

The Optimal Design of Interest Rate Target Changes*

Graeme Guthrie

School of Economics and Finance, Victoria University of Wellington, Wellington, New Zealand

Julian Wright[†]

Department of Economics, University of Auckland, Auckland, New Zealand.

Abstract

Most central banks currently implement monetary policy by targeting a short-term interest rate. This paper asks: “What is the optimal form for such interest rate targeting, given the objectives facing central banks?” We find the optimal rule is for the central bank to change the target rate whenever the deviation between its preferred rate and the current target rate reaches some critical level, and in this case the target rate is changed by a discrete amount in the direction of its preferred rate. Despite the simplicity of this rule, we are able to replicate a number of puzzling features of interest rate targeting observed in practice. Policy implications are drawn for the model calibrated to United States data.

JEL classification: E43; E52; E58

Keywords: Central Banks; Federal Funds Rate; Interest Rate Smoothing.

1 Introduction

Most central banks currently use interest rate targets as their operating objective in the implementation of monetary policy. Central banks in Australia, Canada, Japan, and the United States all target an overnight interbank interest rate, while in most other countries short-term interest rates (tender rates) are targeted. This paper asks: “What is the optimal form for such interest rate targeting, given the objectives facing central banks?” The question is motivated by the practices of the Federal Reserve, which has long targeted the federal funds rate, either directly or indirectly. Goodfriend (1991) suggests the Fed targets the federal funds rate to achieve its ultimate policy objectives, but in doing so it is careful not to “whipsaw the market” and waits till sufficient information has been accumulated before changing the target rate. In particular, he notes that adjustments to the target rate are made at irregular intervals in relatively small steps. Some target changes occur in relatively rapid succession, but when this occurs the changes are in the same direction; target changes are not soon reversed. He also suggests target changes are essentially unpredictable at forecast horizons longer than a month or two. Rudebusch (1995)

*We thank Marvin Goodfriend, Glenn Rudebusch, Carl Walsh, and Michael Woodford for helpful comments. All errors are ours.

[†]Corresponding author: Ph: 64-9-373-7999 extn 7943. Fax: 64-9-373-7427. Email: jk.wright@auckland.ac.nz

has confirmed these features empirically for two periods of explicit funds rate targeting: from September 1974 to September 1979 and from March 1984 to September 1992. He shows that target changes are conducted in small standardized steps with an approximate equality in the size of increases and decreases,¹ that in the first few weeks after a target change the Fed is fairly likely to change the target again in the same direction (but very unlikely to reverse its previous change), while after five weeks without a change there is only a small (but equal) likelihood of an increase or decrease in the target rate.

These practices are not unique to the Federal Reserve. The Bank for International Settlements (1998, p. 68) examines twelve industrial countries that target interest rates and finds that central banks generally move interest rates several times in the same direction before reversing policy, and that the interval between policy adjustments is typically considerably longer when the direction is changed. The Bank for International Settlements' data ends in March 1998, but start dates vary across countries. Using consistent starting dates and longer sample periods for five of these countries, Goodhart (1997) reports similar findings. Given the range of sample periods and countries considered, these findings span a variety of institutional and macroeconomic environments, suggesting a common explanation may be appropriate. Yet, despite the prevalence of interest rate targeting and the important role that it plays in monetary policy implementation, no model exists which explains this set of puzzling facts.² This paper provides such a model.

The model we develop is based on the following assumptions. The central bank at each point in time has a preferred level of some short-term interest rate which is determined by its ultimate objectives (such as inflation and output). We take the simplest case and assume this preferred rate evolves according to driftless Brownian motion. This assumption shows that we do not need to assume mean-reversion in the preferred rate in order to generate positive autocorrelation in target rate changes. The central bank is assumed to suffer flow costs which are quadratic in the difference between the actual interest rate and this preferred level. With just this assumption, the central bank's optimal policy would be to make the target rate exactly track its preferred rate, thus adjusting its target rate by infinitesimal amounts at every point in time. However, we do not observe this in practice. Instead, we observe infrequent discrete changes which suggest

¹Rudebusch also finds empirical evidence of equality in the time between two positive changes and the time between two negative changes, as well as equality in the time between an increase followed by a decrease and the time between a decrease followed by an increase.

²A recent literature, surveyed by Sack and Weiland (1999), explains why central banks may want to smooth interest rates. For instance, Woodford (1998) shows that with precommitment, it is socially optimal to smooth interest rates due to the forward-looking behavior of market participants. This allows the Fed to achieve a desired change in long-term rates, with lower fluctuations in short-term rates. A more gradualist interest rate policy can also be the optimal response to various types of uncertainty facing central banks. Sack (1998) shows that interest rate smoothing is optimal in the presence of uncertainty about the parameters of the central bank's model, while Orphanides (1998) shows that optimal policy involves a less aggressive response to changes in macroeconomic variables when measurement error is taken into account. In our view, while this literature provides a justification for interest rate smoothing, it does not explain the form interest rate adjustments should take — for instance, should interest rates be adjusted in small steps or by a continuous adjustment?

the central bank faces some costs to changing its interest rate target.

We think of these adjustment costs as arising from the central bank's additional role of maintaining orderly financial markets. Target changes may induce adverse market reactions in the form of difficulties in the banking sector (Goodfriend 1987 and Cukierman 1991), unwanted fluctuations in long-term interest rates (Goodfriend 1991), or more general credit market difficulties (for example, dramatic movements in foreign exchange and stock markets).³ The larger the target change the greater the likelihood that financial market stress will arise. However, even a very small target change can give rise to financial market disruptions, especially if markets are sensitive to merely the direction of a target change. Thus we assume both a fixed and proportional adjustment cost.

Given the different types of costs that the central bank faces, we derive the optimal form of the interest rate targeting rule, assuming the central bank chooses the size and timing of target changes to minimize the expected present discounted value of its total costs. The optimal rule is for the central bank to change the target rate whenever the deviation between its preferred rate and the current target rate reaches some critical level, and in this case the target rate is changed by a discrete amount in the direction of its preferred rate. Such a policy economizes on costly adjustments by avoiding changes in target rates that are likely to soon be reversed. This accords well with the observation by Goodfriend (1991) that the Fed waits till sufficient information has been accumulated before changing the target rate, being careful not to whipsaw the market. When the preferred rate is close to the target rate, the Fed will delay a target change, rather than unsettle the market, taking advantage of the fact many times small changes in the preferred rate will be undone by themselves.

Despite the simplicity of the optimal rule, we are able to replicate many observed features of interest rate targeting. Because of the fixed adjustment cost, target changes are discrete and infrequent. The symmetric model set-up gives rise to equality in the size of increases and decreases in the target rate, equality in the time between two positive changes and the time between two negative changes, and equality in the time between an increase followed by a decrease and the time between a decrease followed by an increase. Because of proportional adjustment costs, the target rate will only move part of the way towards the preferred rate. This implies "small" target changes. This also implies that the next target change is likely to be in the same direction (persistence), the average time between reversals is greater than the average time between continuations, and the most likely time for a target change in the same direction as the last one is some short period after the last target change. However, once some time has passed without a subsequent target change, the likelihood that the deviation between the preferred rate and the target rate is still close to the critical level diminishes. Thus, we find reversals become more likely as the time elapsed since the last target change becomes large. In

³In the same spirit, a literature on central bank secrecy suggests that the Fed's desire to smooth interest rates may explain secrecy over its policy direction; see, for example, Goodfriend (1986) and Dotsey (1987). The idea is that the market overreacts when the Fed releases new information and, since the Fed dislikes wild swings in interest rates, it finds it costly to reveal such information.

the limit, reversals and continuations become equally likely, so that the direction of the next target change becomes unpredictable.

We calibrate our model to data on federal funds rate target changes for the United States. Costs implied by the model seem plausible — fixed costs represent less than 3% of the total adjustment costs associated with changing the target rate by its average amount (23.9 basis points). Total adjustment costs from such a change are equivalent to allowing the target rate to deviate from the preferred rate by 50 basis points for around two months. We find that total adjustment costs can proxy the degree of central bank interest rate smoothing, providing us with a novel way to uncover a central bank’s interest rate smoothing tendencies. We also address the popular notion that central banks act “too little, too late” (see for instance, Goodhart 1997, p. 1). In our model, central banks act rationally, given the costs they perceive. However, to the extent central banks over-estimate adjustment costs, our analysis suggests they compromise their ultimate objectives. By moving the target rate to immediately reflect any movement in their preferred rate, they could reduce the average deviation between the preferred and the target rate by 35 basis points.

