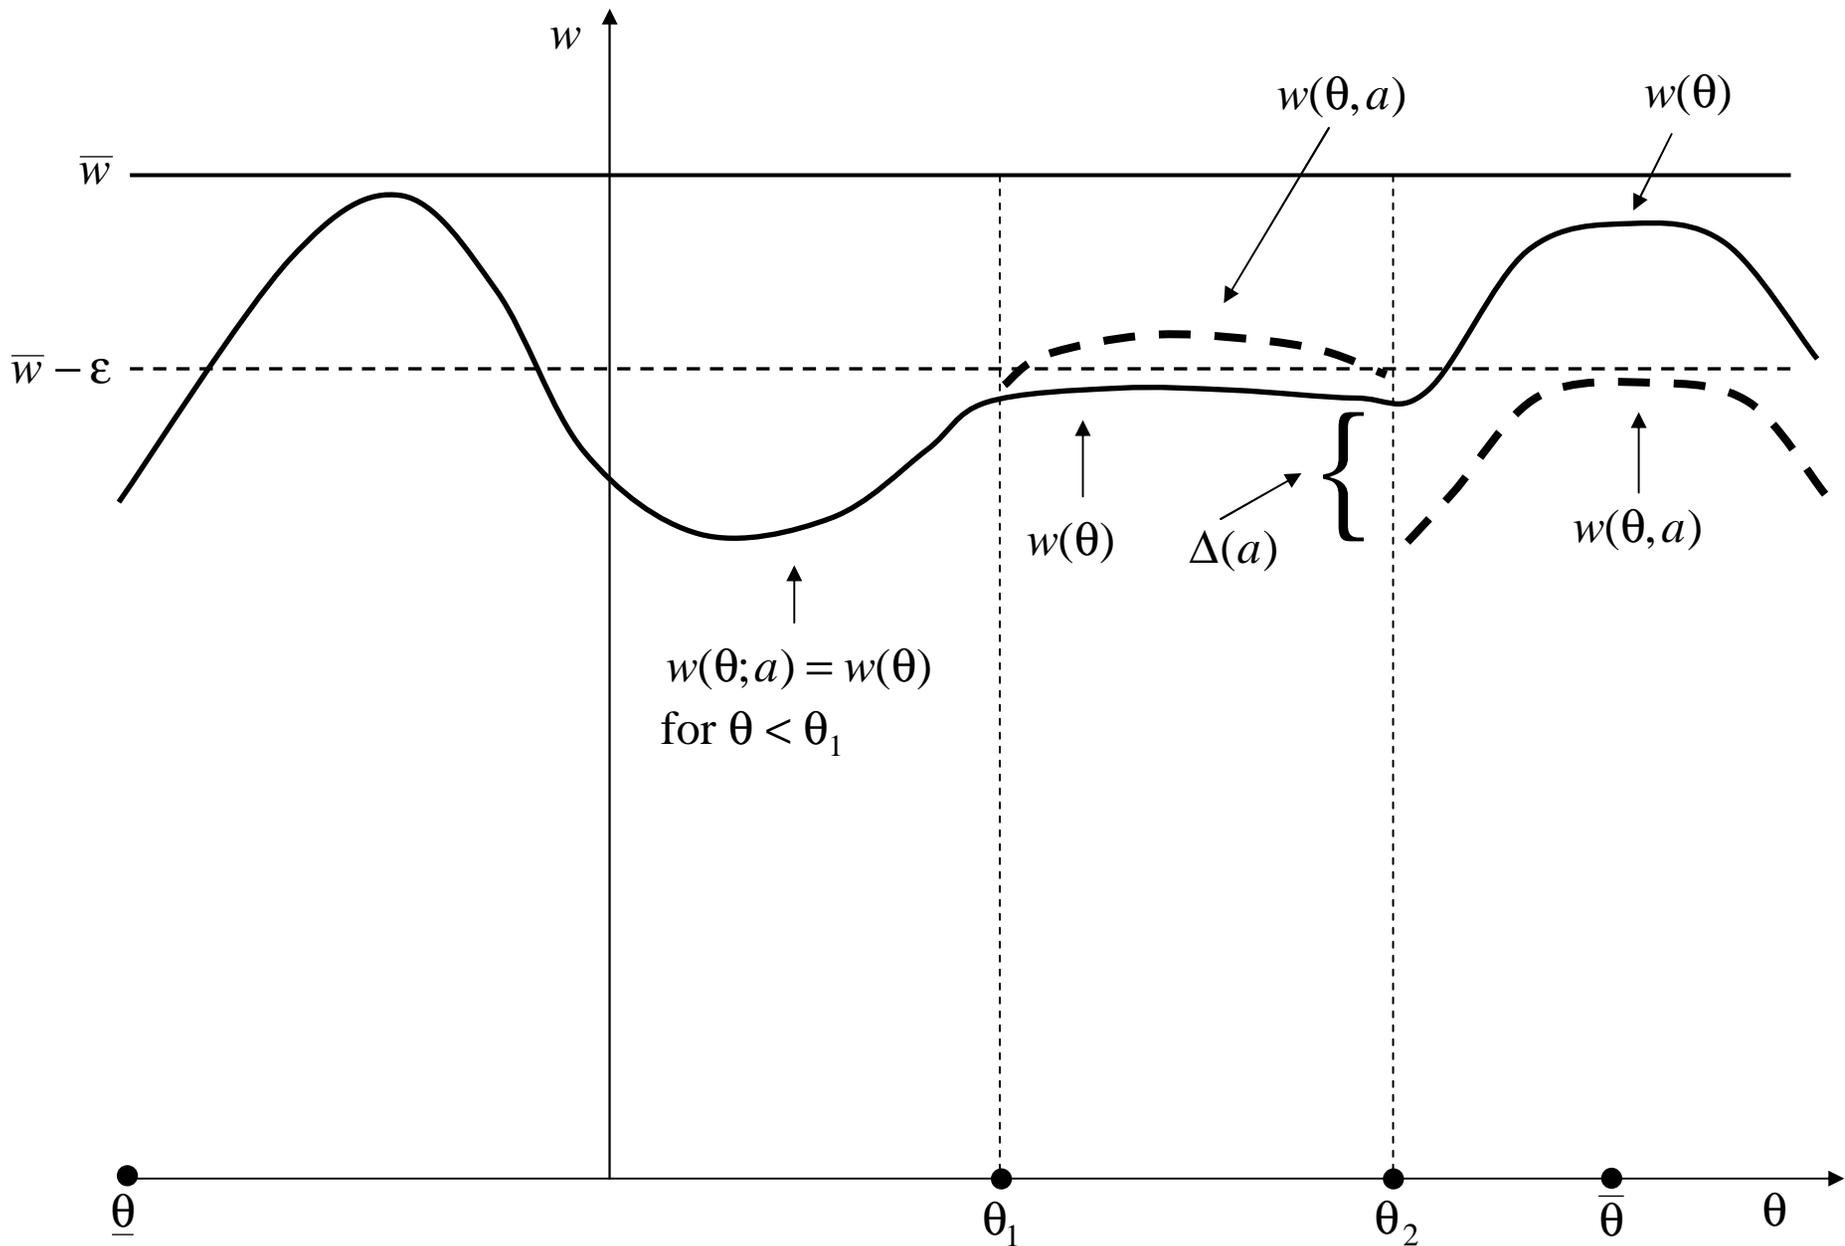


Figure 2: The Continuation Value in the Up Variation

