Federal Reserve Bank of Minneapolis
Research Department

The Optimal Degree of Discretion in Monetary Policy*

Susan Athey, Andrew Atkeson, and Patrick J. Kehoe

Working Paper 626

November 2002

ABSTRACT

How much discretion is it optimal to give the monetary authority in setting its policy? We analyze this mechanism design question 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 flexibility to react to its private information against society's need to guard against the standard time inconsistency problem arising from the temptation to stimulate the economy with unexpected inflation. We find that the optimal degree of monetary policy discretion is decreasing in the severity of the time inconsistency problem. As this problem becomes sufficiently severe, the optimal degree of discretion is none at all. We also find that, despite the apparent complexity of this dynamic mechanism design problem, society can implement the optimal policy simply by legislating an inflation cap that specifies the highest allowable inflation rate.

*Athey, Stanford University and National Bureau of Economic Research; Andrew Atkeson, University of California, Los Angeles and National Bureau of Economic Research; Patrick J. Kehoe, Federal Reserve Bank of Minneapolis, University of Minnesota, and National Bureau of Economic Research. The authors thank the NSF for generous assistance. The views expressed are those of the authors and not necessarily those of the Federal Reserve Bank or the Federal Reserve System.

Kydland and Prescott (1977) began what has become a long debate over the question of how tightly rules should constrain the discretion of the monetary authority in setting monetary policy. Canzoneri (1985) argues that when the monetary authority does not have private information about the state of the economy, the debate is settled: the best outcomes can be achieved by rules which specify the action by the monetary authority as a function of observables. The unsettled question in this debate is the one that Canzoneri posed: What about when the monetary authority does have private information? What, then, is the optimal degree of monetary policy discretion?

Society's two extreme options are obvious: Constrain the decisions of the monetary authority only *loosely*, leaving it the flexibility to change policy in response to its private information, or constrain its decisions *tightly*, leaving it little room to change policy. Canzoneri discusses the pros and cons of several simple types of constraints. These trade off society's desire to give policymakers flexibility to react to their private information against society's need to prevent policymakers from succumbing to the standard time inconsistency problem by giving in to the temptation to stimulate the economy with a surprise inflation.

The purpose of this paper is to answer Canzoneri's question by finding the constraints on monetary policy that, in the presence of private information, optimally trade off society's desires and needs. Canzoneri conjectures that because of the dynamic nature of the problem with private information, the 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 constrains the discretion of the monetary authority by setting an *inflation cap*, an upper limit on the permitted inflation rate. The level of this inflation cap depends on the underlying parameters and is tighter the more severe is the time inconsistency problem. In this sense, the optimal degree of monetary policy discretion is decreasing in the severity of the time inconsistency problem. If this problem is severe enough, then for optimality, the inflation cap must be set so that the monetary authority exercises no discretion at all.

We consider a simple model of monetary policy similar to that of Kydland and Prescott (1977) and Barro and Gordon (1983). The model includes an agreed-upon social welfare function that depends on the random state of the economy. We begin with the simple assumption that the monetary authority observes the state and private agents do not. This assumption creates a tension between flexibility and time inconsistency. Tight constraints on discretion mitigate the time inconsistency problem, in which the monetary authority is tempted to claim repeatedly that its information about 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 observables, but the relevant question is, how tight should they be? How much discretion should be allowed?[1]

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

[1] For some potential empirical support for the idea that the Federal Reserve possesses some nontrivial private information, see Romer and Romer (2000).

when lower values are realized. We set up an infinitely repeated game and use standard recursive techniques to solve for the best equilibrium of this game.

We show that the best equilibrium has one of two forms. Either it has bounded discretion or it has no discretion. Under *bounded discretion,* there is a cutoff state such the for any state less than this cutoff state, the monetary authority chooses its static optimum, 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 rate. Under *no discretion,* the monetary authority chooses some constant rate regardless of its information. In either case, the best dynamic equilibrium ends up being a repetition of the one-period equilibrium.

We show that we can implement the best equilibrium as a repeated static equilibrium in which the monetary authority chooses its optimal policy subject to an inflation cap and in which private 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. As we vary the underlying parameters so that the time inconsistency problem becomes more severe, the inflation cap becomes lower, and the cap is more likely to bind. If the problem is sufficiently severe, then the cap is set sufficiently low that it binds for all realizations of the private information, and the resulting policy is a constant.

Technically, we set up the problem of finding the optimal degree of discretion as a mechanism design problem. To keep the notation simple, we begin with the case in which the state of the economy is never observed by the private agents. We then show that all of our results extend to the case in which this state is observed with noise by the private agents.

One interpretation of this paper is that we solve for the optimal inflation targets. As such, our paper is related to the burgeoning literature on inflation targeting. (See Cukierman and Meltzer (1986), Bernanke and Woodford (1997), and Faust and Svensson (2000), among many others.) In particular, as noted above, our paper is closely related to the work of Canzoneri (1985). Canzoneri advocates a legislative approach to monetary policy with private information. He describes several types of potential inflation targeting rules, including a strict targeting rule, which gives the monetary authority no discretion, and an average targeting rule in which the average inflation rate over two periods has to meet a certain target. He also discusses the merits of solving for the optimal targeting rule.

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. Wider ranges of allowable inflation allow more flexibility while narrow ranges allow less. These ranges or targets for monetary policy, once set, determine the rules of the game played by both the monetary authority and the private agents in the economy.

At a technical level, our paper draws heavily on the literature on recursive approaches to dynamic games. We use the techniques of Abreu, Pearce, and Stachetti (1991) and draw on some of the insights in the work of Athey, Bagwell, and Sanchiricho (2000).

1. The economy

The economy has a government, a monetary authority, and a continuum of agents. The time horizon is infinite with periods indexed $t = 0, 1, \ldots.$.

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. We let x_t denote the average nominal wage. Next, the monetary authority observes the current realization of its private information θ_t. 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 distribution function $P(\theta)$. Given this private information, the monetary authority chooses money growth μ_t in some arbitrarily large compact set $[\underline{\mu}, \bar{\mu}]$

The monetary authority maximizes a social welfare function that depends on unemployment, inflation, and its 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

$$(1) \qquad u_t = U + x_t - \mu_t$$

where x is the average of z across agents and 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 unobserved state θ_t. Our leading example will be the Kydland and Prescott objective function which has the form

$$(2) \qquad -u_t^2/2 - (\pi_t - \theta_t)^2/2.$$

Using (1) and $\pi_t = \mu_t$ we can write this objective function in terms of nominal wage growth x_t, money growth μ_t and the unobserved shock θ_t, as

$$(3) \qquad R(x_t, \mu_t, \theta_t) = -\frac{1}{2}\left[(U + x_t - \mu_t)^2 + (\mu_t - \theta_t)^2\right].$$

In our example the private information is about the inflation target.

We develop our model for general specifications of the social welfare function $R(x_t, \mu_t, \theta_t)$ which imply (3) as a special case. In this general setup, 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. For any x, we define the *static optimum* to be the policy $\mu^*(\theta; x)$ that solves $R_\mu(x, \mu(\theta), \theta) = 0$. We assume that the static optimum is interior in that

(4) $$\underline{\mu} < \mu^*(\theta; x) < \bar{\mu}$$

for all $x \in (\underline{\mu}, \bar{\mu})$. We also assume that if $x = \int \mu(\theta)p(\theta)d\theta$ then

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

This assumption implies that starting from any equilibrium allocation, lowering inflation raises current welfare.

A. Two Ramsey benchmarks

In what follows we will be interested in a game in which the monetary authority cannot commit to its policy. Before analyzing this game it is useful to consider two alternative games with commitment that we think of as benchmarks.

Our first benchmark, *the Ramsey policy*, denoted $\mu^R(\theta)$, yields the highest payoff that can be achieved with commitment. The gap between the associated Ramsey payoff and the payoff in the game without commitment measures the welfare loss of the lack of commitment.

Our second benchmark, *the expected Ramsey policy*, denoted μ^{ER}, is the optimal policy of the monetary authority when it can commit once-and-for-all to a monetary policy that is independent of its private information. This policy is a useful benchmark because in the game without commitment the government can ensure that it implements the expected Ramsey policy as an equilibrium by choosing μ^{ER} as the inflation cap.

For the Ramsey benchmark consider a game with commitment with the following timing scheme. Before the realization of its type, the monetary authority commits to a schedule for money growth rates $\mu(\theta)$ indicating what money growth rate will be implemented once its type is realized. Next, private agents choose their nominal wages z with associated average nominal wages x. Then the government's type θ is privately realized and money growth rate $\mu(\theta)$ is implemented. The equilibrium allocations and policies in this game solve the *Ramsey problem*

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

subject to

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

For our example (3), the Ramsey policy is $\mu^R(\theta) = \theta/2$. Note that the Ramsey policy has the monetary authority choosing a money growth rate that is increasing in its type. Thus, with full commitment, it is optimal to allow the monetary authority flexibility in choosing monetary policy to reflect its private information. This feature of the environment leads to a tension in the repeated game between flexibility and credibility.

For the second benchmark, consider a variant of this game in which the monetary authority is restricted to choosing a money growth μ that does not vary with its type. The equilibrium allocations and policies in this game solve the *expected Ramsey problem*

$$v^{ER} = \max_{x,\mu} \int R(x, \mu, \theta) p(\theta) d\theta \tag{6}$$

subject to $x = \mu$. Let μ^{ER} denote the expected Ramsey policy. For our example (3), the expected Ramsey policy is $\mu^{ER} = 0$. Clearly, it is possible for society to implement the expected Ramsey allocations by legislating constraints on monetary policy which depend only on observables. This policy is analogous to the strict targeting rule discussed in Canzoneri (1985).

Clearly, for our example (3), the Ramsey policy yields strictly higher welfare than the expected Ramsey policy. More generally, when $R_{\mu\theta}(x, \mu, \theta) > 0$, the Ramsey policy $\mu^R(\theta)$ is strictly increasing in θ and the Ramsey policy yields strictly higher welfare than the expected Ramsey policy.

B. The dynamic mechanism design problem

We model the problem of finding the optimal degree of discretion as a dynamic mechanism design problem derived from a direct revelation game. In this problem, society specifies a monetary policy namely, the money growth rate as a function of the history of the monetary authority's reports of its private information. Given the specified rule for money growth rates, the monetary authority chooses a report of its private information. Individual agent's 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 $\left\{\mu_t\left(h_t, \hat{\theta}_t\right)\right\}_{t=0}^{\infty}$ where $\mu_t\left(h_t, \hat{\theta}_t\right)$ specifies the money growth rate that will be chosen in period t following the history $h_t = \left(\hat{\theta}_0, \hat{\theta}_1, \ldots, \hat{\theta}_{t-1}\right)$ of past reports together with the current report $\hat{\theta}_t$. The monetary authority chooses a reporting strategy $\{m_t(h_t, \theta_t)\}_{t=0}^{\infty}$ at 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 at t. As is standard we restrict attention to public strategies, namely strategies that depend only public histories and the current private information, but not on the history of private information.[2] Also, from the Revelation Principle it suffices to 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 private agents are competitive it is we need not record their either their individual past actions or the aggregate of their past actions in the history. See Chari and Kehoe (1992) for details.)

Each agent chooses nominal wage growth equal to expected inflation. Taking monetary policy $\mu_t(h_t, \hat{\theta}_t)$ as given, consumers set $z_t(h_t)$ equal to expected inflation

$$(7) \qquad z_t(h_t) = \int \mu_t(h_t, \theta) p(\theta) d\theta$$

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

The optimal monetary policy maximizes the discounted sum of social welfare

$$(8) \qquad (1-\beta) \sum_{t=0}^{\infty} \int \beta^t R(x_t(h_t), \mu_t(h_t, \theta_t), \theta_t) p(\theta_t) d\theta_t$$

where the future histories h_t are recursively generated from the choice of monetary policy $\mu_t(h_t, \theta_t)$ 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 $\left\{\mu_t\left(h_t, \hat{\theta}_t\right)\right\}_{t=0}^{\infty}$, a reporting strategy $\{m_t(h_t, \theta_t)\}_{t=0}^{\infty}$, a strategy for wage setting by agents $\{z_t(h_t)\}_{t=0}^{\infty}$, and average wages $\{x_t(h_t)\}_{t=0}^{\infty}$ such that (7) is satisfied in every period following every history

[2] See Fudenberg and Tirole (1995), for a discussion of the large class of environments for which this restriction does not alter the set of equilibrium payoffs.