The rest of the paper is organized as follows. Section 2 presents our model of interest rate targeting, deriving the optimal targeting rule, as well as a number of properties of the implied target rate. In Section 3, the hazard function of target changes is constructed, comparative static results evaluated, and some policy implications considered using the model calibrated to United States data. Section 4 concludes with a discussion of some possible modifications to the model and some potential applications.

2 A Model of Interest Rate Targeting

The following sections form the core of the paper. Section 2.1 lays out our model of interest rate targeting, motivating each of the assumptions. Given this model, Section 2.2 derives the optimal form of interest rate targeting. Section 2.3 then analyzes the resulting process for the target rate, showing that the process gives rise to the stylized facts of interest rate targeting.

2.1 The Model Set-Up

Consider the example of a central bank targeting a particular interest rate. We are interested in when and by how much it changes its target rate. We start by supposing the Bank has a preferred level of the target rate at each point in time, which takes into account all relevant factors, except any costs of changing the target rate itself. We assume the process for this preferred level can be approximated by driftless Brownian motion. This assumption is made primarily for mathematical tractability, but it can be motivated in other ways. One motivation of this assumption is Mankiw’s (1987) optimal inflation tax theory. He applies Barro’s (1979) optimal tax-smoothing approach to the situation in which government revenue comes not only from income and sales tax, but also from printing new money (seigniorage). He uses this to explain why the Fed may want to induce a random walk in nominal interest rates. According

to this view, expected deadweight losses are minimized by maintaining expected constancy in both the goods tax rate and the nominal interest rate. Another motivation for our assumption is that it is difficult to reject empirically that short-term interest rates follow a driftless random walk.⁴ This has led others to use the random walk assumption in theoretical work; for instance, Barro (1989) assumes a target interest rate which follows a random walk, noting that “In fact, even for short-term rates, a random walk is a pretty good description of the recent data” (p. 6). An alternative to the random walk approach, motivated by the observation that target changes themselves appear to be positively autocorrelated, would be to allow mean-reversion in the preferred rate around an underlying Brownian motion process. Our approach shows such mean-reversion is not needed to explain the observed persistence in target changes.

We suppose that, by choosing the level of the target rate, the central bank ties down the (instantaneous) market interest rate at this same level. In practice, the market rate will deviate from the target level due to transitory liquidity shocks. However, to the extent these can not be affected by the central bank’s targeting policy, they will not alter the optimal form of interest rate targeting. We denote by ε_t the difference between the central bank’s preferred level and its target level. The central bank has the ability to change ε_t by changing the target rate; an increase in the target rate lowers ε_t . Since we assume the preferred level evolves according to driftless Brownian motion, it follows that as long as the target rate is not changed, ε_t evolves according to the same process. That is, ε_t evolves according to $d\varepsilon_t = \sigma dz_t$, where dz_t is the increment of a Wiener process and σ measures the volatility (assumed constant) of the preferred rate. The central bank is assumed to suffer flow costs $F(\varepsilon_t)dt$ when its preferred level of interest rates deviates from its target level of interest rates by ε_t for a period dt . We assume quadratic flow costs $F(\varepsilon_t) = \varepsilon_t^2$. These can be motivated by a standard loss function in which the central bank faces quadratic loss from deviations from its ultimate objectives.

If the central bank changes the target rate by Δ , it incurs adjustment costs of $C(\Delta)$, where $C(\Delta) = f + c|\Delta|$. Thus, there are two types of adjustment cost — a fixed adjustment cost f that is incurred whenever the target rate is changed, but which does not depend on the size of the target change, and a proportional adjustment cost $c|\Delta|$, which reflects that large target changes will be more costly than small ones. As suggested by Goodfriend (1987), these adjustment costs might arise because of the central bank’s additional role of maintaining orderly financial markets. Target changes may induce disproportionate reactions from financial markets, possibly because financial markets are excessively sensitive to new information from the central bank or because new information may be misinterpreted.⁵ The fixed adjustment cost, which reflects the fact that even a small target change could unsettle markets, is likely to be relatively small. We assume the

⁴Even if we reject empirically that the observed (targeted) interest rate has a unit root, the unobserved preferred rate could still have a unit root.

⁵An alternative view of adjustment costs is that they reflect political and reputational constraints facing the central bank. In this case, there will still be a fixed component to adjustment costs; the central bank may worry that even the direction of a target change could lead to negative comments against the central bank, pressures put on the bank’s chairman (through politicians or from within the bank itself), or changes in beliefs regarding the bank’s competence. Implications of this alternative view are discussed in Section 3.4.

marginal adjustment cost is constant at c . This ensures that the bank cannot reduce adjustment costs simply by dividing a single target change into a series of smaller changes that immediately follow one another; this could be the case with increasing marginal adjustment costs. Absent fixed adjustment costs, it also ensures that the bank cannot reduce adjustment costs simply by combining successive target changes into one; this would be the case with decreasing marginal adjustment costs. However, an increase in the target rate which is immediately offset by a reduction in the target rate (whipsawing the markets) will still be costly with this specification.

2.2 The Optimal Adjustment Rule

This section begins with a formal description of a general interest rate targeting rule and the associated cost function. Rather than study such general policies, we proceed by concentrating on a family of very simple adjustment rules. The expected total cost function associated with such rules is easily calculated. We take advantage of this, and obtain necessary conditions which the best rule from this family must satisfy. We then prove the existence of a unique solution to these necessary conditions. We conclude the section with our main result: no rule for interest rate targeting, even of the general type described at the beginning of the section, achieves a lower total expected cost than a particular simple adjustment rule, which we describe.

The central bank continuously monitors the discrepancy ε and intervenes when necessary to change the target rate. Any given policy will generate an increasing sequence of stopping times $T_1 \leq T_2 \leq \dots \leq T_i \leq \dots$ at which the target rate will be changed. At stopping time T_i , the central bank will change the target rate by some (possibly random) amount, say Δ_i , which can depend only on information available at time T_i . If the central bank adopts the targeting policy P characterized by the stopping times $\{T_i\}$ and target changes $\{\Delta_i\}$, the expected total cost is

$$J(\varepsilon; P) = E \left[\int_0^\infty e^{-\rho t} \varepsilon_t^2 dt + \sum_{i \geq 0} e^{-\rho T_i} (f + c|\Delta_i|) \right],$$

where the deviation between the preferred rate and the target rate is initially ε and $\rho > 0$ is the rate at which future costs are discounted by the central bank. The central bank will adjust the target rate using a rule which minimizes this total cost.

We now describe a very simple adjustment rule and prove that a rule of this form is optimal. The rule we consider is completely described by two constants, b and Δ , satisfying $0 < \Delta < b$. As long as $-b < \varepsilon < b$, the central bank leaves the target rate unchanged. If $\varepsilon \geq b$, the central bank immediately increases the target rate by the amount $\Delta + \varepsilon - b$, resetting the discrepancy to $b - \Delta$. Similarly, if $\varepsilon \leq -b$, the central bank immediately reduces the target rate by the amount $\Delta - \varepsilon - b$, resetting the discrepancy to $-b + \Delta$. Notice that the discrepancy never leaves the interval $[-b, b]$. The target rate behaves in a particularly simple fashion. It is held constant, except for discrete changes, all of the same magnitude.⁶ When the preferred rate

⁶The only exception to this rule involves behavior at time 0. If ε is initially outside the interval $[-b, b]$, the target rate is immediately adjusted in order to bring the discrepancy back to $\pm(b - \Delta)$.

moves sufficiently far away from the target rate, the target rate is adjusted an amount Δ in the direction of the preferred rate.