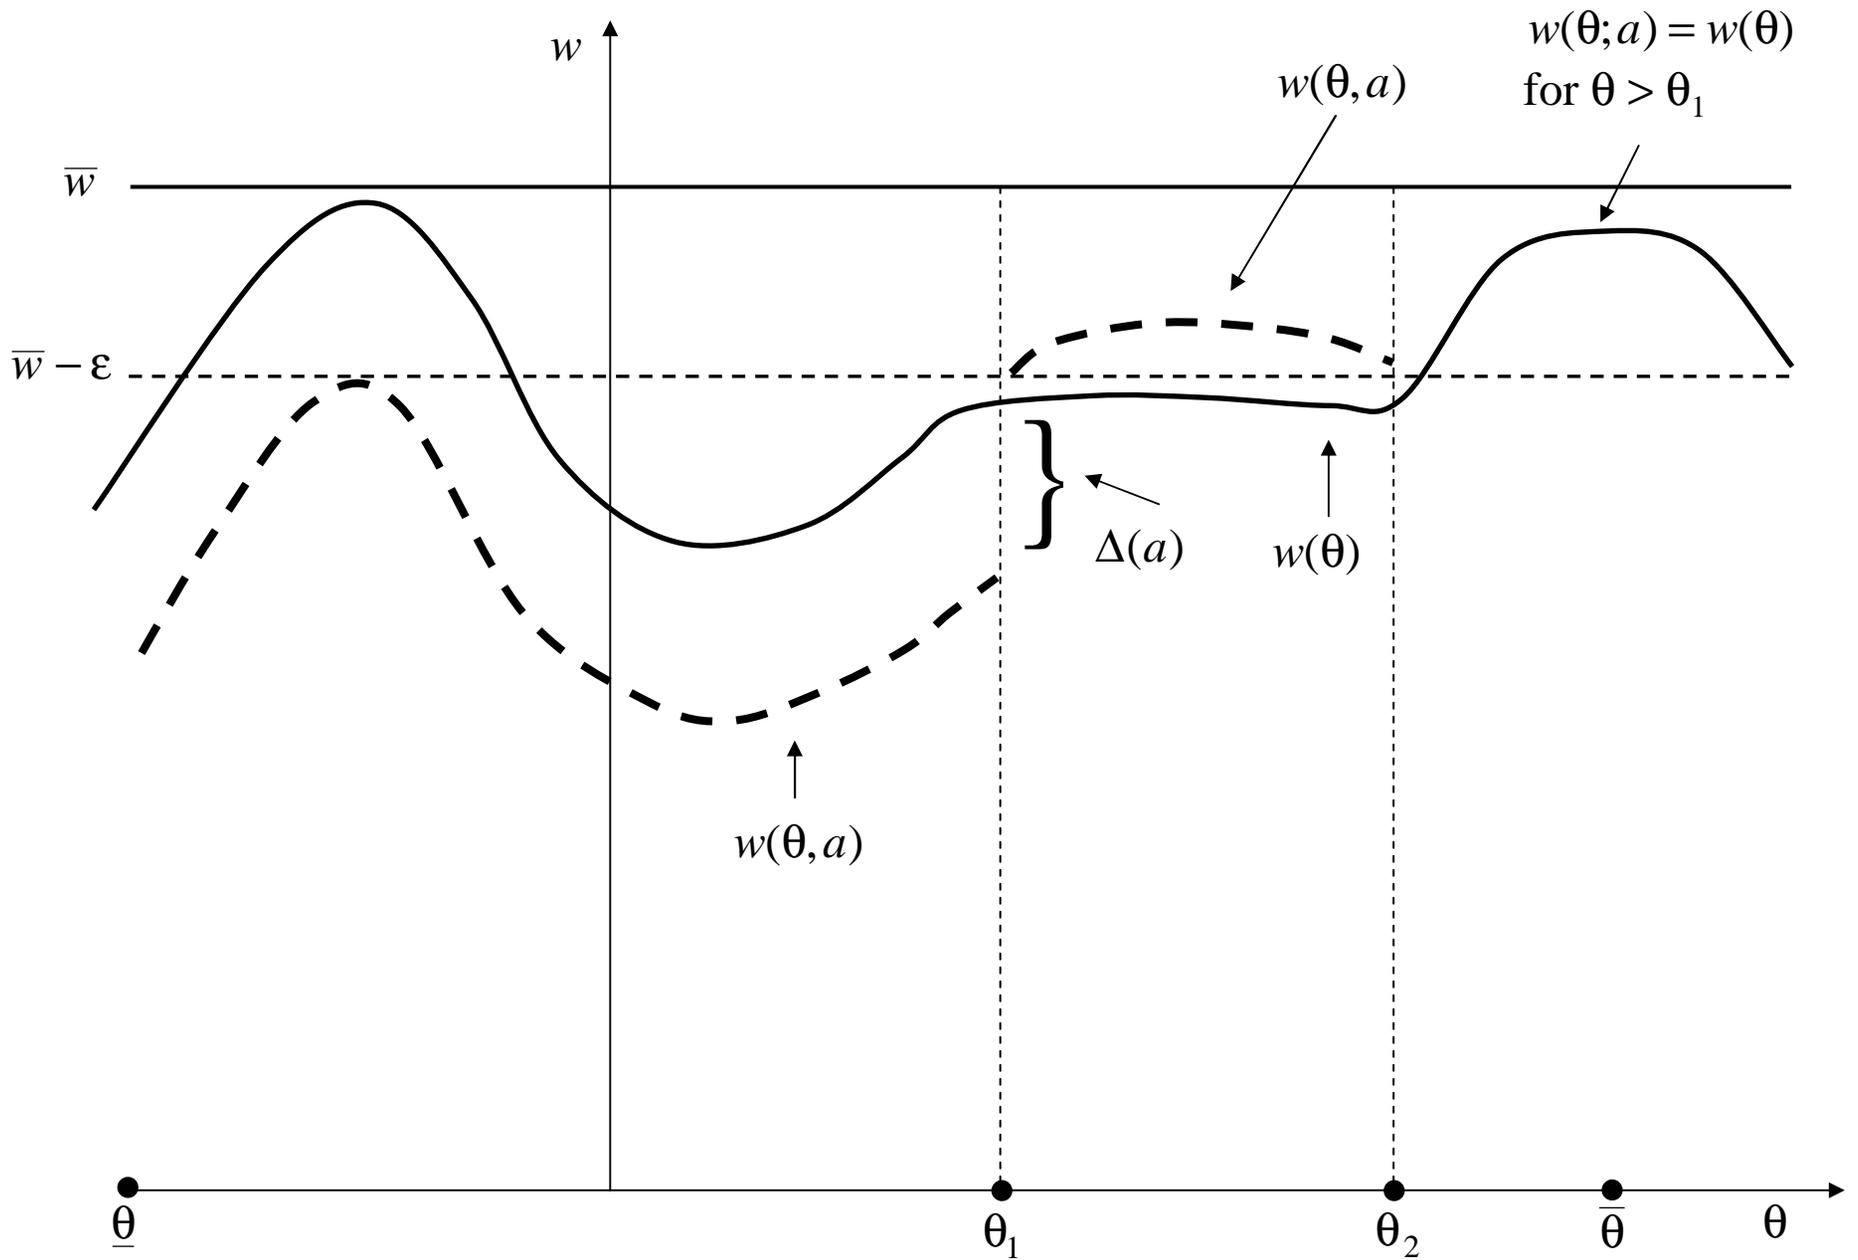


Figure 3: The Continuation Value in the Down Variation

