Inflation Targeting under Political Pressure*

Marina Halac† Pierre Yared‡

June 28, 2019

Abstract

This paper studies optimal monetary policy when the central bank lacks commitment to policies and is subject to privately observed, time-varying political pressure. We characterize the monetary rule that maximizes social welfare subject to the central bank's self-enforcement and private information constraints. Inflation responds to political shocks over time, with extreme shocks potentially triggering a temporary transition from a hawkish low-inflation regime to a dovish high-inflation regime. The threat of transitioning to the dovish regime sustains the hawkish regime, and the promise of returning to the hawkish regime sustains the dovish regime. We examine how inflation dynamics depend on the distribution of political shocks.

Keywords: Rules vs. discretion, time inconsistency, optimal monetary policy, inflation targets

JEL Classification: C73, D82, E52, E58, E6

*This paper was prepared for the 30th anniversary of the Central Bank of Chile conference. We would like to thank Hassan Afrouzi, Rick Mishkin, Jesse Schreger, and Mike Woodford for helpful comments. Miguel Acosta provided excellent research assistance.

†Yale University, Department of Economics and CEPR. Email: marina.halac@yale.edu.

‡Columbia University, Graduate School of Business and NBER. Email: pyared@columbia.edu.

1 Introduction

Historically, many emerging economies, particularly in Latin America, battled persistently high and volatile inflation.[1] Today, emerging economies continue to experience higher inflation than developed ones, and their central banks more frequently deviate from inflation targets.[2] These patterns partly reflect the added political pressure and lower degree of independence faced by central banks in emerging markets. For example, Aisen and Veiga (2006, 2008) find that inflation is higher and more volatile in countries with a lower quality of political institutions and a higher degree of political instability. Using a narrative approach, Binder (2018) finds that on average, 10 percent of central banks face political pressure, and this pressure is associated with higher inflation and inflation persistence.

Motivated by this evidence, this paper studies optimal monetary policy when the central bank lacks commitment to policies and is subject to time-varying political pressure. We characterize the welfare-maximizing policy that can be self-enforced conditional on the degree of central bank independence and the level of political instability. Our analysis elucidates how these political factors affect both average inflation and inflation dynamics.

We cast our model in a Barro and Gordon (1983a,b) framework, in which the central bank's policy determines output and inflation at every date. Following Mishkin and Westelius (2008), we model political pressure as the weight the central bank places on output expansion versus inflation stabilization. A higher weight reflects the increased importance of stimulating output in order to boost an incumbent political party's popularity or to accommodate a fiscal expansion, for example. We take these political shocks to follow an independent and identically distributed (i.i.d.) process. In each period the central bank observes the realized shock prior to its choice of policy. This observation is private, as political pressure cannot be perfectly assessed by an external entity.[3]

[1] Capistrán and Ramos-Francia (2009) provide evidence of these patterns in Latin America.

[2] For a discussion, see Fraga, Goldfajn, and Minella (2003).

[3] In order to focus on the impact of political shocks, we abstract from economic shocks. Under some conditions, observable economic shocks can be introduced without affecting our analysis. Details available upon request.

Due to its lack of commitment, the central bank is inflation-biased when choosing policy. Specifically, the central bank does not internalize the impact of its actions on past inflation expectations, and it thus overweighs the benefit of stimulating output. This bias is increasing in the political shock, which accentuates the focus on output expansion over inflation stabilization. Moreover, since the central bank has full policy discretion, its policy choice must be self-enforcing and can only be disciplined by the policy choices of future central banks, via their effect on the current central bank's continuation value. Such future policies can respond to past policies; however, they cannot depend directly on past political shocks which are privately observed. Hence, monetary policy in our setting can be represented as a rule that assigns the central bank a policy choice and a continuation value for each shock at every date, where this assignment must satisfy the central bank's private information and self-enforcement constraints. Such a rule is optimal if it maximizes social welfare.

To describe the forces underlying our model, suppose first that monetary policy could be perfectly enforced by an external entity. Then political shocks—whether they are publicly observable or not—play no role, and inflation is optimally set at a constant low level. Suppose next that external enforcement is not possible but political shocks are public. Then any deviation by the current central bank (where it does not choose its assigned inflation level) is observable, and it can thus be punished (off path) with the worst continuation value sustained by future central banks' equilibrium behavior. If this punishment is harsh enough, enforcement constraints are non-binding. Otherwise, an optimal rule assigns the lowest level of inflation that is enforceable conditional on the realized shock. Relative to the case of perfect enforcement, inflation under this rule is higher and more volatile, as it responds to political shocks that tighten the enforcement constraint, albeit only temporarily.

An optimal monetary rule in our setting must deal not only with the problem of enforcement, but also with the realistic constraint of private information. Because only the current central bank observes the realized political shock, the rule just described that conditions directly on the shock is not incentive compatible: the central bank can deviate privately from its assigned policy and choose a higher inflation level, making itself strictly better off without being penalized with a lower continuation value. Incentive compatibility requires that, for each political shock, the central bank prefer its assigned inflation level and continuation value to those prescribed for any other shock.

Our main result shows that the optimal monetary rule is characterized by a hawkish low-inflation regime and a dovish high-inflation regime. The threat of transitioning to the dovish regime sustains the hawkish regime, and the promise of returning to the hawkish regime sustains the dovish regime. Moreover, unlike under observable political shocks, a temporary transition from the hawkish regime to the dovish regime may now occur on path, following high enough shocks.

Monetary policy in each regime admits a simple implementation. We show that the hawkish regime takes the form of a *maximally enforced inflation cap*. If the central bank respects the cap, future inflation expectations remain low and the equilibrium restarts in the hawkish regime in the next period. If instead the central bank violates the cap, inflation expectations rise and the equilibrium transitions to the dovish regime. The central bank may not be constrained by the inflation cap when experiencing low political pressure, but it will be constrained under high pressure, and in some cases it will break the cap. Additionally, we show that the dovish regime takes the opposite form of a *maximally enforced inflation floor*. If the central bank respects the floor (by choosing high enough inflation), future inflation expectations decline and the equilibrium returns to the hawkish regime in the next period. If instead the central bank violates the floor, inflation expectations remain high and the equilibrium restarts in the dovish regime. The central bank may not be constrained by the inflation floor when experiencing high political pressure, but it will be constrained under low pressure, and in some cases it will break the floor.