Let $u(\varepsilon)$ denote the expected total cost for such a policy. Suppose that $-b < \varepsilon_t < b$, and the central bank leaves the target rate unchanged for a period of time dt . Its expected total cost equals

$$u(\varepsilon_t) = E_t \left[\int_t^{t+dt} e^{-\rho(s-t)} \varepsilon_s^2 ds + e^{-\rho dt} u(\varepsilon_{t+dt}) \right],$$

comprising the sum of the expected flow cost over the next period of time dt and the expected total cost from time $t + dt$ onwards, appropriately discounted. Using Itô's Lemma,

$$E_t[u(\varepsilon_{t+dt})] = u(\varepsilon_t) + \frac{1}{2}\sigma^2 u''(\varepsilon_t)dt + o(dt).$$

Therefore, the expected total cost is

$$u(\varepsilon_t) = \varepsilon_t^2 dt + u(\varepsilon_t) + \frac{1}{2}\sigma^2 u''(\varepsilon_t)dt - \rho u(\varepsilon_t)dt + o(dt).$$

Taking the limit as $dt \rightarrow 0$, and dropping the time subscript, shows that

$$\varepsilon^2 + \frac{1}{2}\sigma^2 u''(\varepsilon) - \rho u(\varepsilon) = 0, \quad -b < \varepsilon < b. \quad (1)$$

Due to the symmetry of the adjustment policy and both the adjustment and flow cost functions, together with the fact that ε evolves according to a driftless Brownian motion, a discrepancy of $-\varepsilon$ must be exactly as costly as one of ε ; that is, $u(\varepsilon) = u(-\varepsilon)$ for all ε . The general solution to (1) having this property is easily found to be

$$u(\varepsilon) = \frac{\sigma^2}{\rho^2} + \frac{\varepsilon^2}{\rho} + A \cosh(\lambda\varepsilon),$$

where $\lambda^2 = 2\rho/\sigma^2$ and A is an arbitrary constant. The expected total cost if the discrepancy is initially $\varepsilon \geq b$ is

$$u(\varepsilon) = u(b - \Delta) + f + c(\Delta + \varepsilon - b),$$

since in this case the central bank immediately resets the discrepancy to $b - \Delta$ by increasing the target rate by $\Delta + \varepsilon - b$. For similar reasons, the expected total cost if $\varepsilon \leq -b$ is

$$u(\varepsilon) = u(-b + \Delta) + f + c(\Delta - \varepsilon - b).$$


Combining these three results, we see that

$$u(\varepsilon) = \begin{cases} u(-b + \Delta) + f + c(\Delta - \varepsilon - b), & \varepsilon \leq -b, \\ \frac{\sigma^2}{\rho^2} + \frac{\varepsilon^2}{\rho} + A \cosh(\lambda\varepsilon), & -b < \varepsilon < b, \\ u(b - \Delta) + f + c(\Delta + \varepsilon - b), & b \leq \varepsilon. \end{cases} \quad (2)$$

The constant A is determined by the requirement that u is continuous at $\varepsilon = \pm b$:

$$A = \frac{f + c\Delta - b^2/\rho + (b - \Delta)^2/\rho}{\cosh(\lambda b) - \cosh(\lambda(b - \Delta))}. \quad (3)$$

Figure 1: Implementing the optimal adjustment rule



As long as $\varepsilon \in (-b, b)$, the policy parameters b and Δ only influence $u(\varepsilon)$ through A . Since the coefficient on A , $\cosh(\lambda\varepsilon)$, is always positive, the expected total cost for all $\varepsilon \in (-b, b)$ can be minimized by choosing the policy parameters which minimize A . Let the parameters b^* and Δ^* describe such a rule. As shown in Appendix A, necessary conditions for optimality are

$$c = u'(b^* - \Delta^*) \quad (4)$$

and

$$c = u'(b^*). \quad (5)$$

These are the smooth-pasting conditions popularized by Dixit (1991, 1993). Theorem 1 proves the existence of a unique solution to equations (3)–(5), and hence a unique optimal rule of this simple type.⁷

Theorem 1 *There exist constants b^* and Δ^* , satisfying $0 < \Delta^* < b^*$, such that equations (4) and (5) are satisfied by the function u defined by (2) and (3). Moreover, this solution is unique.*

We are now in a position to state our main result: No adjustment policy, no matter how complicated, can achieve a lower expected total cost than the simple rule described in Theorem 1.

Theorem 2 *There exists an optimal adjustment policy of the form: hold the target rate constant if $-b^* < \varepsilon < b^*$; reduce the target rate to reset the discrepancy to $-b^* + \Delta^*$ whenever $\varepsilon \leq -b^*$; raise the target rate to reset the discrepancy to $b^* - \Delta^*$ whenever $\varepsilon \geq b^*$.*

For brevity, we call this rule the (b^*, Δ^*) -rule. Figure 1 illustrates this optimal rule with three cases — with just a proportional adjustment cost; with just a fixed adjustment cost, and with both types of adjustment costs. The thinner of the two lines represents a particular evolution of the preferred rate through time; in each case we use the same path. The solid dark line represents the optimal target rate for each case. The shaded region indicates the band around the target rate, in which the preferred rate can move without provoking a target change.

To understand why the optimal rule has this form, consider first the situation in which there is a proportional, but no fixed, adjustment cost. Suppose the preferred rate has moved some small amount away from the current target level. Should the central bank eliminate this

⁷Proofs of Theorems 1 and 2 can be found in Appendix A.

deviation? Since the flow costs are quadratic in the deviation, the increase in flow costs caused by a small deviation will be very small. Moreover, the deviation will get smaller with probability one-half, in which case a costly adjustment will have been unnecessary. If the deviation does get larger, then the flow costs will rapidly increase, while the cost of adjusting the target rate by a given amount remains the same. This suggests there is some critical point beyond which the central bank will want to act. At such times, the target rate will be adjusted an infinitesimal amount in the direction of the preferred rate. Such a policy economizes on costly adjustment costs by avoiding increases in target rates that are likely to be soon followed by decreases. When a fixed adjustment cost is added to the model, the target will be changed by a more substantial amount to economize on this fixed cost.

2.3 Behavior of the Target Rate

The simple targeting rule described above implies particularly simple behavior for the target rate. Except possibly at time 0, ε never lies outside the interval $[-b^*, b^*]$. As long as $-b^* < \varepsilon < b^*$, the target rate is held constant by the central bank, while, as soon as $\varepsilon = b^*$, the central bank raises the target rate by Δ^* . Similarly, the central bank cuts the target rate by Δ^* as soon as $\varepsilon = -b^*$. Therefore, the behavior of the target rate is completely determined by the behavior of ε (and, of course, the initial level of the target rate). The behavior of regulated stochastic processes, such as the one generating ε , is well understood.⁸ We highlight several properties of the target rate which are implied by this model.

Due to the form of the stochastic process generating the preferred rate, combined with the particular functional forms of the flow and adjustment cost functions, the optimal adjustment policy features a great deal of symmetry — the target rate rises (falls) as soon as $\varepsilon = b^*$ ($\varepsilon = -b^*$). In either case, the target rate changes by Δ^* . Therefore, we have:

Property 1 *The magnitude of a change in the target rate does not depend on the direction of the change.*

This symmetry extends to the analysis of a sequence of target changes. We refer to two possibilities when we talk about a policy continuation: a tightening following a tightening and also a loosening following a loosening. One is the mirror-image of the other. Other than the directions of the changes, the properties of the two types of continuations are identical. In particular

Property 2 *The expected time between one type of continuation is the same as the expected time between the other type.*

Similarly, a policy reversal can take two forms — a tightening following a loosening and also a loosening following a tightening. Again, they are mirror-images of one another. We have

Property 3 *The expected time between one type of reversal is the same as the expected time between the other type.*

⁸Standard references are Cox and Miller (1965, Section 5.10) and Karatzas and Shreve (1988, Section 2.8C).

Rudebusch (1995) could not reject the hypothesis that rises and falls in the target rate had the same magnitude on average. He also found evidence supporting the predictions in Properties 2 and 3.

The following two properties capture the stylized fact that the target rate changes in discrete jumps, separated by discrete time intervals. From Theorem 1, $\Delta^* > 0$, giving

Property 4 *When the target rate changes, it changes by a finite amount.*

The next target change occurs when the preferred rate hits one edge of a band around the current target rate. Since, immediately following a target change, the preferred rate is a discrete distance from the edge of the band (from Property 4), this can only occur after some delay. The expected time between target changes, $\Delta^*(2b^* - \Delta^*)/\sigma^2$ from Proposition B-2 in Appendix B, is positive. That is,

Property 5 *Target changes occur only occasionally.*

Consistent with these two properties, Rudebusch found that the average target change was more than twenty basis points and that target changes occurred on only 199 of the 3427 business days.

The following properties relate to the fact that when the target rate is changed, it is moved only part of the way towards the central bank's preferred level. That is, $\Delta^* < b^*$, as shown in Theorem 1, or

Property 6 *Target rate changes are small relative to the amount needed to move the target rate to the preferred level.*

When the preferred rate hits one edge of the band around the target rate, the whole band moves a small amount (Δ^*) as the target rate is adjusted. After the change, the preferred rate is still closer to that edge of the band (a distance of Δ^*) than it is to the other edge (a distance of $2b^* - \Delta^* > \Delta^*$). It is therefore more likely to hit the same edge, and trigger another target change in the same direction, than it is to hit the opposite edge. That is, the probability of a policy continuation, which is $1 - \Delta^*/2b^*$ from Proposition B-1 in Appendix B, is greater than one-half. Thus we have

Property 7 *At the time of a target change, the next change, whenever it occurs, is more likely to be in the same direction than it is to be in the opposite direction.*

Furthermore, the expected time until the preferred rate moves outside the band, conditional on reaching the closer boundary first, is less than the expected time until it moves outside the band, conditional on reaching the more distant boundary.⁹ We have

Property 8 *The expected time between reversals is greater than the expected time between continuations.*

⁹From Proposition B-2, the expected time between continuations is $\Delta^*(4b^* - \Delta^*)/3\sigma^2$ days and the expected time between reversals is $(2b^* - \Delta^*)(2b^* + \Delta^*)/3\sigma^2$ days.