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 an 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.

C. A recursive formulation

Here we formulate the problem of characterizing the set of equilibrium payoffs of the truth-telling equilibrium of the revelation game recursively along the lines of Abreu, Pearce and Stachetti (1991).

The basic idea is the following. Since there are no physical state variables the set of equilibrium 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 equilibrium strategies 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 equilibrium payoffs. Following this logic, Abreu, Pearce and Stachetti (1991) show that the set of equilibrium payoffs can be found using a recursive method.

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 equilibrium payoffs 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. The continuation payoffs represent the discounted utility for the monetary authority from next period on and are denoted by $w(\hat{\theta})$. These payoffs depend on the publicly observable report $\hat{\theta}$ of the monetary authority. Clearly, these payoffs cannot vary directly with the privately observed state θ.

We say that actions $x, \mu(\theta)$, and function $w(\hat{\theta})$ are *enforceable by* W if,

$$(9) \quad w(\hat{\theta}) \in W$$

$$(10) \quad x = \int \mu(\theta) p(\theta) d\theta$$

and the incentive constraints

$$(11) \quad (1-\beta)R(x, \mu(\theta), \theta) + \beta w(\theta) \geq (1-\beta)R(x, \mu(\hat{\theta}), \theta) + \beta w(\hat{\theta}),$$

for all θ and for all $\hat{\theta}$. Constraint (9) requires that each continuation payoff $w(\hat{\theta})$ be drawn from the set of candidate equilibrium payoffs W while constraint (10) requires that average wages equal expected inflation. Constraint (11) 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 continuation value $w(\theta)$ rather than a money growth rate $\mu(\hat{\theta})$ and corresponding continuation value $w(\hat{\theta})$.

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

$$(12) \quad \Pi(x, \mu(\theta), w(\theta)) = \int \left[(1-\beta)R(x, \mu(\theta), \theta) + \beta w(\theta)\right] p(\theta) d\theta$$

Define the operator T that maps a set of payoffs W into a new set of payoffs $T(W)$ by

$$(13) \quad T(W) = \left\{\Pi(x, \mu(\theta), w(\theta)) \mid x, \mu(\theta), w(\theta) \text{ are enforceable by } W\right\}.$$

As demonstrated by Abreu, Pearce, and Stacchetti, the set of equilibrium payoffs is the largest set W that is a fixed point of this operator, namely

$$(14) \quad W^* = T(W^*).$$

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

$$(15) \quad \bar{v} = \max_{x,\mu(\theta),w(\theta)} \int \left[(1-\beta)R(x,\mu(\theta),\theta) + \beta w(\theta)\right] p(\theta) d\theta$$

subject to constraint that x, $\mu(\theta) \in [\underline{\mu}, \bar{\mu}]$ and $w(\theta)$ are enforceable by W, in that they satisfy (9)-(11). Throughout we assume that $\mu(\theta)$ is a piecewise continuously differentiable function.

The best payoff problem is a mechanism design problem of choosing an incentive compatible allocation $(x, \mu(\theta), w(\theta))$ which maximizes utility. Following the language of mechanism design we refer to θ as the *type* of the government, which in this repeated game changes every period. When we solve this problem with $W = W^*$, (14) implies that the resulting payoff is the highest equilibrium payoff. We refer to this equilibrium as the *best equilibrium* and denote its payoff as $\bar{w}^*$.

2. The Optimal Degree of Discretion

In this section we solve the best payoff problem and use it to characterize the optimal degree of discretion. We show several results. First, the optimal degree of discretion is either bounded discretion or no discretion and the corresponding optimal continuation values $w(\theta)$ are constant at the upper bound $\bar{w}$. Next, the optimal degree of discretion is decreasing in the severity of the time inconsistency problem. Finally, we show that a policy with either bounded discretion or no discretion can be implemented by society setting an upper limit μ^* on the inflation rate the monetary authority is allowed to choose. Subject to this limit each period the monetary authority chooses the inflation rate that is statically optimal and private agents' expectations of future inflation do not vary with the monetary authority's policy choice.

We begin with some definitions. We say that a policy $\mu(\theta)$ has *bounded discretion* if it takes the form

$$(16) \quad \mu(\theta) = \left\{ \begin{array}{c} \mu^*(\theta; x) \text{ if } \theta \in [\underline{\theta}, \theta^*) \\ \\ \mu^* = \mu^*(\theta^*, x) \text{ if } \theta \in (\theta^*, \bar{\theta}] \end{array} \right\}.$$

where $\mu^*(\theta; x)$ is the static optimum inflation rate given wages x. Thus, for $\theta < \theta^*$, the monetary authority chooses the static optimum 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 constant μ^*. Clearly, the best policy with no discretion is the expected Ramsey policy.

The intuition for the 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 either type of policy are ones that raise inflation. Under bounded discretion the types in $[\underline{\theta}, \theta^*)$ are at their static optimum and hence have no incentive to deviate while the types in $(\theta^*, \bar{\theta}]$ have an incentive to deviate to a higher rate than μ^*. Likewise under the expected Ramsey policy of no discretion all types have an incentive to deviate to higher rates of inflation. Hence no other limits on discretion are required to implement such a policy.

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 $(1 - \beta)$ being subsumed in the R function and β being subsumed in the w function.

A. Preliminaries

We assume that the preference satisfy a standard *single-crossing assumption*, namely

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

This implies that higher types have a stronger preference for current inflation. Notice that the single crossing condition implies that the static optimum is strictly increasing in θ, that is

$$(17)\qquad \frac{\partial \mu^*(\theta;x)}{\partial \theta} = -\frac{R_{\mu\theta}(x,\mu(\theta),\theta)}{R_{\mu\mu}(x,\mu(\theta),\theta)} > 0$$

Under the single-crossing assumption a standard lemma allows to replace the global incentive constraints (11) with some local versions of them. We say that an allocation is *locally incentive compatible* if it satisfies three conditions: $\mu(\theta)$ is non-decreasing,

$$(18)\qquad R_\mu(x,\mu(\theta),\theta)\frac{d\mu(\theta)}{d\theta} + \frac{dw(\theta)}{d\theta} = 0$$

wherever $d\mu(\theta)/d\theta$ exists, and for any point θ_i at which this derivative does not exist

$$(19)\qquad \lim_{\theta\nearrow\theta_i} R(x,\mu(\theta),\theta_i) + w(\theta) = \lim_{\theta\searrow\theta_i} R(x,\mu(\theta),\theta_i) + w(\theta).$$

We prove the following lemma in the Appendix.

Lemma 1. Under (A1), $(x,\mu(\theta),w(\theta))$ satisfies the incentive constraints (11) if and only if it is locally incentive compatible.

Given any incentive compatible allocation, we define the *utility of the allocation* as

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

It follows from the lemma that $U(\theta)$ is continuous, differentiable almost everywhere, with derivative $U'(\theta) = R_\theta(x, \mu(\theta), \theta)$. Integrating $U'(\theta)$ implies

$$(20) \quad U(\theta) = U(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} R_\theta(x, \mu(z), z)dz = U(\bar{\theta}) - \int_{\theta}^{\bar{\theta}} R_\theta(x, \mu(z), z)dz.$$

Using integration by parts it is easy to show that for $\theta_2 > \theta_1$,

$$(21) \quad \int_{\theta_1}^{\theta_2} U(\theta)p(\theta)d\theta = P(\theta_2)U(\theta_2) - P(\theta_1)U(\theta_1) - \int_{\theta_1}^{\theta_2} R_\theta(x, \mu(\theta), \theta)P(\theta)d\theta$$

Next we make some joint assumptions on the probability distribution and the return function. Assume for any action profile $(x, \mu(\theta))$ with $\mu(\theta)$ non-decreasing,

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

(A2b) $\dfrac{P(\theta)}{p(\theta)} R_{\theta\mu}(x, \mu(\theta), \theta)$ strictly increasing in θ.

We refer to (A2a) and (A2b) as (A2). In our quadratic case $R_\theta(x, \mu(\theta), \theta) = \mu(\theta) - \theta$ and $R_{\theta\mu}(x, \mu(\theta), \theta) = 1$ so that (A2) reduces to two monotone hazard conditions on probabilities.

B. Characterizing the Optimal Policy

In this section we prove the following proposition with series of lemmas.

Proposition 1. Under (A1) and (A2), the allocation that solves the best payoff problem has the following properties: *i)* $\mu(\theta)$ is continuous, *ii)* $\mu(\theta)$ is either flat in θ, in that $\mu'(\theta) = 0$, or $\mu(\theta)$ is equal to the static optimum $\mu^*(\theta; x)$ so that $R_\mu(x, \mu(\theta), \theta) = 0$, and *iii)* $w(\theta) = \bar{w}$ for all θ.

We begin with a definition. We say that $\mu(\theta)$ *increasing on an interval* (θ_1, θ_2) if 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 $\mu(\theta)$ on this interval, namely

$$\tilde{\mu} = \frac{\int_{\theta_1}^{\theta_2} \mu(\theta)p(\theta)d\theta}{P(\theta_2) - P(\theta_1)}. \tag{22}$$

In words, on this interval $\mu(\theta)$ is strictly below its conditional mean $\tilde{\mu}$ up to $\tilde{\theta}$ and strictly above its conditional mean after $\tilde{\theta}$. This definition of $\mu(\theta)$ increasing subsumes two possibilities that we need to consider below: one with $d\mu(\theta)/d\theta > 0$ for some non-degenerate subinterval of this interval, and another in which $\mu(\theta)$ jumps up at $\tilde{\theta}$.

In the next lemma we show that if the inflation rate is increasing on some interval then it must be that the continuation value $w(\theta)$ is at the highest level $\bar{w}$ on that interval. We prove the lemma by showing that if the continuation value is not at the highest level then we improve welfare by the following variation: marginally move $\mu(\theta)$ towards the conditional mean on this interval and adjust the continuation values to preserve incentive compatibility. It is possible to adjust the continuation values in this way only if $w(\theta)$ is strictly less than $\bar{w}$ on this interval. This flattening of the inflation schedule improves welfare.

Lemma 2. Let $(x, \mu(\theta), w(\theta))$ be an allocation in which $\mu(\theta)$ is increasing on some interval (θ_1, θ_2) and for some $\varepsilon > 0$, $w(\theta) \leq \bar{w} - \varepsilon$ for $\theta \in (\theta_1, \theta_2)$. Then this allocation cannot be optimal.

Proof. Let $(x, \mu(\theta), w(\theta))$ be such an allocation. We show that there is alternative feasible allocation that improves the payoff in the best payoff problem. We must consider two cases. In the first case,

$$\int_{\theta_1}^{\theta_2} R_{\theta\mu}(x, \mu(\theta), \theta) \left[\tilde{\mu} - \mu(\theta)\right] d\theta < 0. \tag{23}$$

In the second case the inequality is reversed.

