Estimating the Implied Distribution of the Future Short-Term Interest Rate Using the Longstaff-Schwartz Model

Peter Hördahl*
Directorate Monetary Policy,
European Central Bank,
D-60311 Frankfurt am Main, Germany

August 31, 2000

Abstract

This paper proposes the use of the two-factor term-structure model of Longstaff and Schwartz (1992a, LS) to estimate the risk-neutral density (RND) of the future short-term interest rate. The resulting RND can be interpreted as the market's estimate of the density of the future short-term interest rate. As such, it provides a useful financial indicator of the perceived uncertainty of future developments in the short-term interest rate. The LS approach used in this paper provides an alternative to option-based estimation procedures, which may be useful in situations where options markets are not sufficiently developed to allow estimation of the implied distribution from observed option prices. A simulation-based comparison of these two approaches reveals that the differences in the results are relatively small in magnitude, at least for short forecast horizons. Furthermore, the LS model is quite successful in capturing the asymmetries of the true distribution. It is therefore concluded that the LS model can be useful for estimating the distribution of future interest rates, when the purpose is to provide a general measure of the market's perceived uncertainty, for example as an indictor for monetary policy purposes.

*This paper was written while the author was at the Monetary and Exchange Rate Policy Department of the Swedish Riksbank. I thank Javiera Aguilar, Tor Jacobson, Anders Vredin, an anonymous referee at the ECB, participants at the European Central Bank workshop on yield curve modelling and seminar participants at Sveriges Riksbank and Lund University for useful comments and suggestions. Views expressed represent exclusively the opinion of the author and do not necessarily correspond to those of the European Central Bank or the Swedish Riksbank.

1 Introduction

The term structure of interest rates is an important source of information for market participants as well as central banks since it provides information on, among other things, the market's expectations concerning future monetary policy. Specifically, the estimated implied instantaneous forward interest rate curve can, given adequate assumptions, be interpreted as the expected short-term interest rate, which is directly or indirectly controlled by the central bank. However, implied forward rates do not provide any information on the uncertainty associated with the expected future short-term interest rate. In particular, at any given time the market's assessment of the degree of uncertainty cannot be inferred from the implied forward rate curve. Furthermore, it does not provide any way of telling whether the risk is perceived by the market to be mainly on the upside, mainly on the downside, or evenly balanced.

Recently, new techniques have been proposed to extract information from option prices in order to address these issues; see e.g. Bahra (1997), Melick and Thomas (1997), Söderlind and Svensson (1997), and Bliss and Panigirtzoglou (1999). The implied risk-neutral density (RND) of the underlying interest rate can be recovered from observed option prices, given a sufficiently developed and liquid market for options written on short-term interest rates. The resulting RND can be interpreted as the market's *ex ante* estimate of the future short-term interest rate density, assuming risk neutrality. The estimated implied RND's are used by a growing number of central banks as a financial indicator of the market's perceived uncertainty of future developments in the short-term interest rate.

Unfortunately, in many countries, including Sweden, the market for options written on short-term interest rates is not sufficiently developed to allow estimation of implied RNDs based on option prices, and it is therefore necessary to resort to alternative approaches in order to estimate the density of the future interest rate. In this paper, the dynamic process that determines the evolution of the short-term interest rate is estimated, and the density for the future interest rate that is implied by the estimated process is subsequently extracted.

An important question is which model to use for the interest rate process. There is a large number of one-factor and multifactor models available that have been suggested in different theoretical and empirical studies. Much of the empirical evidence suggests that one factor is not enough to satisfactorily capture the dynamics of the short-term interest rate. Of the multifactor specifications, the two-factor model of Longstaff and Schwartz (1992a, LS) has become increasingly popular. The two fundamental factors in this model are the level of the short-term interest rate itself, and the volatility of the short rate. This model has several advantages that makes it attractive as a candidate for estimating the interest rate density. First, the two factors in the LS model – the short-term interest rate and its volatility – are intuitively reasonable as determinants of the interest rate process, and they have also been found to be important in empirical work. Second, the LS model is a general equilibrium model, which makes it appealing in a theoretical perspective. Third, from a practical point of view, the model belongs to the affine class of models, and therefore provides closed-form solutions for zero-coupon bond prices, which facilitates estimation of the term structure using cross-sectional data.

Consequently, the LS model is used in this paper to estimate the distribution of the future short-term interest rate. Since this approach is somewhat of a second-best solution given the lack of sufficient option prices, it is interesting to compare the density forecasts obtained using the LS model with more standard option-based density forecasts. Specifically, a relevant question is if, and to what extent, the LS approach yields results that are different from the implied option approach.

The remainder of this paper is organized as follows. Section 2 provides a brief overview of the theoretical background of the LS model and discusses the term structure and the interest rate density that is implied by the model. Section 3 deals with the practical estimation issues of the model as well as the LS density, and applies the model to Swedish data. A number of empirical examples are presented in Section 4, along with a discussion of how the results can be interpreted. Section 5 provides Monte Carlo-based evaluations of the LS density forecasts and comparisons with density estimates obtained from option prices. Finally, Section 6 concludes the paper.

2 The Longstaff-Schwartz Model

2.1 Theoretical Background

Longstaff and Schwartz (1992a) propose a two-factor general equilibrium model for the short-term default-free interest rate and its variance. This section provides a very brief overview of the theoretical foundation of the LS model; for details the reader is referred to the original article.[1] The model is based on a framework similar to the continuous-time economy of Cox, Ingersoll and Ross (1985) (CIR henceforth), but where the number of state variables is two instead of one. These two state variables, X and Y, are assumed to

[1] See also Rebonato (1998), pp.313-340.

have the following mean-reverting dynamics:

$$dX = (a - bX)\,dt + c\sqrt{X}dZ_1, \tag{1}$$
$$dY = (d - eY)\,dt + f\sqrt{Y}dZ_2, \tag{2}$$

where Z_1 and Z_2 are two scalar Wiener processes that are assumed to be uncorrelated. The two state variables, in addition to a stochastic component, govern the return process for investment in the single constant-return-to-scale production technology of the economy. The realized returns on investment follow a stochastic differential equation according to

$$\frac{dQ}{Q} = (\mu X + \theta Y)\,dt + \sigma\sqrt{Y}dZ_3, \tag{3}$$

where Z_3 is a Wiener process assumed to be uncorrelated with Z_1. This means that X can be interpreted as a factor that represents technological changes that are unrelated to uncertainty in the production process, whereas Y represents a factor associated with production uncertainty.

The individuals in the economy are assumed to have identical log-utility preferences, and capital markets are assumed to be perfectly competitive and frictionless. By maximizing expected discounted utility subject to a standard budget constraint, the following equilibrium dynamics are obtained for wealth:

$$dW = (\mu X + \theta Y - \rho)\,W dt + \sigma W\sqrt{Y}dZ_3. \tag{4}$$

The current values of the two state variables and of wealth completely describe the state of the economy and the distribution of future returns on investment.

Values of contingent claims with boundary conditions that do not depend on W can be expressed in terms of the two unobservable state variables X and Y. These values are obtained by solving the fundamental partial differential equation satisfied by all contingent claims. Furthermore, by applying a change of variables, it is possible to express prices of contingent claims in terms of two variables that are observable, or which can be estimated, instead of in terms of the unobservable state variables. This can be done by noting that the instantaneous interest rate is simply equal to the expected return from the production process minus the variance of the production returns. The short-term interest rate is found to be

$$r = \alpha x + \beta y, \tag{5}$$

where $\alpha \equiv \mu c^2$, $\beta \equiv (\theta - \sigma^2)\,f^2$, and where x and y are the rescaled state variables $x \equiv X/c^2$ and $y \equiv Y/f^2$. Similarly, the instantaneous variance of changes in the risk-free rate is given by

$$V = \alpha^2 x + \beta^2 y. \tag{6}$$

Consequently, given that $\alpha \neq \beta$, x and y can be expressed in terms of r and V, which can be used to make a change in variables:

$$x = \frac{\beta r - V}{\alpha (\beta - \alpha)}, \tag{7}$$

$$y = \frac{V - \alpha r}{\beta (\beta - \alpha)}. \tag{8}$$

As a result, the dynamics of the short-term interest rate and the variance of changes in the short rate can be obtained by applying Ito's Lemma to (5) and (6) and using (7) and (8):

$$dr = \left(\alpha\gamma + \beta\eta - \frac{\beta\delta - \alpha\xi}{\beta - \alpha} r - \frac{\xi - \delta}{\beta - \alpha} V \right) dt + \alpha \sqrt{\frac{\beta r - V}{\alpha (\beta - \alpha)}} dZ_1 + \beta \sqrt{\frac{V - \alpha r}{\beta (\beta - \alpha)}} dZ_2, \tag{9}$$

$$dV = \left(\alpha^2\gamma + \beta^2\eta - \frac{\alpha\beta (\delta - \xi)}{\beta - \alpha} r - \frac{\beta\xi - \alpha\delta}{\beta - \alpha} V \right) dt + \alpha^2 \sqrt{\frac{\beta r - V}{\alpha (\beta - \alpha)}} dZ_1 + \beta^2 \sqrt{\frac{V - \alpha r}{\beta (\beta - \alpha)}} dZ_2, \tag{10}$$

where $\gamma \equiv a/c^2$, $\delta \equiv b$, $\eta \equiv d/f^2$, and $\xi \equiv e$.

The parameters α, β, γ, δ, η, and ξ are all assumed to be positive. Hence, equations (9) and (10) imply that both r and V can only assume positive values. Furthermore, in order to avoid complex values of r and V, the arguments under the square root sign must be nonnegative, which requires

$$\alpha r \leq V \leq \beta r. \tag{11}$$

It is clear from equations (9) and (10) that the two processes are interdependent in such a way that a high level of the short rate implies a high value of the current instantaneous volatility, and vice versa. Hence, changes in the short rate and changes in volatility are positively correlated, which is consistent with empirical evidence on short-term interest rates.[2] It is also possible to show that the correlation is bounded between zero and one. The fact that the volatility is not only determined by the level of the short rate displays an important advantage of the two-factor Longstaff-Schwartz model over one-factor models like CIR, where the instantaneous volatility is given by a term of the type σr^{ϕ}.

2.2 The Term Structure and Implied Forward Rates

As noted above, the LS model provides a fundamental partial differential equation (PDE) which is obeyed by all contingent claims. This PDE can be used to obtain the term

[2] See Chan *et al.* (1992) for evidence on U.S. interest rates, and e.g. Dahlquist (1996) for evidence on the behavior of Swedish and other European interest rates.

structure of the economy. The price of a discount bond can be found by solving the PDE, subject to the terminal condition that the value of a bond equals one at maturity. Since the LS model belongs to the class of affine models, a closed-form solution is available for the bond prices. The solution is a two-factor analogue to the closed-form bond prices in the CIR model. In particular, the price of a discount bond τ years before maturity is expressed as:

$$P(r, V, \tau) = A^{2\gamma}(\tau)\, B^{2\eta}(\tau) \exp\left(\kappa\tau + C(\tau)\, r + D(\tau)\, V\right), \tag{12}$$

where

$$\begin{aligned}
A(\tau) &\equiv \frac{2\phi}{(\delta+\phi)\left(\exp(\phi\tau)-1\right)+2\phi}, \\
B(\tau) &\equiv \frac{2\psi}{(v+\psi)\left(\exp(\psi\tau)-1\right)+2\psi}, \\
C(\tau) &\equiv \frac{\alpha\phi\left(\exp(\psi\tau)-1\right)B(\tau)-\beta\psi\left(\exp(\phi\tau)-1\right)A(\tau)}{\phi\psi(\beta-\alpha)}, \\
D(\tau) &\equiv \frac{\psi\left(\exp(\phi\tau)-1\right)A(\tau)-\phi\left(\exp(\psi\tau)-1\right)B(\tau)}{\phi\psi(\beta-\alpha)},
\end{aligned}$$

and

$$\begin{aligned}
\nu &\equiv \xi+\lambda, \\
\phi &\equiv \sqrt{2\alpha+\delta^2}, \\
\psi &\equiv \sqrt{2\beta+\nu^2}, \\
\kappa &\equiv \gamma(\delta+\phi)+\eta(\nu+\psi).
\end{aligned}$$

As displayed by (12), the discount bond prices is a function of the current level of the short-term interest rate, the current instantaneous volatility, and the six parameters α, β, γ, δ, η, and ν. The parameter λ denotes the market price of risk. Since the LS model is an equilibrium model, λ is determined endogenously by the model, which ensures that it is consistent with the absence of arbitrage. Note that λ always appears as a sum with ξ, which means that there exists an infinite number of combinations of λ and ξ that produce the same yield curve. It is therefore not possible to estimate the market price of risk using only data from the term structure of interest rates for a given time. Estimation of λ also requires information about the observed process of the short rate as opposed to the risk-adjusted process. By assuming that the local expectations hypothesis holds, we can set $\lambda = 0$, and express bond prices in terms of ξ instead of ν.

