Interest Rate Dynamics and Interest Rate Targeting*

Graeme Guthrie
Victoria University of Wellington

Julian Wright†
University of Auckland

*We would like to thank Donggyu Sul and Jun Yu for their helpful comments. All errors are ours.

†Corresponding author. Department of Economics, University of Auckland, Auckland, New Zealand. Ph: 64-9-373-7999 extn 7943. Fax: 64-9-373-7427. Email: jk.wright@auckland.ac.nz

Abstract

This paper derives the dynamic properties of interest rates at the short end of the yield curve using a model of central bank interest rate targeting. The model predicts that conditional volatility is persistent and increasing in the spread between the market rate and the central bank's target rate, that there is excess kurtosis in high frequency interest rate movements, and that market rates will revert towards the central bank's target rate. We show these effects are present in a sample of recent daily data on U.S. interest rates. The results suggest that the spread between market rates and the central bank's target rate should be considered in empirical models of interest rate dynamics.

JEL classification: E43, E52, G12

Keywords: interest rates, fed-funds rate, dynamics, volatility

Interest Rate Dynamics and Interest Rate Targeting

Understanding the dynamics of short-term interest rates is fundamental to much of theoretical and empirical finance. Among other things, it is an important ingredient in pricing bonds and bond derivatives, hedging interest rate risk, determining term structure behavior, and forecasting interest rate movements. The main approach adopted by the voluminous literature in finance that models the dynamics of interest rates is to try to estimate the parameters of a discrete approximation to some continuous time stochastic process. Most such models, including those used by Merton (1973), Vasicek (1977), and Cox et al. (1985) can be expressed as special cases of the stochastic differential equation

$$dr_t = (\alpha + \lambda r_t)dt + \sigma r_t^{\gamma} d\xi_t, \tag{1}$$

where r_t is the level of the instantaneous interest rate, ξ_t is a Wiener process, and α, λ, σ, and γ are parameters.

Most empirical studies focus on determining whether there is mean reversion ($\lambda < 0$), as well as the nature of the level effect in conditional volatility (the size of γ). Chan et al. (1992) survey models within the class above. Using GMM estimation on the U.S. one-month Treasury bill yield from June 1964 to December 1989, they find that γ is 1.50 and highly significant, while $\lambda = -0.59$ but not significant. Stronger evidence for mean reversion in interest rates is obtained by Dahlquist (1996) who examines four OECD countries using an approach similar to Chan et al., and by Wu and Zhang (1996) who use a panel approach with twelve OECD countries for the period from January 1978 to September 1994.

Recent work has focused on nonparametric estimation of the drift function $\mu(r_t)$ and the diffusion function $\sigma(r_t)$ in the generalization

$$dr_t = \mu(r_t)dt + \sigma(r_t)d\xi_t$$

of (1). Assuming affine drift, Ait-Sahalia (1996a) finds that a linear specification for σ^2 is a good approximation for interest rates from 0 to 9%, that σ^2 increases faster than r for medium values of interest rates (9% to 14%), and that it flattens and then decreases for higher rates. Stanton (1997) estimates both the drift and the diffusion nonparametrically. He finds that, while the estimated diffusion is similar to that estimated by Chan et al., there is evidence of substantial non-linearity in the drift. In fact, mean reversion only starts to appear once interest rates are well above 15%, but quickly becomes strong for higher interest rates.

In this paper, we explore interest rate dynamics from a new perspective. We consider how central bank interest rate targeting affects interest rate dynamics. Thus, we do not start by assuming some stochastic process which captures the features of drift and the level effect in volatility. Rather we start by assuming the simplest possible specification for the underlying interest rate — a simple Brownian motion. This interest rate, which we call the preferred rate, is interpreted in our model as the instantaneous interest rate which would be preferred by the central bank, absent any costs of actually implementing it. We then characterize the nature of the target rate, the rate the central bank actually sets, by giving a rule which specifies how the preferred rate determines the central bank's target rate. As Guthrie and Wright (1999) show, this rule has three nice properties. Firstly, it is the optimal rule for the central bank to follow, given costs that the central bank is assumed to face in adjusting its target rate. Secondly, it generates properties of the target rate that closely match the dynamic properties of actual target changes by the U.S. Federal Reserve. Finally, its simple structure allows the dynamic behavior of market interest rates to be determined.

When combined with the expectations hypothesis, this framework implies a rich array of dynamic properties for market rates at short horizons. Among other things, the model predicts that conditional volatility is persistent and increasing in the spread between the market rate and the central bank's target rate, that there is excess kurtosis in high frequency interest rate movements, and that market rates will revert towards the central bank's target rate. These results are surprising given our assumption of an underlying simple Brownian motion process. However, not surpisingly given our assumptions, our model of interest rate targeting does not explain the reversion to mean or the level effect in volatility that is discussed above. Thus we see our framework as complementary to standard models of interest rates. By assuming the central bank's preferred rate incorporates mean reversion and a level effect in volatility, we suspect our model will be able to capture these two features, in addition to the other features the model currently explains.

To derive specific predictions, the model is calibrated to data on actual target changes by the Federal Reserve from March 1984 to February 1994, the period of interest rate targeting that most closely matches the assumptions of our model. The calibrated model is used to generate simulated daily observations on the central bank's target rate, as well as corresponding seven-day, one-month and three-month market rates. Some dynamic properties of these rates are then compared to those found using daily data on the Federal Reserve's fed-funds target rate, and the seven-day, one-month, and three-month Eurodollar rates for the same period.

The results are generally supportive of the predictions of our model, and suggest new empirical models of interest rate dynamics should be considered. Our model predicts almost precisely the degree to which the spread between the market rate and target rate declines through time. This occurs not just because the target rate moves to offset large deviations between the market rate and target rate, but more interestingly, absent target changes, market interest rates revert towards the central bank's target rate through time.

Similarly, our model explains persistence in conditional volatility. Volatility clusters natually emerge around times when target changes are likely. According to our model, such times occur when the spread between the target rate and market rates is large. Consistent with the model, we find conditional volatility increases with the magnitude of the deviation between the target rate and market rates. Compared to the size of the level-effect in conditional volatility, the spread-effect is an order of magnitude more important in the actual data.

Because our model implies occasional discrete target changes, excess-kurtosis of interest rate changes is also predicted, especially for high frequency data. The same pattern in results is apparent in the actual data, with daily changes in interest rates exhibiting higher levels of Kurtosis than monthly changes. In fact, for daily interest rate changes, the excess-kurtosis in the actual data is substantially larger than that predicted by our model, suggesting the model is only able to partially explain the patterns in the actual data.

With this in mind, we consider an extension to our model incorporating the fact that the market will not generally know the central bank's preferred rate, but rather will have to infer this from a publicly available variable that is correlated with the central bank's preferred rate. Using this information, as well as observed target changes, investors form rational expectations about future target changes — expectations that determine current market rates. Introducing learning does not change the basic qualitative predictions of the model, but does change some quantitative results. In most cases, the quantitative predictions are closer to those observed empirically. In particular, the extended model does a better job of explaining the high kurtosis of interest rate changes in the actual data. The introduction of private information in the model means that market rates will sometimes react by large amounts to a target change. This shows up most prominently in the excess-kurtosis results for daily changes in market rates, where the kurtosis in the actual data is also at its highest. Nevertheless, the model is still unable to explain a large part of the high kurtosis of daily changes in seven-day and one-month rates.

The rest of paper proceeds as follows. Section 1 reviews our model of central bank interest rate targeting and, in Section 1.1, demonstrates how it can be combined with the expectations hypothesis to generate interest rate dynamics. Section 1.2 enriches the model by allowing the central bank to have private information. The data used to test the model is detailed in Section 2, while Section 3 presents the empirical results. Finally, Section 4 concludes.