In case 1 we proceed as follows. We consider a variation which moves our original allocation marginally towards an allocation with $(\tilde{x}, \tilde{\mu}(\theta))$ defined by

$$(24)\quad \tilde{\mu}(\theta) = \left\{ \begin{array}{c} \tilde{\mu} \text{ if } \theta \in (\theta_1, \theta_2) \\ \\ \mu(\theta) \text{ otherwise} \end{array} \right\}$$

and $\tilde{x} = \int_{\underline{\theta}}^{\bar{\theta}} \tilde{\mu}(\theta) p(\theta) d\theta = x$. Our alternative allocation is defined by $(\tilde{x}(a), \tilde{\mu}(\theta; a), \tilde{w}(\theta; a))$ where

$$(25)\quad \tilde{\mu}(\theta; a) = a\tilde{\mu}(\theta) + (1-a)\mu(\theta)$$

$\tilde{x}(a) = x$ for $a \in [0,1]$. (See Figure 1 for a graph of $\tilde{\mu}(\theta; a)$.) Using (20) the continuation values are then given by

$$(26)\quad \tilde{w}(\theta; a) = U(\underline{\theta}) + \int_{\underline{\theta}}^{\theta} R_\theta(x, a\tilde{\mu}(z) + (1-a)\mu(z), z) dz - R(x, a\tilde{\mu}(\theta) + (1-a)\mu(\theta), \theta).$$

Using (24) we have that for $\theta \leq \theta_1$, $\tilde{w}(\theta; a)$ equals $w(\theta)$, for $\theta \in (\theta_1, \theta_2)$ it equals

$$w(\theta) + \int_{\theta_1}^{\theta} [R_\theta(x, \tilde{\mu}(z; a), z) - R_\theta(x, \mu(z), z)] dz - [R(x, \tilde{\mu}(\theta; a), \theta) - R(x, \mu(\theta), \theta)]$$

and for $\theta \geq \theta_2$ it equals

$$w(\theta) + \int_{\theta_1}^{\theta_2} [R_\theta(x, \tilde{\mu}(z; a), z) - R_\theta(x, \mu(z), z)] dz$$

where $\tilde{\mu}(z; a) = a\tilde{\mu}(z) + (1-a)\mu(z)$. (See Figure 2 for a graph of $\tilde{w}(\theta; a)$ in case 1.)

We first show that for small a this alternative allocation is feasible. It is immediate that our alternative allocation satisfies local incentive compatibility. We need only show that $\tilde{w}(\theta; a) \leq \bar{w}$ for all θ. Clearly, $\tilde{w}(\theta; a) = w(\theta) \leq \bar{w}$ for $\theta \leq \theta_1$ For small a we know that

$\tilde{w}(\theta; a) \le \bar{w}$ for $\theta \in (\theta_1, \theta_2)$ since in this interval $w(\theta) \le \bar{w} - \varepsilon$. To show that for small $a, \tilde{w}(\theta; a) \le \bar{w}$ for $\theta \ge \theta_2$ we note that from (26)

$$\frac{d\tilde{w}(\theta;0)}{da} = \frac{d\tilde{w}(\theta_2;0)}{da} = \int_{\theta_1}^{\theta_2} R_{\theta\mu}(x, \mu(\theta), \theta)\left[\tilde{\mu} - \mu(\theta)\right] d\theta < 0$$

where the last inequality follows directly from (23).

We now show that this alternative allocation improves welfare. From (20) and (21) we can write the value of the objective function as

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

hence under our variation, the value of the alternative allocation is

$$\text{(27)} \quad \tilde{V}(a) = U(\underline{\theta}) + \int_{\underline{\theta}}^{\bar{\theta}} \frac{1 - P(\theta)}{p(\theta)} R_\theta(x, a\tilde{\mu}(\theta) + (1-a)\mu(\theta), \theta)p(\theta)d\theta.$$

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

$$\text{(28)} \quad \frac{d\tilde{V}(0)}{da} = \int_{\theta_1}^{\theta_2} \frac{1 - P(\theta)}{p(\theta)} R_{\theta\mu}(x, \mu(\theta), \theta)\left[\tilde{\mu} - \mu(\theta)\right] p(\theta)d\theta.$$

If we divide (28) by the positive constant $P(\theta_2) - P(\theta_1)$ then we can interpret it to be the expectation of two functions $f(\theta) = \frac{1-P(\theta)}{p(\theta)} R_{\theta\mu}(x, \mu(\theta), \theta)$ and $g(\theta) = \tilde{\mu} - \mu(\theta)$ where the density is $p(\theta)/[P(\theta_2) - P(\theta_1)]$. The function f is strictly decreasing by assumption (A2a) and the function g is decreasing since $\mu(\theta)$ is increasing. By the definition of a covariance, $Efg = cov(f, g) + (Ef)(Eg)$ 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 (22), $Eg = 0$, so $Efg = cov(f, g)$ which is clearly positive since both are strictly decreasing. Thus, (28) is strictly positive. This completes the proof for case 1.

The analysis for case 2 is analogous where $(\tilde{x}, \tilde{\mu}(\theta; a))$ is defined as above and $\tilde{w}(\theta; a)$ is defined by

$$(29) \quad \tilde{w}(\theta; a) = U(\bar{\theta}) - \int_{\theta}^{\bar{\theta}} R_{\theta}(x, a\tilde{\mu}(z) + (1-a)\mu(z), z)dz - R(x, a\tilde{\mu}(\theta) + (1-a)\mu(\theta), \theta).$$

Thus, $\theta \leq \theta_1, \tilde{w}(\theta; a)$ equals

$$w(\theta) - \int_{\theta_1}^{\theta_2} [R_{\theta}(x, \tilde{\mu}(z; a), z) - R_{\theta}(x, \mu(z), z)]dz$$

for $\theta \in (\theta_1, \theta_2)$ it equals

$$w(\theta) - \int_{\theta}^{\theta_2} [R_{\theta}(x, \tilde{\mu}(z; a), z) - R_{\theta}(x, \mu(z), z)]dz - [R(x, \tilde{\mu}(\theta; a), \theta) - R(x, \mu(\theta), \theta)]$$

and for $\theta \geq \theta_2$, it equals $w(\theta)$ where $\tilde{\mu}(z; a) = a\tilde{\mu}(z) + (1-a)\mu(z)$. (See Figure 3 for a graph of $\tilde{w}(\theta; a)$ in case 2.) To check feasibility note that by construction $\tilde{w}(\theta; a) = w(\theta) \leq \bar{w}$ for $\theta \geq \theta_2$. For small a we know that $\tilde{w}(\theta; a) \leq \bar{w}$ for $\theta \in (\theta_1, \theta_2)$ since in this interval $w(\theta) \leq \bar{w} - \varepsilon$. To show that for small a, $\tilde{w}(\theta; a) \leq \bar{w}$ for $\theta \leq \theta_1$ we note that from (29)

$$\frac{d\tilde{w}(\theta; 0)}{da} = \frac{d\tilde{w}(\theta_1; 0)}{da} = -\int_{\theta_1}^{\theta_2} R_{\theta\mu}(x, \mu(\theta), \theta) \left[\tilde{\mu} - \mu(\theta)\right] d\theta < 0$$

where the inequality holds since we are in case 2. Likewise

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

Hence,

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

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

To gain some intuition for how the variation used in Lemma 2 improves welfare consider case 1 and the expression for the change in welfare (28). 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 (A2) 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 effect of increasing a on the current payoff for type θ is

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

The impact of increasing a on the continuation value for this type is given, from Lemma 3, by

$$\frac{d\tilde{w}(\theta;0)}{da}=\int_{\underline{\theta}}^{\theta} R_{\theta\mu}(x, \mu(z), z)\left[\tilde{\mu}(z)-\mu(z)\right]dz-R_{\mu}(x, \mu(\theta), \theta)\left[\tilde{\mu}(\theta)-\mu(\theta)\right]. \tag{31}$$

Hence, the impact on the total utility of type θ is given by

$$\frac{d\tilde{U}(\theta;0)}{da}=\int_{\underline{\theta}}^{\theta} R_{\theta\mu}(x, \mu(z), z)\left[\tilde{\mu}(z)-\mu(z)\right]dz. \tag{32}$$

Note that in case 1 when we change an action for a given type and then adjust the continuation values so as to maintain incentive compatibility, this change has an impact on the total utility of that type and all types above that type. We can decompose the total change in social welfare

$$\int_{\underline{\theta}}^{\theta}\frac{d\tilde{U}(\theta;0)}{da}p(\theta)d\theta=\int_{\underline{\theta}}^{\bar{\theta}}\left[\int_{\underline{\theta}}^{\theta} R_{\theta\mu}(x, \mu(z), z)\left[\tilde{\mu}(z)-\mu(z)\right]dz\right]d\theta$$

as coming from the marginal change $R_{\theta\mu}(x, \mu(\theta), \theta)\left[\tilde{\mu}(\theta)-\mu(\theta)\right]$ resulting from a change in the policy of type θ times the fraction of types at or above type θ. Recall that our single-crossing assumption is that $R_{\theta\mu}(x, \mu(\theta), \theta)>0$ and so the impact of changing the policy at θ depends on the sign of $\tilde{\mu}(\theta)-\mu(\theta)$.

The positive part comes from altering the policy of those types θ below $\tilde{\theta}$. For these types we know that $\mu(\theta) < \tilde{\mu}(\theta)$, so that an increase in a raises $\tilde{\mu}(\theta; a)$ from $\mu(\theta)$ up towards $\tilde{\mu}$ and hence has a positive impact of

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

on the welfare of type θ and all types above θ. There are $(1 - P(\theta))$ such types and thus the marginal contribution to social welfare that we get from raising the inflation rate chosen by this type θ is

$$\frac{(1 - P(\theta))}{p(\theta)} R_{\theta\mu}(x, \mu(\theta), \theta) [\tilde{\mu}(\theta) - \mu(\theta)] \, p(\theta) > 0$$

where we have multiplied and divided by the density $p(\theta)$ for future analytical convenience.

The negative part comes from altering the policy for those types θ above $\tilde{\theta}$. For these types we know that $\mu(\theta) > \tilde{\mu}(\theta)$, so that an increase in a lowers $\tilde{\mu}(\theta; a)$ from $\mu(\theta)$ down towards $\tilde{\mu}$ and hence has a negative impact of

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

on the welfare of type θ and all types above θ. There are $(1 - P(\theta))$ such people, hence the marginal contribution to social welfare that we get from lowering the inflation rate chosen by this type θ is

$$\frac{(1 - P(\theta))}{p(\theta)} R_{\theta\mu}(x, \mu(\theta), \theta) [\tilde{\mu}(\theta) - \mu(\theta)] \, p(\theta) < 0.$$

It follows from our assumption (A2a) that the total positive impact on welfare coming from raising the inflation rate chosen by those $\theta < \tilde{\theta}$ for which $\tilde{\mu}(\theta) > \mu(\theta)$ outweighs the negative impact of lowering inflation chosen by those $\theta > \tilde{\theta}$ for which $\tilde{\mu}(\theta) < \mu(\theta)$.

In case 2 the intuition for the derivative (30) is the same as that for (28), except that in case 2 change in the inflation rate chosen by type θ affects the continuation value of all types below θ.

It is easy to use to use this lemma to prove part ii) of Proposition 1 and that $w(\theta)$ is a step function. We do so in the following lemma.

Lemma 3. Under (A1) and (A2), almost everywhere $\mu(\theta)$ is either flat in θ, in that $\mu'(\theta) = 0$, or $\mu(\theta)$ is equal to the static optimum for type θ in that $R_\mu(x, \mu(\theta), \theta) = 0$. Also, $w(\theta)$ is a step function.

Proof. By way of contradiction, suppose there is some interval (θ_0, θ_2) over which $\mu'(\theta) > 0$. Suppose first that $R_\mu(x, \mu(\theta), \theta) > 0$ over this interval. Then from local incentive compatibility, $w'(\theta) < 0$ and hence there is a subinterval, say (θ_1, θ_2) with $\theta_1 > \theta_0$ over which $\mu(\theta)$ is increasing on this interval and $w(\theta)$ is uniformly bounded below $\bar{w}$. Such an allocation cannot be optimal by Lemma 2. A similar argument applies if $\mu'(\theta) > 0$ and $R_\mu(x, \mu(\theta), \theta) < 0$.