Given values of the two state variables r and V, and estimates of the six parameters $\{\alpha, \beta, \gamma, \delta, \eta, \xi\}$, the yield to maturity on a τ-period zero coupon bond is easily calculated as

$$Y(\tau) = \frac{-(\kappa\tau + 2\gamma \ln A(\tau) + 2\eta \ln B(\tau) + C(\tau) r + D(\tau) V)}{\tau}. \tag{13}$$

Similarly, the implied forward rate with maturity τ is obtained using the standard relation

$$f(\tau) = -\frac{\partial \ln P(r, V, \tau)}{\partial \tau}. \tag{14}$$

For the LS model, the implied forward rate turns out to be a relatively complex expression, which nevertheless can be obtained by calculating the partial derivative in (14). Both the yield and the forward rate converge to a constant value as the maturity τ increases

$$Y(\infty) = \gamma(\phi - \delta) + \eta(\psi - \xi). \tag{15}$$

The fact that yields and forward rates converge to a finite constant is a desirable property for term structure models, since long term yields frequently are interpreted as indicators of the perceived credibility of monetary policy in the long run.

In addition to the term structure of interest rates, the term structure of volatilities can be obtained from the model. The result depends on the maturity, as well as on the two state variables r and V. With fixed initial values of the state variables, the instantaneous bond return variance will converge to zero as the maturity decreases, while the variance converges to a constant value as maturity increases to infinity.

2.3 The Longstaff-Schwartz Density

The dynamics (1) and (2) imply a specific joint density for the two state variables x and y. Longstaff and Schwartz (1992a, b) show that this density is a bivariate noncentral chi-square density with closed form, given initial values of the state variables:

$$\begin{aligned} q(x, y, \tau | x_0, y_0) = {} & \frac{4}{a(\tau) c(\tau)} \left(\frac{x}{b(\tau) x_0}\right)^{\gamma - 1/2} \left(\frac{y}{d(\tau) y_0}\right)^{\eta - 1/2} \\ & \exp\left[\frac{-2}{a(\tau)}(x + b(\tau) x_0)\right] \exp\left[\frac{-2}{c(\tau)}(y + d(\tau) y_0)\right] \\ & I_{2\gamma-1}\left(\frac{4}{a(\tau)}\sqrt{b(\tau) x x_0}\right) I_{2\eta-1}\left(\frac{4}{c(\tau)}\sqrt{d(\tau) y y_0}\right), \end{aligned} \tag{16}$$

where

$$\begin{aligned}
a(\tau) &\equiv \frac{2\phi\left(\exp(\phi\tau)-1\right)}{\phi\left[(\delta+\phi)\left(\exp(\phi\tau)-1\right)+2\phi\right]}, \\
b(\tau) &\equiv \frac{4\phi^{2}\exp(\phi\tau)}{\left[(\delta+\phi)\left(\exp(\phi\tau)-1\right)+2\phi\right]^{2}}, \\
c(\tau) &\equiv \frac{2\psi\left(\exp(\psi\tau)-1\right)}{\psi\left[(\xi+\psi)\left(\exp(\psi\tau)-1\right)+2\psi\right]}, \\
d(\tau) &\equiv \frac{4\psi^{2}\exp(\psi\tau)}{\left[(\xi+\psi)\left(\exp(\psi\tau)-1\right)+2\psi\right]^{2}}, \\
\phi &\equiv \sqrt{2\alpha+\delta^{2}}, \\
\psi &\equiv \sqrt{2\beta+\xi^{2}},
\end{aligned}$$

and where $I_p(\cdot)$ is the modified Bessel function of order p. By using the transformations (5) and (6), we obtain the bivariate noncentral chi-squared density of the two transformed state variables, $q(r, V, \tau \mid r_0, V_0)$. Since the variables r and V are correlated, the transformed density is more elliptical (in the $x-y$ plane) than the density of the two state variables x and y which are uncorrelated by construction.

In principle, it is straightforward to integrate the joint density of r and V over all values of V to obtain the one-dimensional marginal distribution of the short-term interest rate,

$$q(r,\tau \mid r_0) = \int_{\alpha r}^{\beta r} q(r, V, \tau \mid r_0, V_0)\, dV, \tag{17}$$

where the limits of the integral are due to condition (11). Hence, by estimating the parameters of the LS model and then using the parameter estimates to calculate (17), we can obtain an estimate of the future short-term interest rate distribution implied by the underlying process of r and V.

3 Estimation Issues

3.1 Estimating the Longstaff-Schwartz Model

The general equilibrium setup of the LS model implies that the six parameters of the model should be constant over time, which means that the dynamics of the short rate and its volatility are constant as well. The parameters in the model can therefore be estimated using time-series data on a proxy for the short-term interest rate, and an estimated series of volatilities. In practice, however, this kind of model is frequently estimated using cross-sectional data on bills and bonds at some specific time, which is the approach chosen in this paper as well. This results in a new set of parameters each time the model is estimated. An important advantage of using a cross-section approach rather than a time-series approach to estimate the parameters in the model is that by utilizing simultaneous information from the entire term structure, it should be possible to capture changes in the dynamics of the term structure in a much more timely manner. While this time-varying parameter approach violates the equilibrium setup of the model, it is nevertheless used in order to fit the model to observed bond prices as closely as possible. This approach could be compared to the practice of estimating the implied volatility of the Black and Scholes (1973) model using one or several observed option prices, rather than estimating the (assumed constant) volatility using a time-series of returns of the underlying asset.

The estimation procedure relies on (12), which provides a closed form solution for the discount function. Using this expression, the six parameters of the LS model can be estimated with cross-sectional bond price data, given initial values of the two state variables r and V.[3] First, the initial values of r and V are determined. Since we are interested in using the result as a monetary policy indicator, the short rate r is set equal to the official repo rate, which is the Riksbank's key monetary policy instrument.[4]

Next, the initial value of the variance in interest rate changes, V, is estimated. This is done using a simple GARCH(1,1) model, assuming a constant conditional mean:

$$r_t - r_{t-1} = \mu + \varepsilon_t, \qquad \varepsilon_t \mid \Omega_{t-1} \sim N(0, h_t), \tag{18}$$

$$h_t = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta h_{t-1}, \tag{19}$$

where Ω_{t-1} denotes the information set at time $t-1$.[5] In principle, the GARCH model should be fitted to changes in the repo rate, since this is the chosen proxy for r. However, as it is the official rate set by the Riksbank, the repo rate is often constant over long

[3] Alternatively, the values of r and V may be treated as two additional parameters, and be estimated simultaneously with the six parameters of the model. In practice, however, this turns out to be hard to accomplish.

[4] Typically, the repo rate is very close to the interbank overnight interest rate.

[5] This simple specification does not take into account the assumptions in the LS model that both the conditional mean and the conditional variance are dependent on the level of the short-term interest rate as well as the level of the variance. An econometric specification that is more more in line with these

periods of time, with sudden discrete jumps in one direction or the other. It is therefore not suitable to use the repo rate to estimate V. Instead, the variance should be estimated on changes in some market-determined short-term interest rate that evolves continuously over time, such as a treasury bill rate. Due to liquidity problems in one- and two month Swedish t-bills, a three-month bill rate is used in the variance estimations.[6] Specifically, the series consists of daily three-month t-bill yields, expressed as annualized continuously compounded rates, from January 1993 up to the estimation day. Since r is expressed as an annual interest rate, V is annualized by multiplying h_t with 250, the approximate number of trading days per year.

Once the values of r and V have been determined, we are ready to estimate the six parameters in the LS model, $\Theta \equiv \{\alpha, \beta, \gamma, \delta, \eta, \xi\}$. As previously mentioned, the model is estimated using cross-sectional data on Swedish Treasury bills and bonds. Specifically, the official repo rate, all available bills with at least three months to maturity, as well as all benchmark bonds are used in the estimations.[7] Typically, this amounts to around 5 Treasury bills with 3 - 15 months to maturity, and around 10 benchmark bonds with up to 15 years to maturity, in addition to the repo rate; hence, a total of around 16 observations along the yield curve. Next, it is assumed that the observed market prices of these instruments differ from the prices produced by the LS model by an error term with zero expected value. This implies an assumption that the LS model provides the true functional form for pricing bonds, or at least that it is sufficiently flexible to be able

---

features in the model is the following, as suggested by Longstaff and Schwartz (1992a):

$$\begin{aligned} r_t - r_{t-1} &= \mu_0 + \mu_1 r_{t-1} + \mu_2 h_t + \varepsilon_t, \qquad \varepsilon_t \mid \Omega_{t-1} \sim N(0, h_t), \\ h_t &= \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 h_{t-1} + \beta_2 r_{t-1}. \end{aligned}$$

However, estimating this model turned out to be significantly more time-consuming, while it provided conditional variance estimates that were very similar to the ones obtained using (18) - (19). In fact, the correlation coefficient between the conditional variance series produced by the two models was 0.997 (0.995 for the changes in the variances) for daily Swedish 3-month money market rates between January 1996 and December 1998. The simple specification (18) - (19) was therefore chosen for the sake of computational speed.

[6] The use of three-month t-bill rates may also help to avoid the kind of idiosyncratic variation in interest rates of shorter maturity that Duffee (1996) documents for U.S. dollar denominated fixed income securities. While the existence of similar effects in Swedish short-term yields have not been investigated yet, it can not be excluded that they are in fact present.

[7] Benchmark bonds are specific series of government bonds that are selected by the Swedish National Debt Office to act as benchmarks for bonds of different maturities. The market for benchmark bonds is very liquid, since trading tends to be concentrated to these bonds.

to price all bonds correctly. Deviations between observed and model prices may then be seen as reflecting institutional or market microstructural features.[8]

The estimates for the parameters in the LS model are obtained by minimizing the distance between the observed market prices and the model's theoretical prices of bills and bonds, using NLLS:

$$\hat{\Theta} = \arg\min_{\Theta} \sum_{i=1}^{n} \left[P_i - P_i\left(r, V, \Theta\right)\right]^2 . \tag{20}$$

Here, P_i denotes the observed price of bill/bond i among the n available securities with different maturities, while $P_i\left(r, V, \Theta\right)$ is the corresponding LS price given the current values of r, V, and the parameter vector.

3.2 Estimating the Longstaff-Schwartz Density

Given the current values of the state variables r and V, and the parameter estimates of the LS model, the density function of future short rates implied by the LS model can be estimated. In principle, the values of the parameters and the state variables could simply be plugged into the closed form expression for the density (16), and the marginal distribution of r could then be integrated out as in (17). In practice, however, this turns out to be problematic. For some combinations of parameters and state variable values, the numerical evaluation of the modified Bessel function in (16) or the numerical integration of the bivariate density breaks down.[9]

Consequently, the density will have to be estimated using some other approach. The strategy employed in this paper is to use Monte Carlo methods instead of the closed form solution to obtain the density. This is done by using discretized versions of the processes for the short rate and its variance to simulate possible future realizations of r and V. Specifically, by using an Euler approximation and assuming weekly time steps, discrete

[8] See Dahlquist and Svensson (1996) and references therein for a discussion regarding these assumptions, and this approach to estimating the parameters in the model.

