

Nonparametric Estimation of State-Price Densities Implicit in Interest Rate Cap Prices

HAITAO LI AND FENG ZHAO*

July 2007

*Li is from the Stephen M. Ross School of Business, University of Michigan. Zhao is from the Rutgers Business School, Rutgers University. We thank Yacine Aït-Sahalia (the editor), Ivan Brick, Ren-Raw Chen, Jefferson Duarte, Yongmiao Hong, Robert Jarrow, Jin Zhang, an anonymous referee, and seminar participants at Rutgers University, the 17th Derivatives Conference at FDIC, and the 2007 China International Conference in Finance for helpful discussions. We are responsible for any remaining errors.

**Nonparametric Estimation of State-Price Densities Implicit in Interest Rate Cap
Prices**

ABSTRACT

Based on a multivariate extension of the constrained locally polynomial estimator of Aït-Sahalia and Duarte (2003), we provide nonparametric estimates of the probability densities of LIBOR rates under forward martingale measures and the state-price densities (SPDs) implicit in interest rate cap prices conditional on the slope and volatility factors of LIBOR rates. Both the forward densities and the SPDs depend significantly on the volatility of LIBOR rates, and there is a significant impact of mortgage prepayment activities on the forward densities. The SPDs exhibit a pronounced U-shape as a function of future LIBOR rates, suggesting that the state prices are high at both extremely low and high interest rates, which tend to be associated with periods of economic recessions and high inflations, respectively. Our results provide nonparametric evidence of unspanned stochastic volatility and suggest that the unspanned factors could be partly driven by refinancing activities in the mortgage markets.

Over-the-counter interest rate derivatives, such as caps and swaptions, are among the most widely traded interest rate derivatives in the world. According to the Bank for International Settlements, in recent years, the notional value of caps and swaptions exceeds \$ 10 trillion, which is many times larger than that of exchange-traded options. Although the extensive term structure literature of the last decade has mainly focused on explaining bond yields and swap rates,¹ prices of caps and swaptions are likely to contain richer information on term structure dynamics because their payoffs are nonlinear functions of underlying interest rates. As a result, in a recent survey of the term structure literature, Dai and Singleton (2003) suggest that there is an “enormous potential for new insights from using (interest rate) derivatives data in (term structure) model estimations.”

Our paper contributes to the fast-growing literature on interest rate derivatives by providing one of the first nonparametric studies that extracts the rich information on term structure dynamics contained in prices of interest rate caps. Specifically, we provide nonparametric estimates of the probability densities of LIBOR rates under forward martingale measures and the SPDs using caps with a wide range of strike prices and maturities.² The nonparametric forward densities of LIBOR rates can be useful for many purposes. For example, the forward densities estimated using caps, which are among the simplest and the most liquid OTC interest rate derivatives, allow consistent pricing of more exotic and/or less liquid OTC interest rate derivatives based on the forward measure approach.³ The nonparametric forward densities also can reveal potential misspecifications of most existing term structure models, which rely on parametric assumptions to obtain closed-form pricing formulae for interest rate derivatives. As a result, these models are likely to be misspecified and therefore may not be able to fully capture the prices of interest rate derivatives.⁴ Furthermore, by combining the physical and forward densities of LIBOR rates, we can estimate the SPDs or the

¹See Dai and Singleton (2003) and Piazzesi (2003) for excellent surveys of the literature.

²In our paper, state-price density is defined as state price per unit of physical probability, following the definition in Duffie (2001) (see Section 1.F). This is slightly different from that of Ait-Sahalia and Lo (1998, 2000), which is essentially the risk-neutral density.

³Under a risk-neutral measure, the price of any security discounted by the money market account is a martingale. Similarly under a forward measure, the price of any security discounted by a zero-coupon bond associated with the forward measure is a martingale.

⁴For example, the popular LIBOR (Swap) market model assumes that LIBOR forward (swap) rates follow the log-normal distribution and prices caps (swaptions) using the Black (1976) formula. The models of Collin-Dufresne and Goldstein (2002), Han (2002), and Jarrow, Li, and Zhao (2006) rely on the affine jump-diffusion models of Duffie, Pan, and Singleton (2000), while the models of Li and Zhao (2006) rely on the quadratic term structure models of Ahn, Dittmar, and Gallant (2002) and Leippold and Wu (2003).

intertemporal marginal rate of substitutions of the representative investor implicit in cap prices. This allows us to study investor preferences from a perspective that is different from that of most existing studies in the literature, which are mainly based on S&P 500 index options.

The contributions of our paper to the existing literature are both methodological and empirical. Methodologically, we extend the constrained locally polynomial approach of Aït-Sahalia and Duarte (2003) to a multivariate setting. Building on the insights of Breeden and Litzenberger (1978), Aït-Sahalia and Duarte (2003) provide nonparametric estimates of risk-neutral densities using index options. Compared to the nonparametric kernel regression approach of Aït-Sahalia and Lo (1998, 2000) and others, the locally polynomial approach of Aït-Sahalia and Duarte (2003) has superior finite sample performance and guarantees that the estimated risk-neutral densities satisfy the necessary theoretical restrictions. Our multivariate extension of Aït-Sahalia and Duarte (2003) preserves all the advantages of the original method and makes it possible to estimate the forward densities and the SPDs conditional on multiple economic variables.

Based on the newly extended method, we provide nonparametric estimates of the forward LIBOR densities and the SPDs conditional on the slope and volatility factors of LIBOR rates.⁵ The SPDs are estimated as the ratio between the forward and physical densities of LIBOR rates, where the latter is estimated using the kernel method of Aït-Sahalia and Lo (2000). We include the two conditioning variables in our analysis because of the important roles they play for term structure modeling. For example, existing studies, such as Litterman and Scheinkman (1991), have widely documented that the level and slope factors can explain close to 99% of the variations of LIBOR rates. Many studies, such as Han (2002), Jarrow, Li, and Zhao (2007), and Trolle and Schwartz (2007), also have documented the importance of stochastic volatility for pricing interest rate caps. While the level factor is automatically incorporated in existing methods, our new extension of Aït-Sahalia and Duarte (2003) is needed to incorporate the slope and volatility factors in our nonparametric estimation. Interestingly, despite the overwhelming evidence of stochastic volatility in index returns, most existing nonparametric estimates of the SPDs using index options do not allow for stochastic volatility.⁶

Empirically, we find that the forward densities deviate significantly from the log-normal distribution assumed by the standard LIBOR market models and are strongly negatively skewed. More

⁵In this paper, the volatility factor is the first principal component of EGARCH-filtered spot volatilities of LIBOR rates at all maturities. We obtain similar results using GARCH-filtered spot volatilities.

⁶Boes, Drost, and Werker (2005) estimate a two-dimensional SPD as a function of index return and its EGARCH volatility using index options.

important, we find that both the forward densities and the SPDs depend significantly on the slope and volatility factors of LIBOR rates. We also document a pronounced U-shape of the SPDs as a function of future LIBOR rates. This result suggests that investors attach high values to payoffs when interest rates are either extremely high or low. This is consistent with the notion that extremely low interest rates tend to be associated with economic slowdowns or even recessions, while extremely high interest rates tend to be associated with hyper inflations. This pattern differs significantly from that of the SPDs estimated from index options, which is typically a declining function of the level of the equity market. Therefore, unlike most existing studies, our analysis provides new evidence on how the SPDs depend on important term structure factors. Moreover, while the index options used in most existing studies tend to have very short maturities (less than one or two years), the interest rate caps we consider allow us to estimate the SPDs over longer horizons.