1 Modeling Interest Rate Dynamics

The starting point for our model of the term structure of interest rates is the specification of the targeting rule followed by the central bank. We adopt a targeting rule which not only closely matches the dynamic properties of target rates in practice, but is known to be optimal under certain conditions. We first outline these conditions.

We suppose the central bank has a preferred level of the instantaneous rate at each point in time, denoted r_t^*, which evolves according to driftless Brownian motion $dr_t^* = \sigma d\xi_t$, where ξ_t is a Wiener process and σ is a constant. The preferred rate takes into account all the factors that determine the central bank's desired level of the instantaneous rate, except any cost of changing the target rate itself. We suppose that, by choosing the level of the target rate $\hat{r}_t$, the central bank ties down the instantaneous market rate at this same level. In practice, the market rate will deviate from the target level due to transitory liquidity shocks. However, due to their transitory nature, these shocks should have little, if any, impact on the rates we are interested in.[1] The central bank is assumed to suffer flow costs $(r_t^* - \hat{r}_t)^2 dt$ over a period of length dt caused by deviations between its preferred rate and the actual instantaneous rate; deviations

Figure 1: The interest rate targeting rule

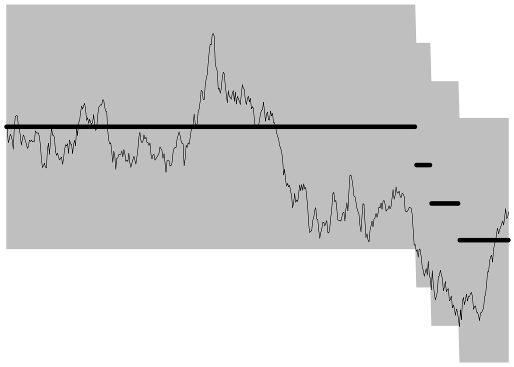

The thin curve plots a particular path of the preferred rate, and the thick curve plots the path of the target rate chosen by the central bank. The shaded region shows the band around the target rate which determines when target changes occur.

are costly since they imply interest rates will be out of line with its objectives. The central bank could eliminate these deviations by having its target rate track movements in its preferred rate. However, doing so is costly. We assume the central bank faces two types of costs of adjusting the target rate by an amount Δ: a fixed adjustment cost f that is incurred whenever the target rate is changed, but which does not depend on the size of the target change, and a proportional adjustment cost $c|\Delta|$, which reflects that large target changes will be more costly than small ones. One motivation for these adjustment costs is that they reflect the risk that markets will react adversely to target changes, even small ones. Other motivations are provided by Guthrie and Wright (1999). Finally, it is assumed the central bank discounts future costs at a rate δ.

Under these assumptions, Guthrie and Wright (1999) derive the optimal targeting rule. The optimal rule is particularly simple, and is completely described by two constants, b and Δ. As long as the preferred rate lies within $(\hat{r}_t - b, \hat{r}_t + b)$, the central bank leaves the target rate fixed. As soon as $r_t^* = \hat{r}_t + b$, the central bank raises the target rate by an amount Δ. Similarly, as soon as $r_t^* = \hat{r}_t - b$, the central bank lowers the target rate by an amount Δ. Figure 1 illustrates the targeting rule for a particular path of the preferred rate. The behavior of the target rate is completely described by b, Δ, and σ. Moreover, these parameters can be calibrated so that the average size of target changes, the probability of policy reversals, and the average time between target changes match their empirical counterparts. Despite the simplicity of the adjustment rule, Guthrie and Wright show that the model replicates a number of puzzling features of interest rate targeting. In particular, the targeting rule leads to occasional discrete target changes which more often than not are in the same direction as previous target changes (continuations) rather than in the opposite direction (reversals). Moreover, the expected time between continuations is shorter than the time between reversals.

To determine the implications of this targeting behavior for the yield curve, the simple targeting rule above is combined with the expectations hypothesis.[2] Due to the continuous time framework we consider, we use the local version of the expectations hypothesis. Thus, the price at time t of any asset with a single payoff at time T equals

$$E_t\left[\exp\left(-\int_t^T \hat{r}_s ds\right) \times \text{payoff}\right], \quad t \leq T, \tag{2}$$

where the expectation is calculated using all information at time t.

We consider two cases. In the first, investors observe the central bank's preferred rate.

1.1 Public Information Case

When the level of the preferred rate is public information, the expectation in (2) is conditional on the level of the preferred rate at time t, as well as the level of the target rate at that time. Denote by $B^o(\hat{r}, r^*, t; T)$ the price at time t of a discount bond paying 1 at time T, where $\hat{r}$ and r^* are the prevailing levels of the target rate and preferred rate respectively. The 'o' is to remind us that the preferred rate is observable. Then

$$B^o(\hat{r}, r^*, t; T) = E\left[\exp\left(-\int_t^T \hat{r}_s ds\right) \Big| \hat{r}_t = \hat{r}, r_t^* = r^*\right]. \tag{3}$$

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 $dr_s^* = \sigma d\xi_s$. By Itô's Lemma, the change in the price of the bond over the next time increment is

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

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

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

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,

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

Consideration of target changes in the opposite direction shows that

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

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

We show in Appendix A.1 that the spread between the yield on a discount bond and the target rate equals $f(\epsilon, \tau)$ for some function f of maturity τ and the current discrepancy $\epsilon = r^* - \hat{r}$ between the central bank's preferred rate and its target rate. Figure 2, which plots the spread

Figure 2: Spreads when the preferred rate is observable

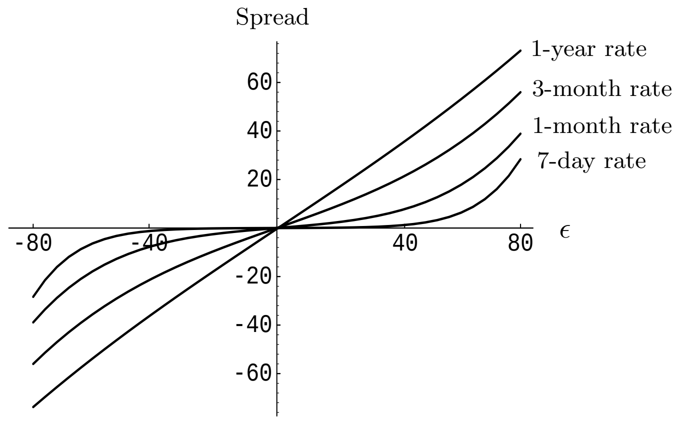

Each curve gives the difference (measured in basis points) between the corresponding interest rate and the target rate as a function of the difference ϵ (measured in basis points) between the preferred rate and the target rate.

for the calibration described in Section 3, demonstrates how market interest rates incorporate anticipated future behavior of the target rate. As $\epsilon \to b$ (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 $\epsilon \to -b$, 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.

Three properties follow immediately from this behavior. Firstly, when the spread between a market rate and the target rate is large, there is a strong tendency for the spread to narrow. For example, suppose the preferred rate is near the top edge of the band. If the preferred rate increases further, a target change will occur: the target rate will climb and, since much of the effect of the target change has already been factored into bond prices, the market rate will climb by a smaller amount, causing the spread to fall. On the other hand, if the preferred rate falls, the market rate will fall with it. Since the target rate will be unchanged, the spread falls. When the spread between the market rate and the target rate is small, the market rate tends to remain close to the target rate because small changes in the preferred rate have little effect on the likelihood of a target change. Thus, the model predicts that the spread will tend to revert back to zero.

Since the preferred rate is assumed to follow a simple Brownian motion, the daily change in this variable is normally distributed. The non-linear relationship between the preferred rate and the market rate evident in Figure 2 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. Thus, the model predicts frequent small daily changes, and occasional large daily changes, in market rates. Kurtosis will be greater for market rates with shorter maturity since, as is clear from Figure 2, the relationship between the preferred rate and such market rates is most non-linear.

