Inflation Targeting in Emerging Markets: A Game-Theoretic Approach

(Preliminary and Incomplete)

José A. Rodrigues-Neto*

September 25, 2006

Abstract

In an inflation targeting regime of an emerging market, we study the strategic interaction of a continuum of myopic market participants with the central bank, modeled as a long-run player. A sufficiently patient central bank can implement its most preferred Nash equilibrium, and can accomplish this without playing the tightest monetary policy in any period. We also study how the structure of market players' payoffs influences optimal monetary policy, and what could be the institutional framework of an optimal target choice process.

Keywords: games, monetary policy, inflation target, optimal target, central banks.

JEL Classification: E50, E52, E58, C72.

*Central Bank of Brazil. E-mail: jose.alvaro@bcb.gov.br. I would like to thank Fabio Araujo, Marta and Waldyr Areosa, Maurício Bugarin, Luciana Fiorini, André Minella, as well as seminar participants at Universidade de Brasília, Universidade de São Paulo, and Central Bank of Brazil. All errors are my own responsibility. The views expressed here are those of the author and not necessarily those of the Central Bank of Brazil.

Contents

1 Introduction

Starting with New Zealand in 1990, a significant number of countries, including several nations in emerging markets, have been adopting inflation targeting. In this context, some natural questions arise. Is the dynamics of this regime the same in developed and developing countries? We know that these two groups of nations differ dramatically in many respects. How these differences are translated in the strategic interaction of the government with private agents when it comes to the benefits and drawbacks of inflation targeting?

Here, we develop a simple game-theoretic model in which the monetary authority, traditionally the country's central bank (CB), publicly announces a target for the future inflation, and then plays an infinitely repeated game with a continuum of anonymous myopic market participants. In this repeated game, the CB is a long-run player.

We find that the CB can always implement its preferred equilibrium of the repeated game. The monetary authority does not need to play the tightest monetary policy in any period to achieve this equilibrium. Its intensity of cooperation should be only sufficiently large to make market participants' best responses be their part in the prescribed equilibrium.[1] The structure of market agents' payoffs determines the minimal level of cooperation that the CB needs to implement its most preferred Nash equilibrium. There is an inflation bias in the choice of the target. We discuss possible solutions.

We assume that the CB wants to decrease inflation as much as possible, but it also

[1] Woodford (1999) has a similar conclusion. In his words, "...some commentators have proposed that U.S. monetary policy has been so successful at inflation stabilization in the 1990s, despite relatively little change in the funds rate for years at a time, because 'the bond market does the Fed's work for it,' responding to disturbances in the way needed to keep inflation stable without the need for large policy adjustments by the Fed. This is exactly what a good policy regime should look like..."

has a short-run incentive to run a loose monetary policy. Independently of its own action, the policy maker always wants to obtain cooperation from the market agents. In order to get this, the CB needs to establish credibility in the sense that market participants must believe that inflation will be sufficiently low in order to cooperate. This credibility of the CB is technically represented by a Nash equilibrium of the repeated game in which the CB sacrifices short-run gains in order to keep obtaining cooperation of market players.

The inflation target choice determines the stage-game payoffs of all players. By choosing a target for future inflation, the CB signals to the market its intensions regarding the conduction of monetary policy, and consequently, the level that future inflation should be. Hence, at first glance, a low target is desired. Under a relatively more ambitious (i.e. lower) target, inflation may in fact become lower, provided that such a target is credible. However, under a smaller target, a defection by the CB is more costly to market agents who trusted in the announced policy than it would be under a relatively larger target. This suggests that the CB may obtain more cooperation when the target is larger (since market agents have less to lose). With more cooperation by market participants the inflation will decrease. Therefore, it is not clear *a priori* what is the optimal target. Furthermore, we also assume that the CB has a significant direct cost for playing a tighter monetary policy, and that this cost is convex and decreasing in the target.

1.1 Brief Literature Review

Many authors, such as Svensson (1996) and (1997), Mishkin and Schmidt-Hebbel (2001), Clarida et al (1999), Giannone and Woodford (2003), and Woodford (2003), have studied inflation targeting. Taylor (2000), Fraga *et al* (2003), Mishkin (2004) and Batini *et al* (2005) bring the particulars of this regime in emerging markets.

Fudenberg *et al* (1990) introduced the study of repeated games with short-run and long-run players.[2] Applications of game theory to monetary policy were made, among others, by Barro and Gordon (1983) and (1986), Barro (1986), Persson and Tabellini (1995) and, more recently, by Carvalho and Bugarin (2006). Observe that, in contrast with most of these papers, we are not using incomplete information in our framework. None of them bring explicitly the target selection as a strategic action in a game, nor they model the CB and market agents as long and short-run players respectively.

Next section describes the game and section 3 calculates the best equilibrium for the CB in the repeated game. Section 5 studies which is the best target for the CB. Section 6 discusses the inflation bias. The appendix has all proofs.

2 The Game

At the beginning of the game, the CB chooses the next inflation target, denoted $\pi^* \in [0, 1]$, and normalized to be inside the unit interval. The target choice defines stage-game payoffs of an infinitely repeated game. In this repeated game, the monetary authority is a long-run player, named player 1. Player 2 is not strategic; it only summarizes the aggregate behavior of a continuum of anonymous and myopic market agents indexed by $j \in (0, 1)$.[3] Describing market participants as anonymous short-run players is a convenient way of modeling agents that cannot be individually punished by the CB because it is too costly to perfectly observe their individual behavior, or because the CB cannot legally punish market agents that are taking some undesirable behavior.

[2] Alternative interpretations of short-run and long-run player models, as well as the interpretation of the anonymity assumption of short-run players are discussed in Mailath and Samuelson (2006).

[3] Alternatively, we can imagine that in each round a market participant is randomly selected to act as player 2. See chapter 2 of Samuelson and Mailath (2006) for more details.

2.1 Actions

Each market agent chooses a pure action and her payoff depends only on her own action and on the action of the CB,[4] while the payoff of player 1 depends on its own action and on the aggregate behavior of all market players.

In the stage-game, the CB chooses a mixture of what we call cooperation (C) and defection (D). Formally, $A_1 = \{C, D\}$ represents the set of pure actions available to the policy maker. Full cooperation by player 1 means that the CB is playing the tightest monetary policy. Mixed actions of the CB are characterized by $x \in [0, 1]$, the intensity of cooperation, that is, the intensity of C.[5]

Each market agent j chooses a pure action, denoted by a_j, in the set $\{C, D\}$. Cooperation by player 2 means that all market agents act in a way that is consistent with the announced inflation target. All actions are observable, including distributions of mixed actions.

2.2 Utility of the Players

In each round $t \in \{0, 1, 2, \cdots\}$, let u_1^t denote the payoff of player 1 in period t. Whenever the context is clear, we will omit the time superscript, and player 1's payoff will be denoted by u_1. The utility function of the CB, denoted by U_1, is defined as:

$$U_1 = (1 - \delta) \sum_{t=0}^{+\infty} \delta^t u_1^t,$$

where $\delta \in (0, 1)$ represents the CB's discount factor.