Next, to see that $w(\theta)$ is a step function note that whenever $\mu'(\theta)$ exists it equals 0. From (18) we know that whenever $\mu'(\theta) = 0$, $w'(\theta) = 0$, so that $w(\theta)$ is a step function.

We now show part i) of Proposition 1, namely that $\mu(\theta)$ is continuous.

Lemma 4. Under (A1) and (A2), $\mu(\theta)$ is continuous and $w(\theta) = \bar{w}$ for all θ.

Proof. It should be clear that once we prove that $\mu(\theta)$ is continuous, (18) implies that $w(\theta)$ is also continuous. Since we know from Lemma 3 that $w(\theta)$ is a step function we conclude that $w(\theta)$ is a constant. Optimality implies that this constant is $\bar{w}$.

Next, we show that $\mu(\theta)$ is continuous. By way of contradiction suppose not, that is suppose that $\mu(\theta)$ jumps up at $\tilde{\theta}$. By Lemma 3, on either side of $\tilde{\theta}$, $\mu(\theta)$ is either the static optimum or flat.

We first rule out that $\mu(\theta)$ is equal to the static optimum on either side of $\tilde{\theta}$. Suppose, for example, that $\mu(\theta)$ equals the static optimum on some interval $(\theta_1, \tilde{\theta})$. Since $\mu(\theta)$ jumps up at $\tilde{\theta}$, it must lie strictly above the static optimum for some interval $(\tilde{\theta}, \theta_2)$ so that

$$\lim_{\theta \nearrow \tilde{\theta}} R(x, \mu(\theta), \tilde{\theta}) > \lim_{\theta \searrow \tilde{\theta}} R(x, \mu(\theta), \tilde{\theta}). \tag{33}$$

Hence from the condition (19) in local incentive compatibility

$$\lim_{\theta \nearrow \tilde{\theta}} w(\theta) < \lim_{\theta \searrow \tilde{\theta}} w(\theta). \tag{34}$$

Thus, for $\theta \in (\theta_1, \tilde{\theta})$, $w(\theta)$ is uniformly bounded below $\bar{w}$ and $\mu(\theta)$ is increasing on the same interval. By Lemma 2 such an allocation can not be optimal. We can rule out the case in which $\mu(\theta)$ equals the static optimum on some interval $(\tilde{\theta}, \theta_2)$ with an analogous argument.

Next we rule out that $\mu(\theta)$ is flat on both sides of $\tilde{\theta}$. Let (μ_1, w_1) denote the allocation immediately below $\tilde{\theta}$ and (μ_2, w_2) denote the allocation immediately above $\tilde{\theta}$. We pick an interval (θ_1, θ_2) that contains $\tilde{\theta}$ and satisfies two conditions. First, the allocations equal (μ_1, w_1) on $(\theta_1, \tilde{\theta})$ and they equal (μ_2, w_2) on $(\tilde{\theta}, \theta_2)$. Second, pick the interval (θ_1, θ_2) small enough so that if $R_\mu(x, \mu_1, \tilde{\theta})$ is strictly positive so is $R_\mu(x, \mu_1, \theta_1)$ and if $R_\mu(x, \mu_2, \tilde{\theta})$ is strictly negative so is $R_\mu(x, \mu_2, \theta_2)$.

We construct feasible allocations that yield higher welfare than the original allocations as in Lemma 2. Suppose that (23) holds. We construct the alternative allocation according to (24)-(26). To show that this allocation is feasible we need to show that for sufficiently small a, $w(\theta; a) \leq \bar{w}$ for all $\theta \in [\underline{\theta}, \bar{\theta}]$.

Clearly, $w(\theta; a) = w(\theta) \leq \bar{w}$ for all $\theta \in [\underline{\theta}, \theta_1)$.

For $\theta \in (\theta_1, \tilde{\theta})$ if $w_1 < \bar{w}$, since a is sufficiently small $w(\theta; a) \leq \bar{w}$. If $w_1 = \bar{w}$ we differentiate (26) to obtain

$$(35) \qquad \frac{\partial \tilde{w}(\theta; 0)}{\partial a} = \int_{\theta_1}^{\theta} R_{\theta\mu}(x, \mu_1, z) [\tilde{\mu} - \mu_1] \, dz - R_\mu(x, \mu_1, \theta) [\tilde{\mu} - \mu_1]$$

We show that (35) is negative as follows. Clearly, $\int_{\theta_1}^{\theta} R_{\theta\mu}(x, \mu_1, z) dz = R_\mu(x, \mu_1, \theta) - R_\mu(x, \mu_1, \theta_1)$ so that (35) equals

$$(36) \qquad [R_\mu(x, \mu_1, \theta) - R_\mu(x, \mu_1, \theta_1)] [\tilde{\mu} - \mu_1] - R_\mu(x, \mu_1, \theta) [\tilde{\mu} - \mu_1] = -R_\mu(x, \mu_1, \theta_1)] [\tilde{\mu} - \mu_1] .$$

To see that (36) is negative note the following. First note that $R_\mu(x, \mu_1, \theta_1)$ is strictly positive. To see this note that since $w_1 = \bar{w}$ and $w_2 \leq \bar{w}$ the incentive constraint $R(x, \mu_1, \tilde{\theta}) + \bar{w} = R(x, \mu_2, \tilde{\theta}) + w_2$ implies that $R(x, \mu_1, \tilde{\theta}) \leq R(x, \mu_2, \tilde{\theta})$. But since $\mu(\theta)$ jumps at $\tilde{\theta}$, $\mu_2 > \mu_1$. Since R is strictly concave it follows that $R_\mu(x, \mu_1, \tilde{\theta}) > 0$. By our construction of the interval, since $R_\mu(x, \mu_1, \tilde{\theta})$ is strictly positive so is $R_\mu(x, \mu_1, \theta_1)$. Next note that since $\mu(\theta)$ is increasing on the interval $\tilde{\mu} > \mu_1$. Thus, $-R_\mu(x, \mu_1, \theta_1)] [\tilde{\mu} - \mu_1] < 0$ so that (36) is negative and hence so is (35).

For $\theta \in (\tilde{\theta}, \theta_2)$ if $w_2 < \bar{w}$, since a is sufficiently small $w(\theta; a) \leq \bar{w}$. If $w_2 = \bar{w}$ we differentiate (26) to obtain

$$\frac{\partial \tilde{w}(\theta; 0)}{\partial a} = \int_{\theta_1}^{\tilde{\theta}} R_{\theta\mu}(x, \mu_1, z) [\tilde{\mu} - \mu_1] \, dz + \int_{\tilde{\theta}}^{\theta_2} R_{\theta\mu}(x, \mu_2, z) [\tilde{\mu} - \mu_2] \, dz - R_\mu(x, \mu_2, \theta) [\tilde{\mu} - \mu_2] .$$

which after some manipulation this can be written as

$$(37) \qquad \frac{\partial \tilde{w}(\theta; 0)}{\partial a} = [R_\mu(x, \mu_1, \tilde{\theta}) - R_\mu(x, \mu_1, \theta_1)] [\tilde{\mu} - \mu_1] - R_\mu(x, \mu_2, \tilde{\theta}) [\tilde{\mu} - \mu_2]$$

We show that (37) is negative in two steps. First we note that $R_\mu(x, \mu_2, \theta_2) < 0$ and $\tilde{\mu} < \mu_2$. To see this note that since $w_2 = \bar{w}$ and $w_1 \leq \bar{w}$ the incentive constraint $R(x, \mu_1, \tilde{\theta}) + \bar{w} =$

$R(x,\mu_2,\tilde{\theta})+w_2$ implies that $R(x,\mu_1,\tilde{\theta}) \geq R(x,\mu_2,\tilde{\theta})$. Since $\mu_2 > \mu_1$ and since R is strictly concave it follows 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 $R_\mu(x,\mu_2,\theta_2)$. Since $\mu(\theta)$ is increasing on the interval $\tilde{\mu} < \mu_2$. Next, we note that since we are in case 1, (23) holds. This inequality is equivalent to

$$(38) \quad \left[R_\mu(x,\mu_1,\tilde{\theta}) - R_\mu(x,\mu_1,\theta_1)\right]\left[\tilde{\mu}-\mu_1\right] + \left[R_\mu(x,\mu_2,\theta_2) - R_\mu(x,\mu_2,\tilde{\theta})\right]\left[\tilde{\mu}-\mu_2\right] < 0.$$

Since $R_\mu(x,\mu_2,\theta_2)\left[\tilde{\mu}-\mu_2\right] > 0$, (38) implies (37) is strictly negative.

Finally, for $\theta \in (\theta_2, \bar{\theta})$ the inequality (23) implies directly that $\partial\tilde{w}(\theta;0)/\partial a < 0$ just as we showed in Lemma 2.

The argument for case 2 is analogous. Thus $\mu(\theta)$ is continuous. *Q.E.D.*

Proving Lemmas 3 and 4 proves Proposition 1.

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

Proof. From Proposition 1, we know that $\mu(\theta)$ continuous and it is either flat or equal to the static optimum $\mu^*(\theta;x)$. Clearly, if it is flat everywhere it equals the expected Ramsey policy. If it is not flat everywhere it must be of the following form. For some θ_1 and θ_2, it must have the form

$$\mu(\theta) = \left\{ \begin{array}{c} \mu_1 = \mu^*(\theta_1;x) \text{ if } \theta \in [\underline{\theta}\ ,\theta_1) \\ \\ \mu^*(\theta,x) \text{ if } \theta \in [\theta_1,\theta_2] \\ \\ \mu_2 = \mu^*(\theta_2;x) \text{ if } \theta \in (\theta_2,\bar{\theta}] \end{array} \right\}.$$

Clearly, it cannot be optimal to have both $\theta_1 > \underline{\theta}$ and $\theta_2 < \bar{\theta}$. To see this observe that if that were so then there would exist an alternative allocation $\tilde{\mu}(\theta)$ of the same form (**??**) with $\tilde{\theta}_1 < \theta_1$ and $\tilde{\theta}_2 > \theta_2$ such that $\int \tilde{\mu}(\theta)p(\theta)d\theta = \int \mu(\theta)p(\theta)d\theta = x$. This alternative allocation would yield a higher payoff since, for each θ, this alternative allocation $\tilde{\mu}(\theta)$ would be closer to $\mu^*(\theta; x)$ wherever it differs from $\mu(\theta)$ and hence strictly preferred by all types θ for which $\tilde{\mu}(\theta)$ differs from $\mu(\theta)$. More formally, observe that the marginal impact on welfare of marginal changes in θ_1 and θ_2 of this type is given by

$$d\tilde{V} = \int_{\underline{\theta}}^{\theta_1} \left\{ R_\mu(x, \mu^*(\theta_1; x), \theta) \frac{\partial \mu^*(\theta_1; x)}{\partial \theta} \Delta\theta_1 \right\} p(\theta) d\theta +$$

$$\int_{\theta_2}^{\bar{\theta}} \left\{ R_\mu(x, \mu^*(\theta_2; x), \theta) \frac{\partial \mu^*(\theta_2; x)}{\partial \theta} \Delta\theta_2 \right\} p(\theta) d\theta > 0$$

where the inequality follows from the facts that $R_\mu(x, \mu^*(\theta_1; x), \theta) < 0$, $\partial\mu^*(\theta_1, x)/\partial\theta > 0$, $\Delta\theta_1 < 0$, $R_\mu(x, \mu^*(\theta_2; x), \theta) > 0$, $\partial\mu^*(\theta_2; x)/\partial\theta > 0$, $\Delta\theta_2 > 0$.