Consistent with Property 6, Rudebusch notes that “target changes occur in small standardized steps.” He also provides evidence that after a target change there is a greater likelihood that the next target change is in the same direction, and that the average time between reversals is greater than that between continuations.

3 Calibrating the Model

In the following sections we calibrate the optimal rule to recent data on federal funds rate target changes for the United States. We use this calibration exercise to derive more properties of the target rate, evaluate the fit of our model, to calculate comparative static results, and to consider some policy implications.

3.1 Calibrating the Adjustment Rule


The behavior of the target rate is completely determined by three parameters: Δ^* , b^* and σ . The information used to calibrate the process for the target level of the federal funds rate is the average absolute value of the change in the target rate ($\hat{\Delta}$), the probability of policy reversals ($\hat{\pi}$) and the average time between target changes (\hat{T}). We set Δ^* equal to $\hat{\Delta}$ and choose b^* and σ so that $\hat{\pi}$ and \hat{T} equal their theoretical counterparts given in equations (10) and (11) in Appendix B, respectively. This calibration is easily shown to be described by

$$\Delta^* = \hat{\Delta}, \quad b^* = \frac{\hat{\Delta}}{2\hat{\pi}}, \quad \sigma = \hat{\Delta} \sqrt{\frac{1 - \hat{\pi}}{\hat{\pi}\hat{T}}}.$$

We calibrate the model to data on federal funds rate target changes for the United States from the Federal Reserve. The data spans March 1, 1984 to March 31, 1998. The period is one where the federal funds rate targets were quite explicit. During this period, the target rate was changed 110 times, with the average change being $\hat{\Delta} = 23.864$ basis points. A proportion $\hat{\pi} = 0.14679$ of those changes were reversals and the average time between changes was $\hat{T} = 30.101$ business days. Therefore, we have $\Delta^* = 23.9$ basis points, $b^* = 81.3$ basis points and $\sigma = 10.5$ basis points, where the unit of time is one business day.

We can check the merits of this calibration by comparing the implied expected time between policy reversals and between policy continuations with the average times observed. From Proposition B-2, conditional on the next change being a reversal, the expected time between target changes is 78.4 days; conditional on it being a continuation, the expected time between target changes is 21.8 days. For the period in question, the average time between reversals was 70.75 business days and between continuations was 23.11 business days. Further encouraging evidence is provided by considering the volatility of the target rate implied by this adjustment rule. We simulate the process for the preferred rate for a period of 10,000 months, approximating Brownian motion by a random walk with 50 draws per day, and calculate the average target rate each month. The standard deviation of the change in this monthly average was found to be 27.10 basis points. Over our sample period, the corresponding standard deviation for the actual federal funds rate was 27.49 basis points.

Figure 2: Hazard functions for changes in the federal funds target rate



The following parameter values were adopted in constructing this figure: $b^* = 81.3$ basis points, $\Delta^* = 23.9$ basis points and $\sigma = 10.5$ basis points. Time is measured in business days.

The dynamic behavior of the target rate predicted by our calibrated model also closely matches observed behavior. Suppose that the Fed raises the target rate at time 0. Define the function $h_c(t)$ such that, conditional on no target changes occurring in the interval $(0, t]$, the Fed will further raise the target rate (a policy continuation) in the interval $(t, t + dt]$ with probability $h_c(t)dt$. Similarly, define the function $h_r(t)$ such that, conditional on no target changes occurring in the interval $(0, t]$, the Fed will reduce the target rate (a policy reversal) in the interval $(t, t + dt]$ with probability $h_r(t)dt$. We plot these hazard functions in Figure 2, using the series expansions given in Appendix B. The bottom curve, which describes $h_r(t)$, shows that in the first few weeks after an increase in the target rate, the Fed is unlikely to reverse its policy. However, the Fed is much more likely to further raise the target rate in the first few weeks after an increase in the target rate, as shown by the behavior of the top curve, which plots $h_c(t)$. After two months, however, the two hazard functions are almost indistinguishable. Therefore, once the time elapsed since the most recent target change becomes large, the direction of the next change is unpredictable.

These properties are likely to be quite general. The likelihood of a target change at any particular time depends on the position of the preferred rate in the band around the target rate at that time. Immediately following a target change, the preferred rate is a distance Δ^* from one boundary and $2b^* - \Delta^* > \Delta^*$ from the other one. If the preferred rate hits the near boundary before it hits the distant one, there will be another target change in the same direction. At all times following a target change, the distribution of the preferred rate will continue to have greater mass close to the near boundary than close to the distant one. Therefore,

Property 9 *At any given time, the probability of a continuation (measured at the time of the last target change) is greater than the probability of a reversal.*

As the time since the most recent target change grows, two things happen to the distribution of the target rate. Firstly, the distribution spreads out, reflecting the fact that, until it reaches one or the other edge of the band around the target rate, the preferred rate evolves according to a Brownian motion. Secondly, conditional on not hitting the edge of the band, the distribution

shifts back towards the middle of the band. This is because it is more likely that the process is in the middle of the band when neither edge has been hit for a long time, than the process is near the edge of the band when neither edge has been hit for a long time. The increased dispersion increases the likelihood of a target change, either a continuation or a reversal, occurring. The shift towards the middle of the band increases the likelihood of a reversal, but reduces the likelihood of a continuation. This explains

Property 10 *The probability of an immediate continuation (measured at the time of the last target change) first increases, then decreases, as the time since the last target change increases.*

and

Property 11 *The probability of an immediate reversal (measured at the time of the last target change) increases as the time since the last target change increases.*

In the limit when the time since the last target change is infinite, the nature of that change (a reduction or increase in the target rate) is irrelevant, and the distribution of the preferred rate is symmetric about the target rate. Therefore,

Property 12 *An immediate continuation and an immediate reversal are equally likely (measured at the time of the last target change) when the time since the most recent target change grows infinitely large.*

3.2 Calibrating the Costs

When constructing the optimal adjustment rule in Section 2.2, we took the cost parameters c and f as given, and solved equations (4) and (5) for b^* and Δ^* . Now, however, we have estimates for b^* and Δ^* and seek cost parameters for which these describe the optimal adjustment rule. This can be achieved by solving equations (4) and (5) for c and f . We obtain

$$c = \frac{2}{\rho} \left(b^* - \Delta^* \frac{\sinh(\lambda b^*)}{\sinh(\lambda b^*) - \sinh(\lambda(b^* - \Delta^*))} \right) \quad (6)$$


and

$$f = \frac{\Delta^*}{\rho} \left(\Delta^* \cdot \frac{\sinh(\lambda b^*) + \sinh(\lambda(b^* - \Delta^*))}{\sinh(\lambda b^*) - \sinh(\lambda(b^* - \Delta^*))} - \frac{2}{\lambda} \cdot \frac{\cosh(\lambda b^*) - \cosh(\lambda(b^* - \Delta^*))}{\sinh(\lambda b^*) - \sinh(\lambda(b^* - \Delta^*))} \right), \quad (7)$$

where $\lambda^2 = 2\rho/\sigma^2$. Substituting $b^* = 81.3$, $\Delta^* = 23.9$ and $\sigma = 10.5$ into these equations, we find that $c = 3908$ and $f = 2859$, where we have assumed that the discount rate is $\rho = 0.01/250$, corresponding to an annual rate of 0.01.¹⁰ Costs implied by the model seem reasonable. For instance, fixed costs represent 2.97% of the total adjustment costs associated with changing the

¹⁰The cost parameters implied by the calibrated adjustment rule are not sensitive to the choice of discount rate. For example, if the Fed is assumed to discount future costs at the rate of $\rho = 0.05/250$, the cost parameters change to $c = 3887$ and $f = 2846$.

Figure 3: Optimal interest rate targeting policy



The top curve in each graph shows the optimal value b^* (measured in basis points) as a function of the indicated variables. The bottom curve shows the optimal value of $b^* - \Delta^*$ (also measured in basis points). The distance between the two curves shows the amount by which the target rate is changed when changes occur. When the difference between the preferred rate and the target rate reaches the level shown by the top curve, it is immediately reset to the level shown by the bottom curve.

target rate by its average amount (23.9 basis points). The average target change incurs total adjustment costs equivalent to a discrepancy of 100 basis points between the preferred and target rates lasting around two weeks, or of 50 basis points lasting around two months.