Most interestingly, we document an important impact of mortgage prepayment activities on the forward LIBOR densities. Specifically, we find that the forward densities at most maturities (especially for 5 and 7 years, which are most relevant for mortgage hedging) are more negatively skewed following a sharp increase in prepayment activities even after controlling for both the slope and volatility factors. In an important study, Duarte (2006) shows that by allowing the volatility of LIBOR rates to be a function of prepayment speed in the string model of Longstaff, Santa-Clara, and Schwartz (2001), one can significantly reduce the pricing errors of at-the-money (ATM) swaptions. Our analysis extends Duarte (2006) in several important aspects. First, while Duarte (2006) focuses on the impact of prepayment on implied volatilities, we focus on its impact on the entire forward densities. It is possible that investors may hedge prepayment options using ATM or OTM interest rate options based on cost-benefit analysis. Therefore, our results show that prepayment affects not only the price of ATM options, but also the relative pricing of interest rate options across moneyness. Second, while the benchmark models in Duarte (2006) have either a constant or CEV volatility, we allow the forward densities to explicitly depend on the slope and volatility factors. While the slope factor is an important determinant of prepayment activities, the volatility factor is important for pricing interest rate options. Therefore, our results further strengthen the findings of Duarte (2006) because we show that even after controlling for stochastic volatility, prepayment still significantly affects the forward densities.

Our empirical results have important implications for one of the most important issues in the current term structure literature, namely the unspanned stochastic volatility (USV) puzzle. Most existing dynamic term structure models (DTSMs) assume that the same set of risk factors drive both

bond yields and interest rate derivative prices. This assumption implies that fixed income markets are complete and that DTSMs can simultaneously price bonds and interest rate options. However, Collin-Dufresne and Goldstein (2002), Heidari and Wu (2003), and Li and Zhao (2006) document the existence of systematic stochastic volatility factors in interest rate derivatives markets that cannot be spanned by bond market factors. These studies demonstrate the limitations of existing DTSMs and suggest that models with USV factors are needed to price interest rate options. The evidence of USV, however, is not without controversy. There are concerns that the results in the above studies, which are mainly based on parametric methods, such as linear regression or model-based hedging, may not be robust to alternative parametric specifications.⁷ Our paper complements the above studies by providing nonparametric and model-independent evidence of USV. More important, our results on the impact of prepayment on forward densities suggest that the unspanned factors could be partially driven by refinancing activities in the mortgage markets.

Our paper is closely related to Aït-Sahalia (1996 a&b), which are among the earliest studies on nonparametric pricing of interest rate derivatives. While Aït-Sahalia (1996 a&b) price interest rate derivatives based on nonparametrically estimated diffusion processes for spot interest rates, we provide nonparametric estimates of the forward densities of LIBOR rates conditional on the slope and volatility factors of LIBOR rates, and thus explicitly allows for USV in interest rate derivatives markets. Beber and Brandt (2006) is one of the few papers that estimates investor preferences using interest rate option prices. Our paper, however, differs from Beber and Brandt (2006) in several important aspects. First, while Beber and Brandt (2006) use the Edgeworth expansion method of Jarrow and Rudd (1982) to estimate the SPDs, we use the extended locally polynomial method of Aït-Sahalia and Duarte (2003). Second, the nonparametric method allows us to estimate the SPDs conditional on important term structure factors, such as the slope and volatility of LIBOR rates. In contrast, Beber and Brandt (2006) do not account for stochastic volatility in interest rates. Third, while Beber and Brandt (2006) consider only short-term interest rate options with maturities typically less than one year, we consider interest rate caps with maturities up to ten years. Most important, the main focus of Beber and Brandt (2006) is on the change of the SPDs around macroeconomic announcements, which is very different from the focus of our paper.

The rest of the paper is organized as follows. In Section 1, we discuss how to obtain the forward

⁷Fan, Gupta, and Ritchken (2003) argue that the linear regression approach of Collin-Dufresne and Goldstein (2002) and Heidari and Wu (2003) cannot fully capture the time-varying hedge ratios of interest rate options.

densities and the SPDs from interest rate cap prices. In Section 2, we discuss our multivariate extension of the constrained locally linear estimator of Aït-Sahalia and Duarte (2003), which is used for estimating the forward densities and the SPDs. In Section 3, we present the data and document a volatility smile in interest rate cap markets. Section 4 reports the empirical findings and Section 5 concludes. The appendix provides the mathematical proof.

1. Forward and State-Price Densities Implicit in Interest Rate Cap Prices

In this section, we discuss the general idea of estimating the forward densities and the SPDs from cap prices. Interest rate caps are portfolios of call options on LIBOR rates. Specifically, a cap gives its holder a series of European call options, called caplets, on LIBOR forward rates. Each caplet has the same strike price as the others, but with different expiration dates. For example, a five-year cap on three-month LIBOR struck at 6% represents a portfolio of 19 separately exercisable caplets with quarterly maturities ranging from six months to five years, where each caplet has a strike price of 6%.

Throughout our analysis, we restrict the cap maturity T to a finite set of dates $0 = T_0 < T_1 < \dots < T_K < T_{K+1}$, and we assume that the intervals $T_{k+1} - T_k$ are equally spaced by δ , a quarter of a year as in the U.S. cap markets. Let $L_k(t) = L(t, T_k)$ be the LIBOR forward rate for the actual period $[T_k, T_{k+1}]$, and let $D_k(t) = D(t, T_k)$ be the price of a zero-coupon bond maturing at T_k . We then have

$$L(t, T_k) = \frac{1}{\delta} \left(\frac{D(t, T_k)}{D(t, T_{k+1})} - 1 \right), \quad \text{for } k = 1, 2, \dots, K. \quad (1)$$

A caplet for the period $[T_k, T_{k+1}]$ struck at X pays $\delta(L_k(T_k) - X)^+$ at T_{k+1} . Although the cash flow of this caplet is received at time T_{k+1} , the LIBOR rate is determined at time T_k and there is no uncertainty about the caplet's cash flow after T_k .

For LIBOR-based instruments such as caps, floors, and swaptions, it is convenient to consider pricing using the forward measure approach. We will therefore focus on the dynamics of LIBOR forward rate $L_k(t)$ under the forward measure \mathbb{Q}^{k+1} , which is essential for pricing caplets maturing at T_{k+1} . Under this measure, the discounted price of any security using $D_{k+1}(t)$ as the numeraire is a martingale. Thus, the time- t price of a caplet maturing at T_{k+1} with a strike price of X is

$$Caplet(L_k(t), X, \tau_k) = \delta D_{k+1}(t) E_t^{\mathbb{Q}^{k+1}} [(L_k(T_k) - X)^+], \quad (2)$$

where $E_t^{\mathbb{Q}^{k+1}}$ is taken with respect to \mathbb{Q}^{k+1} given the information set at t and $\tau_k = T_k - t$, the time horizon over which $L_k(t)$ can randomly fluctuate. The key to valuation is the distribution of $L_k(t)$

under \mathbb{Q}^{k+1} . Once we know this distribution, we can price any security whose payoff on T_{k+1} depends only on $L_k(t)$ by discounting its expected payoff under \mathbb{Q}^{k+1} using $D_{k+1}(t)$.