Next observe that it cannot be optimal for $\mu(\theta)$ to have the form (**??**) with $\theta_1 > \underline{\theta}$ and $\theta_2 = \bar{\theta}$. To see this observe that there would exist an alternative allocation $\tilde{\mu}(\theta)$ of the same form with $\tilde{\theta}_1 < \theta_1$ and $\tilde{\theta}_2 = \theta_2 = \bar{\theta}$. This alternative allocation $\tilde{\mu}(\theta)$ 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 allocation $\tilde{\mu}(\theta)$ would be strictly preferred to $\mu(\theta)$ since the change from $\mu(\theta)$ to $\tilde{\mu}(\theta)$ directly improves welfare for all types $\theta < \theta_1$, holding x fixed, and it also reduces x, which by (5) also contributes to improving total welfare. More formally, observe that the marginal impact on welfare of a marginal reduction in θ_1 is given by

$$d\tilde{V} = \int_{\underline{\theta}}^{\theta_1} \left\{ R_\mu(x, \mu^*(\theta_1; x), \theta) \frac{\partial \mu^*(\theta_1; x)}{\partial \theta} \Delta\theta_1 \right\} p(\theta) d\theta + \int_{\underline{\theta}}^{\bar{\theta}} \left\{ R_x(x, \mu(\theta), \theta) \Delta x \right\} p(\theta) d\theta > 0$$

where the inequality follows from the facts that $R_\mu(x, \mu^*(\theta_1; x), \theta) < 0$, $\partial\mu^*(\theta_1; x)/\partial\theta > 0$, $\Delta\theta_1 < 0$, $\Delta x < 0$, and (5). Q.E.D.

C. The Optimal Degree of Discretion and the Severity of the Time Inconsistency Problem

So far we have shown that the optimal policy either has bounded discretion or no discretion. Here we link the optimal degree of discretion to the size of the time inconsistency problem. We show that the more severe is the time inconsistency problem the smaller is the optimal degree of discretion.

For most of this section we focus on the quadratic objective function (3). For such an objective function we think of the nonnegative parameter U as indexing the severity of the time inconsistency problem. When U equals to zero there is no time inconsistency problem and as U increases from zero the time inconsistency problem gets worse. To see why note that with this objective function the static best response is $\mu^*(\theta, x) = (U + x + \theta)/2$. The static Nash equilibrium inflation rate can be found from solving for the fixed point in x from

$$(39) \quad x = \int \mu^*(\theta, x)p(\theta)d\theta = \frac{U + x}{2} + \frac{1}{2}\int \theta p(\theta)d\theta.$$

Since $\int \theta p(\theta)d\theta = 0$ we have that the Nash inflation rate is $x^N = U$ and the Nash policies are $\mu^*(\theta, U) = U + \theta/2$. The Ramsey inflation rate is $x^R = 0$ and the Ramsey policies are $\mu^R(\theta) = \theta/2$. Thus, for each type θ, the Nash policies are simply the Ramsey policies shifted up by U. This leads to expected inflation in the Nash equilibrium to be higher than that in the Ramsey equilibrium by U. As U gets smaller both the Nash policies and the Nash inflation rate converge to the Ramsey policies and Ramsey inflation rate. When U is zero these policies coincide and thus the Ramsey allocation repeated every period is incentive compatible.

We first state a simple result that that does not depend on the objective function being quadratic. This result follows fairly immediately from Proposition 2.

Proposition 3. Under (A1) and (A2), if $\mu^*(\underline{\theta}, \mu^{ER}) < \mu^{ER}$ then the optimal policy has bounded discretion.

Proof. By way of contradiction, suppose that $\mu^*(\underline{\theta}, \mu^{ER}) < \mu^{ER}$ but the optimal policy is the expected Ramsey policy. Clearly, there is some point $\hat{\theta} \in [\underline{\theta}, \bar{\theta}]$ such that $\mu^*(\hat{\theta}, \mu^{ER}) = \mu^{ER}$. Let θ_1 and θ_2 satisfy $\theta_1 < \hat{\theta} < \theta_2$. Let $\tilde{\mu}(\theta)$ be given by (**??**) with $\tilde{\mu} = \mu^{ER}$ and let $\mu(\theta; a)$ be given by (25). It should be clear that we can improve on the expected Ramsey policy μ^{ER} by using this variation to shift the inflation of the types from $(\theta_1, \hat{\theta})$ down towards their static optimum, shift the types from $(\hat{\theta}, \theta_2)$ up towards their static optimum and do so in a way the keeps expected inflation constant. It is immediate to use arguments like those in Proposition 2 to show that this variation is incentive compatible as well as welfare improving. This gives the contradiction. *Q.E.D.*

When the objective function satisfies (3), the condition $\mu^*(\underline{\theta}, \mu^{ER}) < \mu^{ER}$ in the proposition reduces to $U < -\underline{\theta}$ where $\underline{\theta}$ is a negative number. Proposition 3 thus implies that bounded discretion is optimal when the time inconsistency problem is sufficiently small, in that the static optimum for the low types is below expected Ramsey.

To get a more precise characterization of the link between the optimal degree of discretion and the size of the time inconsistency problem we assume that the distribution satisfies the following condition,

$$\frac{1 - P(\theta)}{p(\theta)}(\bar{\theta} - \theta) \leq 2 - P(\theta) \tag{40}$$

for all θ. This condition is met by the uniform distribution. The degree of discretion is indexed by the cutoff θ^*. When $\theta^* = \underline{\theta}$ there is no discretion. As θ^* increases toward $\bar{\theta}$ the degree of discretion increases. When $\theta^* = \bar{\theta}$ we say there is complete discretion.

Proposition 4. Assume (3) and (40). If $U = 0$, the optimal allocation has complete discretion. If $U \in (0, -\underline{\theta})$, the optimal allocation has bounded discretion with $\theta^* < \bar{\theta}$. The optimal degree of discretion θ^* is decreasing in U. As U approaches $-\underline{\theta}$, the cutoff θ^* approaches $\underline{\theta}$. If $U \geq -\underline{\theta}$, the optimal allocation is the expected Ramsey allocation with no discretion.

Proof. We prove this proposition by computing the optimal cutoff θ^* under bounded discretion as a function of the parameter U. Under the bounded discretion policy $\mu(\theta) = \mu^*(\theta; x)$ for $\theta \leq \theta^*$ and $\mu^*(\theta^*; x)$ for $\theta > \theta^*$ welfare is

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

$$R(x, \mu^*(\underline{\theta}; x), \underline{\theta}) + \int_{\underline{\theta}}^{\theta^*} R_\theta(x, \mu^*(\theta; x), \theta)(1 - P(\theta))d\theta + \int_{\theta^*}^{\bar{\theta}} R_\theta(x, \mu^*(\theta^*; x), \theta)(1 - P(\theta))d\theta. \quad (41)$$

while the expected inflation under this policy is

$$x = \int_{\underline{\theta}}^{\theta^*} \mu^*(\theta, x)p(\theta)d\theta + \int_{\theta^*}^{\bar{\theta}} \mu^*(\theta^*, x)p(\theta)d\theta \quad (42)$$

The first order conditions for the problem of maximizing (41) with respect to θ^* subject to (42) can be reduced to

$$-(1 - P(\theta^*))(U + x) + \int_{\theta^*}^{\bar{\theta}} (1 - P(\theta))d\theta = 0. \quad (43)$$

(See Appendix for details.)

There are multiple solutions to the first order conditions. One solution is at the corner $\theta^* = \bar{\theta}$ for all values of the parameter U. This solution is not optimal if $U > 0$ since it violates the second order condition. To see this differentiate (43) and evaluate it at $\theta^* = \bar{\theta}$ to obtain

$$p(\theta^*)(U + x). \quad (44)$$

At $\theta^* = \bar{\theta}$, the allocation is static Nash, so $x = U$. Since U is strictly positive then so is (44).

Thus, any solution must satisfy $\theta^* < \bar{\theta}$. To evaluate such a solution we first simplify the expression for x in (42) by plugging in the form of the bounded discretion policy and then simplifying to get

$$x = U - \int_{\theta^*}^{\bar{\theta}} (\theta - \theta^*) p(\theta) d\theta. \tag{45}$$

We then use (45) to rewrite the first order condition (43) as

$$(1 - P(\theta^*))(-2U + \int_{\theta^*}^{\bar{\theta}} (\theta - \theta^*) p(\theta) d\theta) + \int_{\theta^*}^{\bar{\theta}} (1 - P(\theta)) d\theta = 0 \tag{46}$$

Consider the region $U < -\underline{\theta}$. We claim the optimal solution satisfies $\theta^* > \underline{\theta}$, so that in this region the solution is strictly interior, $\underline{\theta} < \theta^* < \bar{\theta}$. To see this use the fact that

$$\int_{\underline{\theta}}^{\bar{\theta}} (1 - P(\theta)) d\theta = \int_{\underline{\theta}}^{\bar{\theta}} d\left[\theta(1 - P(\theta))\right] - \int_{\underline{\theta}}^{\bar{\theta}} \theta d(1 - P(\theta)) = -\underline{\theta} + \int_{\underline{\theta}}^{\bar{\theta}} \theta p(\theta) d\theta = -\underline{\theta} \tag{47}$$

implies that note that when $\theta^* = \underline{\theta}$, (46) reduces to $-2(U - \underline{\theta})$ which is strictly positive since $U < -\underline{\theta}$. Next we want to show that for all $U < -\underline{\theta}$ there is at most one interior solution θ^* to (46) and this solution is decreasing in U. To do so rearrange (46) as

$$2U = \frac{\int_{\theta^*}^{\bar{\theta}} (1 - P(\theta)) d\theta}{(1 - P(\theta^*))} + \int_{\theta^*}^{\bar{\theta}} (\theta - \theta^*) p(\theta) d\theta. \tag{48}$$

To show there is at most one solution to (48) which is decreasing in U we show that the right side of this equation is strictly decreasing in θ^*. Differentiating the right hand side (48) with respect to θ^* gives

$$\frac{-(1 - P(\theta^*))^2 + p(\theta^*) \int_{\theta^*}^{\bar{\theta}} (1 - P(\theta)) d\theta}{(1 - P(\theta^*))^2} - (1 - P(\theta^*)) <$$

$$p(\theta^*) \frac{\bar{\theta} - \theta^*}{(1 - P(\theta^*))} - (2 - P(\theta^*)) \leq 0$$

where the first inequality follows from the fact that

$$\frac{\int_{\theta^*}^{\bar{\theta}} (1 - P(\theta)) d\theta}{(1 - P(\theta^*))^2} < \frac{\int_{\theta^*}^{\bar{\theta}} (1 - P(\theta^*)) d\theta}{(1 - P(\theta^*))^2} = \frac{\bar{\theta} - \theta^*}{(1 - P(\theta^*))}$$

and the second inequality follows from our assumption on the distribution (40).

Finally, consider the region with $U \geq -\underline{\theta}$. We want to show that the expected Ramsey allocation is optimal. We know that the right side (48) is decreasing in θ^*. It is minimized at $\theta^* = \underline{\theta}$ and has the value $-2\underline{\theta}$. Consider first $U = -\underline{\theta}$. At this value of U this first order condition holds at $\theta^* = \underline{\theta}$ and at this policy, (45) implies that $x = 0$. Thus, at $U = -\underline{\theta}$ the optimal allocation with bounded discretion is $\theta^* = \underline{\theta}$ and this allocation is the expected Ramsey allocation.

Consider next the region with $U > -\underline{\theta}$. Here there is no interior solution to the first order condition. Within the set of policies with bounded discretion the best on is the lower corner with $\theta^* = \underline{\theta}$. As we showed above the derivative of the objective function at this policy is $-2(U - \underline{\theta})$ which here is positive. But this implies that within the class of bounded discretion policies that best policy specifies a constant inflation rate. Using (45) we know that this constant inflation rate is $x = U + \underline{\theta}$ which is strictly positive when $U > -\underline{\theta}$. But this policy is dominated by the expected Ramsey policy of zero inflation which, by definition, is the best policy among the policies with constant inflation. $Q.E.D.$

D. Implementing the Equilibrium Policies with an Inflation Cap

We have characterized the best equilibrium of a dynamic game. We imagine implementing this equilibrium with an *inflation cap*, namely 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 private agents then simply play the repeated one-shot equilibrium, then the monetary authority will optimally choose the best equilibrium allocation of the mechanism design problem. In this sense, the repeated one-shot game with an inflation cap implements the allocation that solves the best payoff problem.