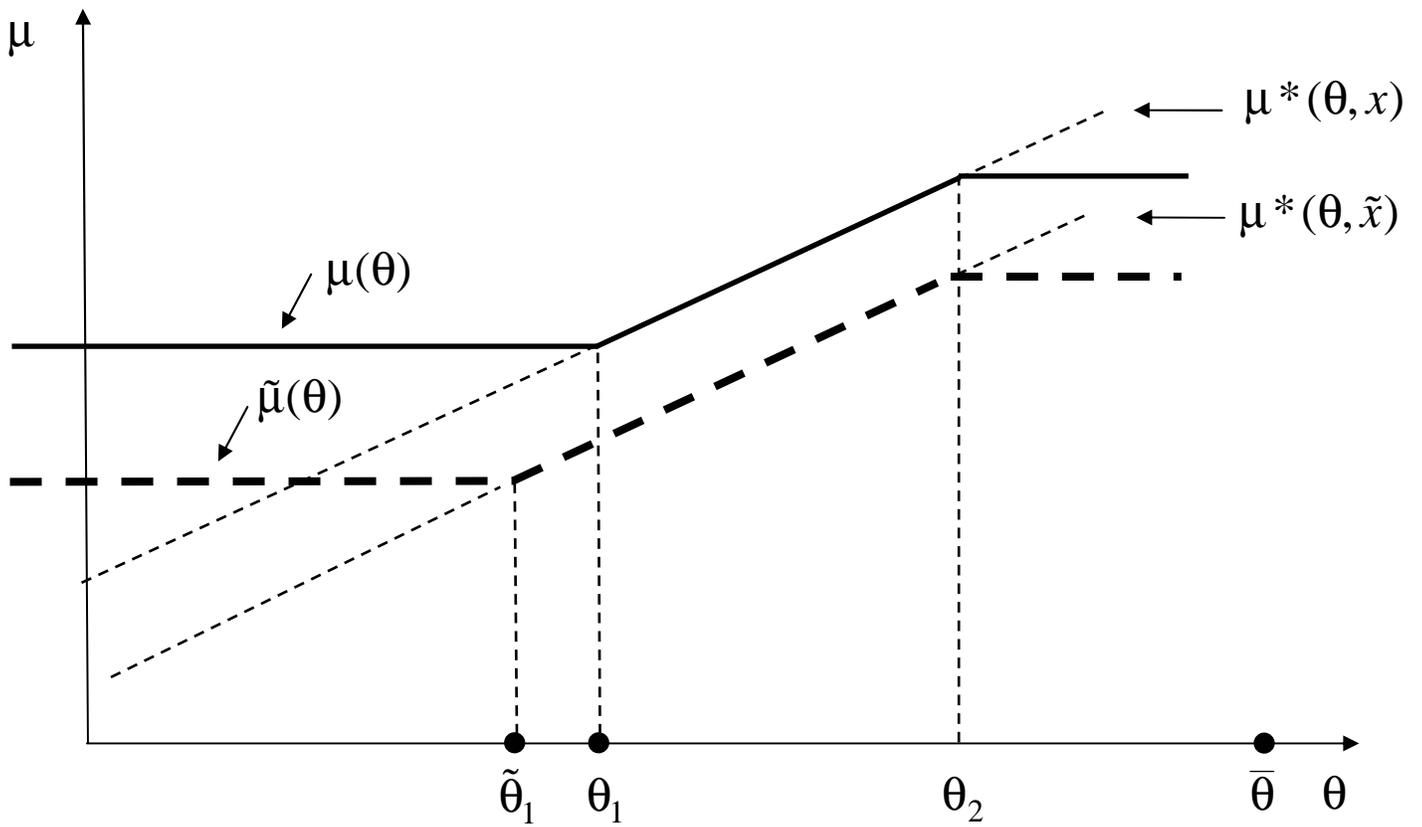


Figure 4: A Welfare-Improving Variation

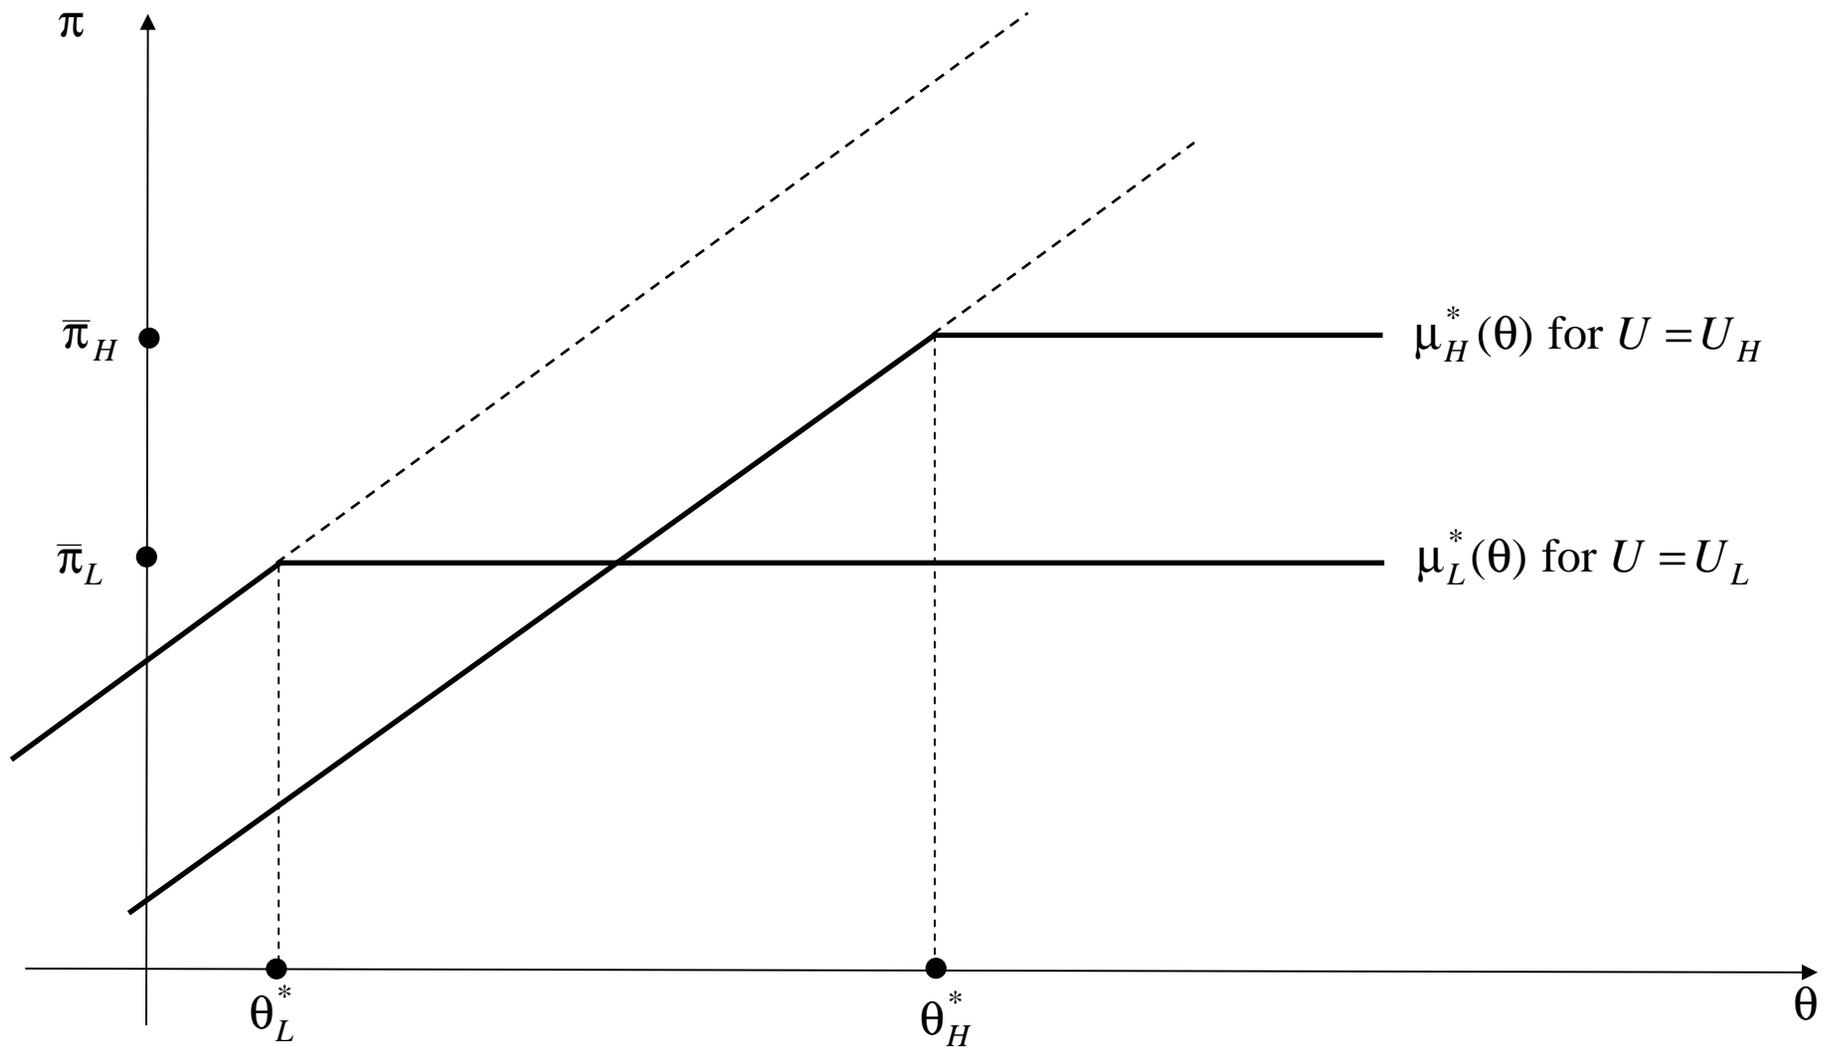


Figure 5: Optimal Discretion with Differing Severity of Time Inconsistency Problems ($U_H > U_L$)

