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, as well as explain some dynamic properties of market interest rates.

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

*We thank Marvin Goodfriend, Glenn Rudebusch, Carl Walsh, Michael Woodford, and two anonymous referees for helpful comments, and Shirley Huang for research assistance. Any errors are ours.

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

are essentially unpredictable at forecast horizons longer than a month or two. Rudebusch (1995) 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, 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 starting 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.[1] 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).[2] We start by taking 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

[1] Eijffinger et al. (1999) use a simpler version of our model in the context of a small macro model to derive implications of discrete interest rate changes for the macro-economy. However, their model cannot explain most of the interesting dynamics of target rates because it has the property that the inflation (or interest rate) gap from the target is reset to zero once it reaches a critical level. A recent literature, surveyed by Sack and Wieland (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.

[2] In the recent inflation targeting literature, the overnight interbank interest rate is often considered the instrument which is used to try to achieve the bank's ultimate targets (or objectives). In this paper, the term 'target' refers to an operating target — the bank sets a target level of interest rates which it attempts to maintain through the use of open market operations.

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 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 when the central bank chooses the size and timing of target changes to minimize the expected 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, because doing so preserves a valuable option, the option to wait and see what happens. Many times small changes in the preferred rate will be undone by themselves.

Despite the simplicity of the optimal rule, the model has a number of predictions for the dynamics of the central bank's target rate and market interest rates. On the one hand, because of fixed costs, the central bank will take advantage of economies of scale by changing the target rate by discrete amounts. On the other hand, because of proportional adjustment costs and quadratic flow costs, the central bank will respond to the diminishing returns to moving the target rate towards the preferred rate by moving the target rate only part of the way towards the preferred rate. The combination of these two features implies that successive target change are likely to be in the same direction (persistence), the average time between reversals is greater

[3]In 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.

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, our model predicts reversals become more likely as the time elapsed since the last target change becomes large. In the limit, reversals and continuations become equally likely, so that the direction of the next target change becomes unpredictable.

When combined with the expectations hypothesis, the optimal targeting rule also has implications for the dynamics of market interest rates. We show that our model implies market rates will revert towards the central bank's target rate as the time since the last target change increases, and that the conditional volatility of market rates is increasing in the spread between the market rate and the central bank's target rate.

We test the different implications of the model using daily data on U.S. interest rates and find considerable support for our model's predictions. We also estimate hazard functions for continuations and reversals, showing they share many of the properties of the hazard functions predicted by our model.

Several extensions to our basic framework are considered. These involve allowing for predictable movements in the central bank's preferred rate, so that there is a business cycle component driving required interest rate adjustments, allowing for the fact that target changes which are in the opposite direction to that anticipated may be especially costly, so that reversals are more costly than continuations, and allowing for regular announcement dates on which the central bank prefers to make its target rate changes, so that surprise announcements cost more.

The paper is organized as follows. Section 2 presents our model of interest rate targeting, deriving the optimal targeting rule. A number of properties of the implied target rate, together with market rates, are derived in Section 3. These implications are tested in Section 4, while Section 5 considers a number of extensions to our basic framework. Conclusions are drawn in Section 6.

2 A Model of Interest Rate Targeting

This section derives the optimal interest rate targeting rule in the context of our benchmark model. Subsequent sections analyze and test the properties of this rule, as well as considering changes to our basic model. Section 2.1 lays out our model of interest rate targeting, motivating each of the assumptions, while Section 2.2 derives the optimal form of interest rate targeting.

2.1 The Model Set-up

Consider a central bank which targets a particular interest rate. We are interested in when and by how much it changes its target rate. We start by supposing the central 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. Over the relatively short time between target changes, the random walk assumption is probably a reasonable modelling approximation for working out the optimal rule. An alternative to the random walk approach, motivated by the observation that target changes themselves appear to be positively autocorrelated, would be to allow gradual mean-reversion in the preferred rate. Our approach shows such mean-reversion is not needed to explain the observed persistence in target changes. However, in Section 5.1 we also consider how trends in the preferred rate affect the optimal rule.

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 cannot 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 the interest rate deviates from its target level 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 $C(\Delta) = f + k|\Delta|$. 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 $k|\Delta|$, 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. 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 marginal adjustment cost is constant at k. 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 $T_1 \leq T_2 \leq \cdots \leq T_i \leq \cdots$ 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

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

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.[4] When the preferred rate 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

[4] The 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 $\pm(b - \Delta)$.

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 $dt \to 0$, and dropping the time subscript, shows that

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

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

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

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

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

Combining these three results, we see that

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

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

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

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

Figure 1: Implementing the optimal adjustment rule

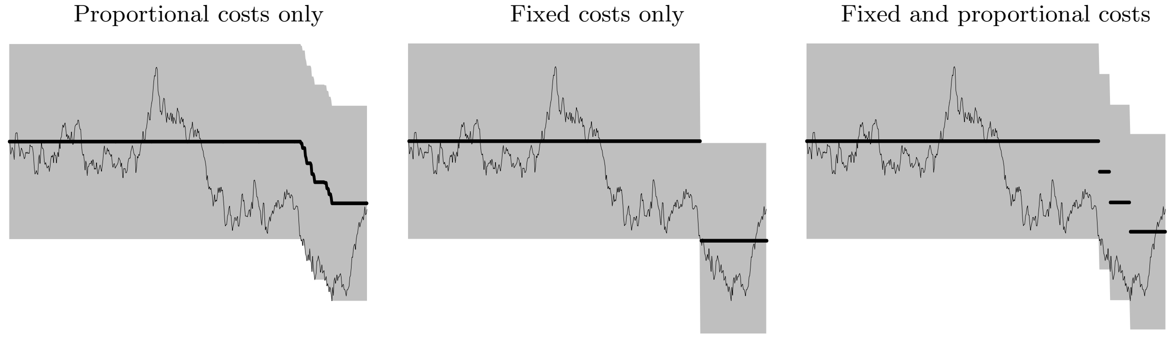

be minimized by choosing the policy parameters which minimize A. Let the parameters b^* and Δ^* describe such a rule. As shown in Appendix A, necessary conditions for optimality are

$$u'(b^* - \Delta^*) = u'(b^*) = k. \tag{4}$$

These are the smooth-pasting conditions popularized by Dixit (1993). Theorem 1 proves the existence of a unique solution to equations (3) and (4), and hence a unique optimal rule of this simple type.[5]

Theorem 1 *There exist constants b^* and Δ^*, satisfying $0 < \Delta^* < b^*$, such that equations (4) 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 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 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

[5]Proofs of Theorems 1 and 2 can be found in Appendix A.

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.

3 Implications of Optimal Targeting

The simple nature of the optimal targeting rule allows us to derive properties for the dynamic behavior of the target rate, as well as market rates. Section 3.1 derives a number of testable properties of the target rate. In Section 3.2 the optimal rule is calibrated to recent data on federal funds rate target changes for the United States and the hazard function for changes in the target rate is derived. Making use of the expectations hypothesis, some implications for market interest rates are derived in Section 3.3.[6]

3.1 Properties 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.[7] 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

[6] In an earlier version of this paper (Guthrie and Wright, 1999) we showed how the optimal targeting rule, the probability of policy continuations, and the expected time between target changes varied with respect to the underlying adjustment costs and volatility parameters.

[7] A standard reference is Karatzas and Shreve (1991, Section 2.8C).

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 taken for one type of continuation is the same as the expected time taken for 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 taken for one type of reversal is the same as the expected time taken for the other type.*

Because $\Delta^* < b^*$, as shown in Theorem 1, after a target 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 4 *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.[8] We have

Property 5 *The expected time taken for reversals is greater than the expected time taken for continuations.*

3.2 The Calibrated Hazard Function

We can derive further properties of the dynamics of the target rate by calibrating the optimal rule to data on federal funds rate target changes for the United States. 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 proportion of policy reversals ($\hat{\pi}$) and the average time between target changes ($\hat{T}$). These three quantities completely determine the three parameters (Δ^*, b^*, and σ) which govern the behavior of the target rate. 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 Propositions B–1 and B–2 in Appendix B, respectively. This calibration is easily shown to be described by

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

[8] From 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.

