

Inflation Target Zones as a Commitment Mechanism

Felipe Schwartzman

Much of the analysis of monetary policy has centered on central banks' alleged inability to commit to an optimal, time-consistent policy. Commitment is important whenever the private sector makes decisions according to what it expects the policymakers to do in the future. By committing to future actions, the central bank can better influence the decisions of the private sector and can generate better policy outcomes.

The literature on monetary policy has identified two types of biases that emerge from lack of commitment. The first is *inflation bias*, as in Kydland and Prescott (1977) and Barro and Gordon (1983) analyses. In that case, equilibrium inflation ends up being, on average, higher than optimal. More recently, the literature has identified a second type of bias, called *stabilization bias*, in which the central bank reacts too much to changes in the state of the economy, leading to excessive volatility in inflation and the output gap as in (Clarida et al. (1999)).

This article aims to analyze the effects of a widely used framework that has been adopted by many central banks seeking to provide a greater degree of predictability in monetary policy—that is, the use of an inflation target zone.¹ Such zones specify a range within which the central bank promises to keep inflation. I interpret these zones as providing constraints on central bank behavior—so that it has to commit to not letting inflation leave that zone—while allowing the central bank to follow, a discretionary policy otherwise. The numerical computations reported in this article indicate that the use of an inflation

■ The original version of this paper was written in 2005 at Princeton. Thanks to Antonio Luis Licha, Lars Svensson, Carlos Viana de Carvalho, Dante Amengual, Paul Ho, Reiko Laski, Pierre-Daniel Sarte, Nico Trachter, Elaine Wissuchek, and John Weinberg for helpful comments and suggestions in different stages of this work.

¹ As Erceg (2002) points out, the definition of a target zone, as opposed to a point target, seems to be adopted by most inflation-targeting countries.

target zone allows for elimination of the inflation bias and, in some cases, substantial correction of the stability bias to a point that will be made explicit shortly. This is because the inflation target zone affects expected inflation even when it is not binding. Furthermore, the existence of a credible target zone dampens the effect of current shocks on expectations because the private agents know that if further shocks occur, the central bank will only react to them so long as the zone is respected. This dampening of the sensitivity of inflation to the underlying state is similar to what Krugman (1991) called a “honeymoon effect” in the context of exchange rate target zones.

At the same time, having the central bank commit to a prespecified inflation rate without any room for variation is hardly optimal either. The reason is that the inflation target zone is determined before all relevant information is revealed. In the presence of supply shocks, it is optimal for a central bank that is not a strict inflation targeter to adjust inflation somewhat in the direction of the shock so that output does not vary too much. This suggests that there is in general a nontrivial optimal width for the inflation target zone. The paper also presents calculations of the optimal zone under different parameter constellations.

Prior literature has concerned itself with the trade-off between greater commitment and discretion to react to incoming information. A classic proposal, Rogoff (1985), to reduce the inflation bias problem is that the central banker should give less weight to output gap variation than society at large, i.e., it should be relatively “weight” conservative. Svensson (1997) proposes that the inflation-targeting framework reduces inflation by forcing central bankers to pursue a lower average rate of inflation than what would be socially optimal, thus compensating for the bias, i.e., he proposes that the central bank should be “inflation target” conservative. The numerical calculations in this article show that the use of a target zone can effectively implement Svensson’s suggestion by setting the center of the zone low enough.

Also, as emphasized by Clarida et al. (1999), Rogoff’s proposal is also able to, at least in part, correct the stability bias by implementing an optimal linear policy rule. There is no way such corrections can implement the first-best policy arrangement, for this is history dependent. The same is true for the target zone arrangement. What the numerical calculations show is that while the target zone mechanism falls short of providing the same gains as assigning a suitable “weight” conservative central banker, at least for some parameter values, it is able to generate very close outcomes.

Inflation target zones can be interpreted as contracts, in which the punishment for the central bank only occurs if inflation exceeds certain

limits. There are a number of papers that try to assess the effects of adopting contracts with central bankers. Walsh (2020) shows that one can use a contract that will fire the central banker if the nominal GDP growth (or some modified version of it) exceeds a certain value to be specified. This contract eliminates the inflation bias, thus implementing what Walsh describes as an optimal policy. However, in Walsh's model, there is no persistence in inflation shocks, so that it does not allow for a discussion of the stability bias in monetary policy.

Another article that addresses the benefits of a target zone is Tetlow (2000). In his model, the central bank follows a Taylor-rule-type policy, but with a randomly varying inflation target, which Tetlow models as following a martingale. The model generates a honeymoon effect, akin to the one observed in Krugman's exchange rate target zone model.

The existence of such an effect was also proposed by Gerlach (1994) and Amano et al. (1997). However, in both cases, the authors retain the flex-price approach of the Krugman model and keep policy exogenous. One of the results presented in this article is that by explicitly modeling the policy, the honeymoon effect can yield an indirect effect on inflation and output gap variability, which is probably more important than the direct impact of nonlinearities implied by the honeymoon effect.

Finally, it is worth mentioning two articles that approach the issue from a different perspective. One is Erceg (2002), who regards the target range as a confidence-interval-type statement of the variance that the central bank wishes inflation to have. In this interpretation, the target range does not have a "hard edge," in the sense that the central bank can allow inflation to obtain values out of the zone, but range should be respected say, 95 percent of the time. However, as Erceg recognizes, target zones seem to be too narrow in that inflation only falls within their bounds less than 95 percent of the time. By allowing the target zone to be more than a simple statement of intentions, the model presented in this paper presents a rationale for why the optimal zone may be "too narrow" by Erceg's criterion.