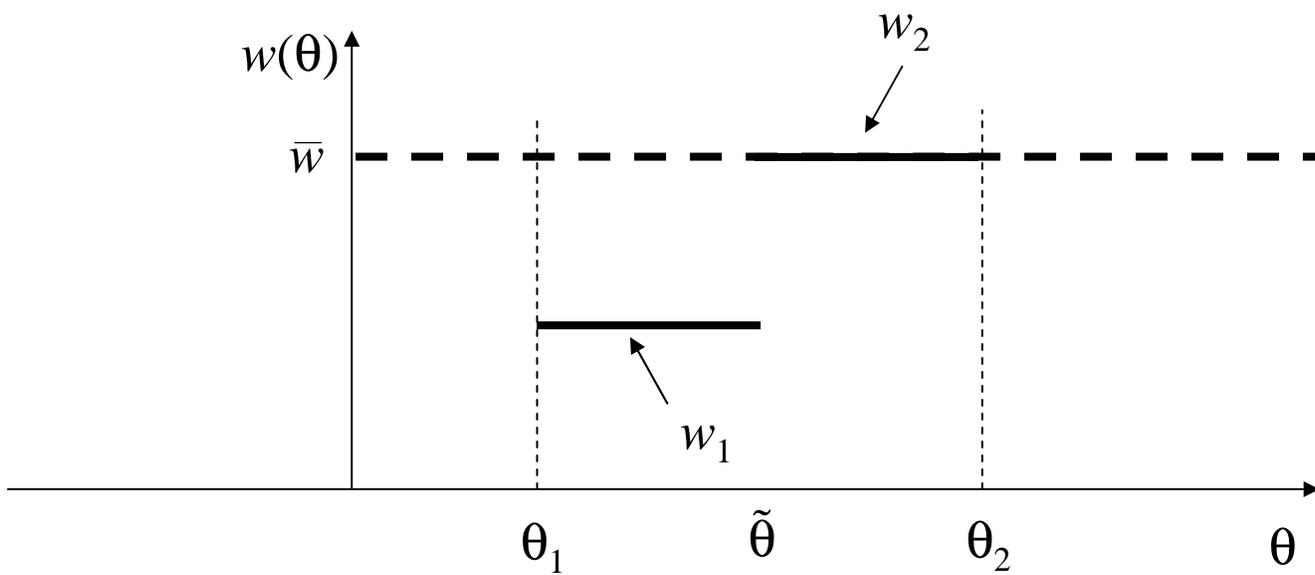
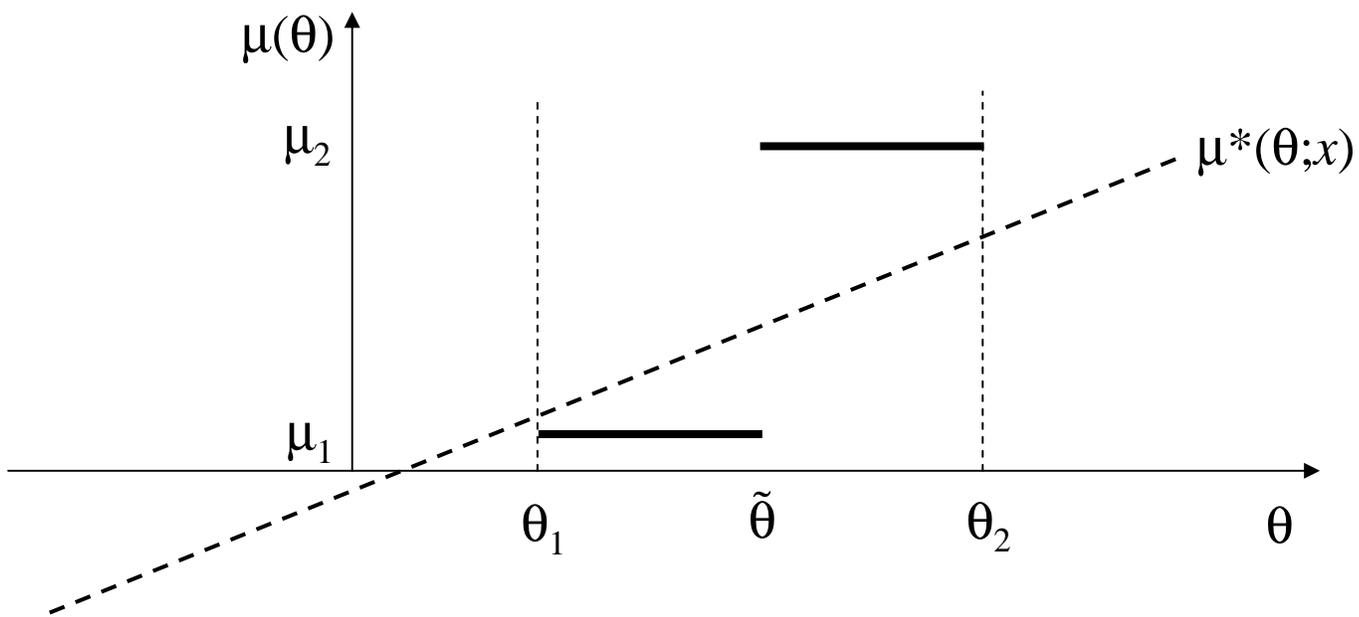