A key feature of our environment is that, ex ante, the central bank shares the same preferences as society for low average inflation. The central bank realizes that private sector expectations are rational, and that future realized inflation will be incorporated into inflation expectations, limiting the benefit of inflation surprises. It is only after private sector expectations are set and the political shock is realized that the central bank sees an added benefit of inflation. Thus, a maximally enforced inflation cap maximizes social welfare by counteracting the political pressure to inflate: the central bank is rewarded for choosing low inflation with a hawkish continuation regime, and it is punished for choosing high inflation with a dovish continuation regime. Analogously, a maximally enforced inflation floor—which serves as a punishment—minimizes social welfare by inducing the central bank to bend to the political pressure. Punishment is always temporary, since a central bank's succumbing to political pressure is rewarded with a transition back to the hawkish regime.

We complete our characterization of inflation dynamics by examining the conditions under which the inflation cap is occasionally broken in the hawkish regime. We find that an optimal rule prescribes on-path violations following high enough shocks only if these shocks are sufficiently unlikely. Intuitively, in this case, the benefit of lowering average inflation by specifying a tight inflation cap exceeds the cost of occasional punishment following extreme (and rare) political shocks.

Our analysis sheds light on the empirical differences in average inflation and inflation volatility in emerging versus developed economies. In our framework, inflation is high and volatile, and temporary political shocks not only impact current inflation but may also persist into the future by changing future inflation expectations. Our results suggest that these patterns, which resemble those in the data, may correspond to the best policy that can be self-enforced when the central bank is subject to time-varying political pressure.

Related literature. Our paper fits into the literature on central bank credibility and reputation, and in particular relates to prior work that examines the role of private information in such a context.[4] We follow Athey, Atkeson, and Kehoe (2005) by taking a mechanism design approach to characterize optimal policy subject to private information constraints.[5] We depart from the literature by combining private information with lack of enforcement, where we show

[4]For example, see Barro and Gordon (1983a,b), Backus and Driffill (1985), Canzoneri (1985), Rogoff (1985), Cukierman and Meltzer (1986), Walsh (1995), and Kocherlakota (2016), among others.

[5]In contrast to Athey, Atkeson, and Kehoe (2005), we study political shocks that are payoff irrelevant for society. Our analysis can be extended to consider shocks to the social cost of inflation, as in their work, without impacting our main results. Details available upon request.

that the latter may lead to transitions between a hawkish and a dovish inflation regime.[6]

Our paper also relates more broadly to the mechanism design literature that studies delegation.[7] Most importantly, our analysis builds on Halac and Yared (2019), which examines optimal fiscal rules under private information and limited enforcement. A main difference is that our current focus is monetary policy, which requires us to incorporate the role of expectations, absent in the context of fiscal policy. Despite this difference, we find that the mathematical arguments developed in Halac and Yared (2019) apply, and thus our results follow from applying the general results in that paper to the present monetary policy application.

Finally, our paper sheds light on the continuing debate about the causes of the rise and fall of inflation in the US and in Latin America in the post-war period (e.g., Sargent, 2001; Sargent, Williams, and Zha, 2009). We find that these persistent regime transitions may reflect the central bank's least socially costly means of responding to temporary political pressure to expand the output gap.[8]

2 Model

Consider an infinite horizon setting with periods $t = \{0, 1, \ldots\}$. At the beginning of each period, an i.i.d. political shock $\theta_t > 0$ is drawn from a bounded set $\Theta \equiv [\underline{\theta}, \overline{\theta}]$, with a continuously differentiable probability density function $f(\theta_t) > 0$ and associated cumulative density function $F(\theta_t)$. The realization of θ_t is

[6] These equilibrium dynamics bear a relationship to the seminal work of Abreu, Pearce, and Stacchetti (1990), who establish the optimality of bang-bang continuation values in a class of repeated games. Their analysis however is constrained to settings with finite actions and a continuous public signal, and thus it does not apply to our environment in which the action is continuous. See Halac and Yared (2019) for a discussion.

[7] The study of delegation in principal-agent settings dates back to Holmström (1977). For recent work, see Amador and Bagwell (2013) and the references cited therein. Yared (2019) discusses fiscal policy applications of delegation theory.

[8] Regime transitions in our setting can also be interpreted as arising from temporary shocks to the central bank's belief about the slope of the Phillips curve (as in Primiceri, 2006, for example). Such shocks would enter the central bank's welfare function in a mathematically identical fashion as our political shocks.

privately observed by the central bank in period t, so we refer to θ_t as the central bank's *type*. We make the following assumption:

Assumption 1. *There exists $\widehat{\theta} \in \Theta$ such that $\theta f'(\theta)/f(\theta) > -2$ if $\theta < \widehat{\theta}$ and $\theta f'(\theta)/f(\theta) < -2$ if $\theta > \widehat{\theta}$.*

Note that this assumption allows for $\theta f'(\theta)/f(\theta)$ to exceed or be below -2 over the whole set Θ; in this case, $\widehat{\theta}$ is defined as either the upper bound or the lower bound of the set Θ. Assumption 1 holds for a broad range of distribution functions, including uniform, exponential, log-normal, gamma, and beta for a subset of its parameters. This assumption is analogous to the distributional assumption used in Halac and Yared (2019).

Following the realization of θ_t, the central bank chooses inflation π_t. Let $\pi_t^e \equiv \mathbb{E}_t[\pi_t]$ be the rational expectation of inflation formed by households at the beginning of the period.[9] The output gap x_t is then determined according to the Phillips curve:

$$x_t = \kappa(\pi_t - \pi_t^e),$$

where $\kappa > 0$ denotes the slope of the Phillips curve.