3.3 Comparative Statics

In this section we consider the effect of changes in cost parameters on the optimal interest rate targeting policy of the Fed and the implications these changes have for the behavior of the target rate.

The first graph in Figure 3 shows the effects of changes in the marginal cost parameter c , holding f and σ constant at their calibrated levels. The top curve plots the optimal bandwidth b^* (measured in basis points) as a function of c . We see that the Fed becomes less tolerant (b^* falls) of deviations between the target rate and the preferred rate as the marginal cost of making target changes falls. The bottom curve plots the optimal level of $b^* - \Delta^*$, the position at which the discrepancy is set when the Fed changes the target rate. Thus, as the marginal cost of target changes falls, the Fed resets the discrepancy closer to zero; equivalently, it moves the target rate closer to the preferred rate. The distance between the two curves shows the amount by which the target rate is adjusted — changes in the target rate increase, but relatively slowly, as the marginal cost falls. In the polar case where the marginal cost of a change in the target rate is zero, it is optimal to set $b^* - \Delta^* = 0$, meaning that the Fed sets the target rate equal to the preferred rate whenever it changes the target rate.

The second graph in Figure 3 repeats this exercise for changes in the fixed cost parameter f , holding c and σ constant at their calibrated levels. As this fixed cost falls, so does the maximum

Figure 4: Probability of policy continuations



Each graph shows the probability of a policy continuation as a function of the indicated variables.

discrepancy allowed by the Fed. The magnitude of target changes falls as the fixed cost falls. Consider the polar case in which there are no fixed costs associated with changing the target rate. In this case, the optimal policy is one of “instantaneous control” (see Harrison, et al. (1983a)), in which the Fed only changes the target rate by the smallest amount necessary to prevent the deviation growing beyond b^* . Most importantly, the target rate will evolve continuously over time. Therefore, if the target rate is to move in discrete jumps, there must be a positive fixed cost associated with every target change.

The third graph in Figure 3 shows that both b^* and Δ^* grow as the preferred rate becomes more volatile.


Figure 4 shows the relative influence of the cost parameters and the volatility of the preferred rate on the degree of persistence evident in the target rate. Each curve represents the probability that the next change in the target rate will be in the same direction as the latest one, calculated at the time of the latest change. As marginal cost falls, so does the probability of policy continuations. In the limiting case when the marginal cost is zero, there is no persistence in the target rate. As the fixed cost falls, the probability of a policy continuation rises, reflecting the fact that lower fixed costs encourage the Fed to make smaller, more frequent, target changes.

In each graph in Figure 5, the three curves indicate the time between different types of target changes. In each case, the lower curve indicates the expected time between continuations and the upper curve represents the expected time between reversals. The middle curve represents the expected time between target changes, regardless of type. In general, we find that the expected time between target changes is increasing in both types of adjustment costs, but is decreasing in the volatility of the preferred rate.¹¹

Figure 5 demonstrates that the extent to which the time between reversals is longer than continuations depends primarily on the marginal cost of target changes. Also of interest is that moderate changes in the volatility of the preferred rate can have a substantial impact on the

¹¹This result, as well as that Δ^* is increasing in volatility, is consistent with the observation of Goodfriend (1991, p. 17), who notes that “...target changes may occur more frequently and step sizes may be bigger in periods of underlying volatility.”

Figure 5: Expected time between target changes



The top curve in each graph shows the expected time (measured in business days) between policy reversals as a function of the indicated variables, while the bottom curve shows the expected time between continuations. The middle curve shows the expected time between target changes, regardless of type.

average time between target changes. As the volatility of the preferred rate decreases, the time between target changes increases sharply.

3.4 Policy Implications

Our comparative statics results allow us to consider the impact of a structural change which shifts one or other of the adjustment costs. As well as affecting the degree of persistence and the time between target changes, adjustment costs also impact on the volatility of changes in the targeted interest rate. In fact, adjustment costs provide a reasonable proxy for the degree of central bank interest rate smoothing.¹² We simulate the process for the preferred rate for a period of 10,000 months, approximating Brownian motion by a random walk with 50 draws per day. We consider 100 different combinations of c and f where fixed costs represent 2.97% of total adjustment costs associated with the optimal adjustment rule. The lowest value of the marginal cost parameter c is set to 100 and the highest is set to 10,000; c increases in increments of 100. For each corresponding optimal rule, we calculate (i) the corresponding path for the target rate; (ii) the total adjustment costs in the central bank’s loss function; (iii) the standard deviation of the change in the monthly average of the target rate. We find the standard deviation in (iii) is monotonically decreasing in the adjustment costs in (ii). Moreover, the correlation coefficient between the standard deviation of interest rates and the adjustment costs is -0.95. This leads to another interpretation of the costs of target changes — they reflect the central bank’s interest rate smoothing motives. To the extent this is true, it suggests a novel way to measure how different central banks compare in terms of their desire to smooth interest rates. Using our

¹²We take interest rate smoothing to mean reducing the volatility of targeted interest rates for a given volatility in the preferred rate.

model, adjustment costs can be backed-out from data on target changes.¹³

Given that adjustment costs are highly correlated with the degree of interest rate smoothing, an interesting question is whether our (b^*, Δ^*) -rule would still be optimal if the costs of adjusting target rates in the central bank's loss function, are replaced by an interest rate smoothing objective. Using the standard deviation of the target rate in the loss function is likely to make the problem of working out the central bank's optimal rule analytically intractable. Instead, starting from the calibrated rule, we examine whether there are other rules of the (b^*, Δ^*) -type which allow for more interest rate smoothing, without increasing flow costs. We find that, at best, the central bank can reduce the standard deviation of the change in the monthly average of the target rate by less than one basis point. Thus, at least within the class of simple rules considered, the rule we found from calibrating our model remains (approximately) optimal even when the bank's loss function is altered to allow for a desire to smooth interest rates (rather than avoid target changes). Interestingly, the rule which minimizes the standard deviation in interest rate changes is the one in which Δ is as close to zero as possible. Thus with an interest rate smoothing objective, target changes should not be discrete jumps. This suggests the observed discreteness in target changes is better explained by an adjustment cost framework, in which there is a fixed cost of each target change, than a pure interest rate smoothing objective.

If we do interpret adjustment costs as proxying interest rate smoothing concerns, one can use our model to ask whether these concerns are justifiable. If the Fed allowed the standard deviation of the change in the monthly average of the target rate to increase by approximately 13 basis points (from 27.1 basis points to 40.3 basis), it could entirely eliminate the discrepancy between preferred and target rates (a reduction of 35 basis points in its average level). On the other hand, to reduce the same standard deviation by 10 basis points, requires an increase in the average deviation between the preferred rate and the target rate of just over 100 basis points. Whether the Fed is behaving optimally, given this trade-off, then depends on a judgement about the relative importance of the flow costs arising from compromising their ultimate objectives and the adjustment costs that proxy interest rate smoothing concerns.

If adjustment costs are not just proxying interest rate smoothing concerns, then policies which can lower the adjustment costs of target changes will be beneficial from the central bank's perspective. One policy which may lower adjustment costs is to specify regular dates when announcements will be made on the desired target rate. Since the timing of these announcements is known in advance, the likelihood of an adverse market reaction on these dates should be less than from an equivalent target change made at an unexpected time. We could model this as a reduction in the adjustment costs of target changes on these days. In this case, it is unlikely our simple rule will still be optimal. Instead, the optimal rule may well involve the band around the target rate becoming wider as the date of each announcement approaches (so a target change is very unlikely when an announcement date is near), and narrowing on the date of the

¹³Simply measuring the volatility of actual interest rates does not suffice. Different countries will face different levels of volatility in macroeconomic variables and so different levels of volatility in their preferred rates. Our model allows this factor to be disentangled from the costs of adjusting interest rates.

announcement (so a target change on an announcement date is much more likely).

An alternative view of adjustment costs is that they reflect political and reputational considerations. In this case, institutional reform (such as operational independence) which reduces these costs could lead to an improvement in the ability of the central bank to meet its ultimate objectives. Goodhart (1997) appears to take this alternative view, suggesting that authorities are not able to vary interest rates sufficiently aggressively, or promptly, to hold inflation to a desired path, because changes in interest rates that preempt future inflationary pressures and that may be quickly reversed are not politically viable. One difficulty with the political costs argument, although not with reputational considerations, is it is difficult to see how this can be reconciled with the symmetry observed in the data on target changes. Political costs are usually associated with constraints to raising interest rates, but not to reducing them. A reputational story for adjustment costs raises the possibility that the central bank will want to be secretive regarding its target changes, or be deliberately vague in justifying them. However, such secrecy may lead to other costs, such as greater uncertainty on the part of investors, which could increase the volatility of market interest rates.