Consider the following setup of a one-shot game. Assume that society legislates a rule that restricts the monetary authority's choices to satisfy $\mu_t \leq \bar{\pi}$. Dropping time subscripts we the problem of the monetary authority at a given θ is as follows. Given aggregate wages x choose money growth for this state θ to solve

$$(49) \quad \max_{\mu(\theta)} R(x, \mu, \theta)$$

subject to

$$(50) \quad \mu(\theta) \leq \bar{\pi}.$$

The private agents decisions on wages are summarized by

$$(51) \quad x = \int \mu(\theta) p(\theta).$$

An equilibrium of this one-shot game consists of aggregate wages x and a money growth policy $\mu(\theta)$ such that *i)* taking x as given, $\mu(\theta)$ satisfies (49) and *ii)* x satisfies (51). We denote the optimal choice of the monetary authority as $\mu^*(\theta; x, \bar{\pi})$. 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 $\bar{\pi}$.

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

$$(52) \quad \bar{\pi} = \mu^{ER}.$$

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

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

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 (52) and (53). Then the equilibrium outcome of the one-shot game with the inflation cap for each period coincides with the best equilibrium outcome of the dynamic game.

Proof. To establish this result we first show that the monetary authority will choose the upper bound $\hat{\pi} = \mu^{ER}$ when the expected Ramsey policy is optimal in the dynamic game. Notice that Proposition 3 implies that whenever the expected Ramsey policy is optimal $\mu^{ER} \leq \mu^*(\underline{\theta}; \mu^{ER})$. Recall also that the single crossing condition (A1) implies that the best response is strictly increasing in θ, so that (17) 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 up by increasing their inflation above μ^{ER} and hence the constraint $\hat{\pi} = \mu^{ER}$ binds and all types choose the expected Ramsey levels.

We next show that when bounded discretion is optimal in the dynamic game, in the associated static game with the inflation caps, 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 response of the θ^* type, namely, $\mu^*(\theta; x^*)$ where x^* is the equilibrium expected inflation rate under bounded discretion. Under (A1), the static best responses are increasing in the type so that the best response of any type

$\theta \geq \theta^*$ is greater $\mu^*(\theta; x^*)$. Thus, in the one-shot game with the inflation cap, the constraint (53) binds for such types. Thus, the equilibrium outcomes of the two games coincide. $Q.E.D.$

Discuss that this is weak implementation in the dynamic game (strong in the static game.)

3. Extensions

So far we have focussed on one extreme in which private agents have no direct information about the state of the economy. Here we discuss extending the analysis to the case in which private agents see a noisy signal of the state. We show that our analysis extends to this case with essentially no modifications.

Consider a version of our model with one change. At the beginning of each period, private agents see a signal s_t about the current state θ_t. This signal is i.i.d. over time and has density $q(s)$. We let $p(\theta|s)$ and $P(\theta|s)$ denote the density and distribution functions over the state θ, conditional on the state.

In the repeated game the history $h_t = (s_0, \hat{\theta}_0, s_1, \hat{\theta}_1, s_1, \ldots, \hat{\theta}_{t-1}, s_t)$ includes the history of signals as well as the history of past reports. Society's objective is

$$(54) \quad (1-\beta)\sum_{t=0}^{\infty}\int\int \beta^t R(x_t(h_t), \mu_t(h_t, \hat{\theta}_t), \theta_t) p(\theta_t|s_t) d\theta_t ds_t.$$

In the recursive formulation, we say that actions $x(s), \mu(\theta, s)$, and function $w(\theta, s)$ are *enforceable by* W if,

$$(55) \quad w(\theta, s) \in W$$

$$(56) \quad x(s) = \int \mu(\theta, s) p(\theta|s) d\theta$$

and the incentive constraints

$$(57)\quad (1-\beta)R(x,\mu(\theta,\ s),\theta)+\beta w(\theta,\ s)\geq(1-\beta)R(x,\mu(\hat{\theta},\ s),\theta)+\beta w(\hat{\theta},s)$$

for all θ and for all $\hat{\theta}$.

The payoff corresponding to $x=(x(s)),\mu(\theta)=(\mu(\theta,s))$, and $w(\theta)=(w(\theta,s))$ is

$$\Pi(x(s),\mu(\theta,s),w(\theta,s))=\int\int\left[(1-\beta)R(x(s),\mu(\theta,s),\theta)+\beta w(\theta,s)\right]p(\theta|s)q(s)d\theta ds$$

Define the operator T that maps sets so payoffs W into new sets of payoffs $T(W)$ according to $T(W)=\{\Pi(x,\mu(\theta,s),w(\theta,s))\mid x(s),\mu(\theta,s),w(\theta,s)$ are enforceable by $W\}$. The best payoff problem is now

$$(58)\quad \bar{v}=\max_{x(s),\mu(\theta,s)\in\left[\underline{\mu},\bar{\mu}\right],w(\theta,s)}\int\int\left[(1-\beta)R(x,\mu(\theta,s),\theta)+\beta w(\theta,s)\right]p(\theta|s)g(s)d\theta ds$$

subject to constraint that for all s, $x(s)$, $\mu(\theta,s)$, and $w(\theta,s))$ are enforceable by W, in that they satisfy (55), (56), (57).

Our previous results carry through in this extension because this best payoff problem can be broken down to a collection of subproblems for each s of the form

$$(59)\quad \bar{v}(s)=\max_{x(s),\mu(\theta,s),w(\theta,s)}\int\left[(1-\beta)R(x,\mu(\theta,s),\theta)+\beta w(\theta,s)\right]p(\theta|s)g(s)d\theta.$$

But each of these subproblems has exactly the same form as the original problem (15) that we have already analyzed. As long the conditional distributions satisfy the analog of (A2) our original results will go through.

4. Conclusion

What is the optimal degree of discretion in monetary policy? For economies with severe time inconsistency problems, it is zero. For economies with less severe time inconsistency problems, it is not zero, but bounded. More generally, the optimal degree of discretion is decreasing in the severity of the time inconsistency problem. And whatever the severity of that problem, the optimal policy can be implemented by enforcing a simple inflation cap.

In our simple model, there is no publicly observed state; hence, the optimal inflation cap is a single number. If the model were extended to have a publicly observed state, then the optimal rule would respond to this state, but not to the private information. To achieve this, the government would specify a rule for setting the inflation cap with public information. We interpret such a rule as a type of inflation targeting.

References

[1] Abreu, Dilip. 1986. Extremal Equilibria of Oligopolistic Supergames, *Journal of Economic Theory*, 39(1), 191-225.

[2] Abreu, Dilip. 1988. On the Theory of Infinitely Repeated Games with Discounting. *Econometrica*, 56(2), 383-96.

[3] Abreu, Dilip, Pearce, David and Stacchetti, Ennio. 1986. Optimal Cartel Equilibria with Imperfect Monitoring *Journal of Economic Theory* 39(1), 251-69.

[4] Abreu, Dilip, Pearce, David and Stacchetti, Ennio. 1990. Toward a Theory of Discounted Repeated Games with Imperfect Monitoring. *Econometrica*, 58(5), 1041-63.

[5] Athey, Susan, Bagwell Kyle, and Sanchirico, Chris. 1998. Collusion and Price Rigidity. MIT Working Paper 98-23.

[6] Barro, Robert J. and Gordon, David. 1983. Rules, discretion and reputation in a model of monetary policy. Journal of Monetary Economics, 12, 101-122.

[7] Canzoneri, Matthew B. 1985. Monetary Policy Games and the Role of Private Information, *American Economic Review*, 75(5), 1056-70.

[8] Chang, Roberto. 1998. Credible Monetary Policy in an Infinite Horizon Model: Recursive Approaches. *Journal of Economic Theory*, 81(2), 431-61.

[9] Chari, V. V., Kehoe, Patrick J. 1990. Sustainable Plans. *Journal of Political Economy*, 98(4), 783-802.

[10] Chari, V. V., Kehoe, Patrick J., and Prescott, Edward C.1989. Time Consistency and Policy, in *Modern business cycle theory*, edited by Robert J. Barro,Cambridge, Mass.: Harvard University Press, 265-305.

[11] Cukierman, Alex, Meltzer, Allan H.A. 1986. A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information. *Econometrica;* 54(5), 1099-1128.

[12] Faust, Jon, Svensson, Lars E. O. 1998. Transparency and Credibility: Monetary Policy with Unobservable Goals, NBER working paper 6452.

[13] Faust, Jon, Svensson, Lars E. O. 1999. The Equilibrium Degree of Transparency and Control in Monetary Policy. NBER working paper 7152.

[14] Goodfriend, Marvin. 1986. Monetary Mystique: Secrecy and Central Banking. *Journal of Monetary Economics*; 17(1), 63-92.

[15] Green, Edward-J., and Porter, Robert H. 1984. Noncooperative Collusion under Imperfect Price Information. *Econometrica*, 52(1), 87-100.

[16] Kydland, Finn, and Prescott, Edward C. 1077. Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy*, 85, 473-491.

[17] Persson, Torsten, and Tabellini, Guido. 1994. *Monetary and Fiscal Policy: Volume 1, Credibility*. MIT Press, Cambridge, MA.

[18] Phelan, Chris, and Stachetti, Ennio. 1999. Sequential Equilibria in a Ramsey Tax Model. Federal Reserve Bank of Minneapolis Staff Report. (Forthcoming *Econometrica*.)

[19] Romer, Christine, and Romer, David. 2000. Federal Reserve Information and the Behavior of Interest Rates. *American Economic Review*, 90, 429-57.

Appendix

Here we derive several of the equations used in the text.

Derivation of (43).

The first order conditions determining the optimal choice of θ^* is given by

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

$$+\int_{\underline{\theta}}^{\theta^*} R_{\theta\mu}(x,\mu^*(\theta;x),\theta)(1-P(\theta))\frac{\partial}{\partial x}\mu^*(\theta;x)\frac{dx}{d\theta^*}d\theta$$

$$+\int_{\underline{\theta}}^{\theta^*} R_{\theta x}(x,\mu^*(\theta;x),\theta)(1-P(\theta))\frac{dx}{d\theta^*}d\theta$$

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

$$+\int_{\theta^*}^{\bar{\theta}} R_{\theta x}(x,\mu^*(\theta^*;x),\theta)(1-P(\theta))\frac{dx}{d\theta^*}d\theta$$

and

$$\frac{dx}{d\theta^*} = \int_{\underline{\theta}}^{\theta^*}\frac{\partial}{\partial x}\mu^*(\theta,x)\frac{dx}{d\theta^*}p(\theta)d\theta + \int_{\theta^*}^{\bar{\theta}}\left[\frac{\partial}{\partial x}\mu^*(\theta^*,x)\frac{dx}{d\theta^*} + \frac{\partial}{\partial\theta^*}\mu^*(\theta^*,x)\right]p(\theta)d\theta.$$