Social welfare at date t is

$$V_t = \mathbb{E}_t \left[\sum_{s=0}^{\infty} \beta^s \left(-\frac{(\pi_{t+s})^2}{2} + \frac{\gamma}{\kappa} x_{t+s} \right) \right], \tag{1}$$

where $\beta \in (0,1)$ is the social discount factor and $\gamma/\kappa > 0$ represents the social weight on output expansion relative to the cost of inflation (normalized by κ to ease the exposition). Note that by the Phillips curve, $\mathbb{E}_t(x_t) = \kappa \mathbb{E}_t(\pi_t - \pi_t^e) = 0$. Substituting into (1), social welfare at t can thus be rewritten recursively as

$$V_t = \mathbb{E}_t \left[-\frac{(\pi_t)^2}{2} + \beta V_{t+1} \right].$$

The central bank's welfare when choosing policy at date t, after expectations

[9]The operator $\mathbb{E}_t$ denotes the expectation at the beginning of period t without knowledge of the realized shock θ_t.

π_t^e have been formed and the shock θ_t has realized, is

$$-\frac{(\pi_t)^2}{2} + \frac{\gamma}{\kappa}\theta_t x_t + \beta V_{t+1}.$$

Substituting with the Phillips curve, this can be rewritten as

$$-\frac{(\pi_t)^2}{2} + \gamma\theta_t\pi_t - \gamma\theta_t\pi_t^e + \beta V_{t+1}. \quad (2)$$

Following Mishkin and Westelius (2008), we model the political shock θ_t as impacting the weight that the central bank places on output expansion versus inflation stabilization. A higher weight reflects the increased importance of stimulating current output in order to boost an incumbent political party's popularity or to accommodate a fiscal expansion, for example.[10]

We require the inflation rate at each date to satisfy $\pi_t \in [\underline{\pi}, \overline{\pi}]$ for finite $\underline{\pi}, \overline{\pi}$, so that welfare is bounded. We take the range $[\underline{\pi}, \overline{\pi}]$ to be wide enough that this constraint is otherwise non-binding.

There are three main features of our environment. First, since $\gamma > 0$, the central bank is time-inconsistent. The central bank at date t shares the same preferences as society from date $t+1$ onward. The reason is that this central bank does not place any weight on future political shocks, and moreover it realizes that future realized inflation will be incorporated into the private sector's rational inflation expectations, limiting the benefit of inflation surprises. Thus, from the perspective of date t, setting $\pi_{t+s} = 0$ for all $s \geq 1$ maximizes both society's and the central bank's welfare. However, the central bank at date $t+s$ is biased relative to society: given a fixed continuation value, its welfare is maximized by setting a strictly positive inflation rate $\pi_{t+s} = \gamma\theta_{t+s} > 0$. The reason is that at the time of choosing policy, the central bank does not internalize the effect of current inflation on past inflation expectations, and it thus underweighs the cost of inflation in its decision-making. This form of time-inconsistency is common to many models of monetary policy (e.g., Barro and Gordon, 1983a,b). In our setting, the degree to which the central bank underweighs the cost of inflation

[10] Our results also apply if the political shock enters additively in the cost of inflation. Under this modification, our results can be extended to a dynamic New Keynesian framework. Details available upon request.

depends on political pressure; specifically, the central bank's bias is increasing in the political shock θ_t. We denote by $\pi^f(\theta_t)$ the statically optimal, or *flexible*, level of inflation for the central bank at date t conditional on θ_t:

$$\pi^f(\theta_t) = \gamma\theta_t. \tag{3}$$

The second feature of our environment is that the political shock θ_t is privately observed by the central bank at date t. This captures the fact that political pressure cannot be perfectly observed or quantified by an external entity, be it an entity in the current period or central banks in future periods.

The third feature of our environment is that the central bank has full discretion when choosing policy. This is a main distinction from previous work, such as Athey, Atkeson, and Kehoe (2005), which assumes that available policies can be restricted arbitrarily and at no cost. Instead, we posit that the central bank can freely choose policy at each date t, and the continuation game following its policy choice serves as reward and punishment for its actions.

3 Equilibrium Definition

We define a self-enforcing rule as a perfect public equilibrium of the interaction between successive central banks. Let $h^{t-1} = \{\pi_0, \ldots, \pi_{t-1}\}$ denote the public history of inflation through time $t-1$ and $\mathcal{H}^{t-1}$ the set of all possible such histories. A public strategy for the central bank in period t is $\sigma_t\left(h^{t-1}, \theta_t\right)$, specifying, for each history $h^{t-1} \in \mathcal{H}^{t-1}$ and current central bank type $\theta_t \in \Theta$, a feasible level of inflation, $\pi_t\left(h^{t-1}, \theta_t\right) \in [\underline{\pi}, \overline{\pi}]$. Expected inflation at h^{t-1}, $\pi_t^e(h^{t-1})$, must be consistent with the central bank's strategy. A perfect public equilibrium is a profile of public strategies $\sigma = \left(\sigma_t\left(h^{t-1}, \theta_t\right)\right)_{t=0}^{\infty}$ such that, for each $t \in \{0, 1, \ldots\}$, $\sigma_t\left(h^{t-1}, \theta_t\right)$ maximizes the t-period central bank's welfare (2) given expectations $\pi_t^e(h^{t-1})$ and the continuation strategies $\left(\sigma_{t+k}\left(h^{t+k-1}, \theta_{t+k}\right)\right)_{k=1}^{\infty}$ of all central banks. We henceforth refer to perfect public equilibria as simply equilibria.