The third property also results from the non-linear 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. From the discussion in the previous paragraph, this is when market rates are most volatile. Furthermore, the magnitude of the spread between the market rate and the target rate will have a greater impact on conditional volatility for very short term interest rates, since the relationship between the preferred rate and these interest rates is highly non-linear.

The quantitative implications of these predictions are determined, and tested, in Section 3.

1.2 Private Information Case

Market participants will not generally know the exact level of the instantaneous rate preferred by the central bank. The best they can do is to use publicly available information to draw inferences about the preferred rate. Interest rates of all maturities will move as more information becomes publicly available, altering the market's beliefs about the central bank's preferred rate. Target changes will lead to discrete jumps in interest rates of all maturities; changes in market rates will be smaller when there are no target changes.

As before, the preferred rate evolves according to $dr_t^* = \sigma d\xi_t$. Now, however, we suppose that the value of r^* is only ever known to the central bank; it cannot be observed by the market. The market can observe the variable y_t, which evolves according to

$$dy_t = dr_t^* + \phi d\eta_t,$$

where ϕ is a constant and $d\eta_t$ is the increment of a Wiener process uncorrelated with $d\xi_t$. That is, the market observes the sum of the actual change in the preferred rate and a noise term. The correlation coefficient of changes in r^* and y,

$$\rho = \frac{\sigma}{\sqrt{\sigma^2 + \phi^2}},$$

is a measure of the central bank's openness — when $\rho = 0$ no useful information about the preferred rate is publicly available, whereas if $\rho = 1$ the central bank's preferred rate is public information and the yield curve behaves in the manner described in Section 1.1.

The exact level of the preferred rate is revealed at the time of a target change: if the target rate is raised at time t, the preferred rate must equal $r_t^* = \hat{r}_t + b - \Delta$, where $\hat{r}_t$ is the new target rate; if the target rate is lowered at time t, the preferred rate must equal $r_t^* = \hat{r}_t - b + \Delta$. Between target changes, the market must rely on public information to update its beliefs. For example, suppose a target change occurs at time t_0. If the market ignores any subsequent public information, its time t_1 estimate of the preferred rate at time t_1 equals $r^*_{t_0}$; the prediction error is normally distributed with variance $\sigma^2(t_1 - t_0)$. (This follows immediately from the process generating the preferred rate.) However, this prediction does not use the information contained in the observable variable y. The change observed in this variable, since the date when the target rate was adjusted, equals

$$y_{t_1} - y_{t_0} = r^*_{t_1} - r^*_{t_0} + \phi(\eta_{t_1} - \eta_{t_0}).$$

Rearranging this expression gives

$$r^*_{t_0} + (y_{t_1} - y_{t_0}) = r^*_{t_1} + \phi(\eta_{t_1} - \eta_{t_0}),$$

where the left hand side is observable at time t_1 and the term in brackets on the right hand side is normally distributed with mean zero and variance $t_1 - t_0$. The market can use this information to refine its beliefs regarding the current level of the preferred rate. Its posterior estimate of the preferred rate equals

$$\frac{\phi^2}{\sigma^2 + \phi^2} r^*_{t_0} + \frac{\sigma^2}{\sigma^2 + \phi^2}(r^*_{t_0} + y_{t_1} - y_{t_0}) = r^*_{t_0} + \frac{\sigma^2}{\sigma^2 + \phi^2}(y_{t_1} - y_{t_0})$$

and the prediction error is normally distributed with variance

$$\frac{\sigma^2\phi^2(t_1 - t_0)}{\sigma^2 + \phi^2} = \sigma^2(1 - \rho^2)(t_1 - t_0). \tag{4}$$

We will use this result to calibrate the model in Section 3.

The path of the target rate contains information which the market can use to further refine its beliefs. For each time t, let $\theta(t)$ denote the date of the most recent target change, measured from time t. If the target rate equals $\hat{r}$ at time t, the market knows that the preferred rate has remained inside the interval $[\hat{r} - b, \hat{r} + b]$ throughout the period $[\theta(t), t]$. If the target rate was raised at time $\theta(t)$, the distribution of r^*_t has density function

$$g^+(r^*, t; \hat{r}) = \frac{\psi^+(r^*, t; \hat{r})}{\displaystyle\int_{\hat{r}-b}^{\hat{r}+b} \psi^+(r', t; \hat{r}) dr'}, \quad \hat{r} - b < r^* < \hat{r} + b, \tag{5}$$

where $\psi^+(v, t; \hat{r})$ is the density function for the variable v_t which evolves according to

$$dv_{t'} = \sigma d\xi_{t'}, \quad \theta(t) \le t' \le t, \quad v_{\theta(t)} = \hat{r} + b - \Delta,$$

with absorbing barriers at $v = \hat{r} \pm b$. The function $g^+(v, t; \hat{r})$ describes the distribution of v_t, conditional on neither barrier having been reached during $[\theta(t), t]$. As demonstrated in Appendix A.2, ψ^+ must satisfy

$$\begin{aligned} \psi^+(v, t'; \hat{r}) &= \psi^+(v, \theta(t); \hat{r}) + \frac{1}{2}\sigma^2 \int_{\theta(t)}^{t'} \frac{\partial^2 \psi^+(v, s; \hat{r})}{\partial v^2} ds \\ &\quad - \rho^2 \int_{\theta(t)}^{t'} \frac{\partial \psi^+(v, s; \hat{r})}{\partial v} dy_s, \quad \hat{r} - b < v < \hat{r} + b, \quad \theta(t) \le t' \le t, \end{aligned} \tag{6}$$

together with the boundary conditions

$$\lim_{v \to \hat{r} \pm b} \psi^+(v, t'; \hat{r}) = 0, \quad \theta(t) \le t' \le t. \tag{7}$$

and the initial condition

$$\psi^+(v, \theta(t); \hat{r}) = \delta(v - (\hat{r} + b - \Delta)), \quad \hat{r} - b < v < \hat{r} + b, \tag{8}$$

where $\delta(\cdot)$ is the Dirac delta function.

If the target was lowered at time $\theta(t)$, we would use $\psi^-(v, t; \hat{r})$, the density function for the variable v_t, but with the initial condition $v_{\theta(t)} = \hat{r} - b + \Delta$, rather than $v_{\theta(t)} = \hat{r} + b - \Delta$. This function also satisfies equations (6) and (7), but the initial condition (8) is replaced with

$$\psi^-(v, \theta(t); \hat{r}) = \delta(v - (\hat{r} - b + \Delta)), \quad \hat{r} - b < v < \hat{r} + b.$$

Let $B(\hat{r}, t; T)$ denote the price at time t of a discount bond paying 1 at time T $(\ge t)$, where $\hat{r}$ is the prevailing target rate. Using the law of iterated expectations, the local expectations hypothesis implies that

$$B(\hat{r}, t; T) = E\left[B^o(\hat{r}, r^*, t; T) \middle| \text{information available at time } t \right].$$

This equals

$$B(\hat{r}, t; T) = \int_{\hat{r}-b}^{\hat{r}+b} g^+(r^*, t; \hat{r}) B^o(\hat{r}, r^*, t; T) dr^* \tag{9}$$

if the last target change was a tightening, and

$$B(\hat{r}, t; T) = \int_{\hat{r}-b}^{\hat{r}+b} g^-(r^*, t; \hat{r}) B^o(\hat{r}, r^*, t; T) dr^* \tag{10}$$

if it was a loosening.

Through equations (9) and (10), any new information which changes the market's view of the current level of the preferred rate will change market interest rates. In particular, market rates of all maturities will jump at the time of a target change, reflecting the fact that target changes reveal the true level of the preferred rate to the market. These discrete interest rate movements will further increase the kurtosis predicted in Section 1.1.