We calibrate the model to data on federal funds rate target changes for the United States from the Federal Reserve from March 1, 1984 to December 31, 1994. This period best matches the assumptions underlying our model. From 1995 onwards, the Federal Reserve adopted a formal policy of announcing at every FOMC meeting its decision concerning interest rates, whether it changed rates or not. As we argue in Section 5.3, this affects the optimal targeting rule, causing most target changes to fall on FOMC meeting dates. Moreover, for several years up to March 1, 1984, the Federal Reserve did not have an interest rate target, but rather targeted monetary aggregates. For this reason, like Rudebusch (1995), we adopt March 1, 1984 as our start date in the 1980s.[9]

During our sample period, the target rate was changed 105 times, with the average change being $\hat{\Delta} = 23.6$ basis points. A proportion $\hat{\pi} = 0.13$ of the changes are reversals and the average time between changes is $\hat{T} = 25.7$ business days. Therefore, we have $\Delta^* = 23.6$ basis points, $b^* = 88.4$ basis points and $\sigma = 11.9$ basis points, where the unit of time is one business day.[10]

Using this calibrated targeting rule we can back out the underlying costs which give rise to it. This can be achieved by solving equations (4) for k and f. We obtain

$$k = \frac{2}{\rho}\left(b^* - \Delta^* \frac{\sinh(\lambda b^*)}{\sinh(\lambda b^*) - \sinh(\lambda(b^* - \Delta^*))}\right) \tag{5}$$

and

$$\begin{aligned} f \quad = \quad & \frac{\Delta^*}{\rho}\left(\Delta^* \cdot \frac{\sinh(\lambda b^*) + \sinh(\lambda(b^* - \Delta^*))}{\sinh(\lambda b^*) - \sinh(\lambda(b^* - \Delta^*))}\right. \\ & \left. - \frac{2}{\lambda} \cdot \frac{\cosh(\lambda b^*) - \cosh(\lambda(b^* - \Delta^*))}{\sinh(\lambda b^*) - \sinh(\lambda(b^* - \Delta^*))}\right), \end{aligned} \tag{6}$$

where $\lambda^2 = 2\rho/\sigma^2$. Substituting $b^* = 88.4$, $\Delta^* = 23.6$ and $\sigma = 11.9$ into these equations, we find that $k = 4160$ and $f = 2378$, where we have assumed that the discount rate is $\rho = 0.01/250$, corresponding to an annual rate of 0.01. Costs implied by the model seem reasonable. For instance, fixed costs represent only 2.37% of the total adjustment costs associated with changing the target rate by its average amount (23.6 basis points). The average target change incurs total adjustment costs equivalent to a discrepancy of 40.6 basis points between the preferred and target rates lasting three months.

[9] Our data on the target rate is from Rudebusch until July 12, 1990, and then from the Federal Reserve Board's website www.federalreserve.gov/fomc/fundsrate.htm

[10] If we take the longer period, March 1, 1984 till March 31, 2001, the calibrated optimal rule does not appear much different: we find $\Delta^* = 24.6$ basis points, $b^* = 78.3$ basis points and $\sigma = 9.6$ basis points. However, the results for this longer period are heavily influenced by the period up till 1995, when 105 out of the 121 target changes occur. When the optimal rule is calibrated separately to the period January 1, 1995 till March 31, 2001, the optimal rule is quite different. Because the remaining 16 target changes have an average magnitude of 31.3 basis points, with 31.3% of the changes being reversals, and an average time between changes of 95.8 business days, the optimal rule for this period is $\Delta^* = 31.3$ basis points, $b^* = 50.0$ basis points and $\sigma = 4.7$ basis points. Despite the different optimal rule, the model still implies the same qualitative properties derived below. Section 5.3 extends our model to capture features of the more recent period.

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

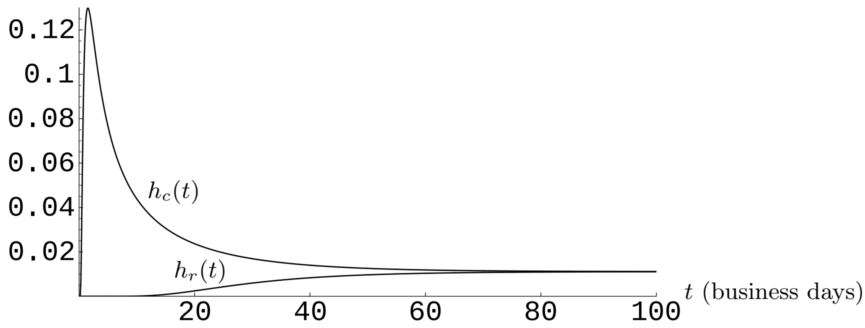

The following parameter values were adopted in constructing this figure: $b^* = 88.4$ basis points, $\Delta^* = 23.6$ basis points and $\sigma = 11.9$ basis points. Time is measured in business days.

The dynamic behavior of the target rate predicted by our calibrated model can be described by hazard functions. 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.

Assuming the preferred rate is unobservable, 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 6 *The probability of an immediate continuation is greater than the probability of an immediate reversal, regardless of how long it has been since the last target change.*

As the time since the most recent target change grows, two things happen to the distribution of the preferred 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 preferred rate is in the middle of the band, rather than near the edges, 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 7 *The probability of an immediate continuation first increases, then decreases, as the time since the last target change increases.*

and

Property 8 *The probability of an immediate reversal 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 9 *An immediate continuation and an immediate reversal are equally likely when the time since the most recent target change grows infinitely large.*

3.3 Behavior of Market Rates

To determine the implications of optimal targeting behavior for market interest rates, the simple targeting rule above is combined with the expectations hypothesis.[11] Due to the continuous time framework we consider, we use the local version of the expectations hypothesis. Thus, the price at time t of a discount bond paying 1 at time T equals

$$B(\hat{r}_t, r_t^*, t; T) = E_t \left[\exp \left(- \int_t^T \hat{r}_s ds \right) \right]. \tag{7}$$

Assuming the level of the central bank's preferred rate is public information, the expectation in (7) is conditional on the level of the preferred rate at time t, as well as the level of the target rate at that time.

If $\hat{r} - b < r^* < \hat{r} + b$, the probability of a target change over the next time increment of length dt is negligible. Thus, $\hat{r}_s$ is constant, while the preferred rate evolves according to $dr_s^* = \sigma d\xi_s$.

[11] Although there is considerable empirical evidence suggesting the expectations hypothesis does not hold for the U.S., recent evidence suggests that for short-maturity interest rates the hypothesis holds up reasonably well. For instance, Hsu and Kugler (1997) find the short-version of the expectations hypothesis cannot be rejected for one- and three-month Eurodollar rates over a period similar to the one we use to test the implications of our model. Their result is robust across different frequencies of interest rates, including the shortest frequency available — daily observations.

Figure 3: Spreads between market rates and the target rate

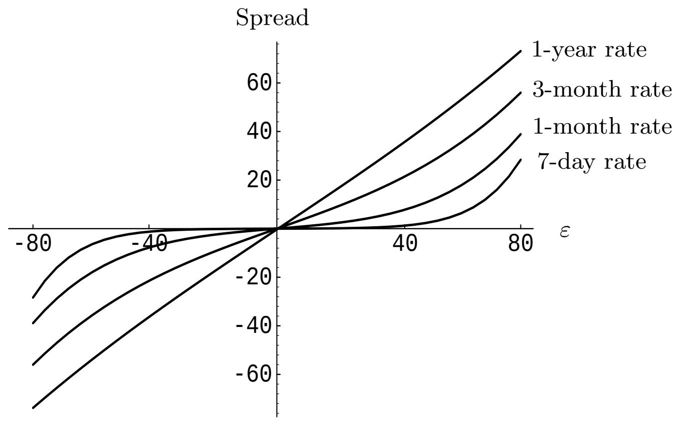

Each curve gives the difference (measured in basis points) between the corresponding interest rate and the target rate as a function of the difference ε (measured in basis points) between the preferred rate and the target rate.

By Itô's Lemma, the change in the price of the bond over the next time increment is

$$dB_t = \left(\frac{\partial B}{\partial t} + \frac{1}{2}\sigma^2 \frac{\partial^2 B}{\partial r^{*2}} \right) dt + \left(\frac{\partial B}{\partial r^*} \right) \sigma d\xi_t.$$

The rate of expected return from holding the discount bond must equal the prevailing target rate. Therefore B must satisfy the partial differential equation

$$\frac{\partial B}{\partial t} + \frac{1}{2}\sigma^2 \frac{\partial^2 B}{\partial r^{*2}} = \hat{r}B, \quad \hat{r} - b < r^* < \hat{r} + b.$$

Boundary conditions are determined by the central bank's targeting policy. If $r^* = \hat{r} + b$, the target rate is immediately raised by Δ. Since observability of the preferred rate makes this change predictable, the bond price must not change. Therefore,

$$B(\hat{r}, \hat{r} + b, t; T) = B(\hat{r} + \Delta, \hat{r} + b, t; T).$$

Consideration of target changes in the opposite direction shows that

$$B(\hat{r}, \hat{r} - b, t; T) = B(\hat{r} - \Delta, \hat{r} - b, t; T).$$

The terminal condition $B(\hat{r}, r^*, T; T) = 1$ reflects the fact that the discount bond pays 1 at maturity.

We show in Appendix C that the spread between the yield on a discount bond and the target rate is a function of maturity and the current discrepancy $\varepsilon = r^* - \hat{r}$ between the central bank's preferred rate and its target rate. Figure 3, which plots the spread for the calibration described in Section 3.2, demonstrates how market interest rates incorporate anticipated future behavior of the target rate. As $\varepsilon \to b$ (that is, as the preferred rate approaches $\hat{r} + b$), all spreads grow larger, reflecting the increased probability of the target rate being raised in the near future. Similarly, as $\varepsilon \to -b$, all spreads become more negative, as the market anticipates lower future levels of the target rate. The spread is most sensitive to changes in ε when the preferred rate is close to either edge of the band around the target rate. In these situations, a small change in the preferred rate greatly alters the likelihood of a target change, leading to a significant revision of the market rate.