Let $V_t(h^{t-1})$ denote the continuation value to the central bank starting from a history h^{t-1}. At any (on- or off-path) history h^{t-1}, the continuation value given the equilibrium strategies can be represented recursively as follows:

$$V_t(h^{t-1}) = \mathbb{E}_t\Big[-\frac{(\pi_t(h^{t-1},\theta_t))^2}{2} + \beta V_{t+1}(h^{t-1},\pi_t(h^{t-1},\theta_t))\Big]. \tag{4}$$

A profile of strategies $(\sigma_t\,(h^{t-1},\theta_t))_{t=0}^{\infty}$ constitutes an equilibrium if and only if, for all $t \in \{0, 1\ldots\}$ and all (on- and off-path) histories h^{t-1}, the following private information and self-enforcement constraints are satisfied:

$$\begin{aligned}&-\frac{(\pi_t(h^{t-1},\theta_t))^2}{2} + \gamma\theta_t\pi_t(h^{t-1},\theta_t) + \beta V_{t+1}(h^{t-1},\pi_t(h^{t-1},\theta_t)) \\ &\quad\geq -\frac{(\pi_t(h^{t-1},\theta_t'))^2}{2} + \gamma\theta_t\pi_t(h^{t-1},\theta_t') + \beta V_{t+1}(h^{t-1},\pi_t(h^{t-1},\theta_t')) \text{ for all } \theta_t,\theta_t' \in \Theta\end{aligned} \tag{5}$$

and

$$\begin{aligned}&-\frac{(\pi_t(h^{t-1},\theta_t))^2}{2} + \gamma\theta_t\pi_t(h^{t-1},\theta_t) + \beta V_{t+1}(h^{t-1},\pi_t(h^{t-1},\theta_t)) \\ &\quad\geq -\frac{(\pi_t')^2}{2} + \gamma\theta_t\pi_t' + \beta V_{t+1}(h^{t-1},\pi_t') \text{ for all } \theta_t \in \Theta \text{ and all } \pi_t' \neq \pi_t(h^{t-1},\theta_t') \text{ for all } \theta_t' \in \Theta.\end{aligned} \tag{6}$$

The private information constraint (5) captures the fact that the central bank at any date t can misrepresent its type. This constraint guarantees that a central bank of type θ_t prefers to pursue its assigned inflation rate rather than that of any other type $\theta_t' \neq \theta_t$. The enforcement constraint (6) captures the fact that the central bank at any date t can freely choose any feasible inflation rate $\pi_t' \in [\underline{\pi}, \overline{\pi}]$, including rates not assigned to any other central bank type. This constraint guarantees that a central bank of type θ_t prefers to pursue its assigned inflation rate rather than any other rate π_t' satisfying $\pi_t' \neq \pi_t(h^{t-1},\theta_t')$ for all $\theta_t' \in \Theta$. Note that in representing both of these constraints, we have ignored inflation expectations $\pi_t^e(h^{t-1})$, as this expectation has no impact on the central bank's strategy at (h^{t-1},θ_t).

Since inflation is bounded and shocks are i.i.d., there exists an upper bound $\overline{V}$ that corresponds to the highest continuation value that can be sustained by equilibrium strategies, with $V_{t+1}(h^{t-1},\pi_t') \leq \overline{V}$ for all h^{t-1} and π_t'. By analogous logic, there also exists a lower bound $\underline{V}$, with $V_{t+1}(h^{t-1},\pi_t') \geq \underline{V}$. Moreover, note that satisfying the enforcement constraint (6) requires that this constraint hold under maximal punishment, namely when $V_{t+1}(h^{t-1}, \pi'_t) = \underline{V}$. In fact, since the inequality must then hold for all $\pi'_t \in [\underline{\pi}, \overline{\pi}]$, it must necessarily hold when $\pi'_t = \pi^f(\theta_t)$. Therefore, a necessary condition for the enforcement constraint to be satisfied is

$$\begin{aligned} -\frac{(\pi_t(h^{t-1},\theta_t))^2}{2} + \gamma\theta_t\pi_t(h^{t-1},\theta_t) + \beta V_{t+1}(h^{t-1},\pi_t(h^{t-1},\theta_t)) \\ \geq -\frac{(\pi^f(\theta_t))^2}{2} + \gamma\theta_t\pi^f(\theta_t) + \beta\underline{V} \quad \text{for all } \theta_t \in \Theta, \end{aligned} \tag{7}$$

where note that the right-hand side is the central bank's minmax payoff.

Constraints (5) and (7) are clearly necessary for a sequence of inflation rates to be supported by equilibrium strategies. Furthermore, these constraints are also sufficient: if a sequence of inflation rates satisfies (5) and (7), then it can be supported by a strategy profile that specifies the worst feasible continuation value following any observable deviation. Since such a deviation is off path, it is without loss to assume that it is maximally punished.

4 Optimal Self-Enforcing Rule

We examine the equilibrium that maximizes social welfare starting from date 0. In what follows, we first consider a recursive representation of the welfare-maximizing equilibrium. We then show that this equilibrium can be characterized by a hawkish low-inflation regime associated with the highest welfare level $\overline{V}$ and a dovish high-inflation regime associated with the lowest welfare level $\underline{V}$. We maintain the assumption that $\overline{V} > \underline{V}$; this inequality is guaranteed to hold provided that $\beta \in (0, 1)$ is sufficiently high.

4.1 Recursive Representation

Let $\{\pi(\theta), V(\theta)\}_{\theta\in\Theta}$ specify the equilibrium inflation rate and continuation value for each type θ at a given date. By our equilibrium definition, this allocation must satisfy the following private information and self-enforcement constraints, analogous to (5) and (7) respectively:

$$-\frac{(\pi(\theta))^2}{2}+\gamma\theta\pi(\theta)+\beta V(\theta)\geq-\frac{(\pi(\theta'))^2}{2}+\gamma\theta\pi(\theta')+\beta V(\theta') \text{ for all } \theta,\theta'\in\Theta, \tag{8}$$

$$-\frac{(\pi(\theta))^2}{2}+\gamma\theta\pi(\theta)+\beta V(\theta)\geq-\frac{(\pi^f(\theta))^2}{2}+\gamma\theta\pi^f(\theta)+\beta\underline{V} \text{ for all } \theta\in\Theta. \tag{9}$$

Additionally, $\{\pi(\theta), V(\theta)\}_{\theta\in\Theta}$ must satisfy the following feasibility constraints:

$$\underline{\pi}\leq\pi(\theta)\leq\overline{\pi} \text{ and } \underline{V}\leq V(\theta)\leq\overline{V} \text{ for all } \theta\in\Theta. \tag{10}$$

A rule is incentive compatible if it satisfies (8)-(9), and it is incentive compatible and feasible if it satisfies (8)-(10).