[9] This seems to be a known problem with evaluating the closed-form solution for the LS density, as pointed out by Rebonato (1998) (page 326). He notes that there are combinations of parameter values that make the probability function tend to infinity in the limit as one of the arguments goes to zero. While the probability distribution always remains integrable, the numerical evaluation of the integral tends to break down in these cases. Furthermore, it appears that this problem occurs relatively frequently when the model is estimated using actual data, which suggests that an alternative, more robust approach to the evaluation of the closed-form density expression is warranted in practice.

versions of the continuous-time dynamics are obtained as follows:

$$\begin{aligned} r_{t+\Delta t} - r_t &= \left(\alpha\gamma + \beta\eta - \frac{\beta\delta - \alpha\xi}{\beta - \alpha} r_t - \frac{\xi - \delta}{\beta - \alpha} V_t\right)\Delta t + \alpha\sqrt{\frac{\beta r_t - V_t}{\alpha(\beta - \alpha)}}\sqrt{\Delta t}\varepsilon_{1,t+\Delta t} \\ &\quad + \beta\sqrt{\frac{V_t - \alpha r_t}{\beta(\beta - \alpha)}}\sqrt{\Delta t}\varepsilon_{2,t+\Delta t}, \end{aligned} \tag{21}$$

$$\begin{aligned} V_{t+\Delta t} - V_t &= \left(\alpha^2\gamma + \beta^2\eta - \frac{\alpha\beta(\delta - \xi)}{\beta - \alpha} r_t - \frac{\beta\xi - \alpha\delta}{\beta - \alpha} V_t\right)\Delta t + \alpha^2\sqrt{\frac{\beta r_t - V_t}{\alpha(\beta - \alpha)}}\sqrt{\Delta t}\varepsilon_{1,t+\Delta t} \\ &\quad + \beta^2\sqrt{\frac{V_t - \alpha r_t}{\beta(\beta - \alpha)}}\sqrt{\Delta t}\varepsilon_{2,t+\Delta t}, \end{aligned} \tag{22}$$

where $\Delta t = 1/52$, while $\varepsilon_{1,t+\Delta t}$ and $\varepsilon_{2,t+\Delta t}$ are drawn from two independent standard normal distributions. Note that since we have assumed that the local expectations hypothesis holds, (21) and (22) are approximations of the risk neutral dynamics of r and V.

The equations for the discretized dynamics, (21) and (22), are used to simulate future values of the short rate and its variance in a recursive manner, starting with the current values of r and V. The simulated sample paths can be extended as far into the future as one likes. In this study, a time horizon of up to one year is chosen, which means that 52 future values of r and V are simulated, given the choice of $\Delta t = 1/52$. This process is then repeated 20,000 times with the same set of parameter values, in order to obtain a relatively large sample of simulated values. Hence, for each of the 52 future weeks following the date of estimation, the described procedure produces a simulated sample consisting of 20,000 r's and 20,000 V's. The next step is to obtain an estimate of the distribution of the future short-term interest rate at each time, which can be done by using e.g. a simple histogram or some kernel estimator. Since the simulations have been performed using the risk neutral dynamics, the resulting density estimate for some forecast horizon is the *risk neutral density*, RND, implied by the observed bond prices, and assuming that the LS model holds. The density obtained using this approach is therefore comparable to the more common option-implied density, which is also risk neutral.

4 RND Estimates

This section presents various ways of displaying the results, discusses how the estimated distributions can be interpreted, and takes a closer look at RND estimates for a few selected dates in the sample. A natural way to show the results is to plot the estimated density itself at a given date for some time horizon that is of interest. Figure 1 displays the estimated RND for the one-month, three-month, six-month, and one-year forecast horizons, implied by the previously discussed set of prices of Swedish bills and bonds on April 20, 1999.[10] This gives a snapshot view of the implied densities at that date. The figure shows that as the forecast horizon grows, the mass of the estimated densities is moved to the right, i.e. to higher interest rate levels. This reflects the fact that the implied forward rate curve was upward sloping at the time.[11] Another feature displayed in the figure is that the densities are more dispersed the longer the horizon. This is natural since there is greater uncertainty about the possible future outcome for longer forecast horizons. However, as described earlier, the short-term interest rate is stationary in the LS model, which implies that the distribution of future rates converges to an unconditional distribution, as the time horizon grows to infinity. Hence, the distribution changes relatively less from one forecast horizon to the next, as the horizon increases.

Figure 1: Estimated RND of future short term interest rate at $t+\tau$ implied by market prices on t = Apr. 20, 1999

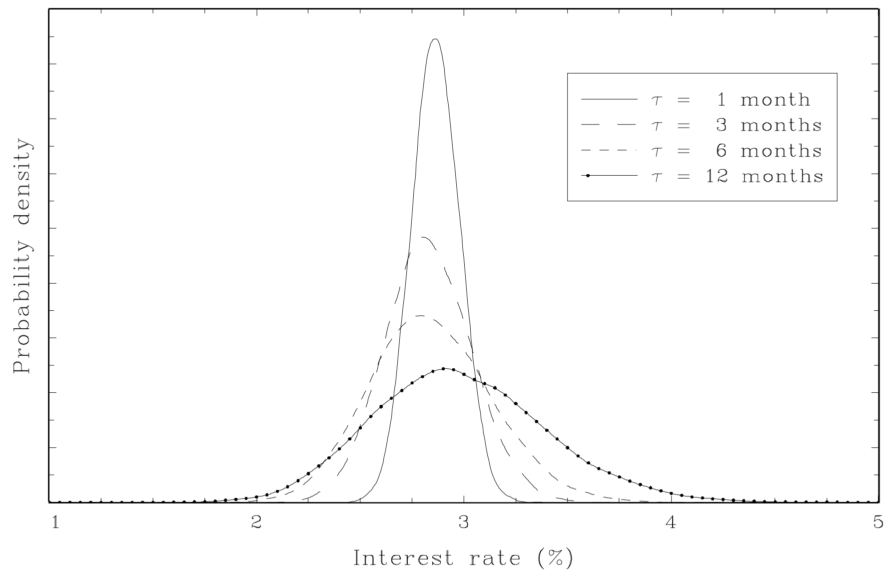

Another way of displaying the results is to plot the PDFs for all horizons up to a year in a three-dimensional diagram, as in Figure 2. Here, the number of weeks ahead in time is displayed on the x-axis, the level of the short-term interest rate is on the y-axis, and

[10] The RNDs in Figure 1 are estimated using a Gaussian kernel with optimal bandwidth, applied to the set of simulated future interest rates. See Silverman (1986), pp. 40-45 for details on kernel density estimation.

[11] On this date, the level of the Riksbank's official repo rate was 2.90%.

the probability density is on the z-axis. This figure clearly shows the convergence of the distribution as the horizon increases. A relevant question is how one should interpret the distributions in figures 1 and 2 from an economic perspective. Clearly, they show the distribution of the future "short-term interest rate". Since the level of the Riksbank's repo rate has been used as the initial short rate in the estimations, and since the repo rate determines the shortest end of the yield curve, it is natural to interpret the results in figures 1 and 2 as the expected distribution of the future repo rate. This is consistent with using the implied forward rate as an estimator of the expected value of future repo rates, since the implied LS forward rate for a given horizon is equal to the expected value of the interest rate distribution.[12]

Figure 2: Estimated RND of future short term interest rate at $t+\tau$ implied by market prices on t = Apr. 20, 1999

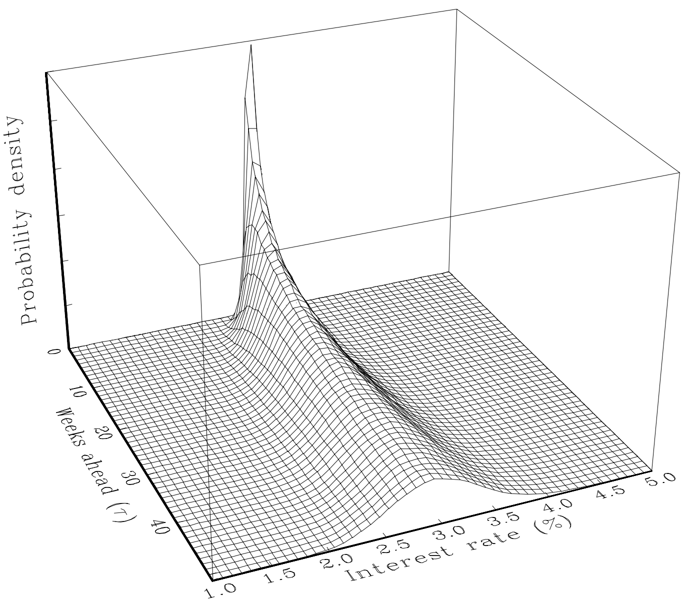

Using this interpretation, the implied PDFs can also be displayed in terms of confidence intervals that show the regions within which the market believes that the repo rate will be, with a certain probability, at different times in the future. Figure 3 shows a time series of the Riksbank's repo rate from 1995 up until April 20 1999, along with various confidence intervals for the repo rate up to a year in the future, estimated using data available on April 20. The darkest central band in the figure covers 10 percent of the probability mass, and each of the wider bands cover an additional 20 percentage points

[12] Naturally, this abstracts from the possibility of any term premia in the yield curve.

of the mass. Hence, the bands display the 10, 30, 50, 70, and 90 percent confidence limits of the implied distribution.

Figure 3: Swedish repo rate and confidence bands for the future rate, estimated on April 20 1999.

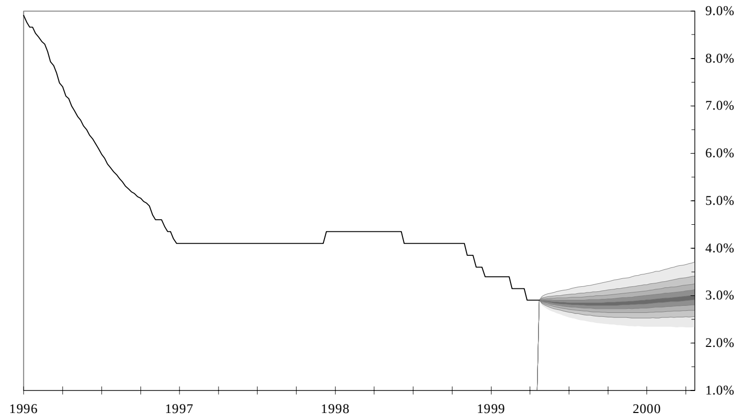

Graphs such as Figure 3 can be useful for e.g. a central bank to get an idea of the markets' perception of the uncertainty concerning the future repo rate. Wide confidence intervals indicate that uncertainty is relatively high as to the direction and the size of repo rate changes. When the distributions are asymmetric, meaning that confidence bands are wider on one side of the central band than on the other, this provides information that the market believes the uncertainty is concentrated mainly on one side. For example, if the bands above the central 10% band are wider than the ones below the central band, then this would suggest that the market considers it more likely that rates will increase more than indicated by the implied forward rates, than that they will increase less or decrease.

It is evident from Figure 3 that the bands for April 20 1999 are approximately symmetric and relatively narrow. Comparing with the PDFs implied by the data on, for example, March 26 1996 and November 11 1997, it is apparent that the distributions have varied considerably over time. Figure 4 shows the estimated distribution obtained on March 26 1996, when the repo rate had been cut on numerous occasions during the months before the estimation date. According to the market prices and the model, there was a strong belief that the Riksbank would continue to lower interest rates in the coming months, but that the easing of monetary policy would end late in 1996. However, the wide confidence intervals show that there was a high degree of uncertainty in the market on the expected size of repo rate changes and also on the timing of a policy change from expansionary to neutral. The asymmetry in the bands also indicates that the risk was perceived by the market to be somewhat higher on the upside than on the downside, compared with the average expected outcome.

Figure 4: Swedish repo rate and confidence bands for the future rate, estimated on March 26 1996.