[4] Allowing the payoff of a market participant to also depend on the aggregate behavior of other market agents does not bring any new issue. See chapter 2 of Mailath and Samuelson (2006). Allowing that each market participant plays a mixed action can only increase the technicalities of the model.

[5] Usually, a mixed action is given by a probability distribution and players try to maximize their expected utility. Our interpretation here is a little different. We assume that player 1 choose its action x in a continuum $[0, 1]$. and all players observe this choice. Mathematically, both situations are equivalent when we assme that, in the former interpretation of the model, all players can observe the distribution with which the CB would be making its action.

All market agents are myopic short-run players. They always play stage-game best responses in every round. The aggregate behavior of a continuum of anonymous market agents is referred to as the mixed action of player 2. This is not a strategic player since we do not assume any coordination among market agents.

2.2.1 Stage-Game Payoffs

Figure 1 brings the stage-game payoffs. We assume that π_H, π_M and π_L represent high (H), medium (M), and low (L) levels of inflation, respectively. The direct cost for the CB of playing the tightest monetary policy is denoted c, and the magnitude of market players' maximum payoff is represented by $g > 0$. We also assume that all these values are smooth functions of the target such that for any $\pi^* \in [0, 1]$:[6]

$$\pi_H > \pi_M > \pi_L > 0 \tag{H1}$$

$$\frac{\partial \pi_L}{\partial \pi^*} > 0, \quad \frac{\partial \pi_M}{\partial \pi^*} > 0, \quad \frac{\partial \pi_H}{\partial \pi^*} > 0, \quad \frac{\partial c}{\partial \pi^*} < 0, \quad \frac{\partial g}{\partial \pi^*} < 0, \tag{H2}$$

$$\frac{\partial^2 \pi_L}{\partial \pi^{*2}} > 0, \quad \frac{\partial^2 \pi_M}{\partial \pi^{*2}} > 0, \quad \frac{\partial^2 \pi_H}{\partial \pi^{*2}} > 0, \quad \frac{\partial^2 c}{\partial \pi^{*2}} > 0, \tag{H3}$$

$$c \in \left(Max\left\{\pi_H - \pi_M, \pi_M - \pi_L\right\}, \; 2\pi_H - \pi_M - \pi_L \right) \tag{H4}$$

We also assume that:

$$\frac{-\partial c}{\partial \pi^*}(0) > \frac{\partial \pi_L}{\partial \pi^*}(0) + \frac{\partial \pi_M}{\partial \pi^*}(0), \qquad \text{and} \qquad \frac{-\partial c}{\partial \pi^*}(1) < \frac{\partial \pi_L}{\partial \pi^*}(1) + \frac{\partial \pi_M}{\partial \pi^*}(1) \tag{H5}$$

Hypothesis (H1) tell us that all inflation levels π_ω, with $\omega \in \{L, M, H\}$, are positive. Hypothesis (H2) says that all π_ω are increasing in the inflation target. Also, (H2) requires that the direct cost c of implementing the tightest monetary policy is decreasing in the target, and that the magnitude g of market agents' possible payoffs decreases with the target.

[6] Observe that hypothesis (H1) implies that both inequalities $2\pi_H - \pi_M - \pi_L > \pi_H - \pi_M$ and $2\pi_H - \pi_M - \pi_L > \pi_M - \pi_L$ are satisfied.

(H4) is the emerging markets hypothesis. It establishes bounds on the direct cost c. The lower bound for c is positive by (H1), and makes the stage-game be such that the CB always prefers to play D, no matter what proportion of market agents play C. In other words, we assume that the cost for the CB of reducing inflation directly is always larger than its direct benefit. Because the stage-game is played repeatedly, reputation effects provide the incentives for the CB to play its part of a Nash equilibrium of the repeated game in which, on the equilibrium path, player 1 plays a mixed action in every round. If the direct cost c were sufficiently small, the outcome (C, C) would be played in every round. As we will see ahead, the upper bound on c is necessary for the implementation of the repeated game Nash equilibrium. If the cost c is too large, the only Nash equilibrium of the repeated game is the perpetual play of (D, D).

Observe that playing D is a dominant strategy for player 1. Thus, the CB is always tempted to run a loose monetary policy. On the other hand, by assumption, $-\pi_L - c > -\pi_M - c$, and $-\pi_M > -\pi_H$. Hence, the CB always benefits if all market participants play C. When all players coordinate their efforts to fight inflation, the outcome is (C, C), and then, inflation is lower the smaller the target is. At the same time, we may observe that the effort to play C may depend on π^*. If the target is relatively smaller, playing C may require the CB to choose a relatively higher interest rate. Technically speaking, we model this by assuming that the function $\pi^* \mapsto c(\pi^*)$ is decreasing. The section that introduces inflation explicitly presents more intuition about the stage game payoffs.

Together with (H2), the role of hypotheses (H3) and (H5) is to provide sufficient conditions for the uniqueness of an interior optimal target for the CB. (H3) says that all π_ω and c are convex functions of the target. Since all inflation functions π_ω are increasing in π^*, their sensibilities are higher at targets close to the maximum. The direct cost c is decreasing, making it more sensible to changes at lower levels of π^*.

		Market Agent	
		C	D
Central Bank	C	$-\pi_L - c$, $+g$	$-\pi_M - c$, $-g$
	D	$-\pi_M$, $-g$	$-\pi_H$, $+g$

Figure 1: Stage-game payoffs.

Observe that this behavior is compatible with (H5).

2.2.2 Market Agents' Possible Payoffs

For markets participants, the goal is to coordinate their action with the CB. In figure 1 we present their possible payoffs. Their payoffs $g > 0$ at (C, C) and (D, D) are decreasing in the target. On the other hand, if a market player misses the CB's action by playing C when player 1 plays D or vice-versa, she obtains only $-g$, which is the lowest feasible payoff. This payoff is an increasing function of the inflation target. The more ambitious the target is (lower π^*), the more costly miscoordination becomes to market players. The intuition behind this is that the lower the inflation target is, the stronger will be the implicit contract between the CB and market players. Therefore, the lower π^* is, the larger will be the degree of the mistake that market agents might be doing by betting on the contrary action to what the CB actually plays.

2.2.3 More Complex Payoffs

In general, we could admit more flexible payoffs with little or no change in our model. For each market agent $j \in (0, 1)$, whenever $u_j(D, D) \neq u_j(D, C)$, define the variable

ξ_j as:

$$\xi_j = \frac{u_j(C,C) - u_j(C,D)}{u_j(D,D) - u_j(D,C)}$$

In the game shown in figure 1, $\xi_j = 1$ for any market player j, regardless of the inflation target. Therefore, regardless of the target, each j prefers to play C when $x > 1/2$, prefers D when $x < 1/2$, and is indifferent when $x = 1/2$. The specific value that ξ_j assumes is not so important. The only essential characteristic of the set of market agents' possible payoffs is given by hypothesis (H6) below:

$$u_j(C,C) > u_j(C,D), \qquad \text{and} \qquad u_j(D,D) > u_j(D,C), \qquad \forall j \in (0,1) \tag{H6}$$

In most of paper, we also assume the following hypotheses:

$$\text{All } \xi_j \text{ , with } j \in (0,1)\text{, are independent of the target } \pi^* \tag{H7}$$

$$\text{All } \xi_j \text{ , with } j \in (0,1)\text{, are constant independently of player } j \tag{H8}$$

There could exist multiple Nash equilibria in the repeated game if we introduce some heterogeneity among market players, that is, if (H8) does not hold. We deal with this issue in a later section.