Existing term structure models rely on parametric assumptions on the distribution of $L_k(t)$ to obtain closed-form pricing formulae for caplets. For example, the standard LIBOR market models of Brace, Gatarek, and Musiela (1997) and Miltersen, Sandmann, and Sondermann (1997) assume that $L_k(t)$ follows a log-normal distribution and price caplet using the Black formula. The models of Jarrow, Li, and Zhao (2007) assume that $L_k(t)$ follows affine jump-diffusions of Duffie, Pan, and Singleton (2000).

In this paper, we estimate the distribution of $L_k(t)$ under \mathbb{Q}^{k+1} using the prices of a cross section of caplets that mature at T_{k+1} and have different strike prices. Breeden and Litzenberger (1979) show that the density of $L_k(t)$ under \mathbb{Q}^{k+1} is proportional to the second derivative of $Caplet(L_k(t), \tau_k, X)$ with respect to X . Specifically, define

$$C(L_k(t), X, \tau_k) = E_t^{\mathbb{Q}^{k+1}} [(L_k(T_k) - X)^+] = \int_X^\infty (y - X) p^{\mathbb{Q}^{k+1}}(L_k(T_k) = y | L_k(t)) dy. \quad (3)$$

Then the conditional density of $L_k(T_k)$ under the forward measure \mathbb{Q}^{k+1} equals

$$p^{\mathbb{Q}^{k+1}}(L_k(T_k) | L_k(t)) = \frac{\partial^2 C(L_k(t), \tau_k, X)}{\partial X^2} \Big|_{X=L_k(T_k)}. \quad (4)$$

We assume that in (3)-(4) the conditional density of $L_k(T_k)$ depends only on the current LIBOR rate, i.e., $p^{\mathbb{Q}^{k+1}}(L_k(T_k) | \mathcal{F}_t) = p^{\mathbb{Q}^{k+1}}(L_k(T_k) | L_k(t))$, where \mathcal{F}_t represents the information set at t . This assumption, however, can be overly restrictive given the multifactor nature of term structure dynamics. For example, while the level factor can explain a large fraction (between 80-90%) of the variations of LIBOR rates, the slope factor still has significant explanatory power of interest rate variations. Moreover, there is overwhelming evidence that the volatility of interest rates are stochastic,⁸ and it has been suggested that interest rate volatility are unspanned in the sense that they can not be fully explained by the yield curve factors such as the level and slope factors.

One important innovation of our study is that we allow the volatility of $L_k(t)$ to be stochastic and the conditional density of $L_k(T_k)$ to depend on not only the level, but also the slope and volatility factors of LIBOR rates. That is, we assume that

$$p^{\mathbb{Q}^{k+1}}(L_k(T_k) | \mathcal{F}_t) = p^{\mathbb{Q}^{k+1}}(L_k(T_k) | L_k(t), Z(t)), \quad (5)$$

⁸See Andersen and Lund (1997), Ball and Torous (1999), Brenner, Harjes, and Kroner (1996), Chen and Scott (2001), and many others.

where $Z(t) = \{s(t), v(t)\}$, $s(t)$ is the slope factor of the LIBOR forward curve, and $v(t)$ is the common volatility factor of all LIBOR rates. Specifically, $s(t)$ is the difference between the 10- and 2-year LIBOR forward rates, and $v(t)$ is the first principal component of EGARCH-filtered spot volatilities of LIBOR rates across all maturities.

Under this generalization, we have

$$C(L_k(t), X, \tau_k, Z(t)) = \int_X^\infty (y - X) p^{\mathbb{Q}^{k+1}}(L_k(T_k) = y | L_k(t), Z(t)) dy. \quad (6)$$

And the conditional density of $L_k(T_k)$ under the forward measure \mathbb{Q}^{k+1} is given by

$$p^{\mathbb{Q}^{k+1}}(L_k(T_k) | L_k(t), Z(t)) = \frac{\partial^2 C(L_k(t), X, \tau_k, Z(t))}{\partial X^2} \Big|_{X=L_k(T_k)}. \quad (7)$$

Next we discuss how to estimate the SPDs by combining the forward and physical densities of LIBOR rates. Given a SPD function H , the price of the caplet can be calculated as

$$Caplet(L_k(t), X, \tau_k, Z(t)) = \delta E_t^{\mathbb{P}} [H \cdot (L_k(T_k) - X)^+], \quad (8)$$

where the expectation is taken under the physical measure. In general, H depends on multiple economic factors, and it is impossible to estimate it using interest rate caps alone. Given the available data, all we can estimate is the projection of H onto the future spot rate $L_k(T_k)$. Define

$$H_k(L_k(T_k); L_k(t), Z(t)) = E_t^{\mathbb{P}} [H | L_k(T_k); L_k(t), Z(t)], \quad (9)$$

then by iterated expectation,

$$\begin{aligned} Caplet(L_k(t), Z(t), \tau_k, X) &= \delta E_t^{\mathbb{P}} [H_k(L_k(T_k); L_k(t), Z(t)) (L_k(T_k) - X)^+] \\ &= \delta \int_X^\infty H_k(y) (y - X) p^{\mathbb{P}}(L_k(T_k) = y | L_k(t), Z(t)) dy. \end{aligned} \quad (10)$$

Comparing the above equation with the forward measure approach, we have

$$H_k(L_k(T_k); L_k(t), Z(t)) = D_{k+1}(t) \frac{p^{\mathbb{Q}^{k+1}}(L_k(T_k) | L_k(t), Z(t))}{p^{\mathbb{P}}(L_k(T_k) | L_k(t), Z(t))}. \quad (11)$$

Therefore, by combining the densities of $L_k(T_k)$ under \mathbb{Q}^{k+1} and \mathbb{P} , we can estimate the projection of H onto $L_k(T_k)$.

Another interpretation of H_k is the intertemporal rate of substitution of consumption of the representative investor. For an economy where the representative agent has a time-additive utility function, we have

$$H_k(L_k(T_k); L_k(t), Z(t)) = E_t^{\mathbb{P}} \left[\frac{U'(c_{T_{k+1}})}{U'(c_t)} \Big| L_k(T_k); L_k(t), Z(t) \right] \quad (12)$$

where $U'(\cdot)$ is the marginal utility of consumption, and $c_{T_{k+1}}$ and c_t are optimal consumptions at T_{k+1} and t , respectively. Therefore, we can estimate the dependence of future consumption $c_{T_{k+1}}$ on future spot interest rate $L_k(T_k)$ using a model-free approach based on the physical and forward densities of $L_{T_k}(T_k)$.⁹ The SPDs contain rich information on how risks are priced in financial markets. While Aït-Sahalia and Lo (1998, 2000), Jackwerth (2000), Rosenberg and Engle (2002), and others estimate the SPDs using index options (i.e., the projection of H onto index returns), our analysis based on interest rate caps documents the dependence of the SPDs on term structure factors.