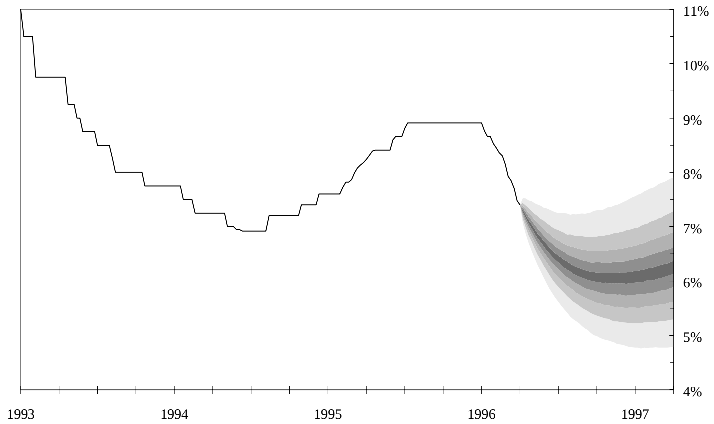

Figure 5 shows the estimated distribution obtained in a very different situation. On November 11 1997, the repo rate had been kept unchanged at 4.10% for almost a year, but the Riksbank had been signalling the need for a somewhat tighter monetary policy for some time. At the same time, the first wave of the Asian crisis had hit the world, leading to increased uncertainty regarding the economic outlook. The sharply upward sloping central band in Figure 5 shows that the Riksbank's signals had reached the market, and that there were expectations of higher interest rates in the near future. Furthermore, the estimated distributions are asymmetric, leading to wider bands above the central band in Figure 5, hence indicating that the risk was perceived as being mainly on the upside at this time.

Figure 5: Swedish repo rate and confidence bands for the future rate, estimated on November 11 1997.

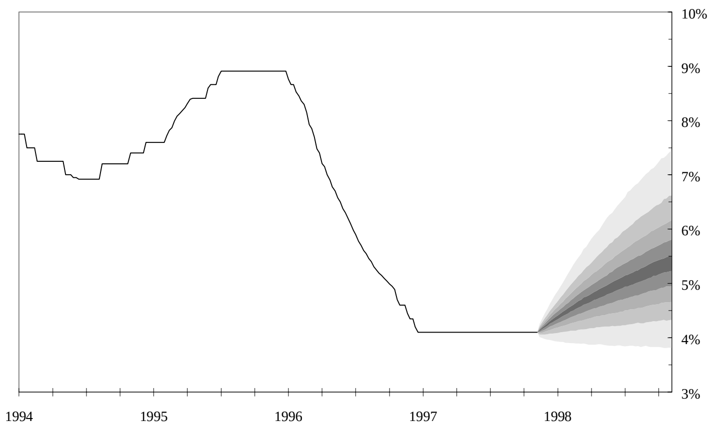

All figures discussed above give snapshot pictures of the estimated uncertainty concerning the future short rate, given prices in the bond market at each specific estimation date. However, these graphs tell us very little about the way the market's perceived uncertainty has changed over time. One illustrative approach that can be used for this purpose is to study time-series of confidence intervals for the expected short rate for some given forecast horizon. Figure 6, for example, displays a time-series of the 6-month implied forward rate, along with a 90% confidence interval obtained from the estimated LS PDF. This graph can be interpreted as displaying the expected repo rate six months ahead, and the interval within which the market believes that this rate will end up with a 90% probability (given the assumptions discussed earlier). The chart clearly shows that the degree of uncertainty has varied over time. For example, the confidence interval widened considerably during the fall/winter of 1997, amid concerns for the Asian crisis and speculations of a rate hike by the Riksbank. It is also interesting to note that the implied distribution at this time became more positively skewed, which resulted in a greater widening of the upper than of the lower bands in Figure 6. Hence, graphs such as Figure 6 can be of assistance in illustrating not only changes in the size of the market's perceived uncertainty, but also changes in the asymmetry of the distribution.

Figure 6: Expected short-term interest rate and 90% confidence interval implied by the LS model, 6 months forecast horizon; January 2 1996 - April 20 1999.

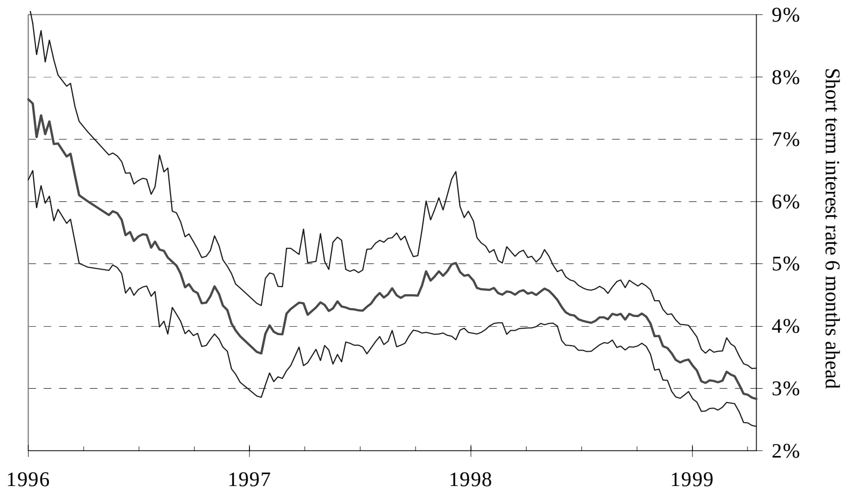

5 Evaluating the LS Density

A natural question to ask is how well the estimated implied LS PDFs reflect the 'true' underlying density of the short-term interest rate. It is also interesting to compare the LS approach with the results obtained by estimating the distribution using option prices. For markets where short-term interest rate options exist that are sufficiently traded to provide contemporaneous option prices over a wide range of exercise prices, it appears that there are few advantages to relying on an approach such as the one described above rather than on an option-implied approach to estimate interest rate PDFs. The option-implied method utilizes information from a large number of option prices that expire at the date for which the density is estimated, whereas the LS method uses bond price data for a number of different maturities, all of which may differ from the forecast horizon for the PDF. Intuitively, this should enable more accurate PDF estimates for those specific horizons that coincide with available maturities of traded option contracts. Furthermore, the option-implied approach is more general in the sense that an estimated terminal PDF can be consistent with a number of different underlying processes for the short-term interest rate, whereas a given process is consistent with only one terminal distribution. This assumes that a suitable estimation technique is used to obtain the option-implied PDF, such as a mixture of lognormal distributions. This is discussed further below.

Since the LS model is an equilibrium model, it could be argued that the LS approach has the advantage of being based on more solid theoretical underpinnings than the option-implied methods commonly used. Specifically, an implied interest rate density obtained using the LS model is always consistent with an economy in which individuals have maximized expected utility and in which asset prices are in equilibrium, etc., whereas this need not be the case for the option-implied alternatives. However, as mentioned before, the equilibrium properties of the LS model are not consistent with calibrating the model to new cross-sectional data at different points in time. Nevertheless, the LS approach does ensure one appealing feature, namely that all implied distributions are consistent with yields observed in the market and, consequently, with the entire term structure of interest rates. On the other hand, one could argue that given a sufficiently flexible method for extracting option-implied PDFs, any distribution obtained using the LS approach could also be obtained using the option-implied approach. From a practical viewpoint, the LS approach has the advantage of allowing the estimation of densities for arbitrarily chosen horizons. This is useful since it provides an easy way of obtaining constant-maturity PDFs, hence facilitating comparisons of various properties of the distribution over time. Option-implied PDFs, on the other hand, can in general only be obtained for fixed expiration dates, which means that the horizon of the implied distributions will vary over time.[13]

For countries like Sweden, where the interest rate option market is not sufficiently liquid and developed to allow estimation of option-implied PDFs, it is necessary to use alternatives such as the LS model. It is therefore of interest to investigate whether this approach yields results that are different from those produced by the option approach, and if so, to what extent. One way of investigating these questions is to estimate and compare the PDFs using the two methods for a country with a sufficiently developed options market. A problem with this approach is that it does not allow us to say if either or both PDFs differ from the *true* PDF, since the true PDF is not observable. A convenient way of dealing with this problem is to use simulated data, for which the true underlying distribution is known. This is the approach taken in this section.

[13] This is the case for exchange-traded options, which have fixed expiration dates. By using OTC-options it is possible to estimate constant-maturity PDFs. In practice, however, data availability appears to limit this possibility to foreign exchange rate options.

5.1 Evaluating Density Forecasts for Different Horizons

In summary, the evaluation procedure is set up as follows. First, a 'realistic' data generating process (DGP) is chosen to generate simulated sample paths for the future short-term interest rate. By generating a large number of paths, it is possible to calculate zero-coupon bond prices, as well as call and put option prices, for arbitrary maturities. Second, the parameters of the LS model are estimated using the cross-section of simulated bond prices, from which the implied LS PDF then can be obtained. Similarly, the option-based implied PDF can be estimated using the simulated option prices. Third, the 'true' PDF can be estimated as the sample distribution of the simulated future short-term interest rates. Finally, the results from the two alternative methods can be evaluated against the 'true' PDF using standard statistical goodness of fit tests, such as the Kolmogorov-Smirnov statistic.

The first issue to deal with is the choice of DGP. It should be able to reproduce observed features of historical short-term interest rates, such as mean-reversion, conditional heteroskedasticity, and level-dependent volatility. A large number of models for the process of the short-term interest rate that capture some or all of these features have been proposed in the term-structure literature. A recent variant was suggested by Koedijk, Nissen, Schotman and Wolff (1997, KNSW), who proposed the following discrete-time specification for the short rate r:

$$r_{t+\Delta t} - r_t = \alpha_0 + \alpha_1 r_t + \alpha_2 r_t^2 + \varepsilon_{t+\Delta t}, \tag{23}$$

$$\varepsilon_{t+\Delta t} = h_{t+\Delta t}\epsilon_{t+\Delta t}, \qquad \epsilon_{t+\Delta t} \sim D(0,1), \tag{24}$$

$$h_{t+\Delta t}^2 = \beta_1 r_t^{2\gamma} + \left(\frac{r_t}{r_{t-\Delta t}}\right)^{2\gamma}\left(\beta_2\varepsilon_t^2 + \beta_3 h_t^2\right). \tag{25}$$

This model encompasses both the level effect of Chan, Karolyi, Longstaff, and Sanders (1992, CKLS) and the conditional heteroskedasticity effect of a GARCH(1,1) model, while at the same time allowing for linear as well as nonlinear mean-reversion in the drift. KNSW find that their model outperforms both the CKLS and a specification where the conditional variance is given by a simple GARCH(1,1) model, since both level effects and conditional heteroskedasticity effects are found to be important features in interest rate data. Consequently, the KNSW model seems to provide a suitable DGP to simulate the short-term interest rate.

Table 1: QML estimates of the KNSW model for weekly German 3-month Eurorates 1983-1998.

Parameter	QML estimate
$\alpha_0 \times 100$	3.452 (3.223)
$\alpha_1 \times 100$	− 1.580 (1.314)
$\alpha_2 \times 100$	0.138 (0.112)
$\beta_1 \times 10^5$	2.505 (4.395)
β_2	0.105 (0.023)
β_3	0.850 (0.062)
γ	1
Log-Likelihood	637.66
Ljung-Box (levels)	0.224
Ljung-Box (squares)	0.968

Figures in parentheses are robust Bollerslev and Wooldridge (1992) standard errors. The Ljung-Box tests for autocorrelation in the levels and the squared standardized residuals are based on 20 lags. The figures given for these tests are p-values.

In order to implement the simulation procedure, parameter values must be chosen for the model above, for example by estimating the model on some time-series of interest rates. For this purpose, a weekly series of German 3-month Eurorates for the period 1983-1998 is chosen.[14] Table 1 displays the parameter values obtained when estimating (23)-(25) with the quasi-maximum likelihood method (i.e. using $\varepsilon_{t+\Delta t} \sim N(0,1)$). In Table 1, the restriction $\gamma = 1$ has been imposed in order to obtain stationary distributions, as discussed by KNSW (1997).[15] The parameters α_0 and α_1 that govern the mean-reversion in the model are of the right sign and have reasonable magnitude, but are not significantly different from zero. This is a common finding in empirical studies dealing with parametric models of the process for the short-term interest rate (see e.g. Aït-Sahalia (1996b)). The GARCH parameters β_1 and β_2 are significantly different from zero, hence

[14] German interest rates are chosen instead of Swedish rates due to the fact that Sweden changed currency regime from fixed to floating in 1992. Hence, it is very likely that there was a regime shift in the Swedish short-term interest rate process as well at this time, which would have made Swedish rates unsuitable as input to the model.