Given this representation, the highest welfare level $\overline{V}$ corresponds to the solution to the following program:

$$\overline{V}=\max_{\{\pi(\theta),V(\theta)\}_{\theta\in\Theta}}\mathbb{E}\Big[-\frac{(\pi(\theta))^2}{2}+\beta V(\theta)\Big]$$
$$\text{subject to (8), (9), and (10).} \tag{11}$$

Analogously, the lowest welfare level $\underline{V}$ is the solution to:

$$\underline{V}=\min_{\{\pi(\theta),V(\theta)\}_{\theta\in\Theta}}\mathbb{E}\Big[-\frac{(\pi(\theta))^2}{2}+\beta V(\theta)\Big]$$
$$\text{subject to (8), (9), and (10).} \tag{12}$$

An optimal self-enforcing rule solves program (11). We assume that the solution admits a sequence of inflation rates that are piecewise continuously differentiable in type. Additionally, if the program admits multiple solutions that differ only on a countable set of types, we select the solution that maximizes social welfare for those types.

4.2 Benchmarks

To understand the role of self-enforcement and private information, it is useful to first consider optimal monetary policy in the absence of these frictions. Suppose that the enforcement constraint (9) in program (11) could be ignored. Then social welfare is maximized by setting $\{\pi(\theta), V(\theta)\} = \{0, \overline{V}\}$ for all $\theta \in \Theta$. That is, in this case inflation can be set at zero in all periods, and the economy does not respond to political shocks. As such, the private information constraint (8) plays no role if the enforcement constraint is never binding.

Suppose next that the enforcement constraint (9) does bind (under some or all $\theta \in \Theta$), but political shocks are observable and thus the private information constraint (8) in program (11) can be ignored. Then using arguments similar to those in Thomas and Worrall (1988), it can be shown that the solution to this program admits $\{\pi(\theta), V(\theta)\} = \{\max\{0, \pi^o(\theta)\}, \overline{V}\}$ for all $\theta \in \Theta$, where $\pi^o(\theta) < \pi^f(\theta)$ satisfies

$$-\frac{(\pi^o(\theta))^2}{2} + \gamma\theta\pi^o(\theta) + \beta\overline{V} = -\frac{(\pi^f(\theta))^2}{2} + \gamma\theta\pi^f(\theta) + \beta\underline{V}.$$

An optimal rule in this case assigns the lowest enforceable level of inflation conditional on the observed political shock. Relative to the case of perfect enforcement, inflation is higher and more volatile, since it directly responds to political shocks that tighten the enforcement constraint. Note however that since the continuation value equals $\overline{V}$ at all dates, the equilibrium restarts in every period, and temporary political shocks only have a temporary impact on inflation.

Our setting incorporates both classes of constraints, due to self-enforcement and private information. A shock-contingent rule as that used in the absence of private information is thus not incentive compatible: the central bank can deviate privately from its assigned policy and choose a higher inflation rate, making itself strictly better off without being penalized with a lower continuation value. Incentive compatibility in our setting requires that, given a realized political shock, the central bank prefer its assigned policy and continuation value to those prescribed for any other shock. We will show that, as a result, temporary political shocks can have persistent effects on inflation under an optimal monetary rule.

4.3 Hawkish Regime

To characterize the solution to program (11), we define the following rule:

Definition 1. $\{\pi(\theta), V(\theta)\}_{\theta\in\Theta}$ *is a maximally enforced inflation cap if there exist* $\theta^* \in [0, \overline{\theta})$ *and finite* $\theta^{**} > \max\{\theta^*, \underline{\theta}\}$ *such that*

$$\{\pi(\theta), V(\theta)\} = \begin{cases} \left\{\pi^f(\theta), \overline{V}\right\} & \text{if } \theta < \theta^* \\ \left\{\pi^f(\theta^*), \overline{V}\right\} & \text{if } \theta \in [\theta^*, \theta^{**}] \\ \left\{\pi^f(\theta), \underline{V}\right\} & \text{if } \theta > \theta^{**} \end{cases} \tag{13}$$

where

$$-\frac{(\pi^f(\theta^*))^2}{2} + \gamma\theta^{**}\pi^f(\theta^*) + \beta\overline{V} = -\frac{(\pi^f(\theta^{**}))^2}{2} + \gamma\theta^{**}\pi^f(\theta^{**}) + \beta\underline{V}. \tag{14}$$

Under this rule, types $\theta \in [\underline{\theta}, \theta^*)$ and $\theta \in \left(\theta^{**}, \overline{\theta}\right]$ choose their flexible inflation level $\pi^f(\theta)$, and types $\theta \in [\theta^*, \theta^{**}]$ choose type θ^*'s flexible inflation level $\pi^f(\theta^*)$. Types $\theta \leq \theta^{**}$ are maximally rewarded with continuation value $\overline{V}$ whereas types $\theta > \theta^{**}$ are maximally punished with continuation value $\underline{V}$. By (14), the enforcement constraint holds with equality for type θ^{**}. This rule can be implemented using an inflation cap $\pi^f(\theta^*)$: if the central bank respects the cap, it receives maximal reward $\overline{V}$; if the central bank breaches the cap, it receives maximal punishment $\underline{V}$.