Two properties follow from this behavior. Firstly, Figure 3 implies the spread between any market rate and the target rate is an increasing function of the spread between the preferred rate and the target rate. When this is combined with Properties 6–9 it suggests that the magnitude of the spread between any market rate and the target rate will decrease as the time since the last target change increases. Immediately after a target change, the preferred rate is still far away from the target rate, since the target rate only moves part of the way towards the preferred rate. If a long time has passed since the last target change, then the distribution of the preferred rate will be centered around the target rate. Thus, the spread between the preferred rate and the target rate will on average be close to zero, reflecting the fact a continuation and a reversal are equally likely. Figure 3 then translates these properties of the spread between the preferred rate and the target rate into spreads between the market rates and the target rate. This implies

Property 10 *The greater the time since a target change has occurred, the smaller is the expected spread between market rates and the target rate.*

Since the preferred rate is assumed to follow a simple Brownian motion, the daily change in this variable is normally distributed. The nonlinear relationship between the preferred rate and the market rate evident in Figure 3 means that when the preferred rate is close to the edge of the band around the target rate, even a small change in the preferred rate will lead to a large change in the market rate. When the preferred rate is close to the target rate, market rates are relatively insensitive to changes in the preferred rate. Because of this nonlinear relationship between the preferred rate and the market rate, the model predicts the volatility of market rates will be high when the spread is large. A large spread between the market rate and the target rate indicates that the preferred rate is near the edge of the band around the target rate and this is when market rates are most volatile. In summary,

Property 11 *The greater the spread between market and target rates, the higher the conditional volatility of market rates.*

4 Testing the Model

In this section we test the properties of target rates and market rates, which were derived in Section 3. We use data for March 1, 1984 to December 31, 1994, for the reasons outlined in Section 3.2. Section 4.1 tests properties on the target rate, while Section 4.2 tests some implications of our model for market rates using data on Eurodollar rates.

4.1 Testing Properties of the Target Rate

Using nonparametric methods, Rudebusch (1995) provides an extensive empirical examination of the properties of changes in the federal funds target rate for the periods September 1974 to September 1979 and March 1984 to September 1992. Rudebusch concludes that "target changes were conducted in small, standardized steps with a rough equality in the size of target increases and decreases" (p. 10), suggesting our Property 1 held for his sample period. Using data on the 105 target changes over the period March 1, 1984 to December 31, 1994 we regress the magnitude of target changes on a constant and a dummy variable for increases in the target rate. The coefficient on the dummy variable is -0.045 with a standard error of 0.027, suggesting that target rate increases are 4.5 basis points smaller on average than decreases. This difference is small, and statistically we cannot reject that target rate increases are the same size as decreases at the 5% significance level (the p-value for the hypothesis in Property 1 is 0.098).

Similarly, Properties 2 and 3 are supported by empirical observation. Using two nonparametric tests, Rudebusch cannot reject the hypothesis that the duration between two consecutive positive target changes has the same distribution as the duration between two consecutive negative changes (with p-values averaging 0.50), and that the duration between a positive change followed by a negative change has the same distribution as a negative change followed by a positive change (with p-values around 0.08). We test Properties 2 and 3 directly by regressing the number of business days between target changes that are continuations (this regression has 91 observations) on a constant and a dummy variable that takes the value one when the continuation refers to two consecutive positive target changes changes and zero when the continuation refers to two consecutive negative target changes. The coefficient on the dummy variable is -4.731 with a standard error of 4.993, suggesting a small and statistically insignificant difference in the duration for the two types of continuations (the p-value for this test is 0.346). The coefficient on the dummy variable in the equivalent regression for reversals is 52.286 with a standard error of 47.798. This implies, the number of business days between target rate decreases which are followed by target rate increases is around 52 days longer than between target rate increases which are followed by target rate decreases. Although this appears to be a large difference, due to the fact there are only 14 such reversals in our sample, and that reversals have durations that vary a great deal, this difference is not statistically significant. Thus, we cannot reject that Property 3 holds (the p-value is 0.295 for this hypothesis test).

A similar method can be used to test Property 5. We regress the number of business days since the last target change on a constant and a dummy variable which takes the value one for continuations and zero for reversals. We find that the duration between target changes is 31.000 days less when they are continuations, with a standard error of 11.180 (the p-value on the hypothesis that the time between two consecutive target changes in the same direction equals the time between two consecutive target changes in opposite directions is 0.007). Further support for Property 5 is the finding by Rudebusch that the p-value on the hypothesis that

Figure 4: Estimated hazard functions

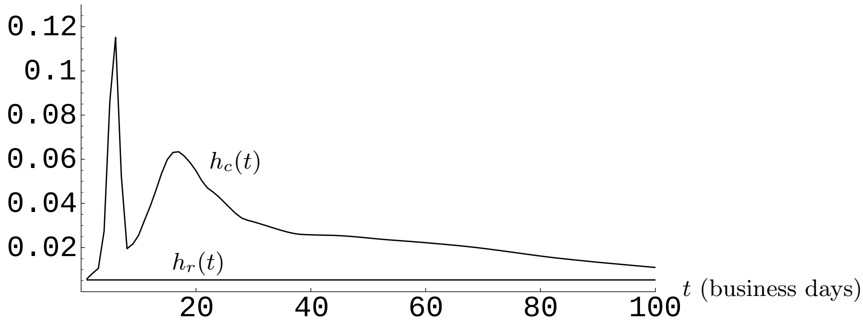

the distribution of durations between consecutive same-sign target changes is the same as the distribution of durations between consecutive different-sign target changes is only 0.015.

Rudebusch also provides evidence which supports Properties 6–9. He estimates the hazard functions for changes in the target rate from the nonparametric hazard rate estimator using kernel functions and quasi-likelihoods, as described in Tanner and Wong (1994). As Figure 3 in his paper demonstrates, the estimated probability of continuations is higher than the probability of reversals, at least for durations up to the 35 business days plotted, (consistent with our Property 6), and the probability of a continuation first increases and then tends to decrease as the time since the last target change increases (consistent with our Property 7). He also finds that the probability of a reversal increases as the time since the last target change increases, for durations up to 25 business days. When this is combined with his finding that after 25 business days the distributions of durations between continuations and reversals are identical (the p-values on two different tests of this hypothesis average 0.15), his results lend support to our Properties 8 and 9. These results also imply we cannot reject Property 4.

We estimate hazard functions for our sample period using two approaches. First, using the same nonparametric kernel estimator as Rudebusch, we obtain the estimated hazard functions in Figure 4.[12] This approach generates an estimated hazard function for continuations which initially increases sharply, suggesting the most likely time for a target change is a short time after the last target change. Apart from a second smaller peak, which is found after a further eleven business days, the probability of a continuation monotonically decreases as the time since the last target change increases. Moreover, for durations without target changes of up to 135 business days, the probability of a target change is more likely to be in the same direction as the last change than in the reverse direction, with the two probabilities converging as the time since the last target change increases. Thus, like Rudebusch, we find some support for Properties 6 and 7. The properties of the hazard functions for very long durations, as well as the properties of the hazard function for reversals, are likely to be dominated by the fact that for our sample period there was one reversal which lasted 349 days. This occurred between September 4, 1992,

[12] We first estimate the optimal smoothing parameters using the HAZRD procedure in IMSL, and then estimate the hazard function using HAZST in IMSL.

Table 1: Estimated hazard functions

	Continuations			Reversals		
Parameter	Estimate	Std error	p-value	Estimate	Std error	p-value
α	1.194	0.127	0.000	1.356	0.354	0.000
γ	0.026	0.011	0.019	0.001	0.002	0.471
θ	-0.007	0.003	0.027	-0.0003	0.002	0.868

when the target rate was lowered to 3% from 3.25% and February 4, 1994, when the target rate was raised from 3% back up to 3.25%. As a result, our estimates of the hazard function for reversals, unlike Rudebusch's estimates, suggest the hazard function is essentially a flat line, showing almost no duration dependence.

Similar results arise when we parametrically estimate the hazard functions for continuations and reversals using the expo-power hazard function. This function is general enough to capture constant, monotonically increasing or decreasing, U-shaped, and inverted U-shaped forms, and nests the Weibull and exponential forms as special cases (see Saha and Hilton, 1997). The hazard function is given by

$$\lambda(t) = \gamma\alpha t^{\alpha-1} e^{\theta t^{\alpha}}.$$

The estimated hazard function for continuations is first increasing and then decreasing, with a peak after 14 days. The estimates are shown in Table 1. We test the hypothesis of no duration dependence (which corresponds to the joint hypothesis that $\alpha = 1$ and $\theta = 0$). The Wald χ^2 test statistic is 30.115 with a p-value of 0.000, suggesting we can clearly reject the null of no duration dependence. Moreover, because $\alpha > 1$, the test of whether the hazard function for continuations has an inverted-U shape or is monotonically increasing is whether $\theta \geq 0$ or $\theta < 0$. If θ is negative then the hazard function is first increasing and then decreasing. Based on the results above, we can reject $\theta \geq 0$ in favour of $\theta < 0$ at the 5% level (the p-value for this one-sided test is 0.013). Thus, our results lend statistical support to Property 7.