By the definition of μ^*, we have

$$R_\mu(x,\mu^*(\underline{\theta};x),\underline{\theta}) = 0.$$

From our quadratic example

$$\mu^*(\theta,x) = \frac{U+x+\theta}{2}$$

so

$$R_x(x,\mu^*(\underline{\theta};x),\underline{\theta}) = -(U+x-\mu^*(\underline{\theta};x)) = -\left(\frac{U+x-\underline{\theta}}{2}\right)$$

$$\frac{\partial}{\partial x}\mu^*(\theta,x) = \frac{1}{2} \text{ and } \frac{\partial}{\partial\theta^*}\mu^*(\theta^*,x) = \frac{1}{2}$$

$$R_{\theta\mu}(x,\mu,\theta) = 1 \text{ and } R_{\theta x}(x,\mu,\theta) = 0$$

Hence, our derivatives come down to

$$\frac{dx}{d\theta^*} = \frac{1}{2}\int_{\underline{\theta}}^{\theta^*} p(\theta)d\theta\frac{dx}{d\theta^*} + \frac{1}{2}\left[\frac{dx}{d\theta^*}+1\right]\int_{\theta^*}^{\bar{\theta}} p(\theta)d\theta$$

or

$$\frac{dx}{d\theta^*} = (1 - P(\theta^*)).$$

And

$$\frac{d}{d\theta^*}\left[\int_{\underline{\theta}}^{\bar{\theta}} U(\theta)p(\theta)d\theta\right] = R_x(x, \mu^*(\underline{\theta}; x), \underline{\theta})\frac{dx}{d\theta^*} + \frac{1}{2}\int_{\underline{\theta}}^{\theta^*}(1 - P(\theta))d\theta\frac{dx}{d\theta^*}$$

$$+\left[\frac{1}{2}\frac{dx}{d\theta^*} + \frac{1}{2}\right]\int_{\theta^*}^{\bar{\theta}}(1 - P(\theta))d\theta = 0.$$

This can be simplified to

$$(1 - P(\theta^*))\left[R_x(x, \mu^*(\underline{\theta}; x), \underline{\theta}) + \frac{1}{2}\int_{\underline{\theta}}^{\bar{\theta}}(1 - P(\theta))d\theta\right] + \frac{1}{2}\int_{\theta^*}^{\bar{\theta}}(1 - P(\theta))d\theta = 0.$$

Note that integration by parts gives

$$\int_{\underline{\theta}}^{\bar{\theta}}(1 - P(\theta))d\theta = \int_{\underline{\theta}}^{\bar{\theta}} d\left[\theta(1 - P(\theta))\right] - \int_{\underline{\theta}}^{\bar{\theta}} \theta d(1 - P(\theta)) = -\underline{\theta} + \int_{\underline{\theta}}^{\bar{\theta}} \theta p(\theta)d\theta$$

Hence our first order condition can be written

$$-(1 - P(\theta^*))\left(U + x\right) + \int_{\theta^*}^{\bar{\theta}}(1 - P(\theta))d\theta = 0.$$

with x given as above.