The following proposition shows that the highest continuation value $\overline{V}$ is sustained by a maximally enforced inflation cap:

Proposition 1 (hawkish regime)**.** *If* $\{\pi(\theta), V(\theta)\}_{\theta\in\Theta}$ *is a solution to* (11) *with* $\pi(\theta) \in (\underline{\pi}, \overline{\pi})$ *for all* $\theta \in \Theta$*, then it satisfies* (13)-(14) *for some* $\theta^* \in [0, \overline{\theta})$ *and finite* $\theta^{**} > \max\{\theta^*, \underline{\theta}\}$*. Hence, any interior solution is a maximally enforced inflation cap.*

The optimal monetary rule therefore consists of an inflation cap that leads to the worst punishment whenever violated. So long as the inflation cap is respected, the economy remains in a hawkish regime that implements this cap in every period.

The proof of Proposition 1 and our other results in the next sections follow from applying the arguments developed in Halac and Yared (2019). We thus omit formal details from this article and refer the reader to the work in Halac and Yared (2019).[11]

As noted in Subsection 4.2, in a setting with perfect enforcement, the optimal rule would set zero inflation at every date. Such a policy corresponds to a maximally enforced inflation cap of 0, associated with type $\theta^* = 0$. Naturally, if this cap can be enforced given $\{\underline{V}, \overline{V}\}$, then it is also optimal under self-enforcement:

Corollary 1. *Suppose*

$$\beta\overline{V} \geq -\frac{(\pi^f(\overline{\theta}))^2}{2} + \gamma\overline{\theta}\pi^f(\overline{\theta}) + \beta\underline{V}. \tag{15}$$

If $\{\pi(\theta), V(\theta)\}_{\theta\in\Theta}$ *is a solution to* (11) *with* $\pi(\theta) \in (\underline{\pi}, \overline{\pi})$ *for all* $\theta \in \Theta$*, then it is the perfect-enforcement inflation cap, with* $\theta^* = 0$ *and* $\theta^{**} \geq \overline{\theta}$*.*

When condition (15) holds, the highest type $\overline{\theta}$, and therefore all types $\theta \in \Theta$, prefer to respect the perfect-enforcement inflation cap of 0 and receive maximal reward $\overline{V}$ rather than inflate above this cap and receive maximal punishment $\underline{V}$. The optimal rule under self-enforcement therefore coincides with that under perfect enforcement and features no on-path punishment. Note that condition (15) trivially holds in a scenario where the central bank is fully independent from political pressure, namely where the political shocks satisfy $\overline{\theta} = \underline{\theta} = 0$.

Our interest is in characterizing the optimal self-enforcing rule when condition (15) does not hold, so the perfect-enforcement inflation cap is not enforceable given $\{\underline{V}, \overline{V}\}$. Proposition 1 implies that this rule takes one of two possible forms. One form is a relaxed inflation cap specifying $\theta^{**} \geq \overline{\theta}$, so that the enforcement constraint is satisfied under all shocks and there are no transitions to

[11] Relative to Halac and Yared (2019), here the bias of the agent (namely, the central bank) takes a different mathematical form, and there is no state variable across periods. Despite these differences, the proof of Proposition 1 follows from analogous arguments, applying Assumption 1 along with the first order approach to simplify the central bank's private information constraints. Details available upon request.

punishment on path. In this case, welfare equals $\overline{V}$ and the economy remains in the hawkish regime at all dates. The second possible form is an inflation cap specifying $\theta^{**} < \overline{\theta}$, so that the enforcement constraint is violated under high enough shocks $\theta > \theta^{**}$ and punishment occurs on path. In this case, the economy remains in the hawkish regime associated with welfare $\overline{V}$ as long as the realized value of θ is below θ^{**}; once a shock $\theta > \theta^{**}$ is realized, the economy transitions to the worst punishment associated with welfare $\underline{V}$. In the next sections, we characterize the equilibrium that sustains $\underline{V}$ and provide a necessary and sufficient condition for on-path punishment to be prescribed by the optimal self-enforcing rule.

4.4 Dovish Regime

In principle, different continuation equilibria could serve as punishment for a central bank violating the inflation cap in the hawkish regime. In fact, Proposition 1 holds independently of the exact structure of punishment. However, the optimal self-enforcing rule requires that the worst punishment be used, as such a punishment maximally relaxes the constraints in program (11) and thus maximizes welfare. We therefore next study the solution to program (12). To characterize this solution, it is useful to define the following rule:

Definition 2. $\{\pi(\theta), V(\theta)\}_{\theta\in\Theta}$ *is a maximally enforced inflation floor if there exist finite* $\theta_n^* > \underline{\theta}$ *and* $\theta_n^{**} \in [\underline{\theta}, \min\{\theta_n^*, \overline{\theta}\}]$ *such that*

$$\{\pi(\theta), V(\theta)\} = \begin{cases} \left\{\pi^f(\theta), \underline{V}\right\} & \text{if } \theta < \theta_n^{**} \\ \left\{\pi^f(\theta_n^*), \overline{V}\right\} & \text{if } \theta \in [\theta_n^{**}, \theta_n^*] \\ \left\{\pi^f(\theta), \overline{V}\right\} & \text{if } \theta > \theta_n^* \end{cases} \quad (16)$$

where

$$-\frac{(\pi^f(\theta_n^*))^2}{2} + \gamma\theta_n^{**}\pi^f(\theta_n^*) + \beta\overline{V} = -\frac{(\pi^f(\theta_n^{**}))^2}{2} + \gamma\theta_n^{**}\pi^f(\theta_n^{**}) + \beta\underline{V}. \quad (17)$$

Under this rule, types $\theta \in [\underline{\theta}, \theta_n^{**})$ and $\theta \in \left(\theta_n^*, \overline{\theta}\right]$ choose their flexible inflation level $\pi^f(\theta)$, and types $\theta \in [\theta_n^{**}, \theta_n^*]$ choose type θ_n^*'s flexible inflation level $\pi^f(\theta_n^*)$. Types $\theta \geq \theta_n^{**}$ are maximally rewarded with continuation value $\overline{V}$

whereas types $\theta < \theta_n^{**}$ are maximally punished with continuation value $\underline{V}$. By (17), the enforcement constraint holds with equality for type θ_n^{**}. This rule can be implemented using an inflation floor $\pi^f(\theta_n^*)$: if the central bank respects the floor, it receives maximal reward $\overline{V}$; if the central bank breaches the floor, it receives maximal punishment $\underline{V}$.