For reversals the estimated hazard function is monotonically increasing up to 123 days, consistent with Property 8. Estimates of the parameters are reported in Table 1. The results show that the positive duration dependence is not statistically significant. The test of no duration dependence for reversals, cannot be rejected. The test statistic for $H_0 : \alpha = 1$ and $\theta = 0$ is 1.021, with a p-value of 0.600. This statistical result is likely to be a function of the low number of reversals in our sample period (14 in total) together with the fact one reversal has a duration of 349 days, much larger than the duration of any other target change. Moreover, similar to our nonparametric estimates, the finding that the probability of a continuation is greater than the probability of a reversal, and that as the time since the most recent target change grows large continuations and reversals become equally likely, holds for durations up to 127 days.

4.2 Testing Properties of Market Rates

To test Properties 10 and 11 we use daily observations on seven-day, one-month and three-month Eurodollar rates, together with the federal funds target rate, from March 1, 1984 to December 31, 1994. The one- and three-month Eurodollar spot rates are bid-side rates quoted in London, collected around 9:30 a.m. Eastern time, and supplied by the Federal Reserve. The seven-day Eurodollar spot rates are midpoint between bid and ask rates, used previously by Ait-Sahalia (1996a, b), and described in detail in his papers. Because of their different sources, the data on the seven-day rates has to be matched to the data on target rates and one- and three-month Eurodollar rates. After eliminating weekends and missing data, the resulting series are available for a common 2,694 observations over our sample.

We use Eurodollar rates since they are comparable to the federal funds target rate. The federal funds target rate refers to a target for the overnight interbank U.S. rate. Eurodollar rates are also private inter-bank U.S. dollar rates from a relatively liquid market. In contrast, Treasury bill rates are for government risk-free securities and so are not directly comparable to the target rate.

According to Property 10, the magnitude of the spread between each of the market rates and the central bank's target rate should be declining as the time since the last target change increases. To test this hypothesis we run the regression

$$|r_t - \hat{r}_t| = \alpha + \frac{\beta}{D_t} + \varepsilon_t,$$

where r_t is a market rate and $\hat{r}_t$ is the target rate at time t, and D_t is the number of days, measured at time t, since the last target change. If the estimate of β is positive, this suggests the magnitude of the spread is declining in the time since the last target change, while a negative value of β implies the magnitude of the spread is increasing in the time since the last target change.[13] We find that β equals 0.044 (0.015), 0.049 (0.018), and 0.016 (0.018) for the seven-day, one-month, and three-month market rates respectively.[14] The reversion of the market rate to the target rate as the number of days since the last target change increases is statistically significant for the seven-day and one-month rates. The results are virtually unchanged when observations where there were target changes are deleted from the sample — these represent only a small fraction of all observations.

It is well known that the conditional volatility of interest rates is persistent and heteroskedastic. Because periods around target changes are times when there is an increased likelihood of large interest rate changes, our model predicts that conditional volatility will be higher during such episodes. These periods are identified as times when the spread between the target rate

[13] A linear specification is not reasonable since a declining spread will eventually imply a negative magnitude of the spread, which is not possible. An exponential specification is also not appropriate, as it suggests the alternative to a geometrically decaying spread is an exponentially exploding spread. Our specification does not suffer from either of these problems.

[14] Standard errors in brackets.

and market rates is high (Property 11). To test this hypothesis we estimate a simple model of interest rate volatility, allowing for the magnitude of the spread between the target and market rates as a determinant of volatility, in addition to the normal level of interest rates. Thus we estimate the following model using FGLS, following the approach of Ait-Sahalia (1996b, p. 409):

$$\begin{aligned} r_t - r_{t-1} &= \alpha + \kappa r_{t-1} + \epsilon_t, \\ E_{t-1}\epsilon_t &= 0, \\ E_{t-1}\epsilon_t^2 &\equiv \sigma_t^2 = \sigma_1^2 |r_{t-1} - \hat{r}_{t-1}|^{2\gamma_1} + \sigma_2^2 r_{t-1}^{2\gamma_2}. \end{aligned}$$

To test Property 11 we estimate these equations using actual data on market and target rates. We find that with seven-day rates, $\hat{\sigma}_1 = 0.438$ (0.022) and $\hat{\gamma}_1 = 1.281$ (0.080); with one-month rates, $\hat{\sigma}_1 = 0.358$ (0.012) and $\hat{\gamma}_1 = 1.154$ (0.035); and with three-month rates, $\hat{\sigma}_1 = 0.175$ (0.007) and $\hat{\gamma}_1 = 1.619$ (0.114). On the other hand, the level-effect is only apparent for three-month rates: with seven-day rates, $\hat{\sigma}_2 = 0.065$ (0.043) and $\hat{\gamma}_2 = 0.237$ (0.327); with one-month rates, $\hat{\sigma}_2 = 0.054$ (0.063) and $\hat{\gamma}_2 = 0.029$ (0.620); with three-month rates, $\hat{\sigma}_2 = 0.027$ (0.007) and $\hat{\gamma}_2 = 0.512$ (0.130). These results demonstrate that the spread term in the conditional volatility equations is highly significant. We call this effect the spread-effect (as opposed to the normal level-effect).

To gauge the economic strength of the spread-effect versus the level-effect, we report the proportion of the total variation in volatility captured by the above explanatory variables, and consider how this R^2 changes when the spread term is eliminated from the estimation. The R^2 from the estimation including both spread- and level-effects is 0.170 for the seven-day Eurodollar rate, 0.259 for the one-month Eurodollar rate, and 0.105 for the three-month Eurodollar rate. When the spread-effect is eliminated, the R^2 drops to just 0.004 with the seven-day rate, just 0.003 with the one-month rate, and 0.025 with the three-month rate. In contrast, when the level-effect is eliminated, the remaining spread-effect implies an R^2 of 0.169 for the seven-day rate, 0.259 for the one-month rate, and 0.081 with the three-month rate. In short, the spread-effect dominates the level-effect on both statistical and economic grounds.

5 Extensions

In this section we present three extensions to our standard model. The first addresses the fact that the central bank will sometimes be able to predict future movements in the preferred rate, and therefore anticipate future target changes. The other two extensions focus on the market's ability to anticipate target changes, both the direction of target changes (Section 5.2), and their timing (Section 5.3).

5.1 Tightening and Loosening Cycles

There are times when there is some deterministic component to economic conditions, so that underlying conditions will call for tighter monetary policy over the medium term, or the reverse.

Table 2: Optimal adjustment rules during a tightening cycle

Annual drift (basis pts)	b_1^*	b_2^*	Δ_1^*	Δ_2^*
0	88.4	88.4	23.6	23.6
50	93.8	83.2	23.0	24.2
100	99.3	78.2	22.5	24.8
150	104.9	73.5	22.0	25.5
200	110.6	69.1	21.6	26.2

The preferred rate has drift μ and volatility $\sigma = 11.9$. Adjustment cost parameters are $f = 2378$ and $k = 4160$. The central bank discounts future costs at the rate $\rho = 0.01/250$.

We model this by introducing a trend into the process for the preferred rate. To be precise, we take the model of Section 2 and change the process for the preferred rate to $dr_t^* = \mu dt + \sigma dz_t$, for constants μ and σ. Thus, the preferred rate is expected to grow by μ basis points each business day. Of course, such a process is unrealistic over long time horizons, but over the sorts of time intervals considered here (25.7 business days between target changes on average), it adequately captures the effects of predictable changes in the preferred rate.

We consider adjustment rules which are natural extensions to the (b, Δ)-rules considered in Section 2.2. The adjustment rule is determined by four parameters $(b_1, b_2, \Delta_1, \Delta_2)$: the target rate is lowered by Δ_1 as soon as ε falls below the threshold $-b_1$; it is raised by Δ_2 as soon as ε climbs above the threshold b_2. We show in Appendix D.1 that the central bank's loss function under such a targeting policy is

$$u(\varepsilon) = \begin{cases} u(-b_1 + \Delta_1) + f + k(-b_1 + \Delta_1 - \varepsilon) & \varepsilon \le -b_1, \\ \frac{\sigma^2}{\rho^2} + \frac{2\mu^2}{\rho^3} + \frac{2\mu\varepsilon}{\rho^2} + \frac{\varepsilon^2}{\rho} + A_1 e^{\lambda_1 \varepsilon} + A_2 e^{\lambda_2 \varepsilon} & -b_1 < \varepsilon < b_2, \\ u(b_2 - \Delta_2) + f + k(\varepsilon - b_2 + \Delta_2) & \varepsilon \ge b_2, \end{cases}$$

where

$$\lambda_1 = \frac{-\mu}{\sigma^2} + \sqrt{\left(\frac{\mu}{\sigma^2}\right)^2 + \frac{2\rho}{\sigma^2}}, \qquad \lambda_2 = \frac{-\mu}{\sigma^2} - \sqrt{\left(\frac{\mu}{\sigma^2}\right)^2 + \frac{2\rho}{\sigma^2}},$$

and the constants A_1 and A_2 are chosen so that u is continuous at $-b_1$ and b_2. A necessary condition for the adjustment policy to be optimal is that $(b_1^*, b_2^*, \Delta_1^*, \Delta_2^*)$ satisfy the smooth-pasting conditions

$$\lim_{\varepsilon \downarrow -b_1^*} u'(\varepsilon) = u'(-b_1^* + \Delta_1^*) = k = u'(b_2^* - \Delta_2^*) = \lim_{\varepsilon \uparrow b_2^*} u'(\varepsilon).$$

This system of equations is straightforward to solve numerically. Results for various levels of drift are reported in Table 2.