[15] The interest rate is nonstationary if $\gamma > 1$; see also Broze *et al.* (1995). The unrestricted estimate of γ is 1.075, which is not significantly different from one.

indicating the presence of conditional heteroskedasticity in the interest rate differentials. The sum of these two parameters is smaller than unity, which implies that the interest rate follows a covariance stationary process (see Engle and Bollerslev (1986)). The Ljung-Box test statistic for the level indicate that the KNSW model is successful in capturing the autocorrelation in the short-term interest rate. The corresponding test for the squared standardized residuals shows that the specification used is able to model the nonlinear dependence in the data as well. In all, the KNSW model seems to be able to capture many of the features of the data, and it should therefore be a suitable choice to use as data generating process.[16] The setup of the Monte Carlo study with the set of parameters in Table 1 is henceforth referred to as Case A.

Next, with the DGP and its parameters chosen, the simulation procedure is implemented by simulating "weekly" data (i.e. Δt is set to one week) recursively for $r_{t+\Delta t}, r_{t+2\Delta t}, ...$ up to five years ahead in time, by drawing values of $\epsilon_{t+\Delta t}$ from the standard normal distribution and substituting them into the model. The procedure is started up by choosing "today's" value of the short rate as $r_t = 3\%$, while the initial conditional variance is set equal to the estimated unconditional variance of the weekly changes in interest rates.[17]

In order to evaluate the LS and the option approach for estimating implied PDFs, bond and option prices have to be calculated. In general, given the existence of a risk neutral probability measure Q, the price of a zero-coupon bond of maturity $(T-t)$ is given by

$$B(t,T) = E_t^Q \left[\exp\left(-\int_t^T r_s ds \right) \right]. \tag{26}$$

Similarly, the price of a call option on the short-term interest rate at a future date τ, with

[16] The fact that the KNSW model appears to do a good job in terms of capturing the salient time-series features of the short-term interest rate raises the question whether this model should not be used instead of the LS model to estimate implied PDFs. While this is possible to do, it would require substantially more effort in practice since there are no analytical expressions available for pricing bonds in the KNSW setup. Consequently, in order to recover the parameters of the model using cross-sectional data, bond prices would have to be obtained numerically as each new set of parameters is evaluated in the NLLS estimation procedure. In practice, this is likely to be too time-consuming to allow the KNSW model to be a realistic alternative for estimating interest rate distributions.

[17] In addition to this, the model requires an initial value for $r_{t-\Delta t}$. This is set equal to $r_t \cdot u$, where u is drawn from normal distribution with mean zero and variance equal to the unconditional variance of the series of first-differenced interest rates.

strike rate X is given by

$$C(t,\tau) = E_t^Q \left[\exp\left(-\int_t^\tau r_s ds \right) \max\left(0, r(\tau) - X\right) \right], \tag{27}$$

while the price of a similar put option can be expressed as

$$P(t,\tau) = E_t^Q \left[\exp\left(-\int_t^\tau r_s ds \right) \max\left(0, X - r(\tau)\right) \right]. \tag{28}$$

Since there are no analytical expressions available for the pricing formulas above in the KNSW setup, the bond and option prices have to be estimated by numerical means. In particular, the prices are obtained using Monte Carlo integration, by calculating the discrete-time counterparts to the expressions above:[18]

$$B(t,T) = \frac{1}{N} \sum_{i=1}^{N} \exp\left(-\sum_{s=t}^{T} r_{i,s} \Delta t \right), \tag{29}$$

$$C(t,\tau) = \frac{1}{N} \sum_{i=1}^{N} \exp\left(-\sum_{s=t}^{\tau-1} r_{i,s} \Delta t \right) \max\left(0, r_{i,\tau} - X\right), \tag{30}$$

$$P(t,\tau) = \frac{1}{N} \sum_{i=1}^{N} \exp\left(-\sum_{s=t}^{\tau-1} r_{i,s} \Delta t \right) \max\left(0, X - r_{i,\tau}\right), \tag{31}$$

where Δt is one week, and the number of simulation runs is set to $N = 20,000$. A total of 14 zero-coupon bond prices with different maturities are calculated, while option prices for calls and puts are obtained for seven different strikes for each of four different maturities.[19] The seven strikes are set equal to the implied forward rate, $f(t,\tau)$, plus three rates above and three below the forward rate, thus resulting in an at-the-money option, three out-of-the-money options, and three in-the-money options. The strikes of the six options that are not in-the-money are equally spaced according to $(1 \pm x) f(t,\tau)$, where $x = 0.1, 0.2,$ and 0.3 respectively.

Using the simulated set of 14 bond prices, the parameters of the LS model can be estimated, and the implied LS density can be obtained in the manner described in earlier sections. The LS PDFs, $q_{LS}(\tau)$, are estimated for the horizons τ corresponding to the four maturities of the options. The option-implied PDFs are estimated using a method proposed by Melick and Thomas (1997). In short, this method assumes that market participants expect future short-term interest rates to be drawn from a mixture of two

[18] Again, it is assumed that the local expectations hypothsis holds.

[19] The maturities selected for the bonds are $\frac{1}{52}, \frac{1}{12}, \frac{1}{6}, \frac{1}{4}, \frac{1}{3}, \frac{5}{12}, \frac{1}{2}, \frac{3}{4}, 1, 1.5, 2, 3, 4,$ and 5 years, while the option maturities are set to $\frac{1}{12}, 0.25, 0.5,$ and 1 year respectively (i.e. 4, 13, 26, and 52 weeks).

lognormal distributions.[20] The specific shape of the terminal distribution is then obtained by minimizing the distance between observed option prices (with identical maturities but different strikes) and the corresponding option prices implied by the mixture of lognormals. This is done with respect to the five parameters to be estimated: a location parameter and a dispersion parameter for each of the lognormal distributions, plus a weighting parameter to determine the relative influence of the two lognormals on the terminal PDF, $q_{OPT}(\tau)$.[21]

Figure 7a: True and implied RNDs, 4−week horizon: Case A

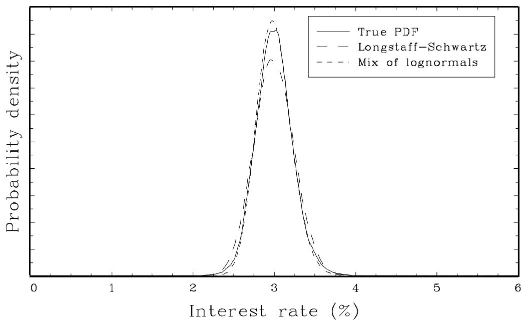

Figure 7b: True and implied RNDs, 13−week horizon: Case A

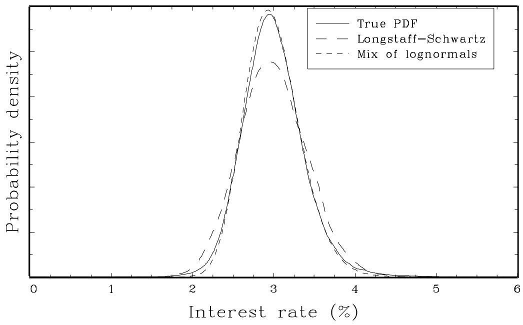

Figure 7c: True and implied RNDs, 26−week horizon: Case A

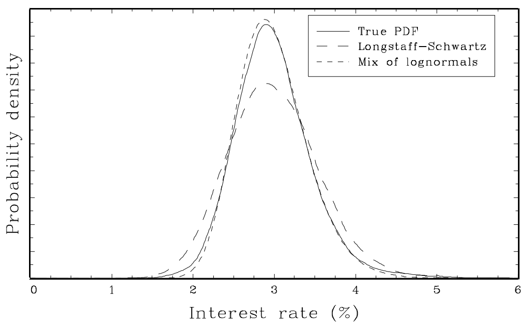

Figure 7d: True and implied RNDs, 52−week horizon: Case A

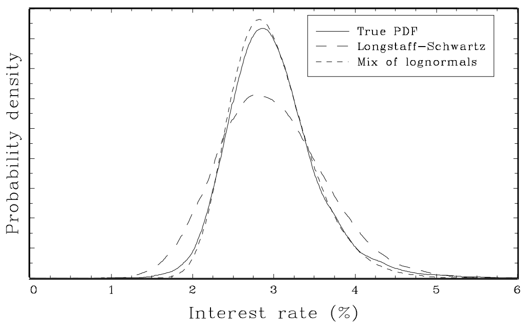

In order to evaluate the density forecasts produced by the two approaches, the "true" PDF, $q_{TRUE}(\tau)$, has to be available. This is obtained by estimating the sample distribution over the 20,000 simulated values of $r_{i,\tau}$, for each of the forecast horizons $\tau = 4$, 13, 26, and 52 weeks. Figures 7a-d displays the estimated implied and "true" distributions (for

[20] In their original application to crude oil options, Melick and Thomas (1997) use a mixture of *three* lognormals. However, Bahra (1997) finds that the use of two lognormals is more tractable from a numerical point of view.

[21] See Appendix 1 for a more detailed description of this estimation method. There are a number of alternative methods that can be used to estimate the implied PDF from option prices. Bahra (1997) discusses these various techniques, and finds the mixture of lognormals approach to be the preferred method.

Case A) for the four horizons. It is clear from the graphs that the mixture of lognormals approach consistently does a better job than the LS method in terms of forecasting the true density. The general picture remains the same when the parameters used in the DGP are changed in various ways to simulate alternative interest rate dynamics, or when an alternative specification is used for the short rate process; see Appendix B. These results are not surprising given that the option-implied approach is more "general" than the LS method. As pointed out by Melick and Thomas (1997), estimation methods based on an assumption about the underlying distribution are more general since any given terminal distribution encompasses many possible underlying processes for the short-term interest rate, whereas a given process is consistent with only one terminal distribution. The option-implied approach enjoys another advantage: for each of the forecast horizons there are 14 option prices available, each carrying information as to the possible outcome of the future interest rate at the end of the forecast horizon. The LS approach, on the other hand, uses at most *one* bond price where the maturity is equal to the forecast horizon.

The question of whether the visible differences between the densities are statistically significant can be addressed by using the Kolmogorov-Smirnov (KS) goodness of fit test. This is based on the following test statistic:

$$D = \sup_{r} \left[\left| Q_a(r) - Q_0(r) \right| \right], \tag{32}$$

where $Q_0(r)$ is the "true" sample distribution function and $Q_a(r)$ is the alternative distribution. In the case where the mixture of lognormals is the alternative distribution, the standard one-sample KS test can be applied, whereas the two-sample KS test (also called the Smirnov test) is used when the LS model provides the alternative, since both the "true" and the LS distributions are based on simulated data in this case.

Table 2 displays the KS test statistics obtained for the two estimation methods and for each of the four forecast horizons. It is clear from the results in Table 2 that we can reject the hypothesis that the distribution of the interest rates generated by the LS model and the distribution obtained using the DGP are the same. However, the KS tests also indicate that there are statistically significant differences between the estimated distribution of the DGP and the implied mix-of-lognormals distribution for all forecast horizons. Hence, it seems that neither of the two estimation methods provide satisfactory results from a statistical point of view.

Table 2: Kolmogorov-Smirnov tests for the goodness of fit between the estimated LS RNDs and the estimated "true" RNDs, and between the RNDs implied by option prices obtained using a mix of lognormals and the "true" RNDs.

Test	Horizon	Test statistic (D)	
$H_0 : Q_{LS}(r) = Q_{TRUE}(r)$	4 weeks	0.030	***
	13 weeks	0.046	***
	26 weeks	0.058	***
	52 weeks	0.078	***
$H_0 : Q_{OPT}(r) = Q_{TRUE}(r)$	4 weeks	0.021	***
	13 weeks	0.011	**
	26 weeks	0.013	***
	52 weeks	0.014	***

$Q_{LS}(r)$, $Q_{OPT}(r)$, and $Q_{TRUE}(r)$ are the estimated distributions using the Longstaff-Schwartz model, the mix of lognormals approach, and the true DGP respectively. The critical values for the two-sample KS test (i.e. for LS vs TRUE) are 0.012 (10%), 0.014 (5%), and 0.016 (1%).The critical values for the one-sample KS test (i.e. for OPT vs TRUE) are 0.009 (10%), 0.010 (5%), and 0.012 (1%). *, **, and *** denote significance at the 10%, 5%, and 1% level respectively.