Figure 6: Ruling Out Discontinuous Continuation Values

Figure 6a: When $w_1 < w_2 \leq \bar{w}$

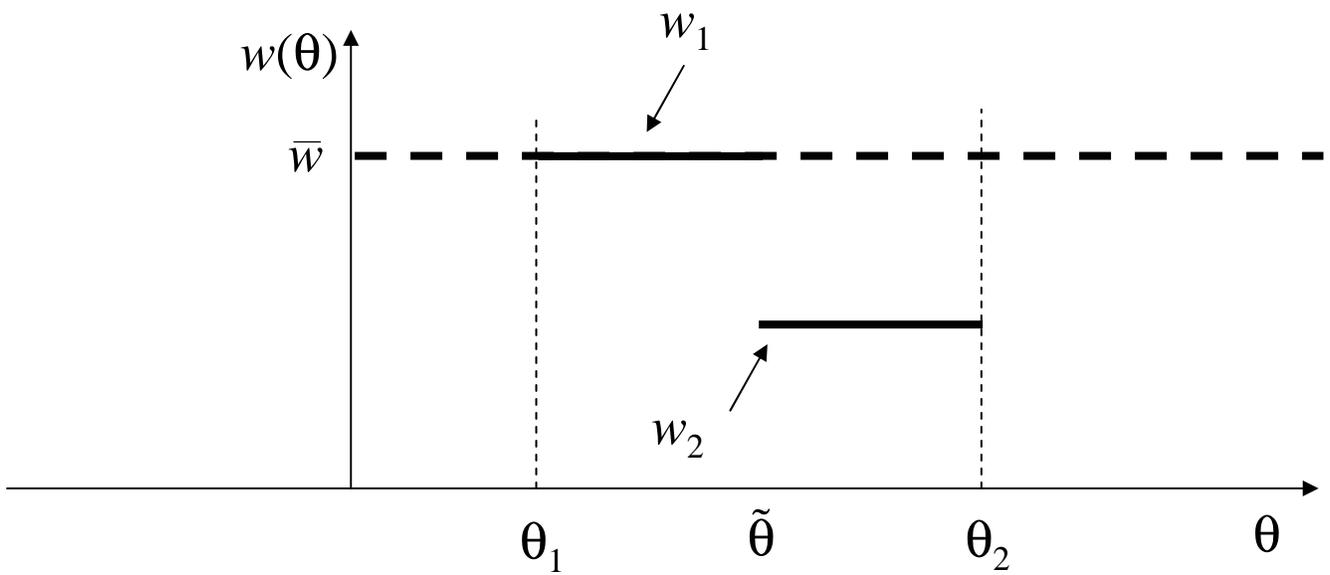