Similar to many existing studies, to reduce the dimensionality of the problem, we assume that the caplet price is homogeneous of degree 1 in the current LIBOR rate. That is,

$$C(L_k(t), X, \tau_k, Z(t)) = L_k(t) C_M(M_k(t), \tau_k, Z(t)), \quad (13)$$

where the moneyness of the caplet $M_k(t) = X/L_k(t)$. Hence, for the rest of the paper we estimate the forward density of $L_k(T_k)/L_k(t)$ as the second derivative of the price function C_M with respect to M :

$$p^{\mathbb{Q}^{k+1}} \left(\frac{L_k(T_k)}{L_k(t)} \middle| Z(t) \right) = \frac{\partial^2 C_M(M_k(t), \tau_k, Z(t))}{\partial M^2} \Big|_{M=L_k(T_k)/L_k(t)}. \quad (14)$$

We also estimate the physical density of $L_k(T_k)/L_k(t)$ using the kernel method of Aït-Sahalia and Lo (2000).

2. Nonparametric Estimation of Forward and State-Price Densities

In this section, we discuss nonparametric estimation of the forward densities and the SPDs implicit in cap prices. We first provide a brief introduction to the locally polynomial approach of Aït-Sahalia and Duarte (2003).¹⁰ Then we discuss how to extend the method of Aït-Sahalia and Duarte (2003) to a multivariate setting to estimate the forward densities conditional on the slope and volatility factors of LIBOR rates. We also discuss how to combine the forward and physical densities of LIBOR rates to estimate the SPDs at different maturities. Finally, we provide simulation evidence on the accuracy of the newly extended nonparametric method.

2.1. A Brief Review of Locally Polynomial Estimation

Most existing nonparametric studies on SPDs typically estimate the option pricing formula $C_M(M_k(t), Z(t), \tau_k)$ nonparametrically, and then differentiate it twice with respect to M to obtain $\partial^2 C/\partial M^2$ and $p^{\mathbb{Q}^{k+1}}(L_k(T_k)|L_k(t), Z(t))$. Nonparametric estimation of the derivatives of a

⁹Here we ignore the small horizon difference between the consumption and spot interest rate and treat them as contemporaneous. The time difference is the tenor of the interest rate caps, which is three months.

¹⁰The discussion here relies heavily on Aït-Sahalia and Duarte (2003), which contains much more detailed descriptions of the local polynomial approach.

regression function, however, requires much more data than estimating the regression function itself. The increase of the dimensionality of the problem due to the conditioning variables $Z(t)$ further worsens the problem. Therefore, for a given sample size, the choice of the estimation method is crucial.

In our analysis, we adopt and extend the constrained locally polynomial method of Ait-Sahalia and Duarte (2003). Compared to the traditional kernel regression method, the locally polynomial approach has several important advantages. First, it has been well-established in the statistics and econometrics literature that the locally polynomial approach is superior to the kernel method in estimating the derivatives of nonlinear functions (See, for example, Fan and Gijbels (1996)). Second, the constrained locally polynomial approach of Ait-Sahalia and Duarte (2003) guarantees that the nonparametric option pricing function $\hat{C}_M(\cdot)$ satisfies the necessary theoretical restrictions. For example, to guarantee the absence of arbitrage across moneyness and the positivity of the density function, we must have $-1 \leq \partial\hat{C}/\partial M \leq 0$ and $\partial^2\hat{C}/\partial M^2 \geq 0$, which we refer to as the shape restrictions on $\hat{C}_M(\cdot)$. Finally, Ait-Sahalia and Duarte (2003) show that the locally polynomial estimation coupled with constrained least square regression also has better finite sample performance than existing models. Below we provide a brief review of the locally polynomial approach.

Suppose we have observations $\{(y_i, x_i)\}_{i=1}^n$ generated from the following relation

$$y = f(x) + \epsilon, \quad (15)$$

where $f(x)$ is an unknown nonlinear function and ϵ is a zero-mean error term. Suppose the $(p+1)^{\text{th}}$ derivative of $f(\cdot)$ at x exists. Then a Taylor expansion gives us an approximation of the unknown function $f(\cdot)$ in a neighborhood of x

$$\begin{aligned} f(z) &\approx f(x) + f'(x)(z-x) + \dots + \frac{f^{(p)}(x)}{p!}(z-x)^p \\ &= \sum_{k=0}^p \beta_{k,p}(x) \times (z-x)^k, \end{aligned} \quad (16)$$

where $\beta_{k,p}(x) = f^{(k)}(x)/k!$ and $f^{(k)}(x) = \frac{\partial f^k(z)}{\partial z^k}|_{z=x}$.

This representation of $f(\cdot)$ suggests that we can model $f(z)$ around x by a polynomial in z , and to use the regression of $f(z)$ on powers of $(z-x)$ to estimate the coefficients $\beta_{k,p}$. To ensure the local nature of the representation, we weight the observations by a kernel $K_h(x_i - x) = K((x_i - x)/h)/h$, where h is a bandwidth. Then the estimates of the coefficients $\hat{\beta}_{k,p}(x)$ are the minimizers of

$$\sum_{i=1}^n \left\{ y_i - \sum_{k=0}^p \beta_{k,p}(x) \times (x_i - x)^k \right\}^2 K_h(x_i - x). \quad (17)$$

At each fixed point x , this is a generalized least squares regression of the y_i s on the powers of $(x_i - x)$ s with diagonal weight matrix formed by the weights $K_h(x_i - x)$. This regression is “local” in the sense that the regression coefficients are only valid in a neighborhood of each point x .

The optimal way to estimate $f^{(k)}(x)$ based on asymptotics is to choose $p = k + 1$ and use the estimator

$$\hat{f}^{(k)}(x) = \hat{f}_p^{(k)}(x) = k! \hat{\beta}_{k,p}(x). \quad (18)$$

For example, a locally linear regression serves to estimate the regression function $\hat{f}_1^{(0)}(x)$, a locally quadratic regression for the first derivative $\hat{f}_2^{(1)}(x)$, and a locally cubic regression for the second derivative $\hat{f}_3^{(2)}(x)$.

Ait-Sahalia and Duarte (2003) show that in small samples the asymptotically optimal solution may not have the best performance. Instead, they consider the following alternative approach, which has superior finite sample performances than the asymptotic approach. Specifically, they estimate $f(\cdot)$ using a locally linear regression

$$\hat{f}(z) = \hat{\beta}_{0,1}(x) + \hat{\beta}_{1,1}(x)(z - x). \quad (19)$$

Hence, the regression function, the first and the second derivatives are estimated as, respectively,

$$\hat{f}(x) = \hat{\beta}_{0,1}(x), \quad \hat{f}^{(1)}(x) = \hat{\beta}_{1,1}(x), \quad \text{and} \quad \hat{f}^{(2)}(x) = \frac{\partial \hat{\beta}_{1,1}(x)}{\partial x} = \hat{\beta}'_{1,1}(x). \quad (20)$$

2.2. A Multivariate Extension of the Constrained Locally Linear Estimator

In this section, we extend the constrained locally polynomial method of Ait-Sahalia and Duarte (2003) to a multivariate setting to incorporate the conditioning variables $Z(t)$ in our analysis. The main objective of Ait-Sahalia and Duarte (2003) is to estimate risk-neutral densities using a small sample of option data, typically one day’s observations. This means that they essentially estimate the option price as a univariate function of the strike price. In our analysis, we would like to estimate the caplet price as a function of its moneyness conditional on the slope and volatility factors of LIBOR rates. The basic idea behind our approach is that we group observations on different dates that share similar values of the conditioning variables, and then within each group we solve a univariate problem similar to that of Ait-Sahalia and Duarte (2003).