In the period following a target change, market beliefs are determined by the nature of the most recent target change (through the initial condition (8)), and the entire path of the observable variable since that target change (through the dy_s term in (6)). Therefore, the model predicts path dependence in interest rate changes.

Moreover, because the change in the preferred rate is unobservable, there is a tendency for the distribution of possible preferred rates to shift back towards the target rate as the time since the last target change grows — if the target rate is unchanged, the preferred rate is more likely to have moved towards the target rate (that is, towards the center of the band) than away from it. Market rates respond by moving closer to the target rate. This general trend compounds the reversion of market rates toward the target rate predicted in Section 1.1.

2 Data

Our data set consists of daily observations on seven-day, one-month and three-month Eurodollar rates, together with the fed-funds target rate, from March 1, 1984 to January 31, 1994. The fed-funds target rate data we use was provided by Brian Sack at the Board of Governors of the Federal Reserve. The one- and three-month Eurodollar spot rates are bid-side rates quoted in London which are collected around 9.30am Eastern time, and which were also provided courtesy of 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 2,463 observations over our sample.

We use Eurodollar rates since they are comparable to the fed-funds target rate. The fed-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 with the market target rate.

The period we use is chosen since it matches the assumptions of our model best. From February 1994 onwards, the Federal Reserve adopted a policy of announcing at every FOMC meeting its decision concerning interest rates.[3] This had the effect of shifting target changes to fall primarily on FOMC meeting dates, the timing of which is announced in advance. Our model does not incorporate this feature, and so only strictly applies to the period up to February 1994.[4] 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, March 1, 1984 is used as our start date.[5]

3 Results

To obtain precise predictions from our theoretical model, we first calibrate it to data on fed-funds target changes. During our sample period, the average size of target changes was 22.48 basis points. Thus we set Δ from our optimal targeting rule discussed in Section 1 to be equal 22.48. For the same rule, Guthrie and Wright (1999) show that the average time between a target change equals $\Delta(2b - \Delta)/\sigma^2$, and the probability of a continuation is $1 - \Delta/2b$. For our sample period, these quantities are 21.43 days and 0.8586 respectively, implying the parameterization $b = 79.47$ basis points and $\sigma = 189.1$ basis points (annualized). The remaining parameter ρ determines how well the market can predict changes in the central bank's preferred rate. We calibrate the model by supposing that the prediction error one month after a target change, assuming that the prediction uses only the path of the observable variable y, equals 25 basis points. Using equation (4), we find that $\rho = 0.8890$.

Using this calibrated model, we generate a series of paths for the preferred rate r^* and the observable variable y using an Euler approximation with each day divided into 24 discrete steps. For each step, we solve equation (6) to obtain the market's beliefs about the level of the preferred rate. We then calculate start of day target rates and corresponding seven-day, one-month, and three-month market rates. We do this for two different assumptions on the observability of the preferred rate ($\rho = 1$ and $\rho = 0.8890$). In each case, we simulate data for 2.463 million days, which corresponds to one-thousand replications of 2,463 observations (the number of observations of actual data from March 1, 1984 to January 31, 1994).

We focus on the three properties of the model discussed in Section 1.1. These are: (1) the excess-kurtosis in high frequency interest rate changes; (2) the reversion of the spread between market rates and the central bank's target rate to zero; and (3) the greater interest rate volatility during times when market rates are far away from the central bank's target rate. In each case, we compare the results using the model-generated data with the result using our actual data. For each property considered, we also consider three hypothesis tests on the model using actual data — the first tests the qualitative predictions of the model discussed in Section 1.1; the second tests the precise predictions of the public information version of the calibrated model; while the third tests the precise predictions of the private information version of the calibrated model. The p-values from these tests are given in the last set of columns of each of the tables below.

Before testing the dynamic properties implied by our model, we first present results on the stationarity, or otherwise, of the variables we use in these tests. Table 1 reports the results on λ from the standard augmented Dickey-Fuller regression,

$$X_t = \alpha + \lambda X_{t-1} + \sum_{j=1}^{q} \theta_{t-j} \Delta X_{t-j} + \epsilon_t, \tag{11}$$

where X_t represents the variable examined. The parameter q is chosen by the AIC2 information criterion, as described in Pantula et al. (1994). We hold q fixed at the level chosen based on the actual data set. The first four rows of Table 1 give results on the level of interest rates. Given the Brownian motion assumption we make, the model implies a unit root in the level of the seven-day, one-month and three-month interest rates, as well as the target rate. The coefficient λ from regressing (11) using 2.463 million observations is equal to 1.0000 (to four decimal places) in both the public information (λ_1) and private information (λ_3) cases. Taking the average of results from 1000 replications, each using 2,463 observations, implies a slightly lower level of λ in both cases (compare λ_2 and λ_4 with λ_1 and λ_3 respectively), suggesting there is some downward bias caused by using a smaller sample. Given this degree of bias, the estimates of λ arising from the actual data (λ_5) are consistent with a unit root in the level of the market and target interest rates. Statistically, this is confirmed by tests of a unit root null (λ equals one) using actual data. The p-values from this test, reported in the last column of the table, are 0.890 for the seven-day rate, 0.798 for the one-month rate, 0.859 for the three-month rate, and 0.556 for the target rate. Thus we cannot reject a unit root for any of the level variables over our sample.[6]

*** Table 1 about here ***

The remaining nine rows of Table 1 present estimates of λ for spreads between market rates and the central bank target rate, and for differenced interest rates.[7] Clearly, even with the full sample of model-generated data, the value of λ is well below unity.[8] The results using actual data paint a similar picture. The highest value of λ is 0.967. In all cases we can easily reject a unit root null, even at the 1% level. This suggests non-stationarity is not a problem for our specifications below, in which we use differenced interest rates and spreads between market and target rates.

3.1 The Kurtosis of Interest Rate Changes

A clear prediction of our model is that there should be excess-kurtosis in daily changes in interest rates. In particular, the model predicts frequent small daily changes, and occasional large daily changes, in market rates. These predictions are strongest when we allow for private information. For lower frequency changes and for longer maturity rates, the excess-kurtosis should become less pronounced. These predictions are evident from Table 2 which presents the kurtosis for the change in interest rates at various frequencies (daily, weekly, and monthly) for both the public and private information versions of the model.

*** Table 2 about here ***

The first column of results in Table 2 reports the kurtosis calculated from the full 2.463 million observations generated from the public information model. Comparing column one with column two confirms there is little bias in calculating kurtosis from a sample of the same size as that available for the actual data. The kurtosis of daily changes in the seven-day rate is more than fifteen, suggesting substantial excess-kurtosis. In contrast, monthly changes in the seven-day rate exhibit a kurtosis of less than six, suggesting excess-kurtosis arises primarily from high frequency changes in rates. Using the private information model, columns three and four show that the kurtosis of daily changes in interest rates is substantially greater compared to the level implied by the public information model.

Column five presents the results on kurtosis using the actual data. The results are consistent with the pattern of results predicted by the private information model. Kurtosis is highest for daily changes in short maturity rates.[9]

The p-values from tests of no excess-kurtosis using actual data are given in the seventh column of results. The standard errors used to construct these p-values are formed using the standard delta method (that is, without assuming normality), and are reported in column six of Table 2. For all three interest rates we can reject that there is no excess-kurtosis at the 5% significance level. At the 1% significance level, we can no longer reject no excess-kurtosis for monthly changes in the one-month and three-month rates. These results seem consistent with the basic qualitative prediction of the model, that excess-kurtosis arises in high frequency interest rate movements.[10] A more exacting test of the model's implications is that the kurtosis in the actual data is equal to the precise value of kurtosis predicted by the model. Columns eight and nine give p-values from tests of the kurtosis values predicted by two versions of the model — the public information version and the private information version. Using a 1% significance level for these tests, we cannot reject the model's implications for ten out of the eighteen cases tested. The model with private information gives a generally better fit of the actual data, although even this version of the model is unable to fully capture the very high levels of kurtosis observed for daily changes in the seven-day and one-month Eurodollar rates.