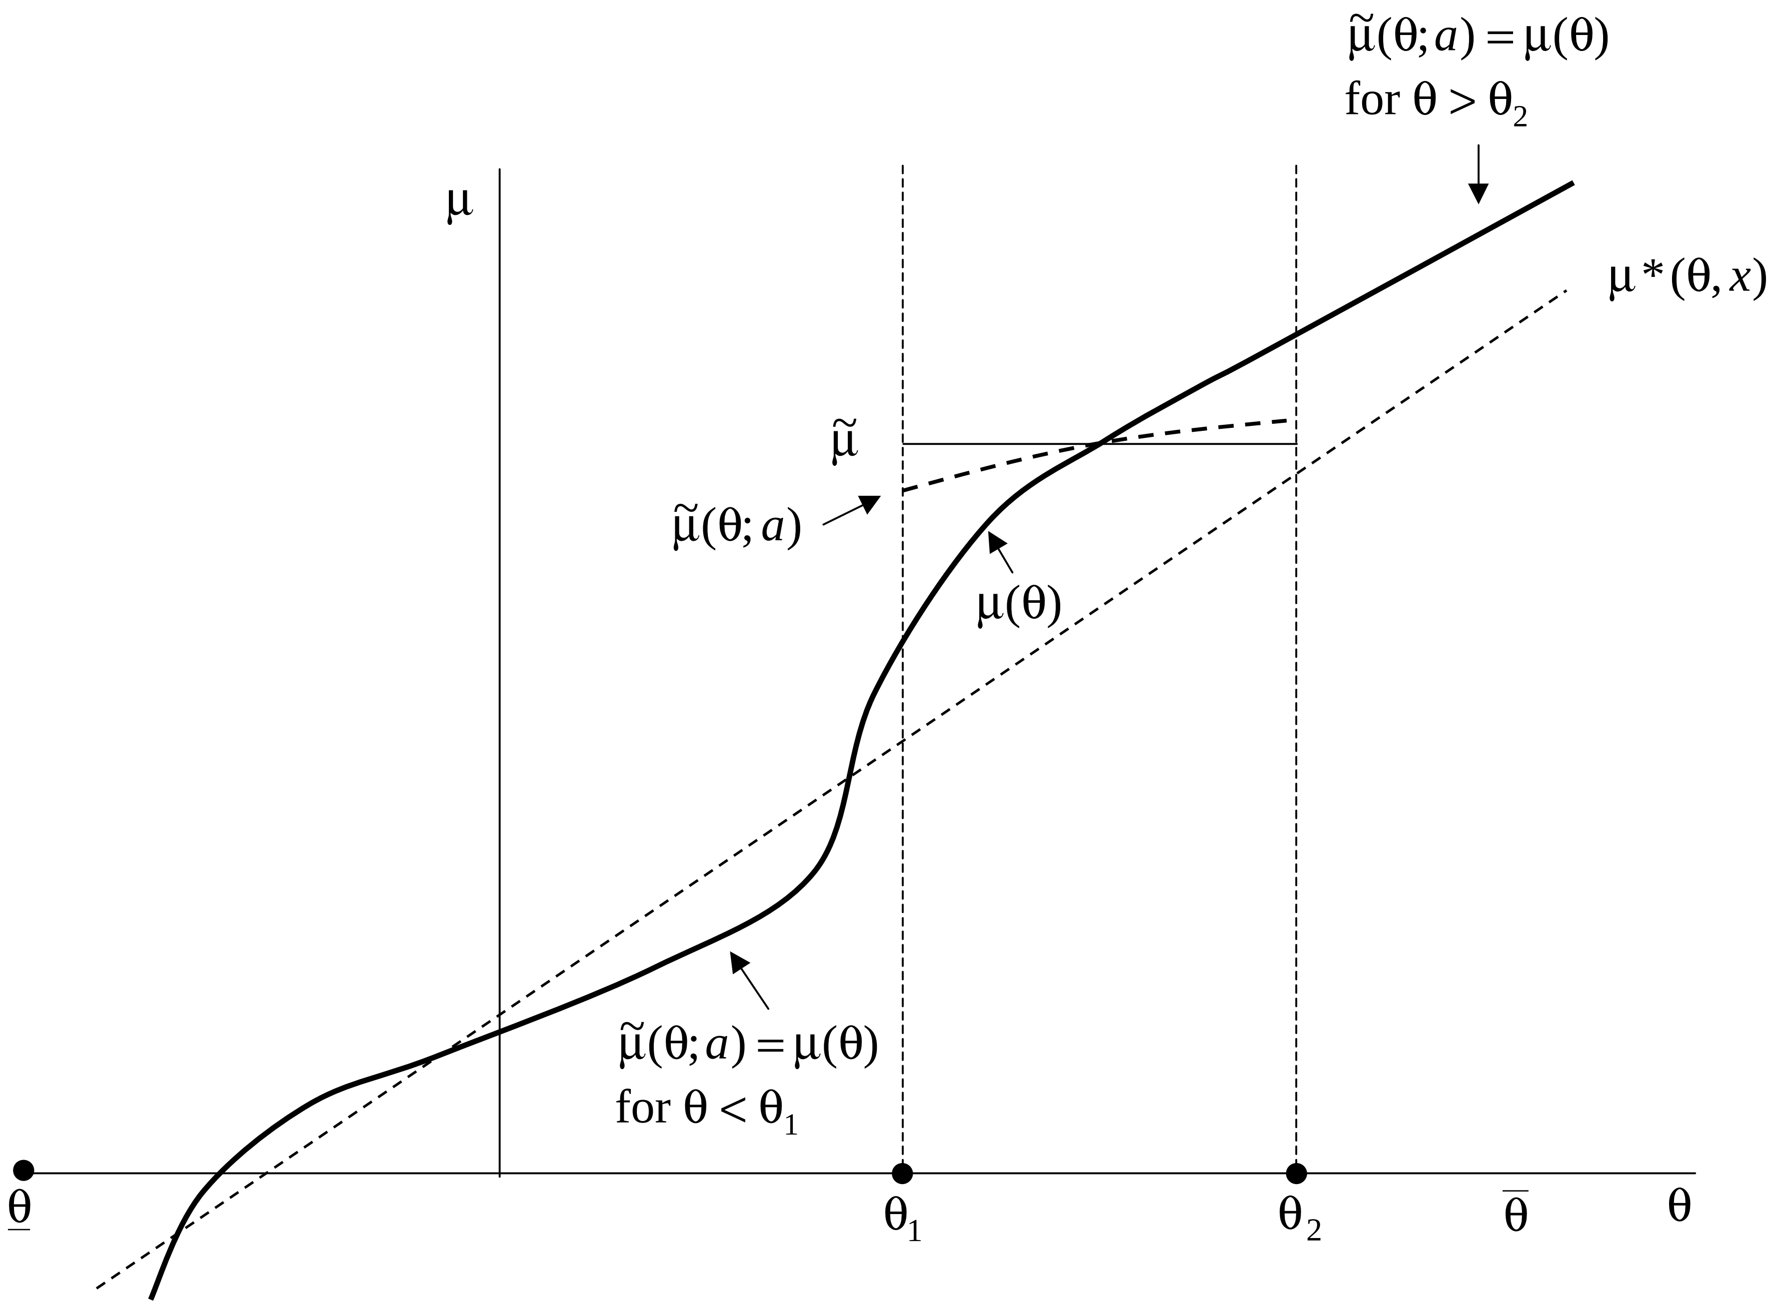

Figure 1: A Welfare Improving Variation

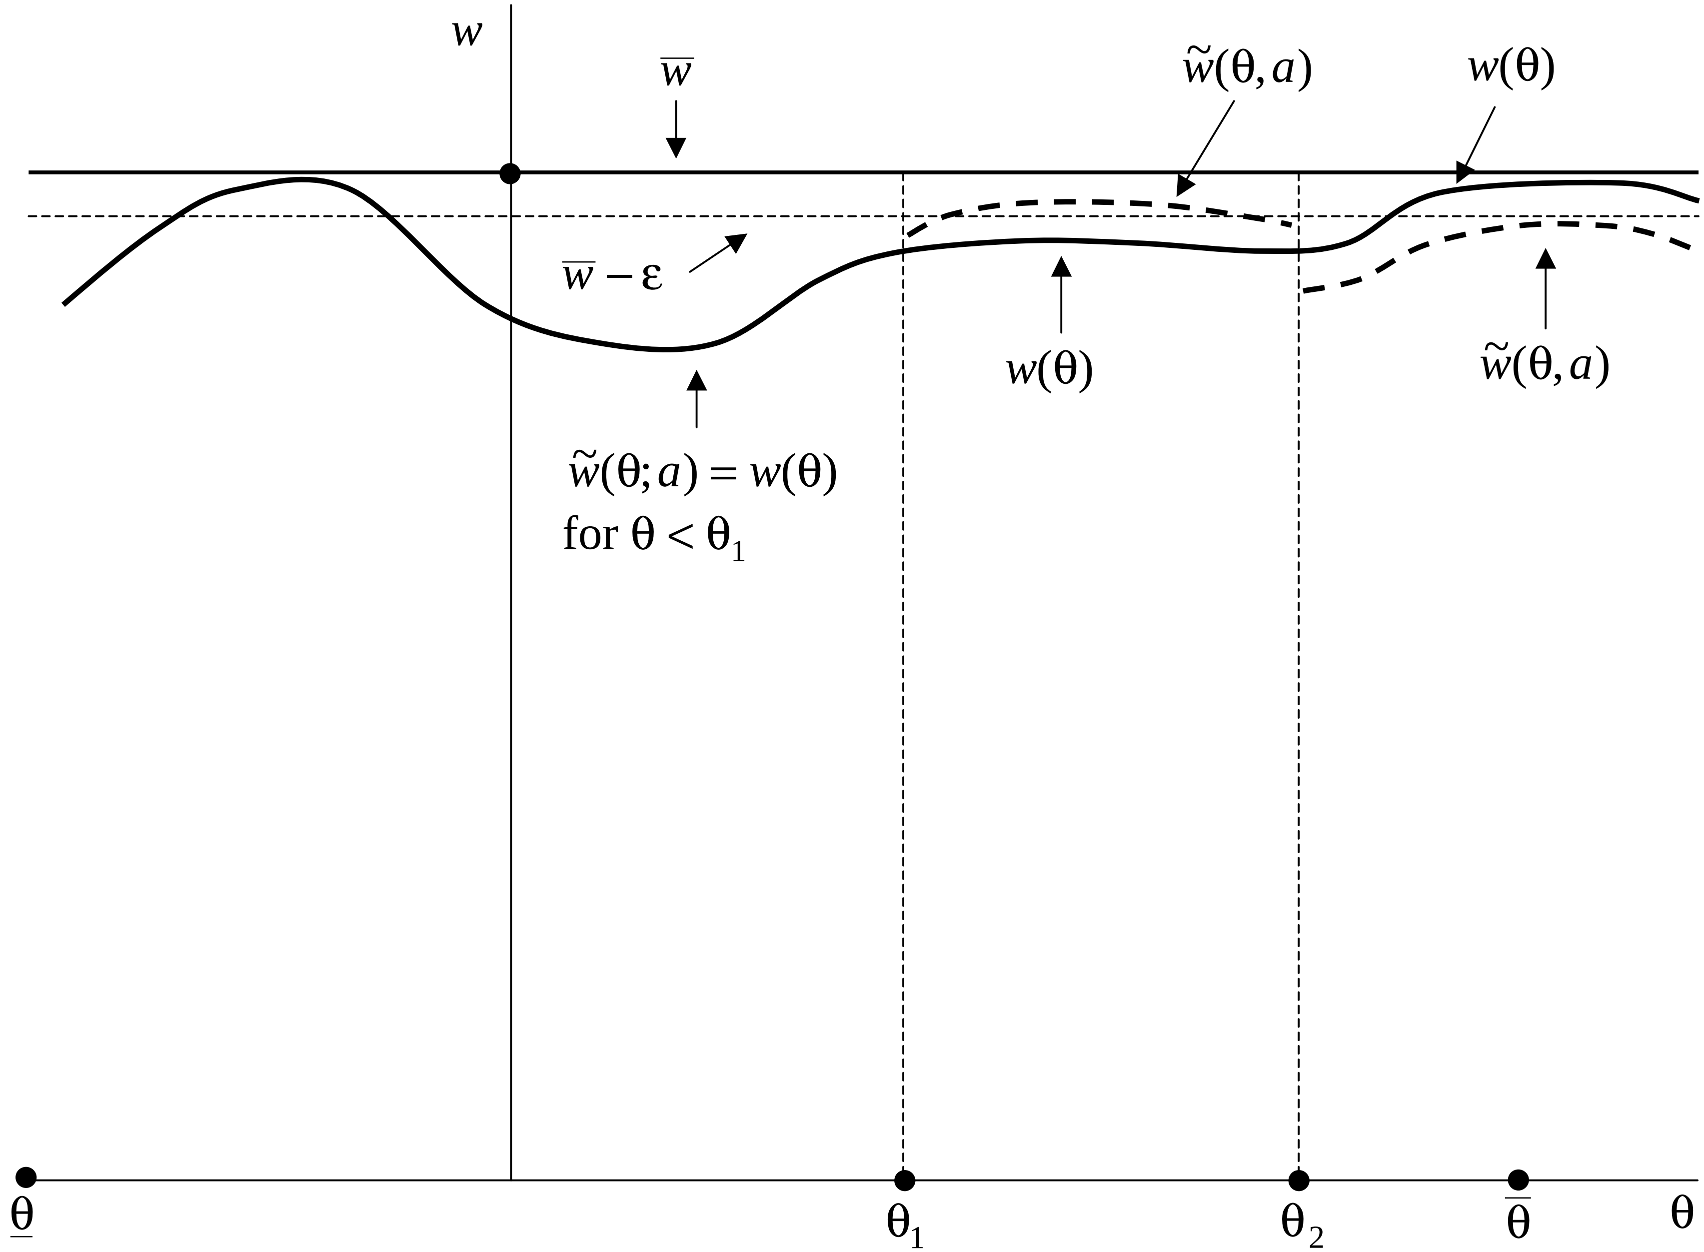

Figure 2: The Associated Continuation Value in Case 1

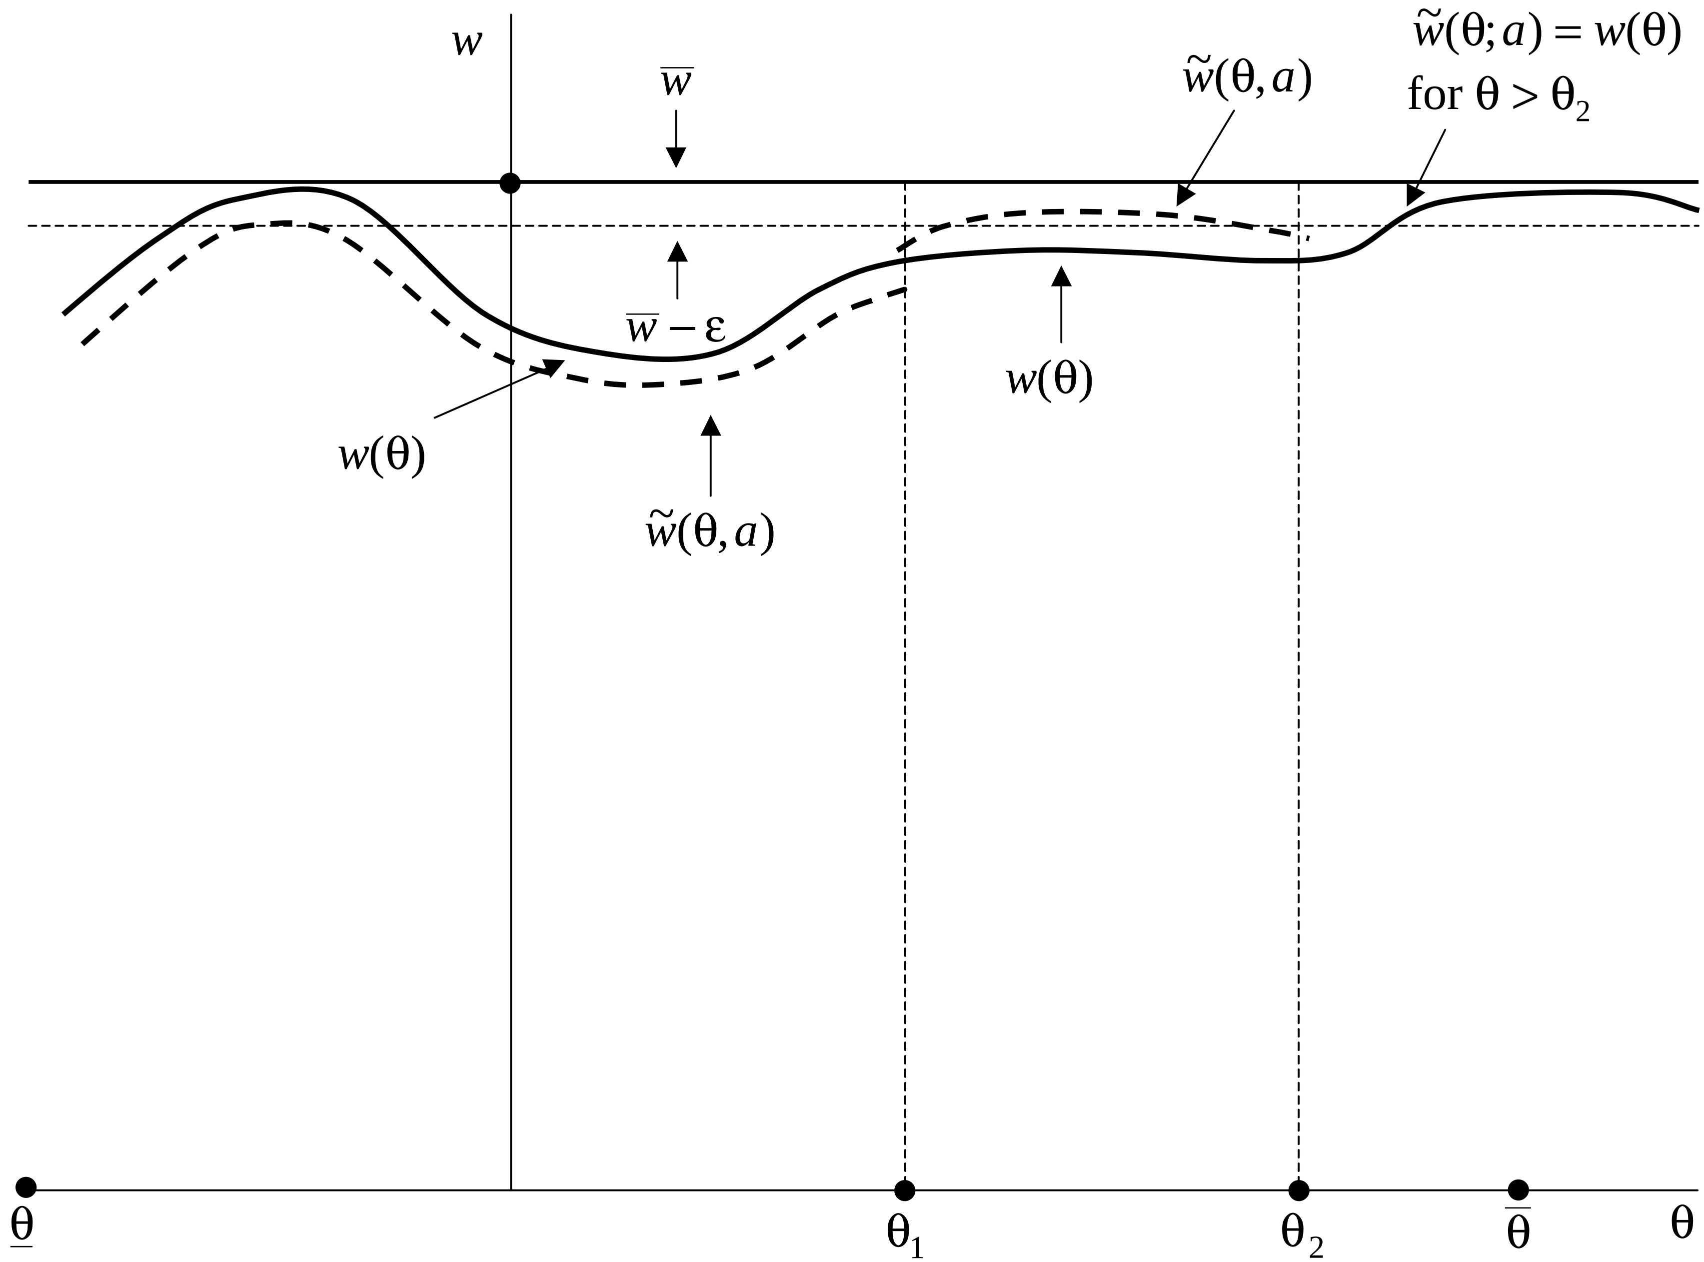

Figure 3: The Associated Continuation Value in Case 2