2.2.4 Payoffs of Mixed Actions

Any mixed action α_1 of the CB can be characterized by $x \in [0,1]$, the intensity that player 1 chooses C. This is consistent with the interpretation of x as an increasing function of the interest rate set by the monetary authority. Similarly, player 2's mixed action α_2 is characterized by $y \in [0,1]$, the proportion of market agents playing C.[7] With this notation, the stage-game payoff of the CB associated with the action profile (x,y) is:

$$u_1(x,y) = x\left[(2\pi_M - \pi_H - \pi_L)\, y + \pi_H - \pi_M - c\right] - \left[y\pi_M + (1-y)\pi_H\right] \tag{1}$$

[7] When the alternative interpretation of footnote 3 is used for player 2, y is the probability that the randomly chosen agent plays C.

For any market player j, possible payoffs in any round are:

$$u_j(x, C) = [2x - 1]g \tag{2a}$$

$$u_j(x, D) = [1 - 2x]g \tag{2b}$$

2.3 Nash Equilibria of the Stage-Game

For the CB, action D is dominant. Hence, in any Nash equilibrium of the stage-game the CB must play D. Every market agent prefers to cooperate (C) whenever the CB puts sufficiently high weight on cooperation, that is, $x > 1/2$. If $x = 1/2$, market players are indifferent between C and D. If $x < 1/2$, the best response for any market participant is defection (D). Notice that the threshold $x = 1/2$ is independent of the target π^*.[8] Therefore, the set of all Nash equilibria of the stage-game is given by the unique action profile that has all players playing D, that is, $(0, 0)$.

3 Nash Equilibria of the Repeated Game

Recall that a strategy σ_1 for the CB in a repeated game is a function assigning a mixed action $x \in [0, 1]$ to any past history of play. Let σ_2 denote the profile of strategies of all short-run players. A strategy profile $\sigma = (\sigma_1, \sigma_2)$ for the repeated game with a long-run and a continuum of short-run players is a Nash equilibrium if and only if:

- Player's 1 strategy σ_1 maximizes $U_1(\sigma_1, \sigma_2)$ over all repeated game strategies, and
- At any history of play that is reached with positive probability under the profile σ, every market player plays a best response against the prescribed action of

[8] If we think that x is an increasing function of the interest rate, then, the minimum level of interest rates that is necessary to achieve cooperation of all market participants is a higher the lower the target is. In other words, a more aggressive target requires more effort from the CB, which is represented by a higher level of interest rate.

player 1 in σ at that history.[9] Next, we will describe the maxmin payoff of the long-run player.

3.1 Market Agents' Best Response and CB's Maxmin Payoff

Recall that player 2 represents the aggregate behavior of all market agents. The best response of player 2 is given by a correspondence from the set of player 1's mixed actions, denoted ΔA_1, to the set ΔA_2 of player 2's mixed actions. The graph of player 2's best response correspondence is denoted $B \subset \Delta A_1 \times \Delta A_2$.

Define player 1's maxmin payoff against a short-lived player, denoted by $\overline{v_1}$, with:

$$\overline{v_1} = \sup_{\alpha \in B} \min_{a_1 \in \text{supp}(\alpha_1)} u_1(a_1, \alpha_2),$$

where $\text{supp}(\alpha_1)$ denotes the support of the mixed action α_1. Fudenberg *et al* (1990) proves that $\overline{v_1}$ is the maximal utility that the long-run player may obtain in any Nash equilibrium of the repeated game. Next, we will describe the equilibrium in which the policy maker has the highest possible utility.

3.2 Best Equilibrium for the CB in the Repeated Game

The next definition describes the proposed Nash equilibrium of the repeated game. For any fixed target level, this candidate equilibrium is such that the CB obtains its best possible payoff of the repeated game.

Definition 1 *(Best Equilibrium for the CB)*

We propose a candidate Nash equilibrium of the repeated game, that is, a strategy profile where, in every round, the action profile $(1/2, 1)$ *is played. Deviations by market agents are ignored, and any deviation by the CB, even if many market players also deviate, triggers perpetual play of mutual defection, i.e. perpetual play of the*

[9] Recall that the actions of other market players do not affect the payoff of a specific market agent.

action profile $(0,0)$. *Deviations by any player, or by any collection of players, in any stage of the punishment phase, do not change the future prescribed play of the punishment phase.*

Now, we will calculate the payoffs in our proposed equilibrium. First, notice that in the stage game, the largest payoff for player 1 occurs when $y = 1$ because, by assumption, $-\pi_L - c > -\pi_M - c$ and $-\pi_M > -\pi_H$. At $(1/2, 1)$, this payoff becomes:

$$u_1(1/2, 1) = \frac{-[\pi_M + \pi_L + c]}{2} \tag{3}$$

If this value is obtained in every period, the CB's utility also becomes $U_1 = u_1(1/2, 1)$. This is the best that the monetary authority could hope for because, for any fixed target $\pi^* \in [0, 1]$, U_1 is equal to its maxmin payoff, that is $U_1 = u_1(1/2, 1) = \overline{v_1}$. From equation (3) we can see that the value $u_1(1/2, 1)$ is decreasing in π_M, π_L and c.

3.2.1 Incentives

When can our proposed strategy profile be sustained as a Nash equilibrium of the repeated game? Well, each market agent is playing a best response in every period. For player 1, a deviation triggers a Nash reversion to the path of mutual defection in every round. The outcome (D, D) is a Nash equilibrium of the stage-game. Thus, the profile (D, D) in every round is a Nash equilibrium of the repeated game. This takes care of the incentives on the punishment path.

Now, we must verify if the CB prefers the proposed path of perpetual play of the action profile $(1/2, 1)$, which implies obtaining a utility of $U_1 = u_1(1/2, 1)$, rather than obtaining $(1 - \delta)(-\pi_M)$ in the first round followed by $-\pi_H$ in every round thereafter. This is the case if and only if:

$$\frac{-[\pi_M + \pi_L + c]}{2} \geq (1 - \delta)(-\pi_M) - \delta\pi_H$$

The inequality above is equivalent to:[10][11]

$$\delta \geq \delta_0, \qquad \text{where} \qquad \delta_0 = \frac{\pi_L - \pi_M + c}{2(\pi_H - \pi_M)} \tag{4}$$

We conclude that the proposed strategy profile is a Nash equilibrium whenever the CB is sufficiently patient. Putting it in another way, a sufficiently impatient CB will not be able to implement a given target. The next proposition summarizes our results so far.