The other article is by Orphanides and Wieland (2000). Their problem is to find what would justify a policy in which the central bank only pays attention to inflation when it is out of the target zone, otherwise focusing the policy on output gap stabilization. They propose a quadratic loss function, which is nearly flat within the target zone, and the possibility of nonlinearities in the Phillips curve so that the inflation-output trade-off only becomes significant for high levels of inflation.

The article proceeds as follows. In the next section, I will lay out the model. Section 2 shows how the existence of the zone affects the discretionary, intrazone policy. Section 3 discusses how the location of

target zone approach can be used to change the average inflation, thus dealing with the inflation bias problem. In Section 4, the effects of different zone widths on the variances of inflation and the output gap are considered. Section 5 provides a discussion of some issues involved in the finding of the optimal target zone as a function of the model parameters. The last section concludes by summarizing policy implications, caveats in the analysis, and directions for future development of the model.

1. THE MODEL

A central bank chooses policy to minimize a quadratic loss function, taking as given a forward-looking Phillips curve, but with no ability to precommit. More specifically, the intertemporal loss function of the central bank is given by

$$L = E_t \left\{ \sum_{s=t}^{\infty} (1 - \delta) \delta^{s-t} \frac{1}{2} \left[(\pi_t)^2 + \alpha (x_t - \bar{x})^2 \right] \right\}, \quad (1)$$

where, as usual, π_t and x_t stand for inflation and output gap, respectively, and \bar{x} stands for the average levels of output gap that the central bank would like to attain in the long run; α is a parameter determining the weight that the central bank gives to deviations of output gap from target relative to inflation; δ is a discount factor; and E_t represents the expectation operator conditional on information available at time t .² Higher values of α imply that the central bank is relatively more concerned with stabilizing output gap close to its target than with stabilizing inflation. The central bank minimizes this loss function subject to a forward-looking Phillips curve:³

$$\pi_t = E_t [\pi_{t+1}] + \lambda x_t + u_t, \quad (2)$$

where $E_t[\cdot]$ is the expectation operator conditional on information available in time t , and u_t is a supply shock, which follows an AR(1) process:

² I multiply the loss function by $(1 - \delta)$ so that it is well-defined in the $\delta \rightarrow 1$ limit.

³ This Phillips curve is similar to the one derived from a Calvo-style, sticky-price model as in Calvo (1983). The difference is that the expected future inflation is not multiplied by the discount rate δ . This simplification is not of great consequence because in any case, δ would be calibrated to be very close to 1. A substantive consequence of this calibration is that the Phillips curve becomes vertical in steady state, so that there is no long-run trade-off between inflation and output gap levels.

$$u_t = \rho u_{t-1} + \eta_t, \quad (3)$$

where ρ is a parameter governing the persistence of u_t and η_t is an i.i.d. $N(0, \sigma^2)$.

A central bank that is fully able to commit chooses inflation and the output gap in each date and given each realization of the supply shock u_t to minimize the discounted losses (1). Given the Phillips curve (equation 2), the central bank understands that by committing to lower inflation at any given date, it will achieve a lower output gap on that same date but will also affect inflation and output gaps on the previous date.⁴ Manipulating first-order conditions, one can verify that the optimal policy with full commitment satisfies:

$$\begin{aligned} \pi_t &= -\frac{\alpha}{\lambda}(x_t - \delta^{-1}x_{t-1}) \quad \text{if } t > 0 \\ \pi_t &= -\frac{\alpha}{\lambda}x_t \quad \text{if } t = 0. \end{aligned}$$

However, when a central bank lacks the ability to precommit, it reoptimizes each period, thus ignoring the effect of its inflation choice on previous variables. Discretion means in this case that the central bank reoptimizes every period, so that the $t = 0$ solution is the only relevant one.

The inflation-targeting framework is interpreted in this article as follows: the central bank is assigned a target zone and is given freedom to act so long as inflation remains within the zone. Therefore, the further constraint added in this model is that the central bank has to keep inflation within a certain zone, i.e.:

$$\pi_t \in [\pi^* - B; \pi^* + B], \quad (4)$$

where π^* denotes the center or the location of the zone, and B denotes its width. Furthermore, for simplicity, it is assumed that the central bank can pick inflation and the output gap at each period, subject to the Phillips curve above, without any error or lags. Therefore, inflation always falls within the zone.⁵

⁴ Put differently, the way in which central banks choose things at any given date has implications for how private agents form their expectations for what is going to happen on that date. The commitment problem emerges exactly because once a given date arrives, the central bank does not have an incentive to take this impact on previously formed expectations into account.

⁵ It would be straightforward to relabel variables in the model so that the central bank controls a component of inflation and commits to keeping that component that it controls within the zone. I do not do that for simplicity of exposition.

I derive the optimal decision of a central bank that does not have the ability to commit to any particular policy rule. To be more specific, I focus on Markovian decision rules, where the central bank only conditions its policy on the current exogenous state of the economy and reoptimizes each period. Under such policy, the way the decision rule affects expectations about future variables is never a consideration, since at the beginning of each period, those are no longer relevant.