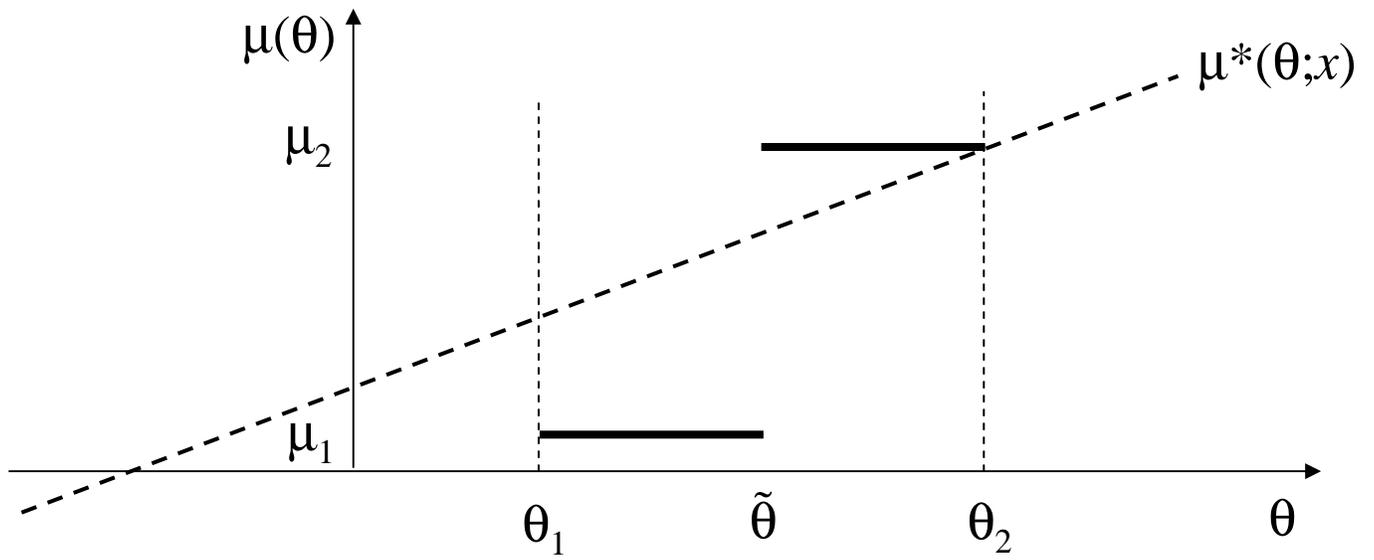
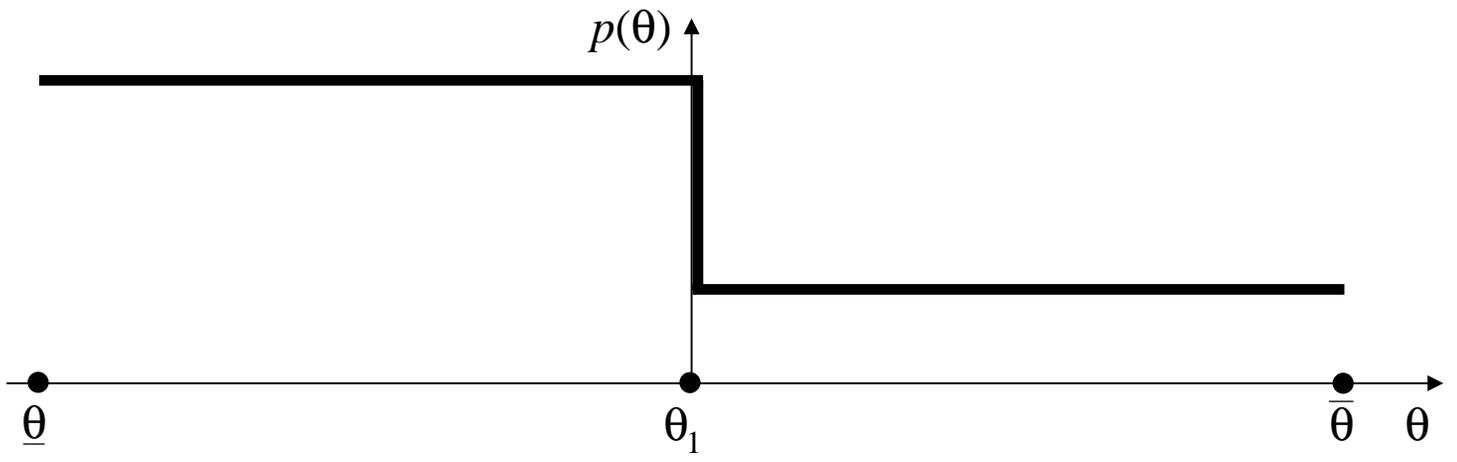
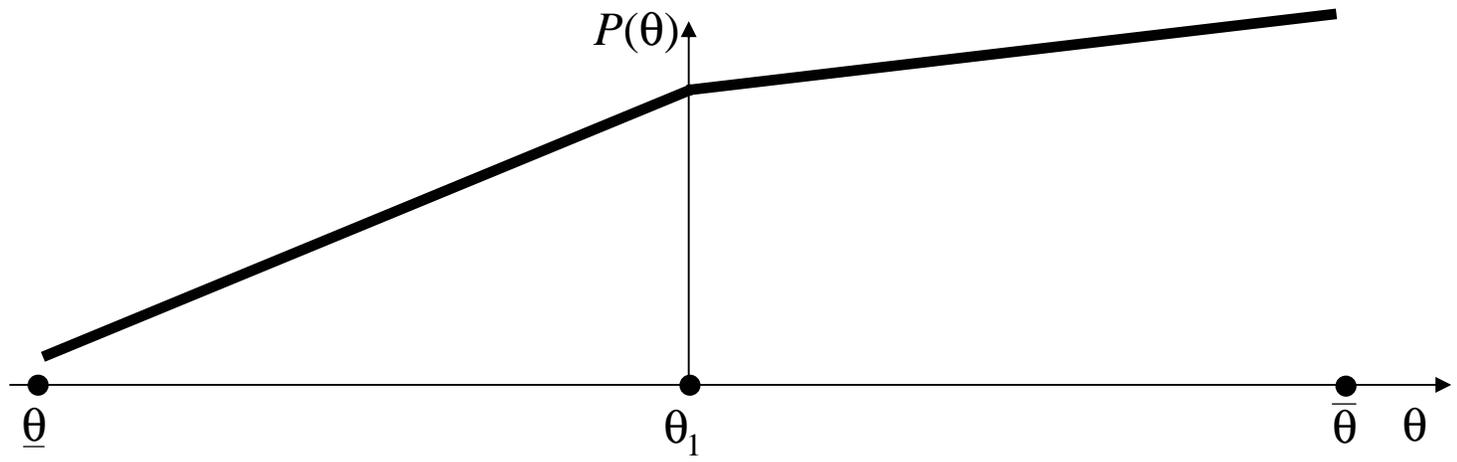