Proposition 1 *(Best Equilibrium for the Central Bank)*

Assume that the hypotheses (H1), (H4), (H6), (H7) and (H8) hold. For any fixed target, if the CB is sufficiently patient, more precisely, if (4) holds, then, it can always implement its most preferred equilibrium of the repeated game, in which the equilibrium path is the perpetual play of the action profile $(1/2, 1)$*. In other words, the CB does not need to play the tightest monetary policy in every round of the repeated game. Its intensity of cooperation,* x*, should be only large enough to make market participants' best responses be their part in the prescribed equilibrium, that is,* $x = 1/2$*.*

Remark 1 *(Market Participants' Payoffs in Equilibrium)*

By playing x *just above* $1/2$*, the CB forces every market participant* $j \in (0, 1)$ *to choose* C*. Applying (2a) and (2b), we may observe that, at any round,* $u_j(1/2, a_j) = 0$*, for any* $a_j \in \{C, D\}$ *and any target* $\pi^* \in [0, 1]$*. It turns out that market players could obtain higher equilibrium payoffs if they were long-run players with sufficiently high discount factor. According to the folk theorem, a long-run market player, having*

[10] Note that $0 < \delta_0 < 1$ because, by hypothesis (H4), $c < 2\pi_H - \pi_M - \pi_L$ and $c > \pi_M - \pi_L$. In other words, if the cost c were sufficiently high, this candidate equilibrium could not be sustained, not even by an arbitrarily patient CB. If the cost were too low, even a myopic CB could implement this equilibrium.

[11] If we had assumed the most general possible payoffs for market players, the expression for δ_0 would depend on ξ.

the same discount factor that the CB has, could obtain a utility arbitrarily close to $+g$ *if the common discount factor were sufficiently close to* 1. *To see this point, observe that the minmax payoff of player* 1 *is* $-\pi_H$. *Hence, the payoff profile* $(-\pi_H, +g)$ *is in the closure of the feasible and individually rational set of payoff profiles.*[12]

4 Complex Payoffs for Market Agents

In this section, the Nash Equilibrium is analyzed when market players have more complex payoff structure. We start with the simple case where all ξ_j are equal to a positive real number, possibly different from 1. Then, we study the case of target dependent payoffs, and finally we let players have heterogeneous ratio ξ_j of possible payoffs.

4.1 Payoff Ratio $\xi = \xi_j \neq 1$

Suppose that all players payoffs are such that hypotheses (H6), (H7) and (H8) are satisfied, but $\xi = \xi_j$ is a positive real number, possibly different from 1. By (1), we have:

$$u_1(x, 1) = x\left[-\pi_L + \pi_M - c\right] - \pi_M$$

Comparing the possible market agents' payoffs, we find that:

$$u_j(x, C) > u_j(x, D) \quad \Leftrightarrow \quad x > \frac{1}{1+\xi}$$

Therefore:

$$u_1\left(\frac{1}{1+\xi}, 1\right) = \frac{-\left(\pi_L - \pi_M + c\right)}{1+\xi} - \pi_M$$

The CB will be willing to play the Nash equilibrium of the repeated game in which the equilibrium path is such that the CB's intensity of cooperation is $(1+\xi)^{-1}$ in

[12] There are many versions of the folk theorem, for instance see Fudenberg and Maskin (1986). In fact, this kind of payoff profile (in the closure of the feasible and individually rational set of payoff profiles) can be obtained in a subgame perfect equilibrium, as long as the common discount factor is sufficiently close to one.

every round and all market agents always cooperate if and only the CB is sufficiently patient, that is:

$$\delta \geq \frac{\pi_L - \pi_M + c}{(1+\xi)(\pi_H - \pi_M)}$$

4.2 Target Dependent Payoffs

So far, we have been assuming that all market players' payoffs are such that all ξ_j are constant among all market players $j \in (0,1)$, and that $\xi = \xi_j$ is independent of the target π^*. In this section we will keep the first two hypotheses, but we will relax the third one by considering that ξ is a monotonic function of the target.

Comparing the possible market agents' payoffs, we find that:

$$u_j(x, C) > u_j(x, D) \quad \Leftrightarrow \quad x > \frac{1}{1+\xi}$$

Hence, if ξ is an increasing function of the target π^*, then the minimal level of cooperation $(1+\xi)^{-1}$ that the CB needs to make the best response of market players be cooperation is a decreasing function of the target. As the target decreases, ξ increases and so does the utility of the CB in equilibrium, namely $u_1(\frac{1}{1+\xi}, 1)$. The minimal patience that the CB needs to implement its most preferred Nash equilibrium is a function of ξ. In other words, the CB's discount factor δ must satisfy:

$$\delta \geq \delta_0(\pi^*), \qquad \text{where:}$$

$$\delta_0(\pi^*) = \frac{\pi_L(\pi^*) - \pi_M(\pi^*) + c(\pi^*)}{[1+\xi(\pi^*)][\pi_H(\pi^*) - \pi_M(\pi^*)]}$$

Proposition 2 *(Necessary Level of Cooperation of the Central Bank)*

Assume that the hypotheses (H1), (H4), (H6), (H8) hold, and that the ratio $\xi = \xi_j$ *increases with the target. Then:*

(i) The minimal level of cooperation that the CB needs in order to obtain cooperation of all market agents is a decreasing function of the target. To obtain cooperation

of all market players when the target is π^*, the CB has to cooperate with an intensity x such that $x \geq (1+\xi(\pi^*))^{-1}$.

(ii) The minimal discount factor of the CB that sustain the proposed Nash equilibrium of the repeated game becomes:

$$\delta_0(\pi^*) = \frac{\pi_L(\pi^*) - \pi_M(\pi^*) + c(\pi^*)}{[1+\xi(\pi^*)][\pi_H(\pi^*) - \pi_M(\pi^*)]}$$

4.3 Heterogeneous Market Agents

Now, suppose that market agents have individual parameters ξ_j that are continuously distributed in an interval $[\underline{\xi}, \overline{\xi}]$, with $0 \leq \underline{\xi} < \overline{\xi}$.