A more relevant question, perhaps, is whether the discrepancies between the true and the estimated distributions are *economically* significant as well. As displayed by Figures 7a-d, the differences between the option-implied and the true distributions are quite small in magnitude. As for the LS RNDs, these also appear to be relatively close to the true distribution, at least for the shorter forecast horizons, and the LS model is quite successful in capturing the main asymmetries of the true distribution. Therefore, it does not seem unreasonable to use the LS model as an alternative to the option-implied approach if option price data is unavailable, when the purpose of the estimated RNDs is to provide a general measure of the market's perceived uncertainty. If, however, the purpose is to use the estimated distributions for e.g. pricing derivatives, the LS model should be used with caution, especially for instruments such as long-maturity out-of-the-money options.

5.2 Evaluating a Sequence of Density Forecasts

In a recent article, Diebold et al. (1998) propose a method for evaluating a sequence of density forecasts. The method is based on the so-called *probability integral transform* of the realized outcomes of the forecasted variable, r_t, with respect to the density forecast, $q_t(r_t)$, defined as

$$\begin{aligned} z_t &= \int_{-\infty}^{r_t} q_t(u)\,du \\ &= Q_t(r_t)\,. \end{aligned} \tag{33}$$

The probability integral transform provides a relationship between the DGP and the sequence of density forecasts, that can be used to evaluate the performance of the forecasts. Diebold et al. (1998) show that if a density forecast coincides with the true density of the DGP, then the density of z_t is simply the uniform density with support over the unit interval. Furthermore, if a sequence of density forecasts, $\{q_t(r_t)\}_{t=1}^{m}$ coincides with the true density sequence, then the sequence of probability integral transforms of $\{r_t\}_{t=1}^{m}$ with respect to $\{q_t(r_t)\}_{t=1}^{m}$ is IID with uniform distribution,

$$\{z_t\}_{t=1}^{m} \overset{IID}{\sim} U(0,1)\,. \tag{34}$$

For practical applications, Diebold et al. (1998) propose testing whether the probability integral transform series is IID $U(0,1)$ using simple visual tools such as histograms and correlograms since these can be helpful in explaining the nature of the violations in case of a rejection.

In this section, the procedure described above is implemented for sequences of interest rate density forecasts. Again, simulated data is used since it allows the construction of a sufficiently long interest rate series to implement the test procedure. Specifically, the same DGP (i.e. Case A) as in the previous section is used to generate a simulated sequence consisting of 14,250 "weekly" short-term interest rate observations, and an equally long sequence of conditional variances. The first 1,000 observations in each series are then discarded to avoid distortions due to initial assumptions used to start up the simulation process. The resulting interest rate series is then taken to be the true sequence of interest rate realizations. The first 250 observations of this series are regarded as the "initially observable" interest rates.

In principle, the density forecasts can be evaluated for any horizon, but in view of the computer-intensive nature of the simulation procedure, the forecast horizon is set equal to three months, i.e. 13 weeks. At every 13th observation along the sequence of interest rate realizations, LS and option-implied densities are estimated, resulting in non-overlapping sequences of 3-month density forecasts. The forecast procedure is started at the last "observable" period, i.e. at observation 250 in the simulated series, and the actual estimations are carried out as follows.

First, theoretical bond and option prices based on the initial interest rate and variance (at $t = 250$) are calculated using (29)-(31) in the same manner as described in the previous section. A total of 14 zero-coupon bond prices, with maturities up to three years, and 14 option prices, with 13 weeks to expiration, are generated using $N = 5,000$ simulation runs. Second, the initial value for the interest rate variance (i.e. V_t) is estimated using the first 250 observations with the simple GARCH(1,1) model in (19). Third, the parameters of the LS model are estimated using the initial interest rate (i.e. r_{250}), the estimated variance, and the 14 generated bond prices. The 13-week LS density is then estimated using 5,000 simulation runs, while the option-implied RND is estimated using the generated option prices. Fourth, the probability integral transform is obtained for each of the two density forecasts, based on the realized outcome 13 weeks later (i.e. r_{263}),

$$z_{LS,263} = Q_{LS,263}\left(r_{263}\right), \tag{35}$$

and

$$z_{OPT,263} = Q_{OPT,263}\left(r_{263}\right). \tag{36}$$

I then move 13 weeks forward, and repeat the steps described above starting at time $t = 263$, to obtain the density estimates for the interest rate 13 weeks later, i.e. for r_{276}, and the resulting set of probability integral transforms, $z_{LS,276}$ and $z_{OPT,276}$. This procedure is repeated a total of 1,000 times in order to generate 1,000 $z-$values for each of the two density forecast methods. As previously mentioned, the sequences $\left\{z_{LS,i}\right\}_{i=1}^{1,000}$ and $\left\{z_{OPT,i}\right\}_{i=1}^{1,000}$ should be IID $U\left(0,1\right)$ if the two approaches provide "good" forecasts of the true density. In order to examine whether the z's are uniformly distributed, Figures 8 and 9 provide histograms (with 20 bins) of the LS and the option-implied $z-$values, along with approximate 95% confidence intervals for the bin heights under the null that z is IID $U\left(0,1\right)$.

Figure 8: Histogram of z (LS)

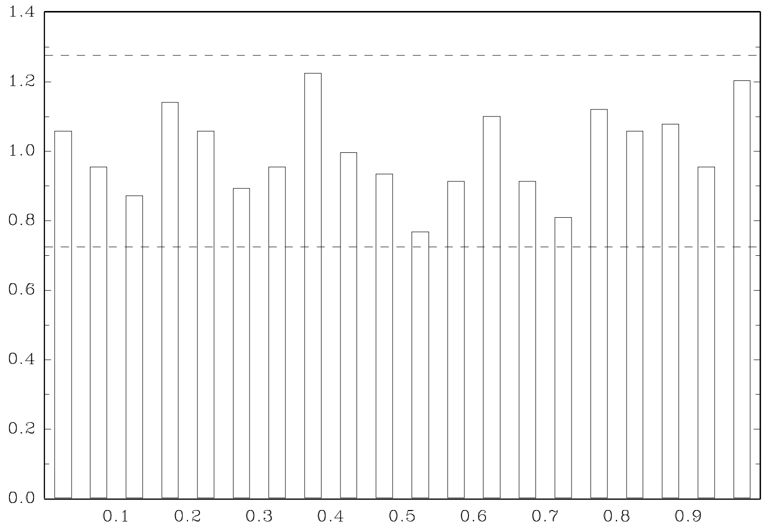

Figure 9: Histogram of z (OPT)

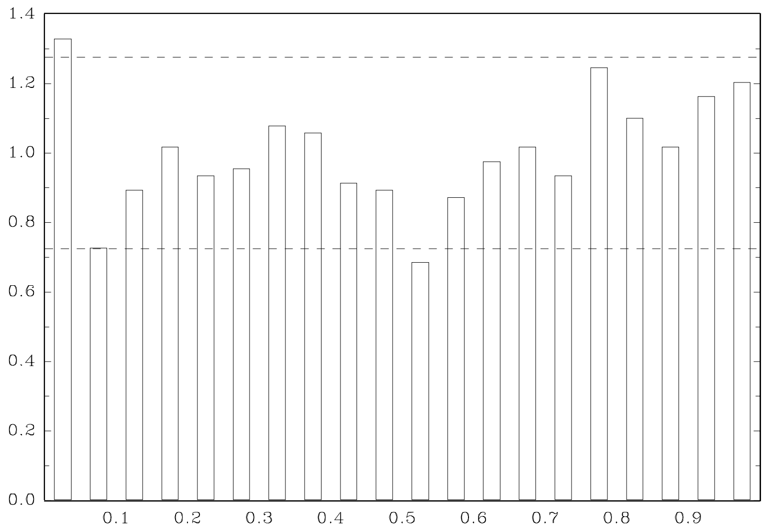

There are no clear visible patterns in the figures that might suggest any systematic violations of the $U(0,1)$ assumption. Furthermore, as shown by Figure 8, the height of all bars in the histogram for the LS density are within the 95% confidence interval, whereas two of the implied-option density bars in Figure 9 end up outside the confidence interval. The question of whether the distribution of the z's really are $U(0,1)$ can be tested more formally using a Kolmogorov-Smirnov test. As Table 3 displays, the null hypothesis that the sequence of probability integral transforms produced by the LS density forecasts is uniformly distributed over the unit interval cannot be rejected at the 5% level (or even at the 10% level). As for the option-implied density forecasts, the uniformity assumption can be rejected at the 10% level, but not at the 5% level. The IID condition can be evaluated using correlograms of $(z - \bar{z})$ and powers of this series. The correlograms reveal no significant serial correlation or any form of systematic nonlinear dependence in the z series of either forecast method; see Appendix C.

Table 3: Kolmogorov-Smirnov tests for the goodness of fit between the empirical distribution of the probability integral transforms (z) and the uniform (0,1) distribution.

Test	Test statistic (D)
$H_0 : Q\left(z_{LS,t}\right) = U\left(0,1\right)$	0.021
$H_0 : Q\left(z_{OPT,t}\right) = U\left(0,1\right)$	0.040 *

$Q\left(z_{LS,t}\right)$ and $Q\left(z_{OPT,t}\right)$ are the distributions of the probability integral transforms (z) for the Longstaff-Schwartz and the mix of lognormals density forecasts respectively. The critical values for the KS test are 0.039 (10%), 0.044 (5%), and 0.052 (1%). *, **, and *** denote significance at the 10%, 5%, and 1% level respectively.

Hence, there is no evidence pointing to a rejection of the hypothesis that z_{LS} is IID $U\left(0,1\right)$, and there is only weak support for rejecting the uniformity assumption for z_{OPT}. It therefore seems that, at least for the three month horizon, both the LS model and the mix of lognormals approach provide adequate density forecasts. In contrast to the results in the previous section, the option-implied density forecasts *do not* seem to outperform the Longstaff-Schwartz forecasts for the three month horizon. This lends support to the conclusion that the LS density forecasts can be useful in providing an indication of the uncertainty in the future short-term interest rate.

6 Conclusions

The possibility to extract information about market expectations from financial asset prices has attracted the interest of market participants and policymakers for a long time. Along with the rapid expansion of derivative markets in recent times, new methods have been developed for exploiting the information content in derivative prices. In particular, a number of different techniques have been proposed for deriving the implied risk-neutral density (RND) from observed prices of options. The implied RND provides valuable information, since it can be interpreted as the markets *ex ante* estimate of the probability density function of the underlying price on the expiration date of the options.

This paper demonstrates an approach for estimating implied distributions of future short-term interest rates, that does not depend on option prices. Instead, the Longstaff-Schwartz (1992a) two-factor model is used, which allows estimation of the implied distribution using cross-sectional data on prices of bills and bonds, conditional on current values of the two state variables: the short-term interest rate and the volatility of the short rate. This method therefore provides a useful alternative when interest rate option data is unavailable for some reason.

In order to test the performance of the LS density estimation approach, the results of the method are compared with those of a more traditional option-implied density estimation method. Furthermore, the implied densities obtained with the two methods are evaluated against the "true" PDF using simulation-based techniques. As might be expected, the results of this exercise show that the option-implied approach performs better than the LS method in the sense that the option-based implied density in general is closer to the "true" PDF. This is not surprising, since the option-implied method utilizes the information in a large number of option prices that expire at the date for which the density is estimated, whereas the LS method uses bond price data for a number of different maturities, all of which may differ from the forecast horizon for the PDF.

However, the density forecast evaluations also show that, at least for the shorter forecast horizons, the difference is relatively small between the estimated densities produced by the two methods, as well as between the estimated and the "true" PDFs. The LS model is also quite successful in capturing the main shape of the true underlying density. Furthermore, an evaluation of a long sequence of three-month density forecasts shows that the LS approach performs well in forecasting the true density, and it is also concluded that the option-implied density forecasts do not outperform the LS forecasts for the chosen forecast horizon. It therefore seems that the LS model is a useful alternative to the option-implied density estimation approach if option price data is unavailable, when the purpose of the estimated RNDs is to provide a general measure of the market's perceived uncertainty, for example as an indicator for monetary policy purposes.