4 Conclusions and Future Directions

This paper was motivated by the observation that interest rate target changes are highly persistent and seldom quickly reversed. According to Goodfriend, this behavior represents a deliberate attempt by the Fed to smooth interest rates and avoid “whipsawing” the market. An alternative view has been put forward by Goodhart (1997), who suggests that central banks “act too little, too late”, thus sacrificing their ultimate objectives. In our model, interest rate targeting does sacrifice a central bank’s ultimate objectives; for the Fed, we found that if it moves the target rate to immediately reflect any movement in its preferred rate, it could reduce the average deviation between the two rates by 35 basis points. However, using a simple model of optimal central bank policy, we showed that given modest costs to adjusting target rates, the persistent behavior of target changes could be rationalized.¹⁴ Interestingly, to get this result we did not need to assume persistence in the underlying shock in our model; nor did we have to assume that reversals were especially costly. Presumably, it is even easier to rationalize the observed behaviour of target rates in the presence of these factors. Instead, in our model the optimal rule follows from uncertainty over future movements in the underlying preferred rate, together with the costs of adjusting the target rate. When the preferred rate is close to the target rate, the Fed is better off to delay a target change, rather than risk unsettling the market, especially since many times small changes in the preferred rate will be undone by themselves.

The starting point for this paper was a simple exogenous stochastic process for the central bank’s preferred rate. In practice, the preferred level of interest rates could itself depend on the target rate policy chosen. Presumably, different interest rate targeting policies will affect the fundamentals of the economy and so feed back into the preferred level of interest rates.

¹⁴Using this model we were also able to explain a number of other puzzling properties of interest rate targeting.

We have assumed away any such feedback effect to make our model tractable. Future work could incorporate such a feedback effect for a specific macroeconomic model of the economy and examine numerically the extent to which the new optimal rule would differ.

Clearly, our assumption that the preferred rate evolves according to driftless Brownian motion is a special case of more general stochastic processes for the preferred rate. However, we believe it is the most interesting case to start with, both because it is difficult to reject a unit root in short-term interest rates and because it shows that one does not need to assume mean-reversion in the central bank's preferred rate to generate persistence in target changes. Nevertheless, it would still be interesting to extend the analysis to consider other stochastic processes for the preferred rate, such as the process for interest rates derived from the policy reaction function of a central bank's model of the economy.

Because our model characterizes the dynamics of the central bank's target rate through time, it also has implications for market interest rates at the short-end of the yield curve. Using the expectations hypothesis, Guthrie and Wright (1999) derive the implications of our calibrated (b^*, Δ^*) -rule for daily observations on one- and three-month market rates. Despite the simplicity of our assumptions, they find a rich array of interest rate dynamics can be derived. Consistent with observed U.S. rates, they show the model implies conditional volatility is persistent and increasing in the spread between the market rate and the central bank's target rate, that there is excess Kurtosis in high frequency interest rate movements, and that market rates will revert towards the central bank's target rate.

Another area where the model seems particularly applicable is in modeling exchange rate targeting. The preferred level of the nominal exchange rate naturally corresponds to the preferred level of interest rates in this paper. The Brownian motion assumption for the preferred level would again be the natural choice. Like interest rate changes, the central bank may view changes in the exchange rate as costly, both in proportion to the size of the exchange rate change, as well as for changes per se. To the extent fixed adjustment costs are more important for exchange rate target changes than for interest rate target changes, our comparative statics results suggest less persistence in, and longer times between, target changes.

A Proofs for Optimal Adjustment Rule

Necessary Conditions for an Optimal Policy


We seek values of b and Δ which minimize A , where A is related to the two choice variables by the requirement that u is continuous at b :

$$u(b) = u(b - \Delta) + f + c\Delta$$

or

$$0 = \frac{b^2}{\rho} + A \cosh(\lambda b) - \frac{(b - \Delta)^2}{\rho} - A \cosh(\lambda(b - \Delta)) - f - c\Delta.$$

Figure 6: The function $v'(\varepsilon; \hat{\varepsilon})$



The Lagrangian for this problem is

$$\mathcal{L} = A - \mu \left(\frac{b^2}{\rho} + A \cosh(\lambda b) - \frac{(b - \Delta)^2}{\rho} - A \cosh(\lambda(b - \Delta)) - f - c\Delta \right),$$

where μ is the Lagrange multiplier, and the appropriate first order conditions are

$$\begin{aligned} 0 &= 1 - \mu (\cosh(\lambda b^*) - \cosh(\lambda(b^* - \Delta^*))), \\ 0 &= -\mu \left(\frac{2b^*}{\rho} + \lambda A \sinh(\lambda b^*) - \frac{2(b^* - \Delta^*)}{\rho} - \lambda A \sinh(\lambda(b^* - \Delta^*)) \right), \\ 0 &= -\mu \left(\frac{2(b^* - \Delta^*)}{\rho} + \lambda A \sinh(\lambda(b^* - \Delta^*)) - c \right). \end{aligned}$$

The second and third conditions become $u'(b^*) = u'(b^* - \Delta^*)$ and $u'(b^* - \Delta^*) = c$, respectively.

Proof of Theorem 1¹⁵

Let $\hat{\varepsilon}$ be an arbitrary positive constant and define the function

$$v(\varepsilon; \hat{\varepsilon}) = \frac{\sigma^2}{\rho^2} + \frac{\varepsilon^2}{\rho} + A(\hat{\varepsilon}) \cosh(\lambda\varepsilon),$$

of ε , where

$$A(\hat{\varepsilon}) = \frac{-\sigma^2}{\rho^2 \cosh(\lambda\hat{\varepsilon})}.$$

Notice that, for a particular value of the integration constant A , v equals the function u , given in (2), on the interval $(-b, b)$, but extends its functional form to the whole real line. The function

$$v'(\varepsilon; \hat{\varepsilon}) = \frac{2\varepsilon}{\rho} + \lambda A(\hat{\varepsilon}) \sinh(\lambda\varepsilon)$$

has turning points at $\varepsilon = \pm\hat{\varepsilon}$. It is drawn in Figure 6. The value of this function at the turning point $\varepsilon = \hat{\varepsilon}$ is

$$\Gamma(\hat{\varepsilon}) = v'(\hat{\varepsilon}, \hat{\varepsilon}) = \frac{2\hat{\varepsilon}}{\rho} + \lambda A(\hat{\varepsilon}) \sinh(\lambda\hat{\varepsilon}).$$

¹⁵The proof of this and the following theorem is based on Constantinides and Richard (1972), who consider a similar problem, with different cost functions, in the context of inventory management.

It is easily shown that Γ is an increasing function of $\hat{\varepsilon}$, with $\Gamma(0) = 0$ and $\Gamma(\hat{\varepsilon}) \rightarrow \infty$ as $\hat{\varepsilon} \rightarrow \infty$. In fact,

$$\Gamma'(\hat{\varepsilon}) = \frac{2}{\rho} \left(1 - \frac{1}{\cosh^2(\lambda\hat{\varepsilon})} \right) > 0.$$

Now, for any value of the parameter $\hat{\varepsilon}$ such that $\Gamma(\hat{\varepsilon}) > c$, the numbers $b(\hat{\varepsilon})$ and $\Delta(\hat{\varepsilon})$ are uniquely determined by the analogs of equations (4) and (5),

$$v'(b(\hat{\varepsilon})) = c \quad \text{and} \quad v'(b(\hat{\varepsilon}) - \Delta(\hat{\varepsilon})) = c,$$

together with the requirement that $\Delta(\hat{\varepsilon}) > 0$. The function u is continuous at $\varepsilon = b^*$ if

$$v(b^*) = v(b^* - \Delta^*) + f + c\Delta^*.$$

Thus, in order to prove existence, we must show that there exists a number $\hat{\varepsilon}^*$ such that $\Psi(\hat{\varepsilon}^*) = f$, where

$$\Psi(\hat{\varepsilon}) = v(b(\hat{\varepsilon})) - v(b(\hat{\varepsilon}) - \Delta(\hat{\varepsilon})) - c\Delta(\hat{\varepsilon}).$$

That is, the shaded region in Figure 6 must have area f . Now, if $\varepsilon \rightarrow \Gamma^{-1}(c)$, then $b(\hat{\varepsilon}) \rightarrow \hat{\varepsilon}$ and $\Delta(\hat{\varepsilon}) \rightarrow 0$, so that $\Psi(\hat{\varepsilon}) \rightarrow 0$ and the shaded region vanishes. On the other hand, as $\hat{\varepsilon} \rightarrow \infty$, we see that $b(\hat{\varepsilon}) \rightarrow \infty$ and $b(\hat{\varepsilon}) - \Delta(\hat{\varepsilon}) \rightarrow \rho c/2$. Since $\Gamma(\hat{\varepsilon}) \rightarrow \infty$, the area of the shaded region becomes infinitely large. Appealing to the Intermediate Value Theorem, we see that there must exist some value of $\hat{\varepsilon}$, call it $\hat{\varepsilon}^*$, for which the shaded region has area f . The required policy parameters are $b^* = b(\hat{\varepsilon}^*)$ and $\Delta^* = \Delta(\hat{\varepsilon}^*)$. It is easy to prove that Ψ is a strictly increasing function of $\hat{\varepsilon}$. These parameters are therefore unique, and the proof is complete.