Comparing the possible market agents' payoffs, we find that player j prefers to cooperate if and only if the CB's intensity of cooperation is sufficiently high. Mathematically:

$$u_j(x, C) > u_j(x, D) \quad \Leftrightarrow \quad x > \frac{1}{1+\xi_j}$$

Hence, agent j will cooperate if and only if $\xi_j > x^{-1} - 1$. Let F represents the distribution of the values ξ_j, and f represent the density function. Hence, a mass of $1 - F(x^{-1} - 1)$ agents cooperate when the CB plays x. Thus:

$$u_1(x, 1 - F(x^{-1}-1)) = x\left[(\pi_H + \pi_L - 2\pi_M)\, F(x^{-1}-1) + \pi_M - \pi_L - c\right] + \\ - (\pi_H - \pi_M) F(x^{-1}-1) - \pi_M$$

By increasing x the CB decreases its payoff from each market player keeping their original action, but more players will cooperate now, and this increases the CB's payoff.

$$\frac{d}{dx}\left[u_1(x, 1 - F(x^{-1}-1))\right] = (\pi_M - \pi_L - c) + (\pi_H + \pi_L - 2\pi_M) F(x^{-1}-1) + \\ + \frac{f(x^{-1}-1)}{x^2}\left[\pi_H - \pi_M + x(2\pi_M - \pi_H - \pi_L)\right]$$

The sum $(\pi_M - \pi_L - c) + (\pi_H + \pi_L - 2\pi_M)F(x^{-1} - 1)$ is always negative and it corresponds to the direct effect of a marginal increase in x, that is, this value is the decrease in CB's utility because player 1 cooperates marginally more now assuming that no market agent switches its action. For each one of the $1-F(x^{-1}-1)$ market players that are cooperating, the marginal increment in CB's utility is equal to $(-\pi_L - c) - (-\pi_M) = \pi_M - \pi_L - c$. For each one of the $F(x^{-1} - 1)$ market players that are not cooperating, the marginal increment in CB's utility is equal to $(-\pi_M - c) - (-\pi_H) = \pi_H - \pi_M - c$. Hence, the direct effect is equal to:

$$\left[1 - F(x^{-1} - 1)\right] [\pi_M - \pi_L - c] + F(x^{-1} - 1) [\pi_H - \pi_M - c] =$$

$$= (\pi_M - \pi_L - c) + (\pi_H + \pi_L - 2\pi_M) F(x^{-1} - 1)$$

On the other hand, the term $f(x^{-1} - 1) [\pi_H - \pi_M + x(2\pi_M - \pi_H - \pi_L)] / x^2$ is always positive and represents the indirect effect, that is, the CB's increase in utility due to the existence of a larger share of market players cooperating now. $f(x^{-1} - 1)/x^2$ is the marginal quantity of market players that switch action because the CB is now cooperating marginally more. The factor $\pi_H - \pi_M + x(2\pi_M - \pi_H - \pi_L)$ is the increase in CB's utility when facing a single market agent that has changed action from D to C. If a market player plays C, the CB playing x obtains $(-\pi_L - c)x + (-\pi_M)(1-x)$ from its interaction with this player. However, if the market agent plays D, the CB playing x obtains only $(-\pi_M - c)x + (-\pi_H)(1 - x)$ from its interaction with this player. Therefore:

$$[-(\pi_L + c)x - \pi_M(1-x)] - [-(\pi_M + c)x - \pi_H(1-x)] = \pi_H - \pi_M + x(2\pi_M - \pi_H - \pi_L)$$

Proposition 3 *(Heterogeneous Market Agents)*

Under hypotheses (H1), (H4) and (H7), the intensities of cooperation where the CB has an incentive to increase cooperation is given by:

$$\frac{d}{dx}\left[u_1(x, 1 - F(x^{-1} - 1))\right] > 0 \quad \Leftrightarrow$$

$$\frac{[\pi_H - \pi_M + x(2\pi_M - \pi_H - \pi_L)]f(x^{-1}-1)}{x^2} > (\pi_L - \pi_M + c) + (2\pi_M - \pi_H - \pi_L)F(x^{-1}-1)$$

5 Target Selection by the CB

The goal of this section is to describe the target choice. To achieve this objective, we need to understand how the CB's utility depends on the inflation target. It turns out that the in our proposed Nash equilibrium, the CB's utility is a single peaked function of the target.

To see this point suppose first that the CB can choose any target $\pi^* \in [0,1]$. The CB maximizes its utility by choosing a target that solves the problem:

$$\underset{\pi^* \in [0,1]}{Max} \left\{ \frac{-[\pi_M + \pi_L + c]}{2} \right\} \tag{5}$$

The first order condition of the CB's maximization problem is:[13]

$$\frac{\partial \pi_M}{\partial \pi^*} + \frac{\partial \pi_L}{\partial \pi^*} = \frac{-\partial c}{\partial \pi^*} \tag{6}$$

The next lemma says that (5) has a unique interior solution, denoted $\hat{\pi}_1^*$.

Lemma 1 *(Uniqueness of the Optimal Target for the Central Bank)*

Assume that the hypotheses (H1), (H2), (H3), (H4) and (H5) hold. Then, there is a unique interior solution, denoted $\hat{\pi}_1^*$, to the CB's target choice problem (5). The objective function of problem (5) is single peaked. At the interval $[0, \hat{\pi}_1^*)$ it is increasing, and, at the interval $(\hat{\pi}_1^*, 1]$ it is decreasing.

5.1 Targets that May be Implemented

We define the set of all targets such that the proposed Nash equilibrium of the repeated game can be sustained. The next definition introduces this concept.

[13] Because the target is inside a compact set, $\pi^* \in [0,1]$, the maximization problem (5) always has a solution. Uniqueness of the solution depends on the behavior of the functions $\frac{\partial \pi_M}{\partial \pi^*}$, $\frac{\partial \pi_L}{\partial \pi^*}$, and $\frac{\partial c}{\partial \pi^*}$.

Definition 2 *(Implementable Targets)*

Let $\{\pi^* \in [0,1] \mid \delta \geq \delta_0(\pi^*)\}$ *be called the set of implementable targets.*

From now on, we assume that the set of implementable target is non-empty because, otherwise, the only Nash equilibrium of the repeated game is the perpetual play of the stage-game Nash equilibrium, and the analysis becomes trivial. Another helpful concept is the smallest implementable target, as we define next.

Definition 3 *(Smallest Implementable Target)*

We define the smallest implementable targets as:

$$\bar{\pi}^* = \inf \{\pi^* \in [0,1] \mid \delta \geq \delta_0(\pi^*)\}$$

5.2 Monotonicity of the Minimal Discount Factor

Consider the derivative of the minimum discount factor function with respect to the target:

$$\frac{\partial \delta_0}{\partial \pi^*} = \frac{(\pi_H - \pi_M)\left(\frac{\partial \pi_L}{\partial \pi^*} - \frac{\partial \pi_M}{\partial \pi^*} + \frac{\partial c}{\partial \pi^*}\right) - (\pi_L - \pi_M + c)\left(\frac{\partial \pi_H}{\partial \pi^*} - \frac{\partial \pi_M}{\partial \pi^*}\right)}{2(\pi_H - \pi_M)^2}$$

A direct calculation reveals that the inequality $\frac{\partial \delta_0}{\partial \pi^*} < 0$ is equivalent to:

$$\frac{-\partial c}{\partial \pi^*} > \left(\frac{\partial \pi_L}{\partial \pi^*} - \frac{\partial \pi_M}{\partial \pi^*}\right) + \frac{(\pi_L - \pi_M + c)\left(\frac{\partial \pi_M}{\partial \pi^*} - \frac{\partial \pi_H}{\partial \pi^*}\right)}{\pi_H - \pi_M} \tag{7}$$

It turns out that, at any target π^* that is smaller or equal to the CB's optimal target $\hat{\pi}_1^*$, condition (7) is automatically satisfied, and then, $\frac{\partial \delta_0}{\partial \pi^*} < 0$. This means that the minimum discount factor of the CB that sustains the proposed Nash equilibrium, namely δ_0, is a decreasing function of the optimal target on the domain $\pi^* \in [0, \hat{\pi}_1^*]$. Formally, we state the next lemma.