Under those assumptions, the central bank will take expectations about future inflation as something it cannot influence within each period—that is, taking them as given. Since the decisions made by the central bank in each period do not affect its loss and hence its action in future periods, the dynamic minimization problem reduces to a sequence of static minimizations of the period loss functions.⁶ If the central bank is allowed to choose output gap and inflation freely, I have in such a framework the usual result that the desired inflation and output gap will satisfy:

$$\pi_t^d = -\frac{\alpha}{\lambda} (x_t^d - \bar{x}), \quad (5)$$

where π_t^d and x_t^d stand for inflation and the output gap desired by the central bank in the absence of the target zone. However, the central bank is constrained by the target zone, so that in effect:

$$\begin{aligned} \pi_t &= \pi_t^d \text{ if } \pi_t^d \in [\pi^* - B, \pi^* + B] \\ \pi_t &= \pi^* - B \text{ if } \pi_t^d \leq \pi^* - B \\ \pi_t &= \pi^* + B \text{ if } \pi_t^d \geq \pi^* + B. \end{aligned} \quad (6)$$

Substituting (5) into (2) yields

$$\pi_t^d = \frac{\alpha}{\alpha + \lambda^2} [E_t[\pi_{t+1}] + \lambda\bar{x} + u_t]. \quad (7)$$

In this simple model, the state of the economy is completely summarized by u_t , so that I should be able to find π_t and π_t^d as functions of u_t .

⁶ More precisely, the problem allows one to look for a Markov equilibrium of the game played between the central bank and its future incarnations, where the actions taken by the central bank are solely a function of the current state of the economy summarized by the shock u_t . This is not the only possible solution for such a game. One could very well look for strategies that take past values of u_t as part of the state. This modification would allow for reputational effects analogous to the ones discussed by Barro and Gordon (1983) or to the implementation of a first-best policy as in Woodford's (2003) "timeless" commitment.

One can find a stationary equilibrium $\pi_t(u_t)$ using the following algorithm:

1. Start with a $\pi^0(u_t)$ function that satisfies (4);
2. Calculate $\pi^{d,1}(u_t)$ by plugging $\pi^0(u_t)$ in the right-hand side of (7);
3. Use equation (6) to find $\pi^1(u_t)$;
4. Iterate until convergence to a fixed point.

As shown in the Appendix, this method is a contraction, so that it necessarily yields a unique reaction function to the central bank. With $\pi(u_t)$ at hand, it is straightforward to find $E[\pi(u_{t+1})|u_t]$ with the use of (3) and then $x(u_t)$ by inverting the Phillips curve (2). With these functions at hand, one can proceed to find the moments of inflation and output gap through numerical integration.

Throughout the article, the resulting policy rule shall be compared with two benchmark cases. One is the optimal policy under discretion in the absence of the target zone.⁷ This is a natural benchmark because it is with respect to this policy that the policy implied by the target zone should improve. The second benchmark is the optimal linear policy, derived by Clarida et al. (1999). The latter is a useful benchmark because it is the optimal discretionary policy for a “conservative” central bank, that is, one that has an inflation target and welfare weights that differ from society at large. These benchmarks are given, respectively, by:⁸

$$\pi^{disc}(u_t) = \alpha\bar{x} + \omega u_t = \alpha\bar{x} + \frac{\alpha}{1 + \alpha(1 - \rho)} u_t$$

and

$$\pi^{opt}(u_t) = \omega^* u_t = \frac{\alpha^c}{1 + \alpha^c(1 - \rho)} u_t,$$

where

⁷ For conciseness, this will be referred to hereafter as simply policy “under discretion” or “discretionary policy.”

⁸ For a derivation, see Clarida et al. (1999). The use of the optimal linear policy as a benchmark, as opposed to the global optimal policy, is justified by the fact that the global optimal policy is history-dependent, and there is no hope that the proposed mechanism can make the central bank adopt such a policy. Also, as seen below, the optimal linear policy does seem to dominate any policy implementable with the target zone.

$$\alpha^c = (1 - \rho) \delta.$$

Note that average inflation in the discretionary case is given by:

$$E \left[\pi^{disc}(u_t) \right] = \alpha \bar{x}.$$

Following the optimal linear policy typically requires a degree of commitment from the central bank, so it is generally not followed. The only circumstance in which it is attained is if the central bank has output gap weight given by α^c and target output gap $\bar{x} = 0$. This amounts to putting in place decision makers in the central bank who are more “conservative” than society in the sense of having a lower output gap weight α and targeting the natural rate of output rather than some other level for the output gap.⁹

The model can be normalized to reduce the number of free parameters without any loss in generality. One normalization is to set $\sigma = \sqrt{1 - \rho^2}$. This sets the unconditional standard deviation of the u_t shocks to 1. The model can be renormalized by multiplying the loss function by $\frac{\sigma^2}{1 - \rho^2}$ and the constraints by $\frac{\sigma}{\sqrt{1 - \rho^2}}$. This means that if the normalized model implies an optimal zone given by (π^*, B) , then the actual optimal zone should be $\left(\frac{\sigma}{\sqrt{1 - \rho^2}} \pi^*, \frac{\sigma}{\sqrt{1 - \rho^2}} B \right)$. Likewise, the resulting variances and squared inflation bias should be all multiplied by $\frac{\sigma^2}{1 - \rho^2}$.

Another normalization is to note that if one can redefine variables so that the welfare weight is reset to $\tilde{\alpha} = \frac{\alpha}{\lambda^2}$ and the pass-through from output gap to inflation set at $\tilde{\lambda} = 1$. Changes in this parameter will indicate both the effects of having a more conservative central bank or less costly short-run trade-off between inflation and the output gap. The required redefinition is that x_t be the direct effect of the output gap on inflation, normally given by λx_t .

2. THE EFFECT OF THE ZONE IN INTRAZONE POLICY

Krugman (1991) identified what he called a “honeymoon effect” in the context of exchange rate target zones. His insight was that because the

⁹ As discussed in Svensson (1997), one can also implement the same policy by instead of requiring the central bank to ignore the target output gap \bar{x} , requesting that the central bank be “inflation conservative” in the sense of desiring a lower inflation that what is socially optimal.

exchange rate is a forward-looking variable, by committing to intervene so that the exchange rate does not get out of certain bounds, the government can in effect make exchange rates less sensitive to underlying fundamentals, thus reducing the variability of the exchange rate even around levels where no intervention takes place. Previous literature, cited above, has already suggested that the framework may be adapted to discuss inflation target zones.