The numbers in Table 2 show that the central bank behaves in quite different ways depending on whether it expects the preferred rate to move towards, or away from, the current level of the target rate. Compared to the case considered in Section 2, the central bank will act much more aggressively when it expects the preferred rate to move further away from the current level of the target rate — it tolerates much smaller deviations before acting, and adjusts the target rate by more when it does act. The greater movement in the target rate is needed to 'catch up' to the preferred rate. Assuming $\mu > 0$, this situation occurs when $\varepsilon > 0$. As soon as ε reaches b_2^*, the central bank raises the target rate by Δ_2^*. Table 2 shows that b_2^* is decreasing, and Δ_2^* is increasing, in drift, so that the bank becomes more aggressive as the trend in the preferred rate becomes stronger. However, the central bank still finds it optimal to act 'too little, too late.' That is, after the target rate is changed, the preferred rate is still expected to move away from the target rate. It is not optimal for the central bank to overshoot by setting the target rate at a level it expects the preferred rate to reach at some point in the future.

The central bank can afford to act more cautiously when it expects the preferred rate to move back towards the current level of the target rate — it tolerates relatively large deviations before acting, and when it does act, adjusts the target rate by a relatively small amount. Again assuming $\mu > 0$, this situation occurs when $\varepsilon < 0$, so that the preferred rate is below the current level of the target rate but is expected to move closer in the future. Table 2 indicates that b_1^* is increasing, and Δ_1^* is decreasing, in drift, so that the central bank is more reluctant to change the target rate, and will change it by a smaller amount, the greater the speed with which the preferred rate is expected to move back towards the target rate.

In summary, target rate changes are relatively large when they are in the same direction as the expected movement of the preferred rate, and relatively small when they are in the opposite direction. It is straightforward to show that the introduction of a trend into the preferred rate also increases the persistence in target changes in the same direction as the trend, and reduces the persistence of changes which move against the trend.

5.2 Anticipated and Unanticipated Target Changes

Given our motivation for adjustment costs, it seems likely that target changes will be more costly when they are in the opposite direction to that expected by the market. We incorporate this possibility by taking the model of Section 2 and setting the fixed and marginal cost parameters at a higher level when the direction of the target change takes the market by surprise. We will find that, provided they cannot observe the central bank's preferred rate, investors will always believe that the next target change is more likely to be a continuation that a reversal. Thus, we set the cost of a target change of magnitude Δ equal to $f_c + k_c\Delta$ if the change is in the same direction as the previous one, and equal to $f_r + k_r\Delta$ if it is in the opposite direction, where $f_r \geq f_c$ and $k_r \geq k_c$. As in the model of Section 2, we suppose that the preferred rate r^* evolves according to the driftless Brownian motion $dr_t^* = \sigma dz_t$.

The analog of the (b, Δ)-rule studied in Section 2 is completely described by four parameters: $(b_c, b_r, \Delta_c, \Delta_r)$.

- Suppose the last target change was a loosening. The central bank leaves the target rate unchanged as long as $-b_c < \varepsilon < b_r$; if $\varepsilon \leq -b_c$, the target rate is lowered (a continuation of the earlier policy) so that ε is reset to $-b_c + \Delta_c$; if $\varepsilon \geq b_r$, the target rate is raised (a reversal of the earlier policy) so that ε is reset to $b_r - \Delta_r$.

- Suppose the last target change was a tightening. The central bank leaves the target rate unchanged as long as $-b_r < \varepsilon < b_c$; if $\varepsilon \leq -b_r$, the target rate is lowered (a reversal of the earlier policy) so that ε is reset to $-b_r + \Delta_r$; if $\varepsilon \geq b_c$, the target rate is raised (a continuation of the earlier policy) so that ε is reset to $b_c - \Delta_c$.

We show in Appendix D.2 that the central bank's loss function under such a targeting policy is $u_L(\varepsilon) = u(\varepsilon)$ if the last target change was a loosening, and $u_T(\varepsilon) = u(-\varepsilon)$ if it was a tightening, where

$$u(\varepsilon) = \begin{cases} u(-b_c + \Delta_c) + f_c + k_c(-b_c + \Delta_c - \varepsilon) & \varepsilon \leq -b_c, \\ \frac{\sigma^2}{\rho^2} + \frac{\varepsilon^2}{\rho} - Ae^{\lambda\varepsilon} - Be^{-\lambda\varepsilon} & -b_c < \varepsilon < b_r, \\ u(-b_r + \Delta_r) + f_r + k_r(\varepsilon - b_r + \Delta_r) & \varepsilon \geq b_r, \end{cases}$$

where $\lambda^2 = 2\rho/\sigma^2$. The constants A and B are chosen so that u is continuous at $-b_c$ and b_r. A necessary condition for the adjustment policy to be optimal is that $(b_c^*, b_r^*, \Delta_c^*, \Delta_r^*)$ satisfy the smooth-pasting conditions

$$\lim_{\varepsilon \downarrow -b_c^*} u'(\varepsilon) = u'(-b_c^* + \Delta_c^*) = -k_c,$$

and

$$\lim_{\varepsilon \uparrow b_r^*} u'(\varepsilon) = -u'(-b_r^* + \Delta_r^*) = k_r.$$

We have found numerical solutions to this system for a range of cost scenarios, and in each case a plot of the hazard function shows that the probability of a continuation is greater than that of a reversal at any time.[15] That is, our assumption that the market always expects a continuation rather than a reversal holds, and the model is internally consistent.

In all the cases we have examined, $\Delta_c^* > \Delta_r^*$, indicating that when the central bank reverses policy, we should see a small reversal, which is usually followed by larger changes in the same direction. The central bank uses a small reversal to advise the market that it is reversing its recent policy. Once the market's expectations have been revised, the central bank can implement a larger target change at relatively low cost. The pattern of small reversals, followed by larger continuations is consistent with those found in the data — over our sample, the average size of a continuation is 24.38 basis points, whereas the average size of a reversal is only 16.64 basis points. Keeping the calibration from Section 3.2 for the volatility of the preferred rate and the adjustment costs of a continuation, while scaling up the fixed and marginal costs of a reversal

[15] This occurs because $b_c^* > b_r^*$. Consider the case where the last target change was a tightening, at which point ε was reset to $b_c^* - \Delta_c^* > 0$. At all times prior to the next target change, the market believes that ε is more likely to be positive than negative. This is the case even when the time since the last target change grows infinitely large, as then the preferred rate is symmetrically distributed around the midpoint of the band, $(b_c^* - b_r^*)/2 > 0$.

by a common factor of 1.02, results in the following optimal adjustment rule:

$$b_c^* = 88.4, \quad b_r^* = 85.8, \quad \Delta_c^* = 23.6, \quad \Delta_r^* = 16.9.$$

As in our sample, reversals are approximately two-thirds the size of continuations.[16]

Another consequence of the additional costs of surprise changes is that a continuation is more likely to occur after a reversal than after another continuation. For example, suppose the last two target changes have been loosenings. At the time of the second target change, $\varepsilon = -b_c^* + \Delta_c^*$. The next target change will be another loosening if ε reaches $-b_c^*$ before it reaches b_r^*. It is straightforward to show that this occurs with probability $1 - \Delta_c^*/(b_c^* + b_r^*)$. It is the probability that one continuation follows another. When the cost parameters for reversals exceed their counterparts for continuations by a factor of 1.02, this probability equals 0.865. Now suppose that the last two target changes have been a tightening followed by a loosening. At the time of the second target change, $\varepsilon = -b_r^* + \Delta_r^*$. The next target change will be another loosening if ε reaches $-b_c^*$ before it reaches b_r^*. This occurs with probability $(2b_r^* - \Delta_r^*)/(b_c^* + b_r^*)$. It is the probability that a continuation follows a reversal. When adjustment costs of reversals are scaled up by a factor of 1.02, this probability equals 0.888. Thus, continuations are somewhat more likely when the last target change was a reversal.[17]

5.3 Announcement Dates

In this section we examine the possibility that by specifying regular announcement dates a central bank can lower adjustment costs when target changes are made on these dates. 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 model this as a reduction in the adjustment costs of target changes on these dates, which we suppose occur at times 0, T, $2T$, and so on. The cost of changing the target rate an amount Δ at time t equals

$$C(\Delta, t) = f(t) + k(t)|\Delta|,$$

for some functions $k(t)$ and $f(t)$.[18] As in the model of Section 2, we suppose that the preferred rate r^* evolves according to the driftless Brownian motion $dr_t^* = \sigma dz_t$.