The following proposition shows that the lowest continuation value $\underline{V}$ is sustained by a maximally enforced inflation floor:

Proposition 2 (dovish regime)**.** *If* $\{\pi(\theta), V(\theta)\}_{\theta\in\Theta}$ *is a solution to* (12) *with* $\pi(\theta) \in (\underline{\pi}, \overline{\pi})$ *for all* $\theta \in \Theta$*, then it satisfies* (16)-(17) *for some finite* $\theta_n^* > \underline{\theta}$ *and* $\theta_n^{**} \in [\underline{\theta}, \min\{\theta_n^*, \overline{\theta}\}]$*. Hence, any interior solution is a maximally enforced inflation floor.*

In the absence of enforcement constraints, the worst punishment would entail forcing all central bank types in all future periods to choose the highest or lowest feasible inflation rates, so as to minimize the value of welfare. However, such a harsh punishment would not be self-enforcing. Proposition 2 shows that the worst punishment that is self-enforcing takes the form of a maximally enforced inflation floor. This floor minimizes welfare by incentivizing overinflation. Intuitively, given a central bank type θ, there are two ways in which (ex-ante) welfare can be reduced: either by inducing too little inflation or by inducing too much inflation. Since the central bank is biased towards overinflating in the present, the latter relaxes enforcement constraints, and it is thus a more efficient means of reducing welfare. As a result, in the worst-punishment allocation, all central bank types choose inflation that is positive and thus socially costly. In fact, inflation is weakly above the flexible level preferred by the central bank. Analogous to Proposition 1, this overinflation is incentivized by maximally rewarding the central bank for respecting the inflation floor and maximally punishing the central bank for violating it.

4.5 Transitions

The results in Proposition 1 and Proposition 2 have important implications for the dynamics of inflation. Starting in a hawkish regime at date t, the central bank is subject to a maximally enforced inflation cap $\pi^f(\theta^*)$. If $\theta_t \leq \theta^{**}$, the central bank respects the cap, future inflation expectations remain low, and the equilibrium restarts in the hawkish regime at $t+1$. If instead $\theta_t > \theta^{**}$, the central bank violates the cap, future inflation expectations rise, and the equilibrium transitions to the dovish regime at $t+1$.

Starting in a dovish regime at date t, the central bank is subject to a maximally enforced inflation floor $\pi^f(\theta_n^*)$. If $\theta_t \geq \theta_n^{**}$, the central bank respects the floor, future inflation expectations decline, and the equilibrium transitions to the hawkish regime at $t+1$. If instead $\theta_t < \theta_n^{**}$, the central bank violates the floor, future inflation expectations remain high, and the equilibrium restarts in the dovish regime at $t+1$.

A maximally enforced inflation cap in the hawkish regime maximizes social welfare by counteracting the political pressure to inflate. A maximally enforced inflation floor in the dovish regime minimizes social welfare by inducing the central bank to bend to the political pressure. The threat of transitioning to the dovish regime sustains the hawkish regime, and the promise of returning to the hawkish regime sustains the dovish regime. Punishment is always temporary, since a central bank's succumbing to political pressure is rewarded with a transition back to the hawkish regime.

A natural question is whether transitions to the dovish regime occur on path (i.e., $\theta^{**} < \overline{\theta}$) or the economy always remains in the hawkish regime (i.e., $\theta^{**} \geq \overline{\theta}$). To answer this question, let $\theta_c < \overline{\theta}$ be the unique type corresponding to the tightest inflation cap that all types $\theta \in \Theta$ would be willing to respect:

$$-\frac{(\pi^f(\theta_c))^2}{2} + \gamma\overline{\theta}\pi^f(\theta_c) + \beta\overline{V} = -\frac{(\pi^f(\overline{\theta}))^2}{2} + \gamma\overline{\theta}\pi^f(\overline{\theta}) + \beta\underline{V}. \tag{18}$$

Note that $\theta_c \leq 0$ whenever the perfect-enforcement inflation cap of 0 is enforceable given $\{\underline{V}, \overline{V}\}$, and $\theta_c > 0$ otherwise. Using this definition of θ_c, the next proposition provides a necessary and sufficient condition for punishment to be optimally imposed along the equilibrium path.

Proposition 3 (use of punishment)**.** *If* $\{\pi(\theta), V(\theta)\}_{\theta\in\Theta}$ *is a solution to* (11) *with* $\pi(\theta) \in (\underline{\pi}, \overline{\pi})$ *for all* $\theta \in \Theta$*, then it is the unique such solution. Moreover, if*

$$(\overline{\theta} - \theta_c)\overline{\theta}\frac{f(\overline{\theta})}{1 - F(\theta_c)} \geq \theta_c, \tag{19}$$

this solution is a maximally enforced inflation cap with $\theta^* = \max\{\theta_c, 0\}$ *and* $\theta^{**} \geq \overline{\theta}$*. Otherwise, this solution is a maximally enforced inflation cap with* $\theta^* \in (0, \theta_c)$ *and* $\theta^{**} < \overline{\theta}$*.*

Whenever the perfect-enforcement cap is enforceable ($\theta_c \leq 0$), condition (19) is satisfied and the optimal rule coincides with that under perfect enforcement. If instead the perfect-enforcement cap is not enforceable ($\theta_c > 0$), then the following tradeoff arises. On the one hand, the monetary rule can raise the value of θ^* to the point that the associated cap $\pi^f(\theta^*)$ satisfies the enforcement constraint of type $\overline{\theta}$ and thus of all types $\theta \in \Theta$. This option entails setting $\theta^* = \theta_c$ and $\theta^{**} = \overline{\theta}$ and has the benefit of avoiding socially costly punishment along the equilibrium path, albeit at the cost of potentially allowing significant overinflation within the relaxed inflation cap. On the other hand, the monetary rule can specify a tighter cap $\pi^f(\theta^*)$ that does not satisfy the enforcement constraint of all types. This option sets $\theta^* < \theta_c$ and $\theta^{**} < \overline{\theta}$ and induces higher discipline on types $\theta \leq \theta^{**}$, but at the cost of transitioning to punishment whenever a shock $\theta > \theta^{**}$ is realized.