I now compare the policy constrained by the zone with the policy adopted under full discretion. This is done in Figure 1. The blue lines represent the policy with the target zone, whereas the red lines represent the policy when there is no such zone. The parameter values used to plot the figures are $\alpha = 3$ and $\bar{x} = 0$. The latter equality implies absence of inflation bias in the discretionary policy. The figures in the left column (1a and 1c) represent inflation as a function of the supply shock u_t , and the figures in the right column its first derivative. For the upper figures (1a and 1b), the shocks were calibrated to be fairly persistent, at $\rho = 0.8$.¹⁰

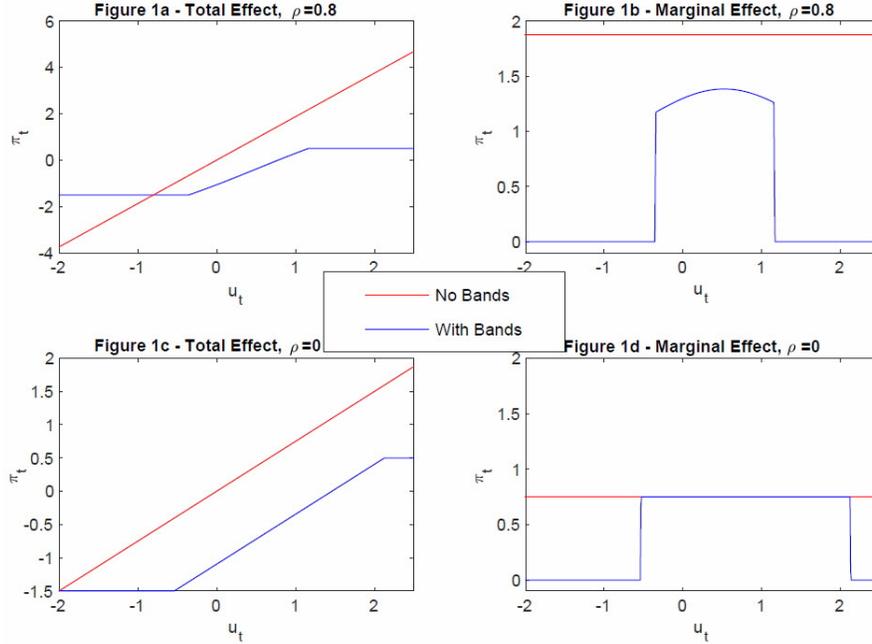
What immediately stands out in figure 1a is that apart from flattening out the reaction of inflation to the shock in the points where inflation is forced against the edges of the zones, the behavior of inflation within the zone is changed. The $\pi(u_t)$ schedule is both translated to the right and rotated downwards. This means that for a given shock, the zone implies a lower inflation level, and for a given change in the shock, the zone implies a weaker reaction.

To gain intuition on the way in which the target bands affect inflation, note that choice of inflation inside the band given by equation (7) does not depend directly on the target band. The only way in which the target band can affect inflation inside the band is, therefore, through its effect on future expected inflation. As the supply shock u_t increases, the probability that the band binds in the next period increases, thus further truncating the right tail of the distribution of inflation outcomes. It follows that future expected inflation increases more slowly with the supply shock than would be the case in the absence of the target band. The narrower the width of the band, the more future inflation is truncated for any current value of u_t . At the same time, the choice of location of the target band affects average inflation by truncating the distribution of future inflation more in one direction than the other.

Figures 1c and 1d are equivalent to 1a and 1b above, but with i.i.d. shocks ($\rho = 0$). By eliminating any source of persistence, the slope

¹⁰ This calibration will be used often in the article, and it corresponds to a half life of approximately three years for the shocks.

Figure 1 Inflation as a Function of the Supply Shock for Different Levels of Persistence



Note: Other parameters are set to $\alpha = 3$, $\lambda = 1$, and $\bar{x} = 0$.

of the inflation schedule does not change with the introduction of the zone, and the nonlinearities within the zone disappear. This is because with $\rho = 0$, inflation expectations do not depend on the current state. It follows that the intrazone policy is described by:

$$\pi_t^d = \frac{\alpha}{\alpha + \lambda^2} [E[\pi_{t+1}] + \lambda \bar{x} + u_t]. \quad (8)$$

With u_t i.i.d., expected future inflation is given by its unconditional expectation and does not depend on current values of u_t . The slope of the intrazone policy rule is given by $\frac{\alpha}{\alpha + \lambda^2}$, which is identically equal to ω when $\rho = 0$. The honeymoon effect disappears. However, the zone may still have an effect on the intercept of the policy by shifting $E[\pi_{t+1}]$ around and thus correcting for eventual inflation biases. In effect, the translation caused by the shifted π^* is, if anything, greater than before.

In the following two sections, it will be shown how the choice of π^* and B can be used to affect the outcome in terms of average inflation and the variances of both inflation and the output gap.