3.2 Target Rate Reversion

The finding that the spread between each of the market rates and the central bank's target rate is stationary suggests that these market rates might revert back to the target rate over time. Table 3 presents results which test this idea. We first regress the spread on a constant and its own lag, showing that indeed the spread between these two variables exhibits reversion to zero. That is, we regress

$$r_t - \hat{r}_t = \alpha + \beta(r_{t-1} - \hat{r}_{t-1}) + \epsilon_t,$$

where r_t is a market rate and $\hat{r}_t$ is the target rate in period t. For both the model-generated and actual data, the estimated value of α is essentially zero. The first, fourth, and seventh rows of results in Table 3 provide information on estimates of β for seven-day, one-month and three-month rates respectively. The public information model predicts β will be approximately 0.72 for seven-day rates, 0.91 for one-month rates and 0.95 for three-month rates, while the private information model predicts β will be 0.82 for seven-day rates, 0.93 for one-month rates and 0.95 for three-month rates.

*** Table 3 about here ***

These implications are borne out by the empirical results. Using actual data, the value of β is 0.80 for seven-day rates, 0.91 for one-month rates and 0.95 for three-month rates. Column seven shows we can easily reject that estimates of β from the actual data are equal to one, confirming statistically there is a tendency for the spread to revert towards zero. Moreover, as columns nine and ten show, we cannot reject that the precise value of β estimated from the actual data is identical to the model-generated data (at the 1% level).

There are two reasons why any spread may be eliminated through time. The market rate might move back towards the target rate or the target rate may be changed in the direction of the existing market rate. Our model of interest rate targeting suggests both types of movements are important. While the movement of target rates to reduce the existing spread between market and target rates simply reflects the fact that market rates anticipate target changes, the reversion of market rates to the target rate is not a well known property of interest rates. To show the latter result arises in our model, we restrict our sample to dates when the target rate was not changed, and run the following regression:

$$r_t - r_{t-1} = \alpha + \beta(r_{t-1} - \hat{r}_{t-1}) + \epsilon_t.$$

As rows two, five and eight of Table 3 show, β is less than zero for both public and private information versions of the model, indicating the market rate reverts towards the target rate in all cases. The private information model implies a substantially higher rate of target-rate reversion compared to the public information model (0.063 compared to 0.026 for one-month rates). With full public information, investors have a better idea of future target changes, and so there is less tendency for market rates to revert to the current target rate, absent target changes, even though the spread reverts to zero more quickly in this case. The model also predicts target rate reversion is much stronger for short maturity rates.

These predictions are again borne out by the actual data. We find significant target rate reversion for each of the Eurodollar rates tested, with the degree of target rate reversion negatively related to the maturity of rates. The hypothesis that $H_0 : \beta_5 = 0$ can be rejected at the 5% significance level, while the precise predictions of the public and private information models cannot be rejected at the 1% level.

An alternative approach to measuring the degree of target rate reversion, based on the cointegration between the market rates and the target rate, is to estimate an Error Correction Model (ECM). If we allow for contemporaneous target changes in the regression, the coefficient on the error correction term will correspond to the rate of reversion of the market rate to the target rate. The model estimated is

$$r_t - r_{t-1} = \alpha + \beta(r_{t-1} - \hat{r}_{t-1}) + \sum_{i=0}^{p} \psi_i(\hat{r}_{t-i} - \hat{r}_{t-i-1}) + \sum_{j=1}^{q} \theta_j(r_{t-j} - r_{t-j-1}) + \epsilon_t,$$

where the parameters p and q are chosen using BIC based on the actual sample and held constant across all other specifications.[11] The results from this specification are very similar to those obtained from dropping target rate observations and not accounting for higher order lag terms. This suggests any autocorrelation in the error terms in the simpler regression does not account for the strong reversion to the target rate observed.

3.3 Conditional Volatility and the Spread-Effect

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 and market rates is high. To test this hypothesis we estimate a standard model of interest rate volatility, replacing the level of interest rates with the magnitude of the spread between the target and market rates. Thus we estimate the following model using FGLS, following the approach of Ait-Sahalia (1996b, p. 409):[12]

$$\begin{aligned} r_t - r_{t-1} &= \alpha + \kappa r_{t-1} + \epsilon_t, && (12) \\ E_{t-1}\epsilon_t &= 0, && (13) \\ E_{t-1}\epsilon_t^2 &\equiv \sigma_t^2 = \sigma^2 |r_{t-1} - \hat{r}_{t-1}|^{2\gamma}. && (14) \end{aligned}$$

The first four columns of Table 4 confirm the predictions from Section 1.1 — the higher is the magnitude of the spread between market and target rates, the higher is the volatility of market interest rates. Columns five, six and seven of Table 4 confirm that the spread is a significant determinant of conditional volatility in actual practice; in all cases, $\hat{\sigma}$ and $\hat{\gamma}$ are statistically different from zero, even at the 1% significance level. We call this effect the spread-effect.

*** Table 4 about here ***

Although the spread-effect is clearly important regardless of whether model-generated or actual data is used, the estimate of γ from the actual data is substantially larger than estimates from our model-generated data. The implication of these differences is that the functional form which captures the spread-effect has greater curvature using actual data. Put differently, volatility is even more sensitive to the spread in reality than in our model.

It is interesting to compare the estimated spread-effect with the level-effect which arises for the same data set.[13] Using r_{t-1} in place of $|r_{t-1} - \hat{r}_{t-1}|$ in (14), we find that for the seven-day Eurodollar rate, $\hat{\sigma} = 0.022$ (0.016) and $\hat{\gamma} = 0.907$ (0.340), where the standard errors of the estimates are given in parentheses. For one-month and three-month Eurodollar rates, the results are similar: using one-month rates, $\hat{\sigma} = 0.033$ (0.021) and $\hat{\gamma} = 0.628$ (0.310); using three-month rates, $\hat{\sigma} = 0.020$ (0.005) and $\hat{\gamma} = 0.744$ (0.122). Thus, there is some evidence of a level-effect in the data, although the t-statistics from tests that the coefficients are zero are considerably smaller than for the spread regressions. To contrast the spread-effect with the level-effect, we allow both to affect conditional volatility at the same time. Thus, we replace (14) with the equation

$$\sigma_t^2 = \sigma_1^2 |r_{t-1} - \hat{r}_{t-1}|^{2\gamma_1} + \sigma_2^2 r_{t-1}^{2\gamma_2}.$$

For each of the three interest rates, the spread-effect remains highly significant: with seven-day rates, $\hat{\sigma}_1 = 0.439$ (0.023) and $\hat{\gamma}_1 = 1.288$ (0.085); with one-month rates, $\hat{\sigma}_1 = 0.360$ (0.012) and $\hat{\gamma}_1 = 1.151$ (0.037); with three-month rates, $\hat{\sigma}_1 = 0.179$ (0.007) and $\hat{\gamma}_1 = 1.490$ (0.111). On the other hand, the level effect is now only apparent for three-month rates: with seven-day rates, $\hat{\sigma}_2 = 0.066$ (0.048) and $\hat{\gamma}_2 = 0.236$ (0.364); with one-month rates, $\hat{\sigma}_2 = 0.063$ (0.073) and $\hat{\gamma}_1 = -0.032$ (0.622); with three-month rates, $\hat{\sigma}_1 = 0.031$ (0.009) and $\hat{\gamma}_1 = 0.454$ (0.143).