For ease of exposition, we describe the extended method in a bivariate setting, although the description can be easily generalized to higher dimensions. Suppose we have a set of n observations y_1, y_2, \dots, y_n and their corresponding explanatory variables $(x_{11}, x_{12}), (x_{21}, x_{22}), \dots, (x_{n1}, x_{n2})$. Throughout the paper, we assume that the observations are ordered by the first explanatory variable,

i.e., $x_{i1} \geq x_{j1}$ if $i > j$, $1 \leq i, j \leq n$. Suppose we want to estimate the y_i 's as a function of x_{i1} 's for a fixed x_2 subject to the necessary shape restrictions. Define

$$D(x_2; \bar{h}) \triangleq \{i | x_{i2} \in [x_2 - \bar{h}, x_2 + \bar{h}], \text{ for } 1 \leq i \leq n\}, \quad (21)$$

where $\bar{h} > 0$. Basically, $D(x_2; \bar{h})$ contains the subsample of observations whose second explanatory variables are grouped around x_2 .

Aït-Sahalia and Duarte (2003) point out that nonparametric estimates of \hat{C}_M using the original data are not guaranteed to be arbitrage-free in finite samples. To address this problem, they first filter the data by solving the following constrained optimization problem:

$$\min_{m \in \mathbb{R}^p} \sum_{i \in D(x_2; \bar{h})} (m_i - y_i)^2 \quad (22)$$

subject to the slope and convexity constraints:

$$-1 \leq \frac{m_{i+1} - m_i}{x_{i+1,1} - x_{i1}} \leq 0 \text{ for all } i = 1, \dots, d-1, \quad (23)$$

$$\frac{m_{i+2} - m_{i+1}}{x_{i+2,1} - x_{i+1,1}} \geq \frac{m_{i+1} - m_i}{x_{i+1,1} - x_{i,1}} \text{ for all } i = 1, \dots, d-2, \quad (24)$$

where ρ is the number of elements in $D(x_2; \bar{h})$. Note that the solution m depends on the fixed value x_2 and the window size \bar{h} . This means that the filtering has to be done for each grid point of x_2 for nonparametric estimation. The basic restriction on \bar{h} is that it should be no smaller than the bandwidth used in the nonparametric smoothing along the second dimension. Aït-Sahalia and Duarte (2003) provide a fast computational algorithm for solving the above constrained optimization problem. The filtered data m closely resembles the original data y and satisfies the shape restrictions imposed by the theory. It is important to note that the filtering is done after obvious data errors have been removed.

To estimate the option pricing function $\hat{m}(x_1, x_2)$, and its first and second partial derivatives, $\frac{\partial \hat{m}(x_1, x_2)}{\partial x_1}$ and $\frac{\partial \hat{m}^2(x_1, x_2)}{\partial x_1^2}$, we minimize the following weighted sum of squared errors

$$\sum_{i=1}^n \{m_i - \beta_0(x_1, x_2) - \beta_1(x_1, x_2) \times (x_{i1} - x_1)\}^2 K_h(x_{i1} - x_1, x_{i2} - x_2), \quad (25)$$

where $K_h(x_{i1} - x_1, x_{i2} - x_2)$ is the joint kernel function. We have the following proposition which extends the analysis of Aït-Sahalia and Duarte (2003) to a bivariate setting.

Proposition 1. *Consider a set of n observations of the dependent variables, y_1, y_2, \dots, y_n and the corresponding bivariate independent variables, $(x_{11}, x_{12}), (x_{21}, x_{22}), \dots, (x_{n1}, x_{n2})$. Without loss*

of generality, let $x_{i1} \geq x_{j1}$ if $i > j$, $1 \leq i, j \leq n$. For any given pair (x_{i1}, x_{i2}) , the estimators $\frac{\partial \widehat{m}(x_1, x_2)}{\partial x_1}$ and $\frac{\partial \widehat{m}^2(x_1, x_2)}{\partial x_1^2}$ satisfy the required constraints in sample: $-1 \leq \frac{\partial \widehat{m}(x_1, x_2)}{\partial x_1} \leq 0$ and $\frac{\partial \widehat{m}^2(x_1, x_2)}{\partial x_1^2} \geq 0$, provided that

1. The transformed data $m_i(x_1, x_2)$, $i = 1, 2, \dots, d$, are obtained through the constrained least squares algorithm based on the original data in $D(x_2; \bar{h})$;
2. The joint kernel function can be written as a product of two univariate kernel functions

$$K_h(x_{i1} - x_1, x_{i2} - x_2) = K_{1,h}(x_{i1} - x_1) K_{2,h}(x_{i2} - x_2);$$

3. The kernel function $K_{1,h}(\cdot)$ is log-concave;
4. The bandwidth for the kernel function $K_{2,h}(\cdot)$, h_2^* , satisfies $h_2^* \leq \bar{h}$, and $K_{2,h}(z) = 0$ for $|z| > \bar{h}$.

Proof. See the appendix.

This proposition can be easily extended to higher dimensions. The key difference from the univariate case is that for the bivariate case we need to apply the constrained least square procedure at each grid point of x_2 . The ideal choice of \bar{h} should be h_2^* . However, the choice of h_2^* often depends on the filtered observations, which depend on the choice of \bar{h} . Therefore, in practice a two-stage procedure should be used in estimation. We can first conduct an unconstrained estimation to obtain the bandwidth h_2^* , which will be used later as the filtering window width in the second-stage constrained estimation.

In our implementation, x_1 represents the moneyness of the option, while x_2 and x_3 represent the slope and volatility factors, respectively. Following Proposition 1, we choose Epanechnikov kernels for the slope and volatility factors and the Gaussian kernel for the moneyness.

While the choices of kernels are relatively straightforward, the choices of bandwidths are much more complicated. The data are relatively evenly distributed along the moneyness dimension. Therefore, a global bandwidth (constant across moneyness) is acceptable in terms of the mean squared error (MSE) criterion. On the other hand, the observations of the slope and volatility factors are rather unevenly distributed. This makes a simple global bandwidth an inappropriate choice, because a global bandwidth will likely be too large and oversmooth the regions with dense observations and too small and undersmooth the regions with sparse observations. To address this problem, we

perform a monotonic transformation of the slope and volatility factors data via their empirical distribution functions and apply a global bandwidth to the transformed data. We first start with a preliminary bandwidth selection from the following relation

$$h_j = c_j \sigma(x_j) n^{-1/(4+d)}, j = 1, 2, 3, \quad (26)$$

where c_j is a constant that is close to one, $\sigma(x_j)$ is the sample standard deviation of x_j , and n is the sample size. Then we fine tune c_j to minimize the MSE via cross validation procedures.

The asymptotic and finite sample distributions of the constrained locally polynomial estimator are not known analytically. We obtain finite sample distribution of the estimator using bootstrap. In our estimation, the sample is divided into cells around certain slope and volatility grids. The estimator for a particular cell does not depend on the data outside that cell. We generate bootstrap samples by randomly drawing observations with replacements from a cell.¹¹ Conducting the estimation on all bootstrap samples, we obtain the finite sample distribution of the estimator.