Proof of Theorem 2

Our proof of the optimality of this simple rule uses the following lemma, which gives sufficient conditions for a policy to be optimal. A proof can be found in Harrison, et al. (1983b).

Lemma 1 *Suppose that u is continuously differentiable, has a bounded derivative, and has a continuous second derivative at all but a finite number of points. If*


$$\begin{aligned} u(\varepsilon) &\leq C(\varepsilon' - \varepsilon) + u(\varepsilon'), && \text{for all } \varepsilon \text{ and } \varepsilon', \\ 0 &\leq F(\varepsilon) + \frac{1}{2}\sigma^2 u''(\varepsilon) - \rho u(\varepsilon), && \text{for almost all } \varepsilon, \end{aligned}$$

then $u(\varepsilon) \leq J(P; \varepsilon)$ for all adjustment policies P and all $\varepsilon \in \mathbb{R}$.

We show that the expected total cost function $u(\varepsilon)$ associated with the adjustment policy constructed in the proof of Theorem 1 satisfies the conditions of Lemma 1. Notice that $u(\varepsilon) = v(\varepsilon; \hat{\varepsilon}^*)$ whenever $-b^* < \varepsilon < b^*$. For other values of ε , $u(\varepsilon)$ is given by the expressions in (2). Figure 7 plots $u'(\varepsilon)$ for the adjustment policy constructed in the proof of Theorem 1.

The function u has a continuous first derivative, which is bounded. Furthermore, it has a continuous second derivative everywhere, except at the points $\varepsilon = \pm b^*$. The regularity conditions of Lemma 1 are therefore satisfied.

Figure 7: The function $u'(\varepsilon)$



If $\varepsilon \leq -b^*$ (respectively $\varepsilon \geq b^*$), it is optimal to change the target in order to bring the discrepancy back to $-b^* + \Delta^*$ (respectively $b^* - \Delta^*$). Therefore, for all ε' ,

$$u(\varepsilon') + C(\varepsilon' - \varepsilon) \geq u(-b^* + \Delta^*) + C(-b^* + \Delta^* - \varepsilon) = u(\varepsilon), \quad \varepsilon \leq -b^*,$$

and

$$u(\varepsilon') + C(\varepsilon' - \varepsilon) \geq u(b^* - \Delta^*) + C(b^* - \Delta^* - \varepsilon) = u(\varepsilon), \quad \varepsilon \geq b^*.$$

If $-b^* < \varepsilon \leq -b^* + \Delta^*$ (respectively $b^* - \Delta^* \leq \varepsilon < b^*$), and the central bank decided to change the target rate, it would do so in such a way that the discrepancy is reset to $-b^* + \Delta^*$ (respectively $b^* - \Delta^*$), since this minimizes the total expected cost after the change. Therefore, for all ε' ,

$$u(\varepsilon') + C(\varepsilon' - \varepsilon) \geq u(-b^* + \Delta^*) + C(-b^* + \Delta^* - \varepsilon) > u(\varepsilon), \quad -b^* < \varepsilon \leq -b^* + \Delta^*,$$

and

$$u(\varepsilon') + C(\varepsilon' - \varepsilon) \geq u(b^* - \Delta^*) + C(b^* - \Delta^* - \varepsilon) > u(\varepsilon), \quad b^* - \Delta^* \leq \varepsilon < b^*.$$

The remaining case to consider is where $-b^* + \Delta^* < \varepsilon < b^* - \Delta^*$. If the central bank decided to change the target rate, the cost-minimizing action would be to change it by zero, since the marginal cost of changing the target rate exceeds the marginal benefit of reducing the discrepancy. Therefore, for all ε' ,

$$u(\varepsilon') + C(\varepsilon' - \varepsilon) \geq u(\varepsilon) + C(0) > u(\varepsilon), \quad -b^* + \Delta^* < \varepsilon < b^* - \Delta^*.$$

Combining these results, we see that u satisfies the first inequality in Lemma 1. Notice, also, that the central bank should only ever change the target rate when the discrepancy is outside the interval $(-b^*, b^*)$.

Let

$$\theta(\varepsilon) = \varepsilon^2 + \frac{1}{2}\sigma^2 u''(\varepsilon) - \rho u(\varepsilon).$$

It is easily confirmed that $\theta(\varepsilon) = 0$ whenever $-b^* < \varepsilon < b^*$. Thus

$$(b^*)^2 - \rho u(b^*-) = \frac{-1}{2}\sigma^2 u''(b^*-) > 0,$$

since $u''(b^*-) < 0$, and, since $u''(b^*+) = 0$, it follows that

$$\theta(b^*+) = (b^*)^2 - \rho u(b^*+) = (b^*)^2 - \rho u(b^*-) > 0.$$

By a similar argument, $\theta(-b^*-) > 0$. Next, notice that whenever $\varepsilon > b^*$ (respectively $\varepsilon < -b^*$), $\theta'(\varepsilon) = 2\varepsilon - \rho c > 2b^* - \rho c > 0$ (respectively $\theta'(\varepsilon) < 0$). It follows that $\theta(\varepsilon) > \theta(b^*+) > 0$ whenever $\varepsilon > b^*$ and that $\theta(\varepsilon) > \theta(-b^*-) > 0$ whenever $\varepsilon < -b^*$. Combining these results, we see that u satisfies the second inequality in Lemma 1 and the proof is complete.

B Proofs for Behavior of the Target Rate

Denote by $g_-(t|\varepsilon_0)$ the probability density of ε reaching $-b^*$ at time t before having reached b^* , conditional on the discrepancy having the value ε_0 at time 0. Similarly, denote by $g_+(t|\varepsilon_0)$ the probability density of passing into b^* at time t before having reached $-b^*$, conditional on the same initial value. Standard techniques can be used to derive the moment-generating functions

$$g_-^*(s|\varepsilon_0) = \int_0^\infty e^{-st} g_-(t|\varepsilon_0) dt = \frac{\sinh(\beta(b^* - \varepsilon_0))}{\sinh(2\beta b^*)} \quad (8)$$

and

$$g_+^*(s|\varepsilon_0) = \int_0^\infty e^{-st} g_+(t|\varepsilon_0) dt = \frac{\sinh(\beta(b^* + \varepsilon_0))}{\sinh(2\beta b^*)}, \quad (9)$$

where $\beta = \sqrt{2s}/\sigma$. These functions enable us to prove the two propositions below.

Proposition B–1 *Measured at the time of a target change, the probability that the next target change, whenever it occurs, is in the same direction equals*

$$1 - \frac{\Delta^*}{2b^*}. \quad (10)$$

PROOF Without loss of generality, suppose that the central bank increases the target rate at time 0; that is, set $\varepsilon_0 = b^* - \Delta^*$. The next target change, whenever it occurs, will be another increase if ε hits b^* before it hits $-b^*$. This occurs with probability

$$\int_0^\infty g_+(t|b^* - \Delta^*) dt = 1 - \frac{\Delta^*}{2b^*},$$

where we have used (9) to calculate

$$\int_0^\infty g_+(t|\varepsilon_0) dt = g_+^*(0|\varepsilon_0) = \frac{b^* + \varepsilon_0}{2b^*}.$$

Since $\Delta^* < b^*$, this probability is greater than 1/2. ■

Proposition B–2 *Consider successive target changes.*

1. Conditional on successive target changes being in the same direction, the expected time between them equals $\Delta^*(4b^* - \Delta^*)/3\sigma^2$ days.

2. Conditional on them being in opposite directions, the expected time is $(2b^* - \Delta^*)(2b^* + \Delta^*)/3\sigma^2$ days.

3. The unconditional mean is

$$\Delta^*(2b^* - \Delta^*)/\sigma^2 \quad (11)$$

days.