Lemma 2 *(Lower Targets Require a More Patient Central Bank)*

Assume that hypotheses (H1), (H2), (H3), (H4) and (H5) are satisfied. Then, in the domain $\pi^* \in [0, \hat{\pi}_1^*]$*, the lower the target is, the more patient the Central Bank needs to be in order to sustain the described equilibrium. Mathematically, the function* $\pi^* \longmapsto \delta_0(\pi^*)$ *is decreasing on* $[0, \hat{\pi}_1^*]$.[14]

5.3 Is the CB's Most Preferred Target Implementable?

If $\bar{\pi}^* \leq \hat{\pi}_1^*$, the CB can implement its optimal target $\hat{\pi}_1^*$ because, by the definition of $\bar{\pi}^*$ and lemma 2, $\delta \geq \delta_0(\bar{\pi}_1^*) \geq \delta_0(\hat{\pi}_1^*)$. Conversely, if the CB can implement its optimal target $\hat{\pi}_1^*$, then $\bar{\pi}^* \leq \hat{\pi}_1^*$. The proof of this fact is that if $\bar{\pi}^* > \hat{\pi}_1^*$, then $\hat{\pi}_1^*$ would not be implementable, violating the hypothesis. The next lemma describes the situation.[15]

Lemma 3 *(CB's Most Preferred Target Implementation)*

The Central Bank can implement its most preferred target $\hat{\pi}_1^*$ *if and only if this value is at least as large as the smallest implementable target* $\bar{\pi}^*$*. Mathematically:*

$$\delta \geq \delta_0(\hat{\pi}_1^*) \quad \Leftrightarrow \quad \bar{\pi}^* \leq \hat{\pi}_1^*$$

5.4 Best Implementable Target

Whenever the CB's most preferred target $\hat{\pi}_1^*$ is not implementable, $\hat{\pi}_1^* < \bar{\pi}^*$. Because the CB's utility is single peaked, then, for the CB, the best implementable target (BIT), denoted π_{BIT}^*, becomes $\bar{\pi}^*$. The next lemma formalizes this point.

Lemma 4 *(Best Implementable Target)*

[14] If $\bar{\pi}^* \leq \hat{\pi}_1^*$, the function δ_0 is weakly decreasing at the point $\bar{\pi}^*$. However, if $\hat{\pi}_1^* < \bar{\pi}^*$, then, in the interval $(\hat{\pi}_1^*, \bar{\pi}^*)$, nothing prevents the function $\frac{\partial \delta_0}{\partial \pi^*}$ from changing its sign multiple times.

[15] Another way of stating this lemma is as follows: the CB can not implement its most preferred target $\hat{\pi}_1^*$ if and only if this value is smaller than the smallest implementable target $\bar{\pi}^*$. Mathematically, $\delta < \delta_0(\hat{\pi}_1^*) \Leftrightarrow \hat{\pi}_1^* < \bar{\pi}^*$.

The best implementable target is given by $\pi^*_{BIT} = \max\{\hat{\pi}^*_1, \bar{\pi}^*\}$*, or equivalently:*

$$\pi^*_{BIT} = \begin{cases} \hat{\pi}^*_1, & \textit{if the CB's most preferred target is implementable, i.e.} \quad \bar{\pi}^* \leq \hat{\pi}^*_1 \\ \bar{\pi}^*, & \textit{if the CB's most preferred target is } \textbf{not} \textit{ implementable, i.e. } \bar{\pi}^* > \hat{\pi}^*_1 \end{cases}$$

5.5 Zero or the Best Implementable Target?

Observe that the CB may prefer the target $\pi^* = 0$ with equilibrium path where all players defect all the time, i.e. $(0, 0)$, rather than our proposed Nash equilibrium, in which the equilibrium path consists of repeated play of $(1/2, 1)$, together with the CB's best implementable target π^*_{BIT}. This is true even when the CB's most preferred target is implementable, that is, $\pi^*_{BIT} = \hat{\pi}^*_1$. Observe that any target choice different from zero generates a lower utility for the CB when the outcome path is the perpetual play of the action profile $(0, 0)$ because π_H is a increasing function.

In order to decide which is its optimal target choice, π^*_{BIT} or $\pi^* = 0$, the CB compares the values $-[\pi_M(\pi^*_{BIT}) + \pi_L(\pi^*_{BIT}) + c(\pi^*_{BIT})]/2$ with $-\pi_H(0)$. The next proposition summarizes our conclusions.

Proposition 4 *(Optimal Target for the Central Bank)*

Assume that hypotheses (H1), (H2), (H3), (H4) and (H5) are satisfied and that the set of implementable targets is non-empty. If

$$\frac{[\pi_M(\pi^*_{BIT}) + \pi_L(\pi^*_{BIT}) + c(\pi^*_{BIT})]}{2} \geq \pi_H(0),$$

then the CB selects the target $\pi^* = 0$*, the Nash equilibrium of the repeated game has an equilibrium path consisting of perpetual play of the stage-game Nash equilibrium action profile* $(0, 0)$*, and consequently:*

$$U_1 = u_1(0, 0) = -\pi_H(0)$$

Otherwise, the BIT is selected, our proposed Nash equilibrium of the repeated game is played, and then:

$$U_1 = u_1(1/2, 1) = \frac{-[\pi_M(\pi^*_{BIT}) + \pi_L(\pi^*_{BIT}) + c(\pi^*_{BIT})]}{2}$$

Remark 2 *Now, suppose that:*

$$\frac{-[\pi_M(\pi^*_{BIT}) + \pi_L(\pi^*_{BIT}) + c(\pi^*_{BIT})]}{2} > -\pi_H(0) \tag{H9}$$

Observe that (H8) is equivalent to $[\pi_M(\pi^*_{BIT}) + \pi_L(\pi^*_{BIT}) + c(\pi^*_{BIT})]/2 < \pi_H(0)$*. By hypothesis (H4), we know that* $[\pi_L(0) + \pi_M(0) + c(0)]/2 < \pi_H(0)$*, so a sufficient condition to obtaining (H8) is that:*

$$\pi_L(\pi^*_{BIT}) + \pi_M(\pi^*_{BIT}) + c(\pi^*_{BIT}) \leq \pi_L(0) + \pi_M(0) + c(0)$$

By the fundamental theorem of calculus, this inequality is equivalent to:

$$\int_{\hat{\pi}^*_1}^{\pi^*_{BIT}} \left[\frac{\partial \pi_L}{\partial \pi^*} + \frac{\partial \pi_M}{\partial \pi^*} + \frac{\partial c}{\partial \pi^*}\right] d\pi^* \leq -\int_0^{\hat{\pi}^*_1} \left[\frac{\partial \pi_L}{\partial \pi^*} + \frac{\partial \pi_M}{\partial \pi^*} + \frac{\partial c}{\partial \pi^*}\right] d\pi^*$$