3. USING THE TARGET ZONE TO CHANGE AVERAGE INFLATION

The original articles on dynamic inconsistency of monetary policy (Kydland and Prescott (1977); Barro and Gordon (1983)) were mainly concerned with the issue of inflation bias. The theory states that the central banks may fall prey to inflation biases because the natural rate of output is above what would be the optimal one. For this reason, central banks would desire to have the output gap permanently at a level above its natural rate, thus pushing steady-state inflation above what would be optimal and, because of the long-term neutrality of the Phillips curve assumed in these models, not achieving their desired goal in terms of the output gap or unemployment.¹¹

As noted above, in the absence of an inflation target zone, the model in this paper also implies the existence of such a bias if $\alpha\bar{x} \neq 0$. In particular, I introduce inflation bias in the exercises shown in figures 2a and 2b by setting $\bar{x} = 1$ (the other parameter values are $\alpha = 3$, $\delta \rightarrow 1$, and $\rho = 0.8$). When plotting these figures, I set the width of the zone to $B = 2$.¹² While the bias under discretion is different for $\alpha = 1$ and $\alpha = 3$ (and given the parameterization, is in fact equal to α), the point that minimizes the loss functions and that sets the bias equal to 0 is close to $\pi^* = -2$ in both cases, so that the upper boundary of the band is close to zero, the optimal inflation. In figure 2b, variation in α affects how quickly inflation increases with π^* , with higher α corresponding to a steeper slope.

It is important to note that in the presence of inflation bias, the center of the band π^* does not coincide with average inflation (the set of points where it does is depicted by the diagonal dashed line in 2b). The loss-minimizing level of inflation target, in fact, implies inflation to be persistently above the center of the band. A simple extension of the model could allow for biases stemming from a preference by the central bank for lower than optimal inflation. In those cases, one might observe inflation to be persistently below the center of the inflation target.

¹¹ Whether this characterization of central banks is realistic is open to disagreement, and at least one insider account, Blinder (1999), disputes this notion.

¹² For comparison, the maximal standard deviation that inflation attains in this class of models (given the normalized variance of the shocks) is the one implied by the discretionary regime and increases with $\omega = \frac{\alpha}{1+\alpha(1-\rho)}$. For the parameterizations used here, this is 5/6 for $\alpha = 1$ and 15/8 for $\alpha = 3$.

4. USING THE TARGET ZONE TO CHANGE INFLATION AND OUTPUT GAP VARIANCES

The monetary policy literature based on new Keynesian type models, such as the one used here, has recognized that even if there is not an inflation bias problem, the fact that price-setting behavior is forward-looking is enough to generate time inconsistency issues in policy making. This appears typically as inflation that is too variable. In particular, it can be the case that both output and inflation are more variable than at the optimum. This behavior is what has been called the “stabilization bias.” By managing the variances of inflation and the output gap, the target zone can potentially substantially ameliorate this issue.

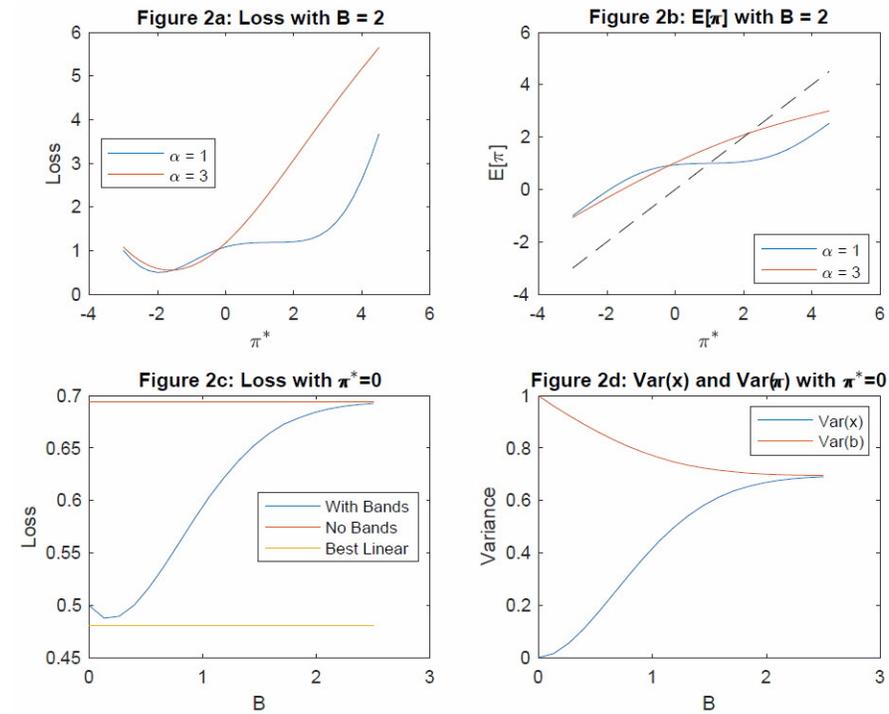
Figures 2c and 2d show how this is done. For plotting these figures, the parameter values were set at $\alpha = 1$, $\delta \rightarrow 1$, $\bar{x} = 0$, and $\rho = 0.8$. Also, in all figures, $\pi^* = 0$. A result that is not shown is that under these conditions, average inflation is not affected by the width of the zone, but the same would not be true if there were some kind of inflation bias problem.

Figure 2c shows how the loss function changes as a function of B and compares it with the loss when there is no zone and when the optimal linear policy is implemented. The loss attains a minimum at a point close to $B = 0$, implying that the chosen parameter values require a fairly tight but nontrivial target zone to achieve the optimal policy. It is not generally the case that the optimal zone is positive and finite. For example, if $\alpha = \delta$ and $\rho = 0$ the optimal zone is to have no zone at all and to let the central bank pursue a discretionary policy, which amounts to the same thing as having $B \rightarrow \infty$.

Another interesting result is that for smaller values of B , the target bands generate substantial gains relative to the fully discretionary policy. At the same time, the loss is always larger than the one implied by the best linear policy, even if at the optimum, it gets very close to that benchmark.¹³

The significance of the fact that loss implied by the optimal linear policy represents a lower bound on the loss function is that, as emphasized by Clarida et al. (1999), in the absence of inflation bias, the optimal linear policy can be interpreted as assigning a central banker with output gap weight given by $(1 - \rho)\delta$. This result implies that the use of the zone is not a substitute for the appointment of a suitably

¹³ However, this result could depend critically on the linear-quadratic nature of the model. For an example of how a departure from such a model can generate target-zone-type behaviors as optimal policies, see Orphanides and Wieland (2000).