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 0.171 for the seven-day Eurodollar rate, 0.260 for the one-month Eurodollar rate, and 0.111 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 0.024 with the three-month rate. In contrast, when the level-effect is eliminated, the remaining spread-effect implies an R^2 of 0.170 for the seven-day rate, 0.258 for the one-month rate, and 0.091 with the three-month rate. In short, the spread-effect dominates the level-effect on both statistical and economic grounds.[14]

4 Conclusion

In this paper we used a simple structural model of interest rate targeting derived from optimal central bank behavior to develop predictions for market interest rate dynamics. To do this we combined the expectations hypothesis with a model of optimal central bank interest rate targeting. We also explicitly accounted for private sector learning about the central bank's intentions regarding target changes, although we find with a few exceptions this plays a relatively minor role in explaining interest rate dynamics. The model captures three particularly prominent features of observed interest rate dynamics.

All three features arise because the magnitude of the response of market interest rates to new information regarding the central bank's preferred level of the target rate varies with the spread between market rates and the target rate. When this spread is large, the market believes a target change is imminent. Any information which delays an imminent target change will have a large effect on short-term rates, since it makes it more likely that the instantaneous rate will remain at its current level over the life of the corresponding bond. Market rates thus respond to the new information by moving rapidly towards the target rate. However, if no target change is believed to be imminent, as will be the case if the spread between market rates and the target rate is small, the effect of the new information on the path of the instantaneous rate over the life of the bond is small, so there is no reason for short-term rates to respond to the new information. In contrast, long-term rates are less sensitive to the exact timing of the next target change; instead, they depend primarily on the long-run behaviour of the target rate. The spread between market rates and the target rate therefore has little effect on the dynamics of long-term interest rates.

This combination of very sensitive short-term rates during some periods, and much more stable short-term rates during other periods, manifests itself as excess-kurtosis in interest rate changes. Since the magnitude of the response of long-term interest rates to new information is less sensitive to the spread, the behaviour of changes in long-term interest rates will be more uniform over time, so that excess-kurtosis is most prominent in short-term interest rate movements. Excess-kurtosis is particularly prominent in high frequency data since over the course of a day the system is unlikely to change from the high-volatility to low-volatility modes, or vice versa. Thus, the periods of high-volatility can easily be distinguished from the low-volatility periods in the data — excess-kurtosis is prominent, especially for short-term rates. However, when observations are made at weekly or monthly intervals, the system will often be in both modes during the course of a single period. This blurs the distinction between the high-volatility and low-volatility interest rate movements. Not surprisingly, excess-kurtosis is not (so) evident.

In addition, the model does a good job of accounting for the target rate reversion which we find in the data. As market rates partially anticipate target changes, when changes do occur, market rates generally jump by a smaller amount than the target rate, narrowing the spread. When target changes are averted, market rates respond by moving back towards the target rate, again narrowing the spread. We find target rate reversion is considerably stronger for short-term rates than for long-term rates, with the movement of the seven-day Eurodollar rate back towards the current fed-funds target, in the absence of target changes, being as much as ten percent of the previous day's spread between the Eurodollar and the target rate.

Finally, the model explains why conditional volatility increases with the magnitude of the spread between the target rate and market rates. This spread is large precisely when the market believes a target change is imminent. In turn, this occurs when market interest rates are most responsive to new information, and hence exhibit the greatest volatility. For our data set, this spread-effect in conditional volatility is found to be an order of magnitude more important than the much studied level-effect.

Despite the success of our model in explaining these features of the data, there remains much that can be done to extend the model, as well as to apply it in new ways. It would be interesting to take a model which allowed for a slow rate of mean reversion in the underlying preferred rate, and see whether this alters the relatively high frequency interest rate dynamics studied in this paper. More generally, does interest rate targeting induce similar interest rate dynamics to those found here when other processes are considered for the central bank's preferred rate?

An interesting implication of our model is that it generates path dependence in interest rate movements between target changes. Path dependence arises from the learning that takes place when market participants have imperfect information about the central bank's preferred interest rate. Our model could thus prove useful in understanding path dependence in interest rates, as well as the non-diffusion nature of interest rates which Ait-Sahalia (1997) has documented. More obviously, the use of our structural model allows us to examine the implications of changes such as increased macroeconomic volatility, increased costs of adjusting target rates, and increased transparency of central bank intentions. These three changes, which are captured by different parameters of our model, are particularly relevant for comparing interest rate targeting in the 1990s with the 1970s, as well as comparing the implications of interest rate targeting across countries.

Perhaps the most important extension of our model is to incorporate the fact that since February 1994 target changes are most likely to coincide with FOMC meeting dates. Filimon (2000) looks at the case that announcements of target changes on dates other than those associated with FOMC meetings are considered more costly by central banks. The optimal rule in this case implies most announcements will be made at the same time as FOMC meeting dates. Future work could model the implications of this change in interest rate targeting for interest rate dynamics. An interesting possibility is then that the time remaining before an FOMC meeting date could become an important determinant of conditional volatility.[15] In our view, empirical modeling of interest rate processes that incorporate the implications of central bank interest rate targeting will prove to be a fruitful area for further research.

Appendix

A.1 Bond Prices With Public Information

The process for the target rate features symmetry which we are able to exploit in simplifying the bond pricing problem. Imagine the target rate and the preferred rate both increase by the same amount, say δ. Then the distribution of all future levels of the target rate shifts to the right by the same amount. (The assumptions crucial to this result are the form of the process generating the preferred rate and the particular adjustment rule adopted by the central bank.) From (3), 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

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

We can exploit a second symmetry. Assuming the current level of the target rate and the preferred rate are unchanged, the price of a discount bond will not change if the current date and the maturity date of the bond both increase by the same amount, say δ. That is, the bond price function satisfies

$$B^o(\hat{r}, r^*, t+\delta; T+\delta) = B^o(\hat{r}, r^*, t; T) \text{ for all } \delta. \tag{A.2}$$

The most general function satisfying (A.1) and (A.2) has the form.

$$B^o(\hat{r}, r^*, t; T) = \exp(-r\tau)u(\epsilon, \tau), \tag{A.3}$$

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

If the function B^o given by (A.3) is to solve the boundary value problem, the function u must solve the simpler boundary value problem comprising the partial differential equation

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

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

$$u(\epsilon, 0) = 1.$$

This problem is readily solved using the Crank-Nicholson finite difference method.

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

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

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.

A.2 Formulating Market Beliefs

We consider the variable which evolves according to $dv_t = \rho^2 dy_t + dz_t$, where z_t is a driftless Brownian motion satisfying $(dy_t)(dz_t) = 0$ and $(dz_t)^2 = (1-\rho^2)\sigma^2 dt$. We approximate the Brownian motions y_t and z_t by dividing time into discrete periods each of length Δt, and supposing that each period y changes by $\pm\Delta y$ (with each possibility equally likely) and z changes by $\pm\Delta z$ (again, with each possibility equally likely). The magnitudes of these changes are chosen such that $(\Delta y)^2 = (\sigma^2+\phi^2)\Delta t$ and $(\Delta z)^2 = (1-\rho^2)\sigma^2\Delta t$.