Applying lemma 1*, we learn that the right-hand side is always positive. If* $\bar{\pi}^* \leq \hat{\pi}^*_1$*, then* $\pi^*_{BIT} = \hat{\pi}^*_1$*, making the left-hand side zero. Therefore, the inequality above is satisfied.*

If $\bar{\pi}^* > \hat{\pi}^*_1$*, then* $\pi^*_{BIT} = \bar{\pi}^* > \hat{\pi}^*_1$*, and the left-hand side is also positive. Hence, the left-hand side tends to be relatively smaller when* $\bar{\pi}^*$ *and* $\hat{\pi}^*_1$ *are relatively closer, while the right-hand side tends to grow with* $\hat{\pi}^*_1$*.*

6 Inflation

The inflation in a given round t, denoted by π^t, is determined by the outcome (x, y) of the stage-game at period t. Formally:

$$\pi^t = xy\pi_L + x(1-y)\pi_M + (1-x)y\pi_M + (1-x)(1-y)\pi_H$$

The inflation accumulated over the repeated game, denoted by Π, is the discounted average of the inflations in each round. Mathematically:[16]

$$\Pi = (1-\delta)\sum_{t=0}^{+\infty}\delta^t\pi^t$$

6.1 Inflation Bias

In the proposed Nash equilibrium of the repeated game, the inflation at every period t, and consequently the accumulated inflation Π, are both given by $(\pi_M + \pi_L)/2$. Note, however, that if there is an implementable target at all, as we will assume from now on, the CB always chooses either $\bar{\pi}^*$ or $\hat{\pi}_1^*$.

Suppose that the socially optimal target is the one that minimizes inflation.[17] Hence, there is an inflation bias in the *choice of the target*. The first best outcome would be a target of zero, together with perpetual play of the action profile $(1, 1)$. This would result in an inflation $\Pi = \pi_L(0)$. Unfortunately, this level of inflation cannot be implemented due to the dynamics and the incentives that players face in the repeated game. The second best involves choosing a target in order to minimize inflation subject to the actual implementation, which comes from the strategic interaction between the CB and market players.

The classic solution for the target bias problem is to have another member of the government, from now on called player 3, choosing the target, taking into account the CB's and market agents' incentives. In other words, this other branch of the government must select its most preferred implementable target. Formally, player 3

[16] In fact, since all Nash equilibria that we describe have only one action profile played on every round in their respective equilibrium paths, then, any sort of average of inflation levels in each period would lead to the same accumulated inflation.

[17] In some models of the literature, the output gap is also considered. However, microfundamented models show that the weight of the output gap term in the respective loss function is so small that we may ignore it without any serious consequence. For instance, see table 5 of Giannone and Woodford (2003), pp. 52.

solves the problem:

$$\underset{\{\pi^* \in [0,1] |\, \delta \geq \delta_0(\pi^*)\}}{Min} \left\{\pi_M(\pi^*) + \pi_L(\pi^*)\right\}, \tag{8}$$

where:

$$\delta_0(\pi^*) = \frac{\pi_L(\pi^*) - \pi_M(\pi^*) + c(\pi^*)}{2\pi_H(\pi^*) - 2\pi_M(\pi^*)}$$

The solution of problem (8) is always $\bar{\pi}^*$ even when $\hat{\pi}_1^*$ is implementable. When the socially optimal implementable target is smaller than the CB's most preferred one, an institutional arrangement that allows player 3 to select the target improves social welfare by decreasing inflation.

Observe that if $\hat{\pi}_1^*$ is not implementable, equivalently, if $\hat{\pi}_1^* < \bar{\pi}^*$, then the preferred implementable targets of players 1 and 3 will coincide. In other words, the government's agency will choose the most preferred implementable target of the CB.

Implicitly in the problem above is the hypothesis that the branch of the government picking the target has the correct incentives to do it without any bias. But what if this other agency has its own bias? Who should select the inflation target?

6.2 Who Should Choose the Target?

If there is a government agency other than the CB that is technically competent and understands well the incentives of the CB, allowing them to select the target may or may not improve the social welfare, depending on its own bias when selecting the target. Hence, the ideal institutional framework will depend on the parameters of the model.

Player 3 always chooses its most preferred implementable target, denoted $\hat{\pi}_3^*$. Interestingly, the lack of patience of the CB helps establishing a balance of power between players 1 and 3. As the CB becomes more impatient, that is, as its discount factor decreases, the set of implementable targets shrinks, making the choice of player 3 more restrict. However, this may or may not benefit the CB. For instance, if players

1 and 3 have the same most preferred target, $\hat{\pi}_1^* = \hat{\pi}_3^*$, then a more restricted set of implementable targets can only hurt the CB. On the other hand, consider the next examples.

Example 1 *Suppose that player* 3*'s preference is represented by a single peaked utility function of the target, with* $\hat{\pi}_3^* < \bar{\pi}^* < \hat{\pi}_1^*$*, that is,* 3*'s preferred target is smaller than the smallest implementable target, which, in turn, is smaller than the CB's favorite one. Because* $\delta_0(\pi^*)$ *is decreasing in the interval* $[0, \hat{\pi}_1^*]$*, the CB would like that player* 3 *believed that its discount factor* δ *was smaller than it actually is.*[18] *This would make player* 3 *think that the value of the smallest implementable target* $\bar{\pi}^*$ *is larger than it really is. Since, given its preference, player* 3 *always chooses what he believes to be* $\bar{\pi}^*$*, the CB's utility would increase (because it is also single peaked, with maximal value at* $\hat{\pi}_1^*$*). A slightly more impatient CB would have a larger utility.*

Example 2 *Another illuminating example occurs when* $\bar{\pi}^* < \hat{\pi}_1^* < \hat{\pi}_3^*$*, player* 3*'s preference is represented by a single peaked utility function of the target, the function* $\delta_0(\pi^*)$ *is decreasing in the interval* $[0, \hat{\pi}_1^* + \varepsilon)$*, for a sufficiently small* ε*, and it increases along the interval* $(\hat{\pi}_1^* + \varepsilon, 1]$*. Suppose that the set of implementable targets is given by* $[\bar{\pi}^*, \hat{\pi}_3^*]$*. Then, the CB would like that player* 3 *believed that its discount factor* δ *was smaller than it actually is. In this case,* 3 *would think that the upper border of the set of implementable targets was smaller than it really is. Thus, 3 would choose a smaller target than what it actually selects when* 3 *knows the truth. Again, by pretending to be more impatient, the CB can improve its utility.*

7 Appendix

We prove the remaining results in this appendix.

[18] The CB would like the most the case in which player 3 believed that $\delta = \delta_0(\hat{\pi}_1^*)$.