Figure 2 Effects of Target Bands

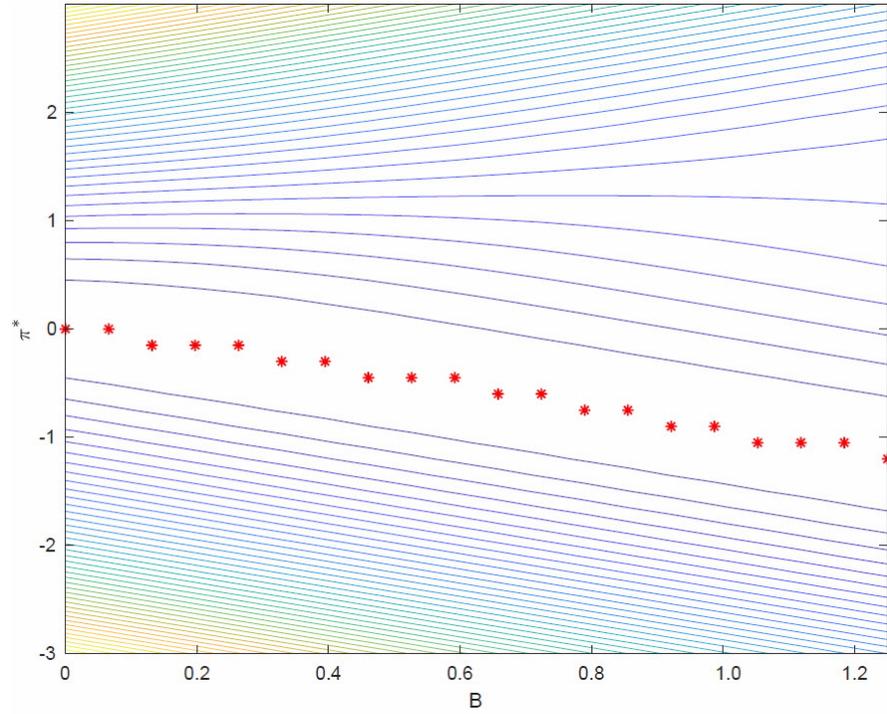
Note: Parameters set to $\alpha = 1$, $\delta \rightarrow 1$, $\bar{x} = 0$, and $\rho = 0.8$.

“weight conservative” central bank, even though in the example shown here, it can approximate the outcome to a large extent.

Figure 2d shows how the variances of inflation and output gap change with the size of the zone. These behave exactly as one would expect, with inflation variance increasing and output gap variance decreasing. However, inflation increases faster than the reduction in the output gap. This illustrates the fundamental trade-off between inflation and the output gap variability faced in the design of the target zone.

5. DEFINING THE OPTIMAL ZONE

The results shown above demonstrate that, at least for certain parameter values, there is, given π^* , an optimal nonzero and finite value for

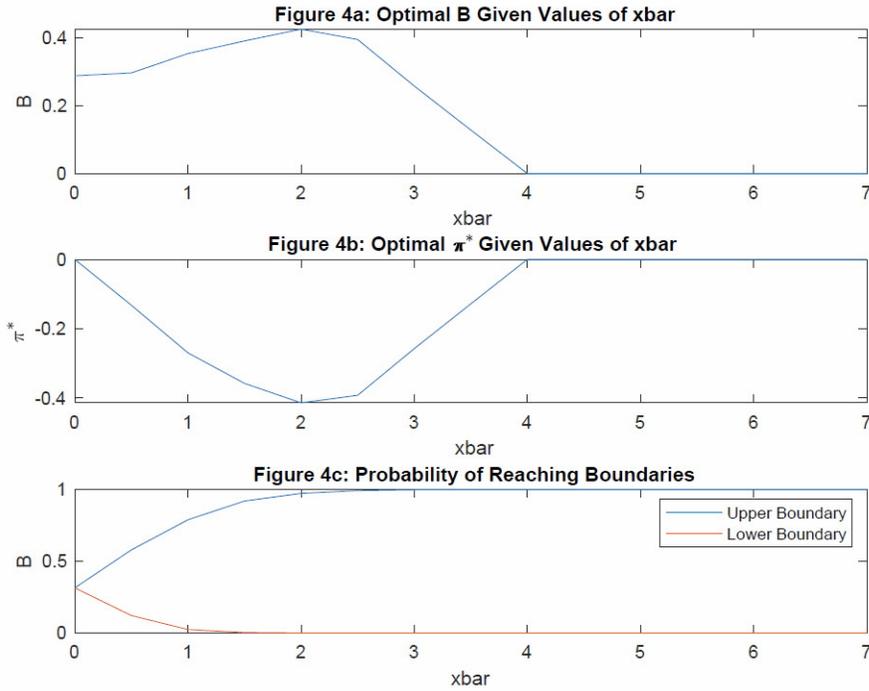
Figure 3 Contour Set of a Loss Function

Notes: The figure shows the loss function for different values of B and π^* . Red dots denote loss-minimizing π^* for different values of B . The figure was plotted with $\rho = 0.8$, $\alpha = 1$, $\delta \rightarrow 1$, and $\bar{x} = 1$. The loss-minimizing B is close to 0.2.

the width of the zone and for a given value of B , an optimal location for π^* . It seems natural to think that there will be in general an optimal combination of B and π^* . The loss associated for different values of B and π^* are shown in Figure 3. The figure was plotted with $\rho = 0.8$, $\alpha = 1$, $\delta \rightarrow 1$, and $\bar{x} = 1$. The optimal B is small, close to 0.2. The red dotted line depicts the optimal π^* for given values of B . Optimal policy design also involves a negative π^* , which is nevertheless smaller in absolute value than the inflation bias.