The adjustment rules we consider, which are the obvious extensions of those considered in Section 2.2, are described by two functions, $b(t)$ and $\Delta(t)$. The central bank will only change

[16] Recall that in our model we are assuming that the market cannot observe the central bank's preferred rate. In practice, much of this information will actually be public, so that we are overstating the market's uncertainty regarding the direction of future target changes. This is one reason why such a small additional cost of reversals in our model can have a significant effect on the optimal size of reversals.

[17] Consistent with these predictions, out of the 14 reversals in our sample, 13 (or 0.929) were followed by a continuation; out of the 91 continuations, 78 (or 0.857) were followed by another continuation.

[18] We assume that $f(t + T) = f(t)$ and $k(t + T) = k(t)$ for all t, so that the cost of a target change depends only on the size of the target change and the time *during the announcement cycle* when the change occurs. The adjustment rules and loss function will share this periodicity.

Figure 5: The optimal adjustment rule with regular announcement dates

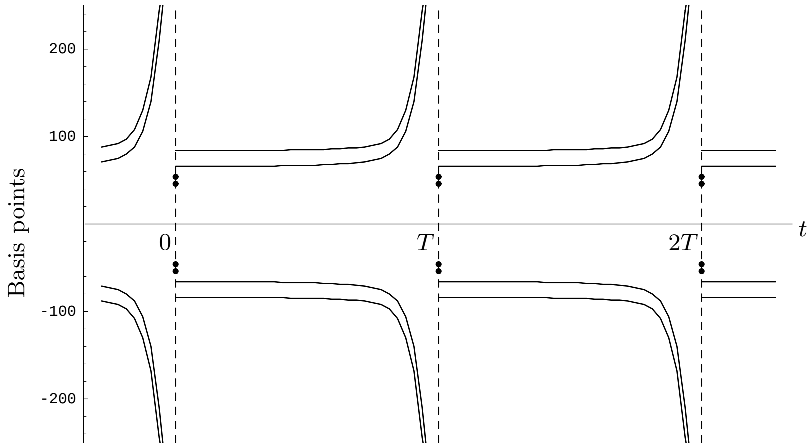

The preferred rate has volatility $\sigma = 11.9$. Adjustment cost parameters are $f = 2378$ and $k = 4160$ between announcement dates, and $0.8f$ and $0.8k$ on announcement dates. The central bank discounts future costs at the rate $\rho = 0.01/250$.

the target rate at time t if $\varepsilon_t \leq -b(t)$, in which case it will lower the target rate in order to reset the discrepancy between the target rate and the preferred rate to $-b(t) + \Delta(t)$, or if $\varepsilon_t \geq b(t)$, in which case it will raise the target rate in order to reset the discrepancy to $b(t) - \Delta(t)$. We show in Appendix D.3 that the central bank's loss function $u(\varepsilon, t)$ satisfies the partial differential equation

$$0 = \frac{\partial u}{\partial t} + \frac{1}{2}\sigma^2 \frac{\partial^2 u}{\partial \varepsilon^2} - \rho u + \varepsilon^2$$

whenever $-b(t) < \varepsilon < b(t)$. For a given $(b(t), \Delta(t))$-rule, the loss function can be calculated by solving the partial differential equation together with the value-matching conditions

$$u(b(t), t) = u(b(t) - \Delta(t), t) + C(\Delta(t), t) \qquad u(-b(t), t) = u(-b(t) + \Delta(t), t) + C(\Delta(t), t),$$

and the requirement that $u(\varepsilon, T) = u(\varepsilon, 0)$. The smooth-pasting conditions

$$\frac{\partial u}{\partial \varepsilon}(b^*(t), t) = \frac{\partial u}{\partial \varepsilon}(b^*(t) - \Delta^*(t), t) = k(t)$$

are necessary for the rule $(b^*(t), \Delta^*(t))$ to be optimal.

Figure 5 shows the optimal adjustment rule in the case where the cost of a target change on an announcement date is 80% of the cost of the same target change at any other time. The outer two curves plot $\pm b^*(t)$ as functions of the time in the announcement cycle, while the inner two curves plot $\pm(b^*(t) - \Delta^*(t))$. Thus, the central bank will only change the target rate between announcement dates when ε_t moves outside one of the outer two curves. It then changes the target rate so that ε_t is brought back to the associated inner curve. The size of the target change equals the distance between the two curves. The situation is slightly different on announcement dates, as the central bank then changes the target rate whenever ε_t lies outside one of the two outer points shown in the figure.[19] When this happens, the target is changed so that ε_t is brought back to the associated inner point. Clearly, the size of such a target change varies with the level of ε_t on the announcement date.[20] The simple targeting rule derived in Section 2.2 is no longer optimal. Instead, the optimal rule involves the band around the target rate becoming wider as the date of each announcement approaches (so that a target change is very unlikely when an announcement date is near), and narrowing on the date of the announcement (so that a target change on an announcement date is much more likely).[21]

Although there is not enough data yet to characterize the properties of targeting under announcement dates with confidence, there does seem to be some evidence that the Fed has a strong preference for changing the target rate only on such dates. In fact, only two of the sixteen target changes over the period January 1, 1995 to March 31, 2001, have been made outside of the FOMC's regular meeting dates. Perhaps the most telling testimony to the Fed's preference for making changes on such dates are comments from the Fed itself, which indicate it has a strong (but not lexiographic) preference towards making changes on FOMC meeting dates. For instance, the following is from the minutes of the FOMC telephone conference on April 11, 2001:

> "In the circumstances, the members could see the need for a further easing of policy at some point, though some had a strong preference for taking such actions at regularly scheduled meetings. They all agreed that an easing on this date would not be advisable, inasmuch as the attendant surprise to most outside observers risked unpredictable reactions in financial markets that had been especially volatile in recent days, and additional important data would become available over the near term."

6 Conclusion

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. 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, as could a number of other puzzling properties of target rates which we observed. We did not need to assume persistence in the underlying shock in our model to get these results; nor did we have to assume that reversals were especially costly. 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

[19] This discontinuity in the width of the band is a consequence of the discontinuity in the adjustment cost functions $f(t)$ and $k(t)$.

[20] Note also that the size of a target change which takes place between announcement dates will generally depend on the time during the announcement cycle at which it occurs.

[21] Filimon (2000) studies the behavior of our model for a large number of different parameter values, and obtains qualitatively similar results for all cases.

target change, rather than risk unsettling the market, especially since many times small changes in the preferred rate will be undone by themselves.

Using data on the Federal Reserve's target rate changes, we tested the dynamics of target changes predicted by our model. Consistent with our model, we found that the magnitude of a change in the target rate does not depend on the direction of the change, that the average time between consecutive target rate increases is the same as the average time between consecutive decreases, that the average time between a target rate increase that is followed by a decrease is the same as the average time between a target rate decrease that is followed by an increase, and that the average time between consecutive changes in the opposite direction is greater than the average time between consecutive changes in the same direction. We estimated parametric and nonparametric hazard functions for our sample of target changes, and found that (except for long durations) they were broadly consistent with the hazard function predicted by our model. Consecutive changes in the same direction are especially likely a short time after the last target change, but for long durations between target changes reversals become almost as likely as continuations.

Because our model characterizes the properties 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, we derived the implications of our calibrated optimal rule for daily observations on seven-day, one-month, and three-month market rates. Despite the simplicity of our assumptions, we found that, consistent with daily observations on Eurodollar rates, conditional volatility is increasing in the magnitude of the spread between the market rate and the central bank's target rate (which we call the spread-effect), and that market rates revert towards the central bank's target rate the longer the time since the last target change.

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 is still interesting to extend the analysis to consider other stochastic processes for the preferred rate.

We considered one such extension by introducing a deterministic component to the central bank's preferred rate. This captured the idea that there can be predictable tightening and loosening cycles to monetary policy. We found that when monetary conditions are expected to move desired interest rates closer to the bank's current target rate, optimal policy requires that the central bank act more cautiously, tolerating relatively large deviations before acting, and only changing the target rate by a small amount if it does act. Conversely, when monetary conditions are expected to move desired interest rates further away from the bank's current target rate, optimal policy requires that the central bank act more aggressively, moving the target rate sooner and by more than would otherwise be the case. Despite the possible need for a 'catching-up' policy, the central bank will still optimally avoid any overshooting of policy, so that the target interest rate is never moved from being considered 'too high' to being 'too low' just because the central bank expects interest rates will need to fall in the future.

Given our motivation for adjustment costs, it seems likely that target changes will be more costly when they are not anticipated by the markets. In a world in which the market does not observe the central bank's preferred rate, we considered two reasons why target changes may be more or less anticipated by investors. We showed that if target changes in the opposite direction to that anticipated by investors are more expensive than target changes in the same direction, a target change which reverses the direction of the previous target change will be more expensive than one which is in the same direction. The central bank's optimal response is to make reversals smaller than continuations. Following a reversal, once the market's expectations have been revised, the central bank can implement a larger target change at relatively low cost. Moreover, this behavior means continuations are somewhat more likely when the last target change was a reversal.