2.3. Nonparametric Estimation of State-Price Densities

While we estimate $p^{\mathbb{Q}^{k+1}}(L_k(T_k) | L_k(t), Z(t))$ from caplet prices, we estimate $p^{\mathbb{P}}(L_k(T_k) | L_k(t), Z(t))$ using the underlying LIBOR rates based on the kernel method of Aït-Sahalia and Lo (2000). Let $L(t, T)$ be the time- t three-month LIBOR forward rate with a maturity of T , and $L(T, T)$ be the three-month LIBOR spot rate at T . Suppose we have the following time series observations $\{L(t_i, T_i), L(T_i, T_i)\}_{i=1}^n$, where $T_i - t_i = T - t$. We define the log-return of the LIBOR rates as

$$u_{t_i, T_i} = \log(L(T_i, T_i)) - \log(L(t_i, T_i)). \quad (27)$$

We use Z_{t_i} to denote the conditioning variables at t_i . The joint distribution of the log-return and the conditioning variables under the physical measure, $p_{u, Z}^{\mathbb{P}}(u_{t, T}, Z_t)$, can be estimated as

$$\hat{p}_{u, Z}^{\mathbb{P}}(u, z) = \frac{1}{n} \sum_{i=1}^n K_{h_u} \left(\frac{u_{t_i, T_i} - u}{h_u} \right) K_{h_z} \left(\frac{Z_{t_i} - z}{h_z} \right), \quad (28)$$

where $K_h(\cdot) = K(\cdot)/h$ and $K(\cdot)$ is a kernel function. In our estimation, we use the Gaussian kernel and assume the bivariate kernel is a product of one-dimensional kernels. The joint density of the conditioning variables under the physical measure can be estimated as

$$\hat{p}_Z^{\mathbb{P}}(z) = \frac{1}{n} \sum_{i=1}^n K_{h_z} \left(\frac{Z_{t_i} - z}{h_z} \right), \quad (29)$$

¹¹For identical observations, our estimator essentially keeps one among them.

and the density of the log-return conditional on $Z(t)$ is given by

$$\widehat{p}_{u|Z}^{\mathbb{P}}(u|Z) = \frac{\widehat{p}_{u,Z}^{\mathbb{P}}(u, z)}{\widehat{p}_Z^{\mathbb{P}}(z)}. \quad (30)$$

The bandwidths are chosen according to the following relation

$$h_j = c_j \sigma_j n^{-1/(4+d)}, j = u, z \quad (31)$$

where σ_j is the unconditional standard deviation of the data, c_j is a constant, and d represents the dimension of the problem ($d = 3$ for $\widehat{p}_{u,Z}^{\mathbb{P}}(u, z)$ and $d = 2$ for $\widehat{p}_Z^{\mathbb{P}}(z)$). Since the joint density estimator converges slower than the marginal one, the speed of convergence for the conditional density estimator is the same as that of the joint distribution, i.e., $O(n^{-2/(4+d)})$. In practice, similar to Aït-Sahalia and Lo (1998), we let the bandwidth converge to zero slightly faster than the MSE-optimal rate stated above in order to eliminate the asymptotic bias. This makes the asymptotic variance slightly larger and the MSE convergence speed slightly slower. The asymptotic distribution for the kernel estimator based on this bandwidth selection method is provided below.

$$n^{1/2} \prod_{j=1}^d h_j^{1/2} \left[\widehat{p}_{u|Z}^{\mathbb{P}}(u|Z) - p_{u|Z}^{\mathbb{P}}(u|Z) \right] \xrightarrow{d} \text{Normal} \left(0, \frac{p_{u|Z}^{\mathbb{P}}(u|Z)}{p_Z^{\mathbb{P}}(z)} \left(\int K(z)^2 dz \right)^d \right). \quad (32)$$

The estimator of the SPD equals

$$\widehat{H}_k(u|Z) = D_{k+1}(t) \frac{\widehat{p}_{u|Z}^{\mathbb{Q}^{k+1}}(u|Z)}{\widehat{p}_{u|Z}^{\mathbb{P}}(u|Z)}. \quad (33)$$

Note that $\widehat{H}_k(u|Z)$ is the projection of H onto the log ratio between the future spot rate and the current forward rate. Based on similar arguments in Aït-Sahalia and Lo (2000), we obtain the distribution of the pricing kernel estimator. Intuitively, the estimator of the physical density $\widehat{p}_{u|Z}^{\mathbb{P}}(u|Z)$ converges faster than that of the forward density $\widehat{p}_{u|Z}^{\mathbb{Q}^{k+1}}(u|Z)$, since the latter involves estimation of second-order derivatives with the same dimensionality. So, the asymptotic distribution of $\frac{\widehat{p}_{u|Z}^{\mathbb{Q}^{k+1}}(u|Z)}{\widehat{p}_{u|Z}^{\mathbb{P}}(u|Z)}$ is identical to that of $\frac{\widehat{p}_{u|Z}^{\mathbb{Q}^{k+1}}(u|Z)}{p_{u|Z}^{\mathbb{P}}(u|Z)}$, where we replace the true physical density with the estimate. Since we do not have the asymptotic distribution of the estimator $\widehat{p}_{u|Z}^{\mathbb{Q}^{k+1}}(u|Z)$, we obtain finite sample distribution using bootstrap, which tends to give a larger confidence interval.

2.4. Simulation Evidence

In this section, we provide simulation evidence on the accuracy of the multivariate extension of the constrained locally linear estimator. The basic idea of the new approach is to group observations on different dates whose conditioning variables are within a certain range and to apply the original

constrained locally linear estimator to the grouped data. Therefore, the data used in estimation within each group are likely to be more than just one day’s data. As a result, we expect the new method to work well in practice given the excellent finite sample performance of the original method of Ait-Sahalia and Duarte (2003) when only one day’s observations are used.

Our simulation design is similar to that of Ait-Sahalia and Duarte (2003). That is, we assume that caplet prices are determined by the Black (1976) formula. To capture the observed volatility skew in the data and the dependence of forward densities on the conditioning variables, we further assume that the log of the volatility is a linear function of moneyness, slope, and volatility of LIBOR rates with observation errors. The parameters used in the following specification are chosen so that the effects of the conditioning variables can be better illustrated:

$$\log \sigma (M_k (t), Z (t)) = -0.1 - 1.5M_k (t) - 0.75s (t) + 0.75v (t) + \epsilon (t), \quad (34)$$

where $\epsilon (t) \sim Uniform [-0.5\%, +0.5\%]$.¹²

We generate 500 random samples of $\sigma (M_k (t), Z (t))$, each with 2,000 observations, by producing random draws of $M_k (t)$, $s (t)$, and $v (t)$ from three independent uniform distributions with 20, 10, and 10 grid points, respectively. While we choose fixed grid points for $s (t)$ and $v (t)$, we allow the grid points of $M_k (t)$ to be random to avoid observations with identical moneyness. Based on the randomly generated volatilities $\sigma (M_k (t), Z (t))$, we calculate the prices of two-year caplets based on the Black (1976) formula. Ait-Sahalia and Duarte (2003) consider fixed moneyness grids in their simulations because they only use one day’s observations. Since we have to group caplet prices on different dates with similar conditioning variables, our simulation design captures the fact that moneyness of caplets change over time in our sample. While the actual slope and volatility factors do not follow uniform distributions, in our estimation we work with the transformed conditioning variables via their empirical distributions, which do follow uniform distributions. Thus, our simulation design is consistent with this estimation approach. We choose the sample size to be 2,000, because given our bandwidth, there are roughly 100 observations within each slope-volatility cell, which are similar to that in our empirical analysis.