A Estimating Implied PDFs from Option Prices[22]

Assuming the existence of an equivalent martingale, or risk-neutral, probability measure, the price at time t of a European call option expiring at time T, with strike price X, written on an underlying asset S, can be expressed as the present value of the option's expected payoff at the maturity date, where expectations are taken with respect to the risk neutral measure. Hence, we can write

$$c = e^{-r(T-t)} \int_{X}^{\infty} q\left(S_T\right)\left(S_T - X\right) dS_T, \tag{37}$$

where r is the risk-free interest rate during the option's life, and $q\left(S_T\right)$ is the risk-neutral density (RND) for the price of the underlying asset at the expiration date. Similarly, the price of a European put option can be written as

$$p = e^{-r(T-t)} \int_{-\infty}^{X} q\left(S_T\right)\left(X - S_T\right) dS_T. \tag{38}$$

It is possible to estimate the RND since the option prices above are functions of the density of the underlying asset. Assume that a weighted average of two lognormal densities is a suitable candidate for the underlying density of S_T :

$$q\left(S_T\right) = \theta L\left(\alpha_1, \beta_1, S_T\right) + (1-\theta) L\left(\alpha_2, \beta_2, S_T\right), \tag{39}$$

where

$$L\left(\alpha_i, \beta_i, S_T\right) = \frac{1}{S_T \beta_i \sqrt{2\pi}} \exp\left[-\frac{\left(\ln S_T - \alpha_i\right)^2}{2\beta_i^2}\right], \quad i = 1, 2, \tag{40}$$

and where α_i and β_i are location and dispersion parameters for each of the two lognormal distributions, while θ is the weighting parameter that determines the relative influence of the two lognormals on the terminal distribution.

By using at least five simultaneously observed option prices with the same expiration date, but with different strikes, the parameter vector $\Phi = \{\alpha_1, \alpha_2, \beta_1, \beta_2, \theta\}$ can be estimated using NLLS:

$$\min_{\Phi} \sum_{j=1}^{m} \left(c_j(\Phi) - c_j^*\right)^2 + \sum_{j=1}^{n} \left(p_j(\Phi) - p_j^*\right)^2, \tag{41}$$

where * denotes observed values.

[22] For further details about the estimation method, the reader is referred to the original article by Melick and Thomas (1997). See also Bahra (1997) and Söderlind and Svensson (1997).

B True and Implied RNDs for Alternative DGPs

Figure B1a: True and implied RNDs, 4−week horizon: Case B

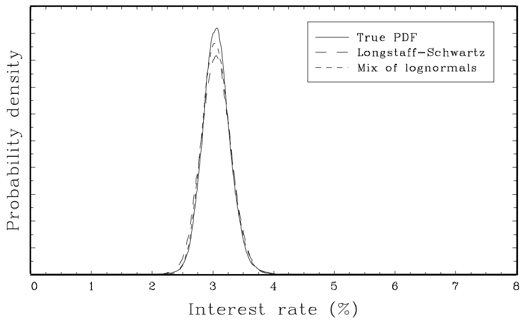

Figure B1b: True and implied RNDs, 13−week horizon: Case B

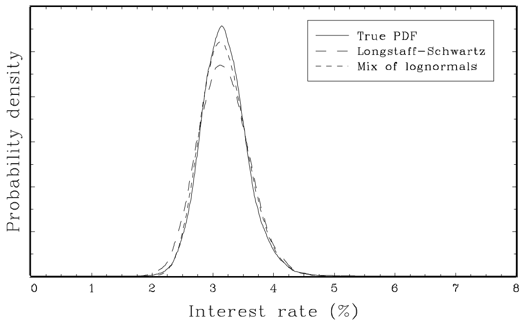

Figure B1c: True and implied RNDs, 26−week horizon: Case B

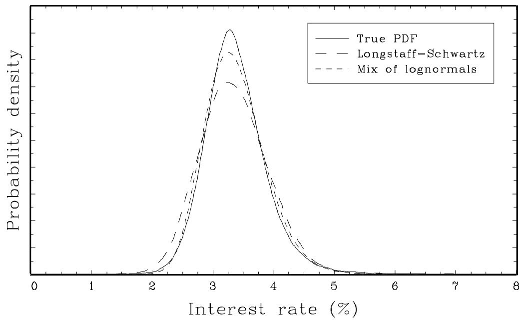

Figure B1d: True and implied RNDs, 52−week horizon: Case B

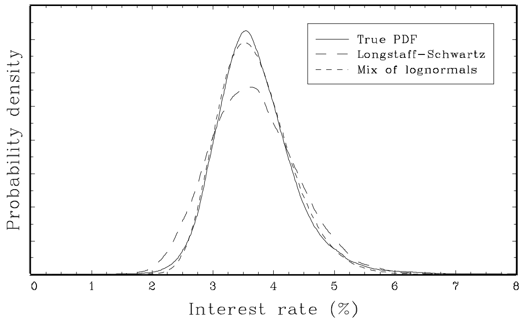

The DGP for Case B is identical to the one for Case A, except for the parameter α_0 which has been increased to 0.05, therefore producing an upward sloping implied forward rate curve as a result of the larger drift.

Table B1: Kolmogorov-Smirnov tests for the goodness of fit between the estimated and the "true" RNDs: Case B.

Test	Horizon	Test statistic (D)	
$H_0 : Q_{LS}(r) = Q_{TRUE}(r)$	4 weeks	0.040	***
	13 weeks	0.049	***
	26 weeks	0.061	***
	52 weeks	0.068	***
$H_0 : Q_{OPT}(r) = Q_{TRUE}(r)$	4 weeks	0.028	***
	13 weeks	0.020	***
	26 weeks	0.026	***
	52 weeks	0.012	***

*, **, and *** denote significance at the 10%, 5%, and 1% level respectively.

Figure B2a: True and implied RNDs, 4−week horizon: Case C

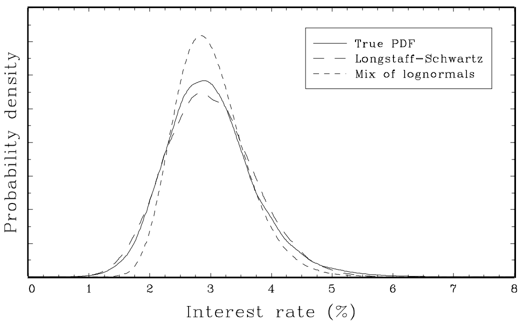

Figure B2b: True and implied RNDs, 13−week horizon: Case C

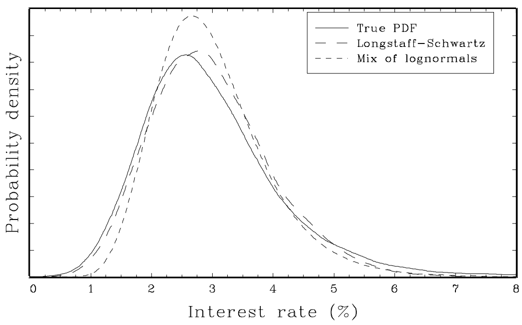

Figure B2c: True and implied RNDs, 26−week horizon: Case C

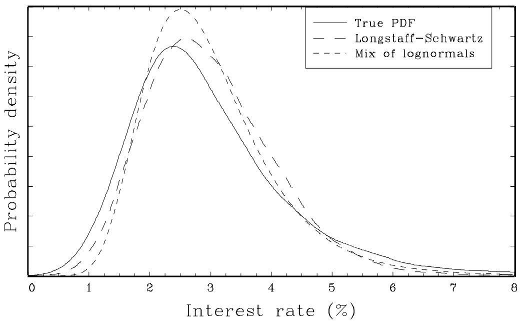

Figure B2d: True and implied RNDs, 52−week horizon: Case C

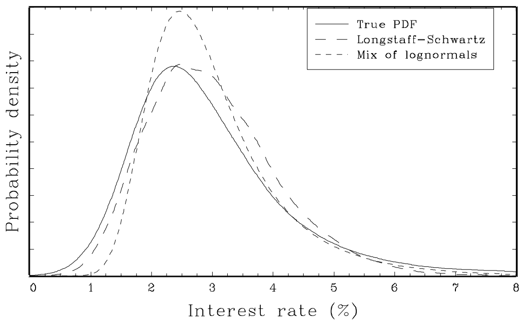

The DGP for Case C is identical to the one for Case A, but the initial variance has been increased by a factor of 10, thereby simulating turbulent market conditions at the estimation time.

Table B2: Kolmogorov-Smirnov tests for the goodness of fit between the estimated and the "true" RNDs: Case C.

Test	Horizon	Test statistic (D)
$H_0 : Q_{LS}(r) = Q_{TRUE}(r)$	4 weeks	0.021 ***
	13 weeks	0.035 ***
	26 weeks	0.062 ***
	52 weeks	0.062 ***
$H_0 : Q_{OPT}(r) = Q_{TRUE}(r)$	4 weeks	0.054 ***
	13 weeks	0.057 ***
	26 weeks	0.060 ***
	52 weeks	0.070 ***

*, **, and *** denote significance at the 10%, 5%, and 1% level respectively.

Figure B3a: True and implied RNDs, 4−week horizon: Case D

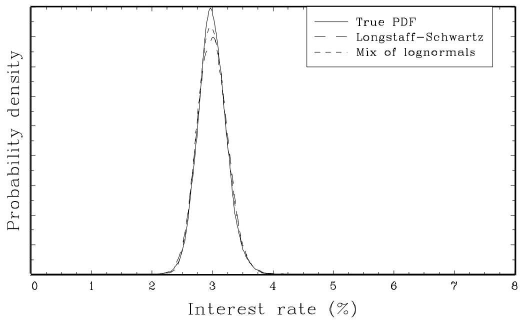

Figure B3b: True and implied RNDs, 13−week horizon: Case D

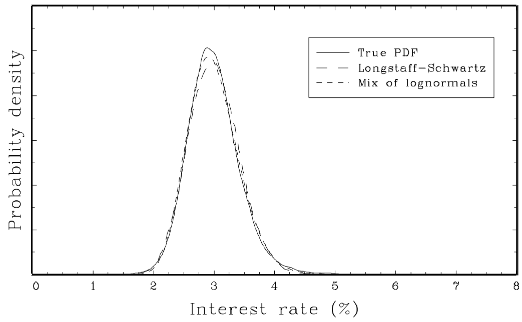

Figure B3c: True and implied RNDs, 26−week horizon: Case D

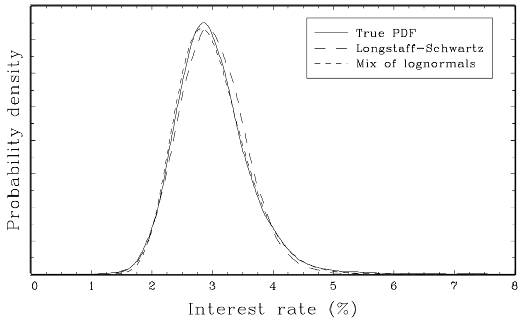

Figure B3d: True and implied RNDs, 52−week horizon: Case D

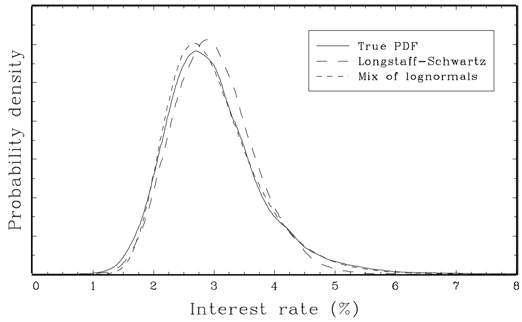

The DGP for Case D is identical to the one for Case A, except for the parameter β_1 which has been increased by a factor of 2, therefore producing a more volatile process for the short rate. (Compared with Case A, the variance in the simulated outcomes one year from the starting date is increased by about 45%).

Table B3: Kolmogorov-Smirnov tests for the goodness of fit between the estimated and the "true" RNDs: Case D.