Suppose that between times $t-\Delta t$ and t, the public variable changes by Δy. Consequently, if $v_t = v$, then either $v_{t-\Delta t} = v - (\rho^2\Delta y + \Delta z)$ or $v_{t-\Delta t} = v-(\rho^2\Delta y - \Delta z)$, each with probability 1/2. Let $\psi(v,t)$ denote the density function of v_t. Then

$$\begin{aligned}
\psi(v,t) &= \frac{1}{2}\psi(v-(\rho^2\Delta y+\Delta z), t-\Delta t) + \frac{1}{2}\psi(v-(\rho^2\Delta y - \Delta z), t-\Delta t) \\
&= \frac{1}{2}\left(\psi(v,t) - (\rho^2\Delta y + \Delta z)\psi_1(v,t) - \Delta t\psi_2(v,t) + \frac{1}{2}(\rho^2\Delta y+\Delta z)^2\psi_{11}(v,t)\right) \\
&\quad + \frac{1}{2}\left(\psi(v,t) - (\rho^2\Delta y - \Delta z)\psi_1(v,t) - \Delta t\psi_2(v,t) + \frac{1}{2}(\rho^2\Delta y-\Delta z)^2\psi_{11}(v,t)\right) \\
&= \psi(v,t) - \rho^2\Delta y\psi_1(v,t) - \Delta t\psi_2(v,t) + \frac{1}{2}((\rho^2\Delta y)^2 + (\Delta z)^2)\psi_{11}(v,t)
\end{aligned}$$

Since $(\rho^2\Delta y)^2 = \rho^4(\sigma^2+\phi^2)\Delta t = \rho^2\sigma^2\Delta t$ and $(\Delta z)^2 = (1-\rho^2)\sigma^2\Delta t$, this equation reduces to

$$\psi_2(v,t)\Delta t = \frac{1}{2}\sigma^2\psi_{11}(v,t)\Delta t - \rho^2\psi_1(v,t)\Delta y.$$

In the limit as $\Delta t \to 0$, we have

$$\psi_2(v,t)dt = \frac{1}{2}\sigma^2\psi_{11}(v,t)dt - \rho^2\psi_1(v,t)dy_t.$$

Integrating with respect to t results in equation (6).

A.3 Bond Prices With Private Information

It is obvious from the symmetry of the problem that the density function has the following symmetry:

$$g^+(r^*, t; \hat{r}) = g^+(r^*+\delta, t; \hat{r}+\delta) \text{ for all } \delta.$$

Therefore,

$$g^+(r^*, t; \hat{r}) = g^+(r^*-\hat{r}, t; 0).$$

Combining this result with equation (A.3) shows that

$$B(\hat{r}, t; T) = \int_{\hat{r}-b}^{\hat{r}+b} g^+(r^*, t; \hat{r}) B^o(\hat{r}, r^*, t; T)dr^*$$

can be rewritten as

$$B(\hat{r}, t; T) = e^{-\hat{r}(T-t)}\int_{-b}^{b} g^+(\epsilon, t, 0)u(\epsilon, T-t)d\epsilon.$$

That is,

$$B(\hat{r}, t; T) = e^{-\hat{r}(T-t)}v(t, T-t),$$

where

$$v^{+}(t_1,t_2) = \int_{-b}^{b} g^{+}(\epsilon, t_1; 0) u(\epsilon, t_2) d\epsilon.$$

This functional form has the same advantages as that given in equation (A.3). In particular, the yield on a bond with time τ until maturity equals

$$\hat{r} + \frac{-1}{\tau} \log v^{+}(t,\tau),$$

so that the current target rate determines the level of the yield curve, while the market's beliefs regarding the discrepancy between the target rate and the preferred rate determine the spread, and hence the shape of the yield curve.

References

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

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

Ait-Sahalia, Y. "Do Interest Rates Really Follow Continuous-Time Markov Diffusions?" Working Paper. University of Princeton (1997).

Balduzzi, P., G. Bertola, and S. Foresi. "A Model of Target Changes and the Term Structure of Interest Rates." *Journal of Monetary Economics* 39 (1997), 223–249.

Chan, K. C., G. A. Karolyi, F. A. Longstaff, and A. B. Sanders. "An Empirical Comparison of Alternative Models of the Short-Term Interest Rate." *Journal of Finance* 47 (July 1992), 1209–1227.

Cox, J. C., J. E. Ingersoll, and S. A. Ross. "A Theory of the Term Structure of Interest Rates." *Econometrica* 53 (March 1985), 363–384.

Dahlquist, M. "On Alternative Interest Rate Processes." *Journal of Banking and Finance* 20 (1996), 1093–1119.

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

Guthrie, G., and J. Wright. "The Optimal Design of Interest Rate Target Changes." Mimeo, University of Auckland (December 1999).

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

Merton, R. "Theory of Rational Option Pricing." *Bell Journal of Economics and Management Science* 4 (1973), 141–183.

Pantula, S. G., G. Gonzalez-Farias, and W. A. Fuller. "A Comparison of Unit-Root Test Criteria." *Journal of Business and Economic Statistics* 12(4) (October 1994), 449–459.

Rudebusch, G. D. "Federal Reserve Interest Rate Targeting, Rational Expectations, and the Term Structure." *Journal of Monetary Economics* 35(2) (1995), 245–274.

Stanton, R. "A Nonparametric Model of Term Structure Dynamics and the Market Price of Interest Rate Risk." *Journal of Finance* 52 (December 1997), 1973–2002.

Vasicek, O. "An Equilibrium Characterization of the Term Structure." *Journal of Financial Economics* 5 (1977), 177–188.

Wu, Y. and H. Zhang. "Mean Reversion in Interest Rates: New Evidence from a Panel of OECD countries." *Journal of Money, Credit, and Banking* 28(4) (November 1996), 604–621.

Table 1: Unit root tests

Description	Model with $\rho = 1$		Model with $\rho = 0.889$		Actual data	p-value
	Full sample	2463 obs	Full sample	2463 obs	Estimate	$H_0 : \lambda_5 = 1$
	λ_1	λ_2	λ_3	λ_4	λ_5	
r_t^{7d}	1.000	0.999	1.000	0.999	0.999	0.890
r_t^{1m}	1.000	0.999	1.000	0.999	0.999	0.798
r_t^{3m}	1.000	0.999	1.000	0.999	1.000	0.859
$\hat{r}_t$	1.000	0.999	1.000	0.999	0.999	0.556
$r_t^{7d} - r_{t-1}^{7d}$	0.460	0.449	0.408	0.396	-0.697	0.000
$r_t^{7d} - \hat{r}_t$	0.865	0.861	0.908	0.905	0.862	0.000
$\lvert r_t^{7d} - \hat{r}_t \rvert$	0.711	0.692	0.767	0.758	0.737	0.000
$r_t^{1m} - r_{t-1}^{1m}$	0.415	0.397	0.369	0.351	-0.178	0.000
$r_t^{1m} - \hat{r}_t$	0.940	0.938	0.949	0.947	0.921	0.000
$\lvert r_t^{1m} - \hat{r}_t \rvert$	0.822	0.817	0.850	0.845	0.869	0.000
$r_t^{3m} - r_{t-1}^{3m}$	0.061	0.057	0.056	0.052	-0.013	0.000
$r_t^{3m} - \hat{r}_t$	0.956	0.954	0.958	0.956	0.967	0.000
$\lvert r_t^{3m} - \hat{r}_t \rvert$	0.846	0.845	0.866	0.865	0.926	0.000

NOTES: The columns labeled 'Full sample' give results using all 2.463 million simulated observations. The columns labeled '2463 obs' report the average of 1000 estimates each based on 2463 simulated observations. The columns labeled 'Actual data' use 2463 daily observations from March 1, 1984 to January 31, 1994.