Based on the simulated caplet prices, Figure 1 provides nonparametric estimates of the caplet price function, the first derivative of the price function, and the forward density conditional on the slope and volatility factors using the multivariate constrained locally linear estimator. The two levels of the slope and volatility factors used in Figure 1 are sufficiently different from each other to better

¹²The bid-ask spread of ATM caps quoted in implied volatility is 1%.

illustrate the effects of the conditioning variables, although we obtain similar results for conditioning variables at other levels. For each of the functions that we are trying to estimate, Figure 1 reports the true function, the average estimate over the 500 random samples, and the 95% confidence band over the 500 estimates. It is clear that our method can accurately recover the forward densities from the simulated caplet prices even when the two conditioning variables are involved.

3. The Data

In this section, we introduce the two datasets used in our empirical analysis. The first one contains daily LIBOR and swap rates, and the second one contains daily prices of interest rate caps with different strike prices and maturities.

We obtain daily LIBOR rates with maturities of three, six, and twelve months, as well as daily two-, three-, four-, five-, seven- and ten-year swap rates between August 13, 1990 and December 8, 2005 from Datastream. We bootstrap the swap rates to obtain daily three-month LIBOR forward rates with maturities beyond one year. Figure 2.A contains the term structure of LIBOR forward rates during our sample period. Both the level and shape of the LIBOR forward curve have exhibited rich variations during the 15-year time period, which spans the longest economic boom in U.S. history as well as the spectacular crash of the technology bubble.

One of the conditioning variables we use is the slope of the term structure, which is defined as the difference between the 10- and 2-year LIBOR forward rates. Besides the level factor, the slope factor has the biggest explanatory power of the variations of LIBOR rates. Figure 2.B contains the time series plot of the slope factor during the sample period. It is obvious that the slope of the forward curve has changed quite dramatically during our sample period, producing from very flat to quite steep forward curves. Figure 2.C reports the distribution of the slope factor in terms of both histogram and nonparametric kernel density estimator. For simplicity, we report the slope factor in percentage terms in Figure 2.C. In a later part of the paper, we report our empirical results based on three typical levels of the slope factor (0.6, 1.6, and 2.4), which represent flat, average, and steep forward curves, respectively.

We also obtain daily prices of interest rate caps between August 1, 2000 and July 26, 2004 from SwapPX. Jointly developed by GovPX and Garban-ICAP, SwapPX is the first widely distributed service delivering 24-hour real-time rates, data, and analytics for the world-wide interest rate swaps market. GovPX, established in the early 1990s by the major U.S. fixed-income dealers in a response to regulators' demands for increased transparency in the fixed-income markets, aggregates quotes from most of the largest fixed-income dealers in the world. Garban-ICAP is the world's leading swap

broker specializing in trades between dealers and trades between dealers and large customers. The data are collected every day the market is open between 3:30 and 4 p.m. Our data set is one of the most comprehensive ones available for caps written on dollar LIBOR rates. Other studies in this area include Gupta and Subrahmanyam (2005), Deuskar, Gupta, and Subrahmanyam (2003), and Troller and Schwartz (2007).

One advantage of our data is that we observe prices of caps over a wide range of strike prices and maturities. For example, every day for each maturity, there are 10 different strike prices: 4.0, 4.5, 5.0, 5.5, 6.0, 6.5, 7.0, 8.0, 9.0, and 10.0% between August 1, 2000 and October 17, 2001; 1.0, 1.5, 2.0, 2.5, 3.0, 3.5, 4.0, 4.5, 5.0 and 5.5% between October 18 and November 1, 2001; 2.5, 3.0, 3.5, 4.0, 4.5, 5.0, 5.5, 6.0, 6.5, and 7.0% between November 2, 2001 and July 15, 2002; 2.0, 2.5, 3.0, 3.5, 4.0, 4.5, 5.0, 5.5, 6.0, and 6.5% between July 16, 2002 and April 14, 2003; and 1.5, 2.0, 2.5, 3.0, 3.5, 4.0, 4.5, 5.0, 5.5 and 6.0% between April 15, 2003 and July 26, 2004. Moreover, caps have 15 different maturities throughout the whole sample period: 0.5, 1.0, 1.5, 2.0, 2.5, 3.0, 3.5, 4.0, 4.5, 5.0, 6.0, 7.0, 8.0, 9.0, and 10.0 years.

Our analysis uses prices of caplets, although we only observe cap prices. To obtain caplet prices, we consider the difference between the prices of caps with the same strike and adjacent maturities, which we refer to as *difference caps*. A difference cap includes a few caplets between two neighboring maturities with the same strike. For example, 1.5-year difference caps with a specific strike represent the sum of the 1.25-year and 1.5-year caplets with the same strike. We assume all individual caplets of a difference cap share the same Black implied volatility and calculate the price of each individual caplet using the Black formula.

Due to daily changes in LIBOR rates, caplets realize different moneyness (defined as the ratio between the strike price and the LIBOR forward rate underlying the caplet) each day. Therefore, throughout our analysis, we focus on the prices of caplets at given fixed moneyness. That is, each day we interpolate caplet prices with respect to the strike price to obtain prices at fixed moneyness. Specifically, we use locally cubic polynomials to preserve the shape of the original curves while smoothing over the grid points. We refrain from extrapolation and interpolation over grid points without nearby observations, and we eliminate all observations that violate various arbitrage restrictions. We also eliminate observations with zero prices, and observations that violate either monotonicity or convexity with respect to the strikes.

Figure 3.A contains the average Black implied volatilities of caplets across moneyness and maturity, while Figure 3.B plots the average implied volatilities of ATM caplets over the whole sample

period. Consistent with the existing literature, the implied volatilities of caplets with a moneyness between 0.8 and 1.2 have a humped shape with a peak at around a maturity of two years. However, the implied volatilities of all other caplets decline with maturity. There is also a pronounced volatility skew for caplets at all maturities, with the skew being stronger for short-term caplets. The pattern is similar to that of equity options: In-the-money (ITM) caplets have higher implied volatilities than do out-of-the-money (OTM) caplets. The implied volatilities of the very short-term caplets are more like a symmetric smile than a skew. Figure 3.C plots the time series of ATM implied volatilities of 1-, 3-, and 5-year caplets. It is clear that the implied volatilities are time varying and have increased dramatically during our sample period.

Both the volatility skew and the time series of ATM implied volatilities strongly suggest that LIBOR volatilities are time-varying and stochastic. Due to the evidence of USV and the important role of stochastic volatility for cap pricing demonstrated by existing studies, such as Jarrow, Li, and Zhao (2007), we include volatility as the second conditioning variable in our nonparametric estimation of the forward densities and the SPDs.

We construct the spot volatility factor in the following way. We first filter out spot volatilities of changes of three-month LIBOR forward rates at different maturities using an EGARCH model. Then we conduct principal component analysis of the spot volatilities and use the first principal component, which captures most of the variations of the spot volatilities, as the volatility factor. Figure 3.D reports the first three principal components of the spot volatilities, which explain 91.53%, 6.22%, and 1.33% of the variations of spot volatilities, respectively. Interestingly, these principal components also have an interpretation as the level, slope, and curvature of the spot volatilities. Figure 3.E contains the time series of the volatility factor, which has been normalized to have a mean that equals to one, while Figure 3.F reports the distribution of the volatility factor in terms of both histogram and kernel density estimator. The volatility factor has fluctuated within a certain range and is strongly mean reverting. In later part of the paper, we report our empirical results based on three typical levels of the volatility factor, which are below, at, and above the sample mean.