Proof. Proposition 1 (Best Equilibrium for the Central Bank)

The proof is in the main text. ■

Proof. Proposition 2 (Necessary Level of Cooperation of the Central Bank)

Suppose that $\xi = \xi(\pi^*)$ is an increasing function of the target. Comparing the possible market agents' payoffs:

$$u_j(x, C) > u_j(x, D) \quad \Leftrightarrow \quad x > \frac{1}{1+\xi}$$

Hence, the minimal level of cooperation, namely $(1+\xi(\pi^*))^{-1}$, that the CB needs to make the best response of all market players be cooperation is a decreasing function of the target π^*. This concludes this proof. ■

Proof. Proposition 3 (Heterogeneous Market Agents)

A direct calculation reveals this result. This concludes this proof. ■

Proof. Lemma 1 (Uniqueness of the Optimal Target for the Central Bank)

Consider the following continuous functions of the inflation target $\pi^* \in [0, 1]$: $h = -\frac{1}{2}(\pi_M + \pi_L + c)$ and $g = \frac{\partial h}{\partial \pi^*}$. By (H3), $\frac{\partial g}{\partial \pi^*}$ is negative on the interval $[0, 1]$. Thus, g is always decreasing. By the first part of hypothesis (H5), $g(0)$ is positive. By the second part of (H5), $g(1)$ is negative. Therefore, there is a unique value $\hat{\pi}_1^* \in (0, 1)$ such that $g(\hat{\pi}_1^*) = 0$. If $\pi^* < \hat{\pi}_1^*$, $g(\pi^*)$ is positive, and then, h is increasing. If $\pi^* > g(\pi^*)$ is negative and h is decreasing. This proves this result. ■

Proof. Lemma 2 ((Lower Targets Require a More Patient Central Bank)

To prove that $\frac{\partial \delta_0}{\partial \pi^*} < 0$ in $[0, \hat{\pi}_1^*]$ is equivalent to show that inequality (7) holds. By lemma 1 we know that $\frac{-\partial c}{\partial \pi^*} \geq \frac{\partial \pi_M}{\partial \pi^*} + \frac{\partial \pi_L}{\partial \pi^*}$ on $[0, \hat{\pi}_1^*]$. Hence, to prove (7) it is enough to show that for any π^* on $[0, \hat{\pi}_1^*]$:

$$\frac{\partial \pi_M}{\partial \pi^*} + \frac{\partial \pi_L}{\partial \pi^*} > \left(\frac{\partial \pi_L}{\partial \pi^*} - \frac{\partial \pi_M}{\partial \pi^*}\right) + \frac{(\pi_L - \pi_M + c)\left(\frac{\partial \pi_M}{\partial \pi^*} - \frac{\partial \pi_H}{\partial \pi^*}\right)}{\pi_H - \pi_M}$$

After some algebra, we find that this inequality is equivalent to:

$$\frac{\partial \pi_M}{\partial \pi^*}(2\pi_H - \pi_M - \pi_L - c) > -\frac{\partial \pi_H}{\partial \pi^*}(\pi_L - \pi_M + c)$$

But this is clearly true because the left-hand side is positive by (H2) and (H4), while the right-hand side is negative also because of (H2) and (H4). ■

Proof. Lemma 3 (CB's Most Preferred Target Implementation)

The proof is in the main text. ■

Proof. Lemma 4 (Best Implementable Target)

The proof is in the main text. ■

Proof. Proposition 4 (Optimal Target for the Central Bank)

The proof is in the main text. ■

References

[1] Barro, Robert J. "Reputation in a Model of Monetary Policy with Incomplete Information," *Journal of Monetary Economics* , Vol. 17 (1986), pp. 3-20.

[2] Barro, Robert J. and Gordon, David B., "A Positive Theory of Monetary Policy in a Natural Rate Model," *The Journal of Political Economy*, Vol. 91, No. 4. (Aug., 1983), pp. 589-610.

[3] Barro, Robert J. and Gordon, David B., "Rules, Discretion and Reputation in a Model of Monetary Policy," *Journal of Monetary Economics*, Vol. 12 (1986), pp. 101-121.

[4] Batini, N., Kuttner, K., Laxton, D., "Does Inflation Targeting Work in Emerging Market?" World Economic Outlook, chapter 4 (September, 2005), pp. 161-186.

[5] Canzoneri, Matthew B., "Monetary Policy Games and the Role of Private Information," *The American Economic Review* , Vol. 75, No.5 (December, 1985), pp. 1056-1070.

[6] Carvalho, Fabia A. and Bugarin, M., "Heterogeneity of Central Bankers and Inflationary Pressure," working paper, 2006.

[7] Clarida, Richard; Gali, Jordi and Gertler, Mark, "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, Vol. 37, No. 4. (Dec., 1999), pp. 1661-1707.

[8] Giannone, Marc and Woodford, M., "Optimal Inflation Targeting Rules," working paper, 2003.

[9] Fraga, A., Goldfajn, I., and Minella, A., "Inflation Targeting in Emerging Market Economies," NBER Working Paper No. 10019 (October, 2003).

[10] Fudenberg, D.; Kreps, D. and Maskin, E., "Repeated Games with Long-run and Short-run Players," *Review of Economic Studies*, Vol. 57 (1990), pp. 555-573.

[11] Fudenberg, Drew and Maskin, Eric "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, Vol. 54, No. 3 (May, 1986), pp. 533-554.

[12] Kydland, Finn E. and Prescott, Edward C., "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *The Journal of Political Economy*, Vol. 85, No. 3 (June, 1977), pp. 473-492.

[13] Mailath, George J. and Samuelson, Larry, "Repeated Games and Reputations: Long-Run Relationships," Oxford University Press, 2006.

[14] Mishkin, Frederic S., "Can Inflation Targeting Work in Emerging Market Countries?" NBER Working Paper No. 10646 (July, 2004).

[15] Mishkin, Frederic S. and Schmidt-Hebbel, Klaus, "One Decade of Inflation Targeting in the World: What Do We Know and What Do We Need to Know?" NBER Working Paper No. 8397 (July, 2001).

[16] Persson, Torsten and Tabellini, Guido, editors "Monetary and Fiscal Policy, Volume 1: Credibility", Second Printing, The MIT Press, 1995.

[17] Svensson, L. "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets," Institute for International Economics Studies, Stockholm University; CEPR and NBER (1996).

[18] Svensson, L., "Optimal Inflation Targeting, 'Conservative' Central Banks, and Linear Inflation Contracts," *The American Economic Review*, Vol. 87 No. 1 (Mar., 1997), pp. 98-114.

[19] Taylor, John B., "Using Monetary Policy Rules in Emerging Market Economies," Stabilization and Monetary Policy: The International Experience. Bank Mexico (2000), pp. 441-457.

[20] Walsh, Carl E. "Optimal Contracts for Central Bankers," *The American Economic Review*, Vol. 85, No.1 (March, 1995), pp. 150-167.

[21] Woodford, M., "Optimal Monetary Policy Inertia," working paper, 1999.

[22] Woodford, M., "Interest and Prices," Princeton University Press, Princeton, 2003.