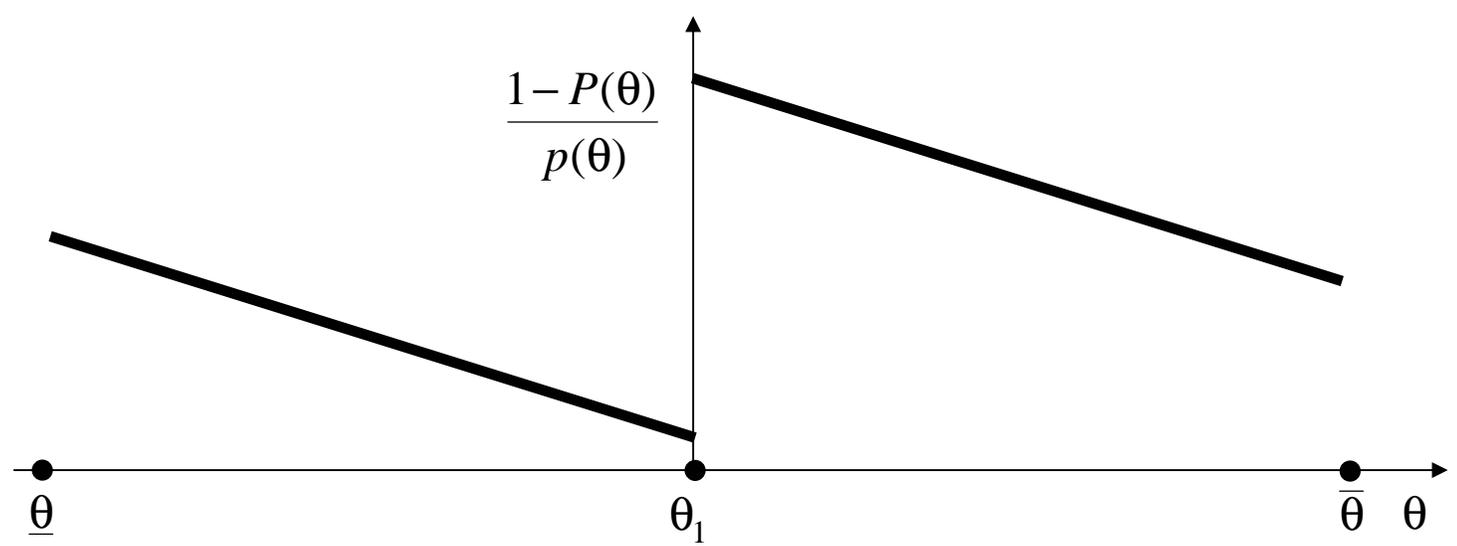
Figure 6b: When $\bar{w} \geq w_1 > w_2$



The Density



The Corresponding Distribution



The Corresponding Hazard

Figure 7: A Distribution With a Nonmonotone Hazard