PROOF From (8) and (9), we find that

$$\begin{aligned} \int_0^\infty tg_-(t|\varepsilon_0)dt &= -\frac{d}{ds}g_-^*(s|\varepsilon_0)\Big|_{s=0} = \frac{(b^* - \varepsilon_0)(b^* + \varepsilon_0)(3b^* - \varepsilon_0)}{6b^*\sigma^2}, \\ \int_0^\infty tg_+(t|\varepsilon_0)dt &= -\frac{d}{ds}g_+^*(s|\varepsilon_0)\Big|_{s=0} = \frac{(b^* - \varepsilon_0)(b^* + \varepsilon_0)(3b^* + \varepsilon_0)}{6b^*\sigma^2}. \end{aligned}$$

Without loss of generality, suppose that the central bank increases the target rate at time 0; that is, set $\varepsilon_0 = b^* - \Delta^*$. Then:

1. The expected time until the next target change, conditional on that change being another increase in the target rate, equals

$$\frac{\int_0^\infty tg_+(t|b^* - \Delta^*)dt}{\int_0^\infty g_+(t|b^* - \Delta^*)dt} = \frac{\Delta^*(4b^* - \Delta^*)}{3\sigma^2}$$

days.

2. The expected time until the next target change, conditional on that change being a reduction in the target rate, equals

$$\frac{\int_0^\infty tg_-(t|b^* - \Delta^*)dt}{\int_0^\infty g_-(t|b^* - \Delta^*)dt} = \frac{(2b^* - \Delta^*)(2b^* + \Delta^*)}{3\sigma^2}$$

days.

3. The unconditional mean time between target changes equals

$$\int_0^\infty t(g_-(t|b^* - \Delta^*) + g_+(t|b^* - \Delta^*))dt = \frac{\Delta^*(2b^* - \Delta^*)}{\sigma^2}$$

days. ■

Suppose that the central bank raises the target rate at time 0. Define the function $h_c(t)$ such that, conditional on no target changes occurring in the interval $(0, t]$, the central bank will further raise the target rate (a policy continuation) in the interval $(t, t + dt]$ with probability $h_c(t)dt$. Similarly, define the function $h_r(t)$ such that, conditional on no target changes occurring in the interval $(0, t]$, the central bank will reduce the target rate (a policy reversal) in the interval $(t, t + dt]$ with probability $h_r(t)dt$. These so-called hazard functions are

$$h_c(t) = \frac{g_+(t|b^* - \Delta^*)}{1 - \int_0^t (g_-(t'|b^* - \Delta^*) + g_+(t'|b^* - \Delta^*))dt'}$$

$$h_r(t) = \frac{g_-(t|b^* - \Delta^*)}{1 - \int_0^t (g_-(t'|b^* - \Delta^*) + g_+(t'|b^* - \Delta^*)) dt'}.$$

The probability densities $g_-(t|\varepsilon_0)$ and $g_+(t|\varepsilon_0)$ can be found by calculating the inverse Laplace transforms of (8) and (9). They are

$$g_-(t|\varepsilon_0) = \frac{1}{\sqrt{2\pi\sigma^2 t^3}} \sum_{n=-\infty}^{\infty} (4nb^* + b^* + \varepsilon_0) \exp \left\{ -\frac{(4nb^* + b^* + \varepsilon_0)^2}{2\sigma^2 t} \right\}$$

and

$$g_+(t|\varepsilon_0) = \frac{1}{\sqrt{2\pi\sigma^2 t^3}} \sum_{n=-\infty}^{\infty} (4nb^* + b^* - \varepsilon_0) \exp \left\{ -\frac{(4nb^* + b^* - \varepsilon_0)^2}{2\sigma^2 t} \right\}$$

respectively.

Equations (10) and (11) are used to calibrate the model. From equation (10), the probability of a policy reversal is

$$\pi = \frac{\Delta^*}{2b^*},$$

while, from (11), the expected time between target changes is

$$T = \frac{\Delta^*(2b^* - \Delta^*)}{\sigma^2}.$$

Solving these two equations for b^* and σ gives

$$b^* = \frac{\Delta^*}{2\pi}, \quad \sigma = \Delta^* \sqrt{\frac{1-\pi}{\pi T}}.$$

Replacing Δ^* with the average absolute value of the change in the target rate ($\hat{\Delta}$), π with the proportion of target changes which are policy reversals ($\hat{\pi}$), and T with the average time between target changes (\hat{T}), gives the calibration used in Section 3.1.

References

- BANK FOR INTERNATIONAL SETTLEMENTS. *68th Annual Report*. Basle, Switzerland: Bank for International Settlements, 1998.
- BARRO, ROBERT. "On the Determination of the Public Debt," *Journal of Political Economy*, 1979, 87, pp. 940–971.
- _____. "Interest-Rate Targeting," *Journal of Monetary Economics*, 1989, 23, pp. 3–30.
- CONSTANTINIDES, GEORGE M. AND RICHARD, SCOTT F. "Existence of Optimal Simple Policies for Discounted-Cost Inventory and Cash Management in Continuous Time," *Operations Research*, 1978, 26, pp. 620–636.
- COX, D. R. AND MILLER, H. D. *The theory of stochastic processes*. London: Chapman and Hall, 1965.

- CUKIERMAN, ALEX. “Why does the Fed Smooth Interest Rates?” in *Monetary Policy on the 75th Anniversary of the Federal Reserve System*. Ed. MICHAEL BELONGIA. Norwell, Mass. and Dordrecht: Kluwer Academic Publishers, 1991, pp. 111–147.
- DIXIT, AVINASH. “A Simplified Treatment of the Theory of Optimal Regulation of Brownian Motion,” *Journal of Economic Dynamics and Control*, 1991, 15, pp. 657–673.
- _____. *The art of smooth pasting*, Vol. 55 in *Fundamentals of pure and applied economics*. Eds. JACQUES LESOURNE AND HUGO SONNENSCHNEIN. Chur, Switzerland: Harwood Academic Publishers, 1993.
- DOTSEY, MICHAEL. “Monetary Policy, Secrecy, and Federal Funds Rate Behavior,” *Journal of Monetary Economics*, 1987, 20, pp. 463–474.
- GOODFRIEND, MARVIN. “Monetary Mystique: Secrecy and Central Banking,” *Journal of Monetary Economics*, 1986, 17, pp. 63–92.
- _____. “Interest Rate Smoothing and Price Level Trend-Stationarity,” *Journal of Monetary Economics*, 1987, 19, pp. 335–348.
- _____. “Interest Rates and the Conduct of Monetary Policy,” *Carnegie-Rochester Conference Series on Public Policy*, 1991, 34, pp. 7–30.
- GOODHART, CHARLES. “Why do the Monetary Authorities Smooth Interest Rates?” in *European Monetary Policy*. Ed. S. COLLIGNON. London, 1997, Chapter 8, pp. 119–174.
- GUTHRIE, GRAEME AND WRIGHT, JULIAN. “Interest Rate Dynamics Under Fed-Funds Rate Targeting,” 1999, mimeo, University of Auckland.
- HARRISON, J. MICHAEL; SELLKE, THOMAS M. AND TAYLOR, ALLISON J. “Instantaneous Control of Brownian Motion,” *Mathematics of Operations Research*, 1983a, 8(3), pp. 439–453.
- _____. “Impulse Control of Brownian Motion,” *Mathematics of Operations Research*, 1983b, 8(3), pp. 454–466.
- KARATZAS, IOANNIS AND SHREVE, STEVEN E. *Brownian motion and stochastic calculus*. New York: Springer-Verlag, 1988.
- MANKIW, GREGORY N. “The Optimal Collection of Seigniorage: Theory and Evidence,” *Journal of Monetary Economics*, 1987, 20, pp. 327–341.
- ORPHANIDES, ATHANASIOS. “Monetary Policy Evaluation with Noisy Information,” *Finance and Economics Discussion Series*, 1998, No. 98-3, Board of Governors of the Federal Reserve System.
- RUDEBUSCH, GLENN D. “Federal Reserve Interest Rate Targeting, Rational Expectations, and the Term Structure,” *Journal of Monetary Economics*, 1995, 35 (2), pp. 245–274.

SACK, BRIAN. "Uncertainty, Learning and Gradual Monetary Policy," *Finance and Economics Discussion Series*, 1998, No. 98-34, Board of Governors of the Federal Reserve System.

SACK, BRIAN AND WIELAND, VOLKER. "Interest-Rate Smoothing and Optimal Monetary Policy: A Review of Recent Empirical Evidence," *Finance and Economics Discussion Series* 1999, No. 99-39, Board of Governors of the Federal Reserve System.

WOODFORD, MICHAEL. "Optimal Monetary Policy Inertia," 1999, mimeo, Princeton University.