Test	Horizon	Test statistic (D)	
$H_0 : Q_{LS}(r) = Q_{TRUE}(r)$	4 weeks	0.023	***
	13 weeks	0.026	***
	26 weeks	0.030	***
	52 weeks	0.047	***
$H_0 : Q_{OPT}(r) = Q_{TRUE}(r)$	4 weeks	0.024	***
	13 weeks	0.017	***
	26 weeks	0.014	***
	52 weeks	0.010	**

*, **, and *** denote significance at the 10%, 5%, and 1% level respectively.

Figure B4a: True and implied RNDs, 4−week horizon: Case E

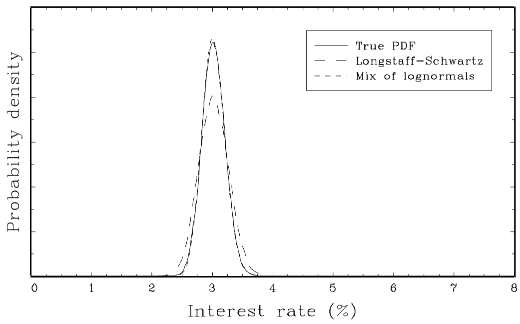

Figure B4b: True and implied RNDs, 13−week horizon: Case E

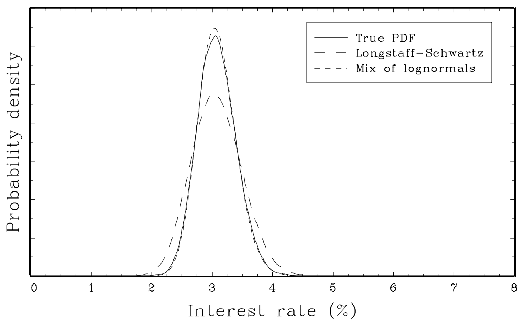

Figure B4c: True and implied RNDs, 26−week horizon: Case E

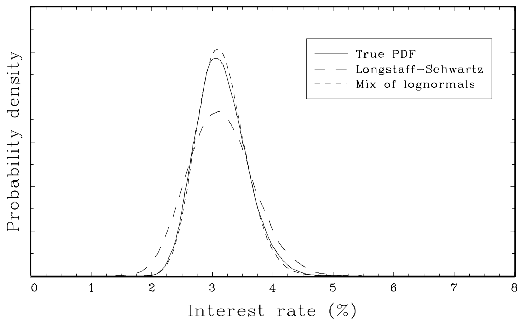

Figure B4d: True and implied RNDs, 52−week horizon: Case E

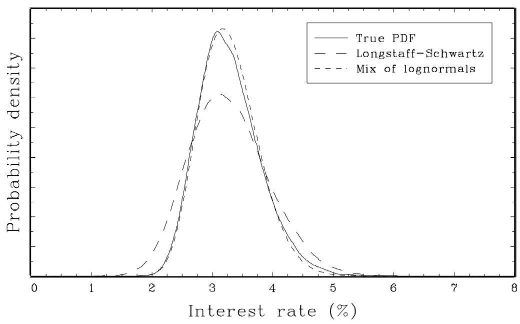

The DGP for Case E is the "general parametric model" of Aït-Sahalia (1996b), where the short term interest rate evolves according to $dr = \left(\alpha_0 + \alpha_1 r + \alpha_2 r^2 + \frac{\alpha_3}{r}\right) dt + \sqrt{\beta_0 + \beta_1 r + \beta_2 r^{\beta_3}} dZ$. This is the only parametric specification that was not rejected by Aït-Sahalia (1996b) when testing a number of different specifications using the distance between the parametric and a nonparametric density estimator. The parameters used for the DGP were obtained by estimating the model on the German 3-month Eurorates discussed in Section 5. The actual estimation was performed using FGLS on a discretized version of the model. The initial short-term interest rate is set to 3% in the simulations.

Table B4: Kolmogorov-Smirnov tests for the goodness of fit between the estimated and the "true" RNDs: Case E.

Test	Horizon	Test statistic (D)	
$H_0 : Q_{LS}(r) = Q_{TRUE}(r)$	4 weeks	0.067	***
	13 weeks	0.073	***
	26 weeks	0.077	***
	52 weeks	0.080	***
$H_0 : Q_{OPT}(r) = Q_{TRUE}(r)$	4 weeks	0.017	***
	13 weeks	0.010	**
	26 weeks	0.012	***
	52 weeks	0.015	***

*, **, and *** denote significance at the 10%, 5%, and 1% level.

C Correlograms of Powers of $(z - \bar{z})$

Figure C1: Correlograms of $(z - \bar{z})$ (upper left figure), $(z - \bar{z})^2$ (upper right figure), $(z - \bar{z})^3$ (lower left figure), and $(z - \bar{z})^4$ (lower right figure) for the Longstaff-Schwartz density forecasts. Dashed lines are approximate 95% confidence intervals.

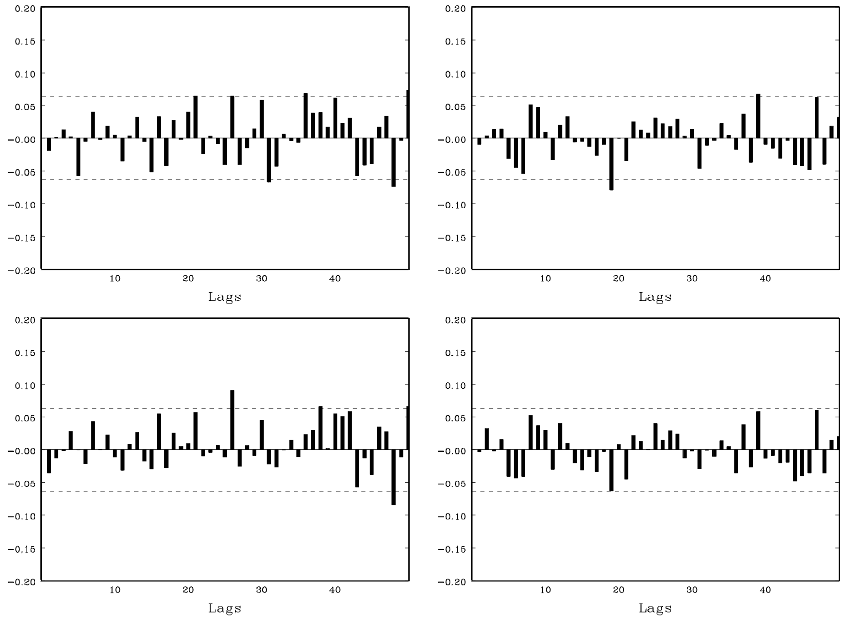

Figure C2: Correlograms of $(z - \bar{z})$ (upper left figure), $(z - \bar{z})^2$ (upper right figure), $(z - \bar{z})^3$ (lower left figure), and $(z - \bar{z})^4$ (lower right figure) for the mix of lognormals density forecasts. Dashed lines are approximate 95% confidence intervals.

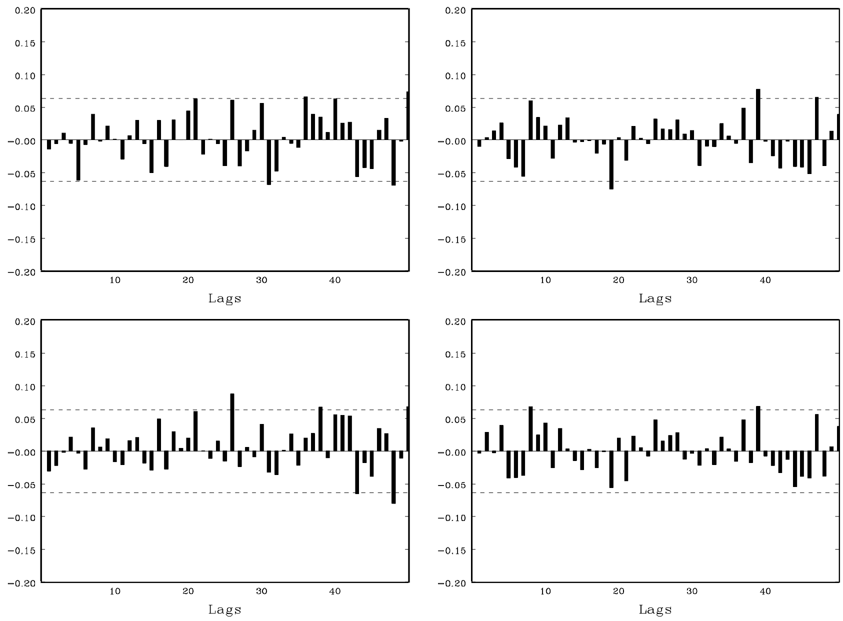

References

Aït-Sahalia, Y. (1996a): "Nonparametric Pricing of Interest Rate Derivative Securities," *Econometrica*, 64, 527-560.

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

Bahra, B. (1997): "Implied Risk-Neutral Probability Density Functions from Option Prices: Theory and Application," Working Paper, Bank of England.

Black, F., and M.S. Scholes (1973): "The Pricing of Options on Corporate Liabilities," *Journal of Political Economy*, 3, 637-654.

Bliss, R.R., and N. Panigirtzoglou (1999): "Testing the Stability of Implied Probability Density Functions," working paper Federal Reserve Bank of Chicago and Bank of England.

Bollerslev, T. (1986): "Generalized Autoregressive Conditional Heteroscedasticity," *Journal of Econometrics*, 31, 307-327.

Bollerslev, T., and J.M. Wooldridge (1992): "Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time-Varying Covariances," *Econometric Reviews*, 11, 143-172.

Brenner, R.J., R.H. Harjes, and K.F. Kroner (1996): "Another Look at Models of the Short-Term Interest Rate," *Journal of Financial and Quantitative Analysis*, 31, 85-107.

Broze, L., O. Scaillet, and J.-M. Zakoian (1995): "Testing for Continuous-Time Models of the Short-Term Interest Rate," *Journal of Empirical Finance*, 2, 199-223.

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

Cox, J.C., J.E. Ingersoll, and S.A. Ross (1980): "An Analysis of Variable Rate Loan Contracts," *Journal of Finance*, 35, 389-403.

Cox, J.C., J.E. Ingersoll, and S.A. Ross (1985): "A Theory of the Term Structure of Interest Rates," *Econometrica*, 53, 385-407.

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

Dahlquist, M., and L.E.O. Svensson (1996): "Estimating the Term Structure of Interest Rates for Monetary Policy Analysis," *Scandinavian Journal of Economics*, 98, 163-183.

Diebold, F.X., T.A. Gunther, and A.S. Tay (1998): "Evaluating Density Forecasts, with Applications to Financial Risk Management," *International Economic Review*, 39, 863-883.

Duffee, G.R. (1996): "Idiosyncratic Variation of Treasury Bill Yields," *Journal of Finance*, 51, 527-551.

Engle, R.F. (1982): "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of the United Kingdom inflation," *Econometrica*, 50, 987-1007.

Engle, R.F., and T. Bollerslev (1986): "Modelling the Persistence of Conditional Variances," *Econometric Reviews*, 5, 1-50.

Koedijk, K.G., F.G.J.A. Nissen, P.C. Schotman, and C.C.P. Wolff (1997): "The Dynamics of Short-Term Interest Rate Volatility Reconsidered," *European Finance Review*, 1, 105-130.

Longstaff, F.A., and E.S. Schwartz (1992a): "Interest-Rate Volatility and the Term Structure: A Two Factor General Equilibrium Model," *Journal of Finance*, 47, 1259-1282.

Longstaff, F.A., and E.S. Schwartz (1992b): "A Two-Factor Interest Rate Model and Contingent Claims Valuation," *Journal of Fixed Income*, 3, December, 16-23.

Melick, W.R., and C.P. Thomas (1997): "Recovering an Asset's Implied PDF from Option Prices: An Application to Crude Oil during the Gulf Crisis," *Journal of Financial and Quantitative Analysis*, 32, 91-115.

Rebonato, R. (1998): *Interest Rate Option Models*, John Wiley & Sons Ltd., Chichester.

Söderlind, P., and L.E.O. Svensson (1997): "New Techniques to Extract Market Expectations from Financial Instruments," *Journal of Monetary Economics*, 40, 383-429.

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