Also, for small-to-medium values of B , the level curves tend to slope downward, with higher B s being associated with more negative values of π^* . The intuition is straightforward. If, for example, B is set at, say, 0.5 and π^* at -1 , as is optimal when B is very large, this

Figure 4 Optimal Band Design Given Different Degrees of Inflation Bias



Note: Parameters set to $\alpha = 1$, $\delta \rightarrow 1$, and $\rho = 0.8$.

means that inflation will always be lower than its optimal long-term level, which is equal to zero. It follows that if the width of the target zone is constrained to be very wide, one may want to set the center of that zone at a lower level than if the width is fairly tight.

Figures 4a through 4c show how the optimal policy changes with \bar{x} given $\alpha = 1$, $\delta \rightarrow 1$, and $\rho = 0.8$. First, note that for $\bar{x} = 0$, the optimal zone is centered at 0, as one might expect. As \bar{x} increases, the optimal width of the zone increases slowly at first and then, at around $\bar{x} = 1.5$, it starts to grow faster. This behavior is mirrored by the optimal location of the zone, which at first departs only slowly from its initial value of π^* and then, from $\bar{x} = 1.5$, it starts to decrease at roughly one for one with the increase in \bar{x} . More interestingly, at some value of \bar{x} close to 4, the optimal zone collapses to include only the point

$\pi = 0$. It is as if the mechanism is able to deal with moderate values of inflation bias, but after a point, it doesn't have any option better than just telling the central bank to stick to zero inflation whatever happens. Figure 4c shows how in the absence of bias, the probabilities of inflation reaching either side of the zone are equal and sizable. As \bar{x} increases, the probability of inflation reaching the upper part of the zone increases, and the probability of reaching the lower part decreases. The optimal zone in the presence of a moderate inflation bias will tend to be asymmetric not only in the sense that average inflation will exceed the center of the zone, but also in the sense that it will reach the upper edge of the zone more often than its lower edge.

6. CONCLUSIONS, CAVEATS, AND DIRECTIONS FOR FUTURE WORK

Blinder (1999) commented about how economists often behave as in the joke in which one of them is confronted with some real-world fact and asks himself or herself about whether it is also true in theory. This is in some sense exactly what was done in this paper. The use of inflation target zones is fairly common among inflation-targeting countries. Not only their location, but also their width varies from country to country, and the latter is often regarded as being too "tight" in the sense that inflation hits its edges too often.

The model presented in this paper provides a theoretical rationale for these real-world facts. The use of inflation target zones is shown to be able to eliminate the inflation bias (which leads to inflation that is too high) and, at least in some circumstances, significantly reduce the stability bias (which leads to inflation that is too volatile). Also, in this model, the zone's optimal width and location depend in complicated ways on the underlying parameters of the economy. Finally, the tightness of the zone emerges endogenously from the optimization of the policy designer, so that it is to be expected that whenever one thinks that the zone can be used to deal with the stability bias issue, one would see inflation reaching its edges fairly often.

The ability of the inflation target band to approximate optimal policy under commitment is limited because it cannot make policy history dependent. At the same time, the numerical simulations imply that it can approximate the best linear, non-history-dependent policy.

In terms of empirical applications, the model implies there is an inflation bias, and if the zone is being set optimally, inflation should be on average above the center of the target zone, and one should see it touch the upper edge of the zone more often than the lower edge. This would serve as a test both of the adequacy of this model, in case one is

willing to maintain the hypothesis that the inflation bias exists, and as a test of the existence of an inflation bias if one is willing to maintain the hypothesis that the model is a valid description of reality.

It would be interesting to evaluate what changes in the model could allow it to account for the inflation persistently below target observed in the United States in the past years. Those changes are unlikely to stem from a central bank that pursues a negative output gap, while economic distortions such as market power would typically lead the natural rate of output to be below the first-best. More likely, the model would suggest inflation may fall persistently below target if central bank staff has a preference for inflation that is below the one that the Federal Reserve communicates to the public. For example, to the extent that the experience of run-away inflation in the 1970s still looms large as a cautionary tale for the current generation of central bank leadership, this could lead to a negative inflation bias.

The main caveat to the findings shown here is that the results came from a very simplified model. Future work should try to evaluate which of these results are robust to the introduction of a more realistic utility function that gives weight to other features of the distribution of inflation and output gap apart from their mean and variance or nonlinearities in the Phillips curve, along the lines of what has been done by Orphanides and Wieland (2000). More importantly, there should be some investigation into how well the target zone fares when there is persistence not only in the shocks, but in the inflation and output gap processes as well.

REFERENCES

- Amano, Robert, Richard Black, and Marcel Kasumovich. 1997. "A Band-Aid Solution to Inflation Targeting." Bank of Canada Staff Working Paper 97-11 (May).
- Barro, Robert J., and David B. Gordon. 1983. "A Positive Theory of Monetary Policy in a Natural Rate Model." *Journal of Political Economy* 91 (August): 589–610.
- Blinder, Alan S. 1998. *Central Banking in Theory and Practice*. Cambridge, Mass: MIT Press.
- Calvo, Guillermo A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (September):

383–398.