Proposition 3 tells us that which of these two options is optimal depends on whether the inequality in (19) holds or not. To analyze this condition, keep fixed the support $[\underline{\theta}, \overline{\theta}]$ and the value of θ_c. Then condition (19) shows how the use of punishment depends on the distribution of political shocks. The condition implies that punishment is not imposed on path if high political shocks are relatively likely, namely if $f(\overline{\theta})/(1 - F(\theta_c))$ is sufficiently high. This situation arises, for example, under a uniform distribution of shocks. In this case, it is optimal to set a relaxed inflation cap that is never violated, as punishing the central bank following high shocks would be too costly.

In contrast, punishment is optimally imposed on path if high political shocks are relatively unlikely, namely if $f(\overline{\theta})/(1 - F(\theta_c))$ is sufficiently low. This situation arises whenever the perfect-enforcement inflation cap is not enforceable

$(\theta_c > 0)$ and $\lim_{\theta \to \bar{\theta}} f(\theta) = 0$, as is true for example under a beta distribution with beta shape parameter greater than one. In this case, it is optimal to set a tight inflation cap that is violated following high enough political shocks, as such events are sufficiently rare that the cost of punishing the central bank is then relatively low.

5 Concluding Remarks

This paper has studied optimal monetary policy when the central bank lacks commitment to policies and is subject to privately observed, time-varying political pressure. We showed that a maximally enforced inflation cap mitigates the political pressure to inflate in a hawkish regime. A maximally enforced inflation floor instead accommodates the political pressure in a temporary dovish regime, serving as a punishment for violations of the inflation cap in the hawkish regime. We examined the conditions for regime transitions to occur on path and how they depend on the distribution of political shocks.

Our analysis takes political pressure as given and explores how central banks optimally respond to it. A remaining question of interest concerns the underlying nature of political shocks, and the extent to which central bank policy can endogenously affect their distribution. In fact, political shocks may themselves be endogenous to economic shocks that also impact the central bank's welfare. Studying how central bank behavior can affect the nature of political and economic shocks jointly may be an interesting direction for future research.

References

Abreu, D., D. Pearce, and E. Stacchetti (1990): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58, 1041–63.

Aisen, A. and F. J. Veiga (2006): "Does Political Instability Lead to Higher Inflation? A Panel Data Analysis," *Journal of Money, Credit, and Banking*, 38, 1379–89.

——— (2008): "Political Instability and Inflation Volatility," *Public Choice*, 135, 207–23.

Amador, M. and K. Bagwell (2013): "The Theory of Delegation with an Application to Tariff Caps," *Econometrica*, 81, 1541–600.

Athey, S., A. Atkeson, and P. J. Kehoe (2005): "The Optimal Degree of Discretion in Monetary Policy," *Econometrica*, 73, 1431–75.

Backus, D. and J. Driffill (1985): "Rational Expectations and Policy Credibility Following a Change in Regime," *Review of Economic Studies*, 52, 211–21.

Barro, R. J. and D. B. Gordon (1983a): "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Political Economy*, 91, 589–610.

——— (1983b): "Rules, Discretion and Reputation in a Model of Monetary Policy," *Journal of Monetary Economics*, 12, 101–121.

Binder, C. C. (2018): "Political Pressure on Central Banks," Working Paper.

Canzoneri, M. B. (1985): "Monetary Policy Games and the Role of Private Information," *American Economic Review*, 75, 1056–70.

Capistrán, C. and M. Ramos-Francia (2009): "Inflation Dynamics in Latin America," *Contemporary Economic Policy*, 27, 349–62.

Cukierman, A. and A. H. Meltzer (1986): "A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information," *Econometrica*, 54, 1099–128.

Fraga, A., I. Goldfajn, and A. Minella (2003): "Inflation Targeting in Emerging Market Economies," *NBER Macroeconomics Annual*, 18, 365–416.

Halac, M. and P. Yared (2019): "Fiscal Rules and Discretion under Limited Enforcement," Working Paper.

Holmström, B. (1977): "On Incentives and Control in Organizations," Ph.D. Thesis, Stanford University.

Kocherlakota, N. (2016): "Rules versus Discretion: A Reconsideration," *Brookings Papers on Economic Activity*, Fall.

Mishkin, F. S. and N. J. Westelius (2008): "Inflation Band Targeting and Optimal Inflation Contracts," *Journal of Money, Credit and Banking*, 40, 557–82.

Primiceri, G. E. (2006): "Why Inflation Rose and Fell: Policy-Makers' Beliefs and U.S. Postwar Stabilization Policy," *Quarterly Journal of Economics*, 121, 867–901.

Rogoff, K. (1985): "The Optimal Degree of Commitment to an Intermediate Monetary Target," *Quarterly Journal of Economics*, 100, 1169–1189.

Sargent, T. J. (2001): *The Conquest of American Inflation*, Princeton University Press, Princeton, New Jersey.

Sargent, T. J., N. Williams, and T. Zha (2009): "The Conquest of South American Inflation," *Journal of Political Economy*, 117, 211–56.

Thomas, J. and T. Worrall (1988): "Self-Enforcing Wage Contracts," *Review of Economic Studies*, 55, 541–53.

Walsh, C. (1995): "Optimal Contracts for Central Bankers," *American Economic Review*, 85, 150–67.

Yared, P. (2019): "Rising Government Debt: Causes and Solutions for a Decades-Old Trend," *Journal of Economic Perspectives*, 33, 115–140.