Another reason why target changes may be more or less of a surprise to the markets is the use of regular announcement dates. By specifying regular announcement dates a central bank can reduce the surprise element of target changes that are made on these dates, thus lowering the associated adjustment costs. In this case the optimal rule involves 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 announcement (so a target change on an announcement date is much more likely). We believe that when combined with a reduction in underlying volatility, this extension of the model may be capable of explaining the changes in dynamics of target rates observed since the Federal Reserve shifted to formal announcements on FOMC meeting dates. Future research, armed with more observations on target changes outside of meeting dates, could test the implications of our model in this new regime.

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. 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.

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:

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

or

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

The Lagrangian for this problem is

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

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

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

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

Proof of Theorem 1[22]

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

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

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

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

[22] The proof of this and the following theorem is based on Constantinides and Richard (1978), who consider a similar problem, with different cost functions, in the context of inventory management.

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

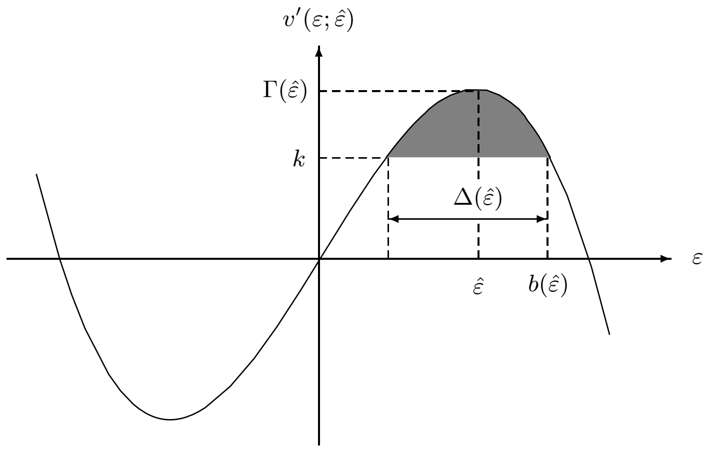

Now, for any value of the parameter $\hat{\varepsilon}$ such that $\Gamma(\hat{\varepsilon}) > k$, the numbers $b(\hat{\varepsilon})$ and $\Delta(\hat{\varepsilon})$ are uniquely determined by the analogs of equations (4),

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

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

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

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

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

That is, the shaded region in Figure 6 must have area f. Now, if $\varepsilon \to \Gamma^{-1}(k)$, then $b(\hat{\varepsilon}) \to \hat{\varepsilon}$ and $\Delta(\hat{\varepsilon}) \to 0$, so that $\Psi(\hat{\varepsilon}) \to 0$ and the shaded region vanishes. On the other hand, as $\hat{\varepsilon} \to \infty$, we see that $b(\hat{\varepsilon}) \to \infty$ and $b(\hat{\varepsilon}) - \Delta(\hat{\varepsilon}) \to \rho k/2$. Since $\Gamma(\hat{\varepsilon}) \to \infty$, 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. (1983).

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*

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

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

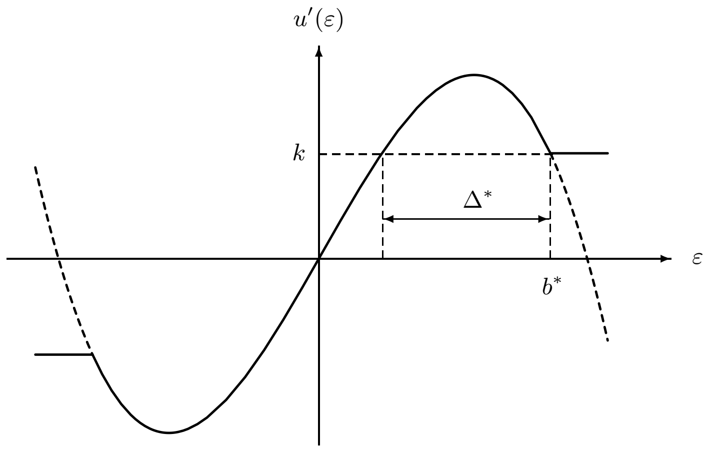

then $u(\varepsilon) \leq J(P;\varepsilon)$ *for all adjustment policies* P *and all* $\varepsilon \in \mathbb{R}$.

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.

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 $\varepsilon < -b^*$), $\theta'(\varepsilon) = 2\varepsilon - \rho k > 2b^* - \rho k > 0$ (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

Proposition B–1 *The probability that the next target change, whenever it occurs, is in the same direction equals* $1 - \Delta^*/(2b^*)$.

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 $-b^*$. From Karatzas and Shreve (1991, p. 100), this occurs with probability $(2b^* - \Delta^*)/(2b^*) = 1 - \Delta^*/(2b^*)$. ■

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* $\Delta^*(2b^* - \Delta^*)/\sigma^2$ *days.*

PROOF 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. From Karatzas and Shreve (1991, p. 100),

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

and

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

where $\beta = \sqrt{2s}/\sigma$. It follows that

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

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

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

days.

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

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

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

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

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

When plotting the hazard functions, we use the following series expansions from Karatzas and Shreve (1991, p. 100):

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

and

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

C Determining Market Rates

The bond pricing problem is greatly simplified by noting two particular properties of bond prices in our model. If the target rate and the preferred rate both increase by δ, the distribution of all future levels of the target rate shifts to the right by the same amount. From (7), the effect is to scale the time t price of a discount bond maturing at time T by the factor $\exp(-\delta(T-t))$. That is, the bond price function satisfies

$$B(\hat{r} + \delta, r^* + \delta, t; T) = e^{-\delta(T-t)} B(\hat{r}, r^*, t; T) \text{ for all } \delta. \tag{C-1}$$

If the current date and the maturity date of the bond both increase by s, the price of a discount bond will not change. That is, the bond price function satisfies

$$B(\hat{r}, r^*, t+s; T+s) = B(\hat{r}, r^*, t; T) \text{ for all } s. \tag{C-2}$$

The most general function satisfying (C-1) and (C-2) has the form

$$B(\hat{r}, r^*, t; T) = \exp(-r\tau) u(\varepsilon, \tau), \tag{C-3}$$

where $\tau = T - t$ is the time remaining until the bond matures and $\varepsilon = r^* - \hat{r}$ is the discrepancy between the preferred rate and the target rate.

If the function B given by (C-3) is to solve the boundary value problem, the function u must solve the simpler boundary value problem comprising the partial differential equation

$$\frac{\partial u}{\partial \tau} = \frac{1}{2}\sigma^2 \frac{\partial^2 u}{\partial \varepsilon^2}, \quad -b < \varepsilon < b,$$

together with the boundary conditions

$$u(b, \tau) = e^{-\Delta\tau} u(b - \Delta, \tau), \quad u(-b, \tau) = e^{\Delta\tau} u(-b + \Delta, \tau),$$

and the initial condition $u(\varepsilon, 0) = 1$. This problem is readily solved using the Crank-Nicholson finite difference method.

This functional form has a natural interpretation involving the yield on the discount bond:

$$\frac{-1}{T-t} \log B(\hat{r}, r^*, t; T) = \hat{r} + \frac{-1}{\tau} \log u(\varepsilon, \tau).$$

The yield on a discount bond with time τ until maturity equals the target rate plus an amount which depends on the bond's maturity and the extent to which the target rate deviates from the preferred rate. Thus, the current target rate determines the level of the yield curve, while the discrepancy between the target rate and the preferred rate determines the spread, and hence the shape of the yield curve.

D Appendix for Section 5

D.1 Tightening and Loosening Cycles

It is easily shown, by mimicking the derivation of equation (1), that whenever $-b_1 < \varepsilon < b_2$, the central bank's loss function satisfies the differential equation

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

The solution is

$$u(\varepsilon) = \frac{\sigma^2}{\rho^2} + \frac{2\mu^2}{\rho^3} + \frac{2\mu\varepsilon}{\rho^2} + \frac{\varepsilon^2}{\rho} + A_1 e^{\lambda_1 \varepsilon} + A_2 e^{\lambda_2 \varepsilon},$$

where

$$\lambda_1 = \frac{-\mu}{\sigma^2} + \sqrt{\left(\frac{\mu}{\sigma^2}\right)^2 + \frac{2\rho}{\sigma^2}}, \qquad \lambda_2 = \frac{-\mu}{\sigma^2} - \sqrt{\left(\frac{\mu}{\sigma^2}\right)^2 + \frac{2\rho}{\sigma^2}},$$

and A_1 and A_2 are constants to be determined. If $\varepsilon \leq -b_1$, the central bank lowers the target rate, resetting the discrepancy to $-b_1 + \Delta_1$, so that

$$u(\varepsilon) = u(-b_1 + \Delta_1) + f + k(-b_1 + \Delta_1 - \varepsilon).$$

If $\varepsilon \geq b_2$, the central bank raises the target rate, resetting the discrepancy to $b_2 - \Delta_2$, so that

$$u(\varepsilon) = u(b_2 - \Delta_2) + f + k(\varepsilon - b_2 + \Delta_2).$$

The constants A_1 and A_2 are determined by the condition that u is continuous at $-b_1$ and b_2.