- Clarida, Richard, Jordi Gali, and Mark Gertler. (1999). “The Science of Monetary Policy: A New Keynesian Perspective.” *Journal of Economic Literature* 37 (December): 1661–1707.
- Erceg, Christopher J. 2002. “The Choice of an Inflation Target Range in a Small Open Economy.” *American Economic Review* 92 (May): 85–89.
- Gerlach, Stefan. 1994. “On the Symmetry between Inflation and Exchange Rate Targets.” *Economics Letters* 44 (February): 133–137.
- Krugman, Paul R. 1991. “Target Zones and Exchange Rate Dynamics.” *Quarterly Journal of Economics* 106 (August): 669–682.
- Kydland, Finn E., and Edward C. Prescott. 1977. “Rules Rather than Discretion: The Inconsistency of Optimal Plans.” *Journal of Political Economy* 85 (June): 473–491.
- Orphanides, Athanasios, and Volker Wieland. 2000. “Inflation Zone Targeting.” *European Economic Review* 44 (June): 1351–1387.
- Rogoff, Kenneth. 1985. “The Optimal Degree of Commitment to an Intermediate Monetary Target.” *Quarterly Journal of Economics* 100 (November): 1169–1189.
- Svensson, Lars E. O. 1997. “Optimal Inflation Targets, ‘Conservative’ Central Banks, and Linear Inflation Contracts.” *American Economic Review* 87 (March): 98–114.
- Tetlow, Robert J. 2000. “Inflation Targeting and Target Instability.” Finance and Economics Discussion Series 2000–01, Board of Governors of the Federal Reserve System, revised 2000.
- Walsh, Carl E. 2002. When Should Central Bankers Be Fired? *Economics of Governance* 3 (March): 1–21.
- Woodford, Michael. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, N.J.: Princeton University Press.

APPENDIX

Proposition 1 *Let the operator T be given by the algorithm presented in Section 2. Then for $\lambda = 1$, T is a contraction with modulus $\frac{\alpha}{\alpha+1}$.*

I will show that T is a contraction by showing that it satisfies the Blackwell sufficient conditions. First, note that there is no loss in generality in focusing our attention on bounded $\pi(\cdot)$ as for given finite B and π^* , $\pi(\cdot)$ is necessarily bounded. To show that T is a contraction, one has to show, furthermore, that it satisfies the following two conditions:

i) T is monotone:

$$\pi \leq \pi' \Rightarrow T\pi \leq T\pi'.$$

I have after setting $\lambda = 1$ that

$$\pi^{d,n+1}(u_t) = \frac{\alpha}{\alpha+1} (E_t[\pi^n(u_{t+1})] + \bar{x} + u_t).$$

By the properties of the expectations operator, if $\pi^{n'}(u_{t+1}) \geq \pi^n(u_{t+1}) \forall u_{t+1}$, then it must be the case that $\pi^{d,n+1'}(u_t) \geq \pi^{d,n+1}(u_t) \forall u_t$. As $\pi^{n+1}(u_t)$ is such that if $\pi^{d,n+1}(u) \geq \pi^{d,n+1}(u')$, then $\pi^{n+1}(u) \geq \pi^{n+1}(u')$. It must therefore be the case that $\pi^{d,n+1'}(u_t) \geq \pi^{d,n+1}(u_t) \forall u_t$, establishing the result.

ii) Discounting:

$$\exists \delta \in (0, 1) \text{ s.t. } T(\pi + c)(u_t) \leq T(\pi) + \delta c, c \geq 0$$

Let $\pi^{n'}(u_t) = \pi^n(u_t) + c$ with $c > 0$. Then

...

$$\begin{aligned} \pi^{d,n+1'}(u_t) &= \frac{\alpha}{\alpha+1} (E_t[\pi^n(u_{t+1}) + c] + \bar{x} + u_t) \\ &= \pi^{d,n+1}(u_t) + \frac{\alpha}{\alpha+1}c. \end{aligned} \tag{9}$$

To see that this implies that $\pi^{n+1'}(u_t) \leq \pi^{n+1}(u_t) + \frac{\alpha}{\alpha+1}c$, note that

$$\begin{aligned}\pi^{n+1} &= \max\left(\min\left(\pi^{d,n+1}, \pi^* + B\right), \pi^* - B\right) \\ \pi^{n+1'} &= \max\left(\min\left(\pi^{d,n+1'}, \pi^* + B\right), \pi^* - B\right).\end{aligned}$$

If $2B < \frac{\alpha}{\alpha+1}c$, then we are done. Otherwise, there are five relevant cases:

- 1 : $\pi^{d,n+1'}(u_t) < \pi^{d,n+1}(u_t) \leq \pi^* - B$, and
- 2 : $\pi^* + B \leq \pi^{d,n+1'}(u_t) < \pi^{n+1}(u_t)$.

In these two cases, $\pi^{n+1}(u_t) = \pi^{n+1'}$. The third case is

$$3 : \pi^* - B < \pi^{d,n+1'}(u_t) < \pi^{d,n+1}(u_t) < \pi^* + B.$$

In this case, $\pi^{n+1}(u_t) = \pi^{d,n+1}(u_t)$, and $\pi^{n+1'}(u_t) = \pi^{d,n+1'}(u_t)$, so that equation (9) holds. The last two cases are:

- 4 : $\pi^{d,n+1'}(u_t) \leq \pi^* - B < \pi^{d,n+1}(u_t)$.
- 5 : $\pi^{d,n+1'}(u_t) \leq \pi^* + B < \pi^{d,n+1}(u_t)$.

In both of these cases, $\pi^{n+1}(u_t) - \pi^{n+1'}(u_t) \leq \pi^{d,n+1}(u_t) - \pi^{d,n+1'}(u_t)$, establishing the result. So with $\delta = \frac{\alpha}{\alpha+1} < 1$, the discounting property is established.