Table 2: Kurtosis

Frequency	Model with $\rho = 1$		Model with $\rho = 0.889$		Actual data		p-values		
	Full sample	2463 obs	Full sample	2463 obs	Estimate	Std error	$H_0 : K_5 \leq 3$	$H_0 : K_5 = K_2$	$H_0 : K_5 = K_4$
	K_1	K_2	K_3	K_4	K_5				
$r_t^{7d} - r_{t-1}^{7d}$									
daily	15.135	15.163	21.752	21.858	59.800	10.859	0.000	0.000	0.000
weekly	10.084	9.898	10.355	10.233	23.688	7.959	0.005	0.083	0.091
monthly	5.584	5.310	5.575	5.294	11.477	1.931	0.000	0.001	0.001
$r_t^{1m} - r_{t-1}^{1m}$									
daily	7.376	7.340	15.903	15.871	64.672	9.895	0.000	0.000	0.000
weekly	6.913	6.740	8.034	7.932	17.883	5.337	0.003	0.037	0.062
monthly	4.992	4.735	5.026	4.782	5.076	0.949	0.015	0.719	0.757
$r_t^{3m} - r_{t-1}^{3m}$									
daily	3.734	3.720	10.887	10.797	10.899	1.566	0.000	0.000	0.948
weekly	3.920	3.840	5.657	5.583	8.134	1.588	0.001	0.007	0.108
monthly	3.810	3.642	4.125	3.955	3.764	0.421	0.036	0.772	0.650

Table 3: Target rate reversion

Description			Model with $\rho = 1$		Model with $\rho = 0.889$		Actual data		p-values			
			Full sample	2463 obs	Full sample	2463 obs	Estimate	Std error	$H_0 : \beta_5 = 1$	$H_0 : \beta_5 = 0$	$H_0 : \beta_5 = \beta_2$	$H_0 : \beta_5 = \beta_4$
	y	x	β_1	β_2	β_3	β_4	β_5					
7-day rate												
all obs	$r_t - \hat{r}_t$	$r_{t-1} - \hat{r}_{t-1}$	0.716	0.712	0.823	0.820	0.799	0.040	0.000		0.029	0.599
$\Delta\hat{r}_t = 0$	Δr_t	$r_{t-1} - \hat{r}_{t-1}$	−0.091	−0.094	−0.121	−0.124	−0.183	0.044		0.000	0.043	0.183
all obs (ECM)	Δr_t	x_t^1	−0.121	−0.124	−0.163	−0.166	−0.169	0.038		0.000	0.235	0.934
1-month rate												
all obs	$r_t - \hat{r}_t$	$r_{t-1} - \hat{r}_{t-1}$	0.908	0.906	0.930	0.928	0.907	0.031	0.003		0.980	0.499
$\Delta\hat{r}_t = 0$	Δr_t	$r_{t-1} - \hat{r}_{t-1}$	−0.024	−0.026	−0.061	−0.063	−0.074	0.034		0.028	0.156	0.746
all obs (ECM)	Δr_t	x_t^2	−0.026	−0.027	−0.077	−0.079	−0.075	0.035		0.031	0.171	0.900
3-month rate												
all obs	$r_t - \hat{r}_t$	$r_{t-1} - \hat{r}_{t-1}$	0.946	0.945	0.951	0.950	0.951	0.010	0.000		0.585	0.929
$\Delta\hat{r}_t = 0$	Δr_t	$r_{t-1} - \hat{r}_{t-1}$	−0.023	−0.024	−0.050	−0.052	−0.026	0.010		0.012	0.832	0.013
all obs (ECM)	Δr_t	x_t^3	−0.020	−0.022	−0.058	−0.060	−0.026	0.010		0.010	0.638	0.001

NOTES:

1. x_t^1 is the vector containing $r_{t-1} - \hat{r}_{t-1}$, $\Delta\hat{r}_t$, $\Delta\hat{r}_{t-1}$, Δr_{t-1}.
2. x_t^2 is the vector containing $r_{t-1} - \hat{r}_{t-1}$, $\Delta\hat{r}_t$, $\Delta\hat{r}_{t-1}$, $\Delta\hat{r}_{t-2}$, Δr_{t-1}.
3. x_t^3 is the vector containing $r_{t-1} - \hat{r}_{t-1}$, $\Delta\hat{r}_t$, $\Delta\hat{r}_{t-1}$, $\Delta\hat{r}_{t-2}$, Δr_{t-1}, Δr_{t-2}.
4. All regressions include a constant.

Table 4: Heteroskedasticity in volatility

Description	Model with $\rho = 1$		Model with $\rho = 0.889$		Actual data		p-values		
	Full sample	2463 obs	Full sample	2463 obs	Estimate	Std error	$H_0 : e_5 = 0$	$H_0 : e_5 = e_2$	$H_0 : e_5 = e_4$
	e_1	e_2	e_3	e_4	e_5				
7-day rate									
σ	0.241	0.252	0.281	0.351	0.477	0.018	0.000	0.000	0.000
γ	0.410	0.430	0.455	0.460	1.160	0.064	0.000	0.000	0.000
1-month rate									
σ	0.222	0.221	0.198	0.195	0.367	0.011	0.000	0.000	0.000
γ	0.565	0.561	0.458	0.444	1.132	0.033	0.000	0.000	0.000
3-month rate									
σ	0.135	0.133	0.123	0.122	0.200	0.005	0.000	0.000	0.000
γ	0.286	0.272	0.177	0.171	0.728	0.045	0.000	0.000	0.000

Notes

[1]This is consistent with the findings of Rudebusch (1995), who finds these liquidity shocks are quickly eliminated, and Balduzzi et al. (1997), who show that spreads between short-term rates and the overnight target rate are mainly driven by expectations of changes in the target and not by the transitory dynamics of the overnight rate around the target rate.

[2]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 calibrate our model. Their result is robust across different frequencies of interest rates, including the frequency we focus on — daily observations.

[3]We have also estimated the models discussed below using actual data on one-month and three-month Eurodollar rates for the longer period March 1984 to June 1999, and obtained very similar results to those reported. Seven-day Eurodollar rates are not available for this later period.

[4]Filimon (2000) derives the optimal target rate adjustment rule when announcements on dates other than those associated with FOMC meetings are more costly. The optimal rule implies almost all announcements will be made on FOMC meeting dates. Future work could model the implications of this form of interest rate targeting for interest rate dynamics.

[5]This corresponds to the start date used by Rudebusch (1995) to model the Federal Reserve's interest rate targeting regime.

[6]Of course, whether one finds evidence of a unit root in interest rates, or not, very much depends on the sample chosen. We interpret the results in the first four rows of Table 1 as saying that over our relatively short sample period, interest rates behave as if they have a unit root.

[7]Before taking the absolute value of the spread between a market rate and a target rate, we first subtract the mean level of this spread over the sample, so that on average the level of the spread is zero. This ensures that a positive number for the absolute value of the spread is not simply capturing a difference in the definition of the Eurodollar rates compared to the target rate such that the two differ on average, or any constant term-premium between the Eurodollar rates and the fed-funds target rate. The absolute value of the spread is used only in Section 3.3 below.

[8]Note, a value of λ of 0.95 with daily data is equivalent to a λ of 0.77 when using weekly data and 0.36 when using monthly data.

[9]The difference between kurtosis for daily changes and monthly changes in interest rates is not driven by autocorrelation in the interest rate changes. For both the model-generated data

and the actual data, we regressed interest rate changes on the optimal number of own lags implied using the Akaike information criterion, calculating the kurtosis of the residual. Similar results emerge to those presented in Table 2, although the differences in kurtosis between daily and monthly interest rate changes was in some cases fractionally reduced.

[10]We also tested the null of no-skewness in the actual data, and could not reject this hypothesis in interest rate changes, even at the 10% level. This is consistent with the symmetry assumed in our model.

[11]For the seven-day interest rate, $p = 1$ and $q = 1$, for the one-month interest rate, $p = 2$ and $q = 1$, while for the three-month interest rate, $p = 2$ and $q = 2$.

[12]We use the lagged level of interest rates in equation (12) to maintain consistency with the rest of the literature. We have also re-estimated the model, setting κ to zero in equation (12), as well as using $r_{t-1} - \hat{r}_{t-1}$ in place of r_{t-1} in equation (12). Neither of these changes alters the results on $\hat{\sigma}$ and $\hat{\gamma}$ in any appreciable way.

[13]Because of the underlying simple Brownian motion assumption, the model does not imply any level-effect.

[14]These results also confirm the greater impact of the spread-effect with seven-day and one-month rates, relative to three-month rates, as suggested by the analysis in Section 1.1.

[15]Along the same lines, the time remaining before an impending economic data release could also be useful in predicting heightened interest rate volatility.