D.2 Anticipated and Unanticipated Target Changes

The problem is simplified by noting its inherent symmetry: the situation when the last target change was a loosening and the preferred rate currently exceeds the target rate by ε is equivalent to the situation when the last target change was a tightening and the target rate currently exceeds the preferred rate by ε. If the optimal policy in the first situation is to raise the target rate by Δ, the optimal policy in the second situation must be to lower the target rate by Δ. It follows that $u_L(\varepsilon)$, the central bank's loss function when the last target change was a loosening, and $u_T(\varepsilon)$, the loss function when the last target change was a tightening, are related by $u_L(\varepsilon) = u_T(-\varepsilon)$.

We consider the situation in which the last target change was a loosening in more detail.

- If $-b_c < \varepsilon < b_r$, the target rate will not be changed. As in the standard model, the loss function must satisfy the ordinary differential equation

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

  The general solution is

$$u(\varepsilon) = \frac{\sigma^2}{\rho^2} + \frac{\varepsilon^2}{\rho} - Ae^{\lambda\varepsilon} - Be^{-\lambda\varepsilon},$$

  where A and B are arbitrary integration constants and $\lambda^2 = 2\rho/\sigma^2$.

- If $\varepsilon \leq -b_c$, a further loosening will be triggered. The central bank will lower the target rate, resetting ε to $-b_c + \Delta_c$. The loss function must satisfy

$$u(\varepsilon) = u(-b_c + \Delta_c) + f_c + k_c(-b_c + \Delta_c - \varepsilon).$$

- If $\varepsilon \geq b_r$, a tightening will be triggered. The central bank will raise the target rate, resetting ε to $b_r - \Delta_r$. Since this constitutes a reversal of the central bank's previous policy, the loss function must satisfy

$$\begin{aligned} u(\varepsilon) &= u_T(b_r - \Delta_r) + f_r + k_r(\varepsilon - b_r + \Delta_r) \\ &= u(-b_r + \Delta_r) + f_r + k_r(\varepsilon - b_r + \Delta_r). \end{aligned}$$

Now, the requirement that $\lim_{\varepsilon \downarrow -b_c} u(\varepsilon) = u(-b_c)$ implies that

$$\begin{aligned} \frac{(b_c)^2}{\rho} - Ae^{-\lambda b_c} - Be^{\lambda b_c} &= \frac{(b_c - \Delta_c)^2}{\rho} \\ &\quad - Ae^{-\lambda(b_c - \Delta_c)} - Be^{\lambda(b_c - \Delta_c)} + f_c + k_c\Delta_c. \end{aligned} \tag{D-1}$$

Similarly, the requirement that $\lim_{\varepsilon \uparrow b_r} u(\varepsilon) = u(b_r)$ implies that

$$\begin{aligned} \frac{(b_r)^2}{\rho} - Ae^{\lambda b_r} - Be^{-\lambda b_r} &= \frac{(b_r - \Delta_r)^2}{\rho} \\ &\quad - Ae^{-\lambda(b_r - \Delta_r)} - Be^{\lambda(b_r - \Delta_r)} + f_r + k_r\Delta_r. \end{aligned} \tag{D-2}$$

These two equations give A and B implicitly as functions of $(b_c, b_r, \Delta_c, \Delta_r)$.

Differentiating equation (D-1) with respect to b_r and rearranging shows that

$$\frac{\partial A}{\partial b_r} e^{-\lambda b_c}(e^{\lambda\Delta_c} - 1) + \frac{\partial B}{\partial b_r} e^{\lambda b_c^*}(e^{-\lambda\Delta_c} - 1) = 0.$$

Since the coefficients of the two partial derivatives have opposite signs, it must be the case that the two partial derivatives share the same sign. Differentiating equation (D-1) with respect to Δ_r shows that the Δ_r-derivatives of A and B also have the same sign. If we differentiate equation (D-2) with respect to b_c, a similar argument shows that $\partial A/\partial b_c$ and $\partial B/\partial b_c$ have the same sign. Likewise, differentiating that equation with respect to Δ_c, a similar argument shows that $\partial A/\partial A_c$ and $\partial B/\partial \Delta_c$ have the same sign. It follows that for any change in $(b_c, b_r, \Delta_c, \Delta_r)$,

A and B will either both increase in value, or they will both fall in value. They must both have (local) maxima at the same points.

If we differentiate equation (D-1) with respect to b_c, and set $\partial A/\partial b_c = \partial B/\partial b_c = 0$, we obtain an equation which can be written in the form

$$u'(-b_c) = u'(-b_c + \Delta_c).$$

If we differentiate with respect to Δ_c, instead, and set $\partial A/\partial \Delta_c = \partial B/\partial \Delta_c = 0$, we obtain an equation which can be written in the form

$$0 = u'(-b_c + \Delta_c) + k_c.$$

Now we turn to equation (D-2). Differentiating it with respect to b_r, and setting $\partial A/\partial b_r = \partial B/\partial b_r = 0$, leads to the condition

$$u'(b_r) = -u'(-b_r + \Delta_r).$$

If, instead, we differentiate with respect to Δ_r, and set $\partial A/\partial \Delta_r = \partial B/\partial \Delta_r = 0$, we obtain

$$0 = u'(-b_r + \Delta_r) + k._r$$

These four equations comprise the smooth-pasting conditions for this problem.

D.3 Announcement Dates

If $-b(t) < \varepsilon_t < b(t)$ and the central bank leaves the target rate unchanged for a period of time dt, then Itô's Lemma implies that

$$E_t[u(\varepsilon_{t+dt}, t + dt)] = u(\varepsilon_t, t) + \left(\frac{1}{2}\sigma^2 \frac{\partial^2 u}{\partial \varepsilon^2}(\varepsilon_t, t) + \frac{\partial u}{\partial t}(\varepsilon_t, t)\right) dt + o(dt).$$

Since the central bank discounts future costs at rate ρ,

$$\begin{aligned} u(\varepsilon_t, t) &= \varepsilon_t^2 dt + e^{-\rho\, dt} E_t[u(\varepsilon_{t+dt}, t + dt)] + o(dt) \\ &= u(\varepsilon_t, t) + \left(\varepsilon_t^2 + \frac{1}{2}\sigma^2 \frac{\partial^2 u}{\partial \varepsilon^2}(\varepsilon_t, t) + \frac{\partial u}{\partial t}(\varepsilon_t, t) - \rho u(\varepsilon_t, t)\right) dt + o(dt). \end{aligned}$$

Taking the limit as $dt \to 0$, and rearranging, gives us the following partial differential equation for the loss function:

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

References

Ait-Sahalia, Y. "Nonparametric Pricing of Interest Rate Derivative Securities," *Econometrica*, 1996a, 64(3), pp. 527–560.

AIT-SAHALIA, Y. "Testing Continuous-Time Models of the Spot Interest Rate," *Review of Financial Studies*, 1996b, 9(2), pp. 385–426.

BANK FOR INTERNATIONAL SETTLEMENTS. *68th Annual Report.* Basle, Switzerland: Bank for International Settlements, 1998.

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.

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. *The Art of Smooth Pasting*, Vol. 55 in *Fundamentals of Pure and Applied Economics.* Eds. JACQUES LESOURNE AND HUGO SONNENSCHEIN. 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.

EIJFFINGER, SYLVESTER, SCHALING, ERIC, AND VERHAGEN, WILLEM. "A Theory of Interest Rate Stepping: Inflation Targeting in a Dynamic Menu Cost Model," CEPR Discussion Paper No. 2168, June 1999.

FILIMON, ARTEMIZA. "The Optimal Form of Interest Rate Targeting with Announcement Dates," Master of Science thesis, University of Auckland (April 2000).

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. "The Optimal Design of Interest Rate Target Changes," 1999, mimeo, University of Auckland.

HARRISON, J. MICHAEL; SELLKE, THOMAS M. AND TAYLOR, ALLISON J. "Impulse Control of Brownian Motion," *Mathematics of Operations Research*, 1983, 8(3), pp. 454–466.

Hsu, C. and Kugler, P. "The Revival of the Expectations Hypothesis of the U.S. Term Structure of Interest Rates." *Economic Letters*, 1997, 55, pp. 115–120.

Karatzas, Ioannis and Shreve, Steven E. *Brownian Motion and Stochastic Calculus* (2nd ed.). New York: Springer-Verlag, 1991.

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.

Saha, Atanu and Hilton, Lynette. "Expo-Power: A Flexible Hazard Function for Duration Data Models," *Economics Letters*, 1997, 54, pp. 227–233.

Tanner, Martin and Wong, Wing. "Data-Based Nonparametric Estimation of the Hazard Function with Applications to Model Diagnostics and Exploratory Analysis," *Journal of the American Statistical Association*, 1984, 79, pp. 174–182.

Woodford, Michael. "Optimal Monetary Policy Inertia," 1999, mimeo, Princeton University.