Duarte (2006) argues that mortgage refinancing activities have both an “actual volatility effect” and an “implied volatility effect.” Dynamic hedging of MBS using either Treasury securities or interest rate swaps affects time series volatility of LIBOR rates, which is the “actual volatility effect.” Static hedging of MBS using interest rate options affects the implied volatility of LIBOR rates, which is the “implied volatility effect.” While we could use either spot or implied volatility as the volatility factor, we choose spot volatility mainly because we are interested in testing the “implied volatility

effect” of Duarte (2006). As long as the two effects are not exactly the same, our approach makes it possible to test the effects of mortgage hedging on LIBOR forward densities.

4. Empirical Results

4.1 Nonparametric Estimates of Forward Densities of LIBOR Rates

The volatility smile or skew documented in the cap markets strongly suggests that the log-normal assumption of the standard LIBOR market models is violated in the data. Instead of considering parametric extensions of the log-normal model, we provide nonparametric estimates of the probability densities of LIBOR rates under forward martingale measures conditional on the slope and volatility factors of LIBOR rates.

As a preliminary evidence of the importance of the slope and volatility factors for the forward densities, Figure 4 provides nonparametric estimates of the Black implied volatilities as a function of moneyness and the slope factor.¹³ We see clearly that both the level and skewness of implied volatilities increase with the slope of the term structure. That is, when the slope is flat, the implied volatility curve is rather flat for caplets with maturities longer than 3 years. On the other hand, when the yield curve is steep, the implied volatility curve exhibits a strong volatility skew, with ITM caplets having much higher implied volatilities than ATM and OTM caplets.

We further plot the implied volatilities as a function of moneyness at three levels of the slope and volatility factors in Figure 5. The three levels of the slope factor correspond to flat, average, and steep yield curves, while the three levels of the volatility factor correspond to low, medium, and high levels of volatility. Again, we see clear dependence of the implied volatilities on the volatility factor after controlling for the slope effect. For a flat term structure, we see the strongest volatility skew at the medium level of the volatility factor. We see a similar pattern of volatility skew for an average-sloped term structure for caplets with maturities less than 3 years. When the slope is steep, the skewness increases with spot volatility for shorter maturities and remains the same for longer maturities. Across maturities, the implied volatilities of the 7- and 10-year caplets exhibit similar behaviors as those of the 2- to 5-year caplets, although the general levels of the implied volatilities are lower and the volatility skews are not as dramatic. This is consistent with the fact that the longer maturity LIBOR rates have lower spot volatilities and the longer maturity caplets have less significant volatility skews.

Figure 6 provides three-dimensional plots of nonparametric forward densities of the log-returns

¹³While we only consider caplets with maturities of 1, 3, 5, and 7 years, we obtain similar results for all other maturities.

of LIBOR rates conditional on the slope factor. Figure 7 plots the same forward densities at three different levels of the slope and volatility factors. Again we only consider caplets with 1, 3, 5, and 7 year maturities. The 95% confidence intervals are obtained through bootstrap. Under the forward measures, the LIBOR rates should be a martingale and the forward densities should have a mean that is close to zero. The expected log-returns of the LIBOR rates are slightly negative due to an adjustment from the Jensen's inequality.

Figures 6 and 7 show that one common feature of the forward densities is that the log-normal assumption underlying the popular LIBOR market models is grossly violated in the data, and the forward densities across all maturities are significantly negatively skewed. Later results show that the physical densities of the LIBOR rates are not as negatively skewed. Therefore, the negative skewness in the forward densities is mostly due to the negative skewness in the SPDs. Another interesting finding is that all the forward densities depend significantly on both the slope and volatility factors. For example, when the slope of the term structure is steeper than average, the forward densities across all maturities become much more dispersed. With a steep term structure, the current spot rate is low and is expected to rise in the future. This coincides with periods when the Fed lowers the short rate to spur economic growth. This result reveals a positive relation between the volatility of future spot rate and the slope of the term structure.

Though the dependence of the forward densities on the volatility factor is not very transparent when the term structure is very steep, the volatility effect is very significant when the slope is around the average level: When the spot volatility is low, the forward densities are compact with high peaks; as the spot volatility rises, the forward densities become much more dispersed and negatively skewed. This pattern holds for all maturities, although the effect becomes weaker for longer maturities. This result suggests that the volatility process is very persistent because current high spot volatility leads to high future spot volatility. When the term structure is relatively flat, the effects of the volatility factor vary across maturities. The 2- and 3-year forward densities are more negatively skewed when the spot volatility is high. The seven- and ten-year forward densities, however, become more dispersed when the spot volatility is either below or above the sample mean.

Our nonparametric analysis reveals significant nonlinear dependence of the forward densities on both the slope and volatility factors of LIBOR rates. These results have important implications for one of the most important and controversial topics in the current term structure literature, namely the USV puzzle. While existing studies on USV mainly rely on parametric methods, our results provide nonparametric evidence on the importance of USV: Even after controlling for important bond market

factors, such as level and slope, the volatility factor still significantly affects the forward densities of LIBOR rates. Our results also reveal the challenges in modeling volatility dynamics due to their nonlinear impacts on the forward densities.

4.2 Impact of Mortgage Prepayment on LIBOR Forward Densities

Some recent studies have documented close connections between activities in mortgage and interest rate derivatives markets. For example, in an interesting study, Duarte (2006) shows that ATM swaption implied volatilities are highly correlated with prepayment activities in the mortgage market. Duarte (2006) extends the string model of Longstaff, Santa-Clara, and Schwartz (2001) by allowing the volatility of LIBOR rates to be a function of the prepayment speed in the mortgage market. He shows that the new model has much smaller pricing errors for ATM swaptions than the original model with a constant volatility or a CEV model. Duarte’s findings suggest that if activities in the mortgage market, notably the hedging activities of government sponsored enterprises, such as Fannie Mae and Freddie Mac, affect the supply/demand of interest rate derivatives, then this source of risk may not be fully spanned by the factors driving the evolution of the term structure.

In this section, we examine the impact of mortgage prepayment on cap prices using our nonparametric method. Our analysis extends Duarte (2006) in several important dimensions. First, while Duarte focuses on the impact of prepayment on ATM swaption implied volatilities, we focus on the entire forward densities, and thus the pricing of caps/floors across moneyness. Second, while the benchmark models of Duarte (2006) only allow a constant or CEV volatility of LIBOR rates, we explicitly allow LIBOR forward densities to depend on the slope and volatility factors of LIBOR rates. While the slope factor can have nontrivial impact on prepayment behavior, the volatility factor is crucial for pricing interest rate options. Therefore, in the presence of these two factors, it is not clear whether prepayment still has incremental contributions in explaining interest rate option prices.

Our measure of prepayment activities is the weekly refinancing index, which is based on the number of mortgage loan applications from Mortgage Bankers Association of America (MBAA) during the same sample period of our caplet data. We take first-order difference of the refinancing index. We denote weeks with changes that are in the top 20 percentile of all weekly changes and the subsequent three weeks after each spike as high prepayment periods. As indicated by Duarte (2003) it usually takes four weeks for the loans to be approved. We denote the rest of the sample as low prepayment periods. Figure 8 reports the log of the MBAA index, where the shaded areas denote periods of high prepayments defined above. Interestingly, we find that most high (low) prepayment periods are indeed associated with steep (flat) term structures. This illustrates the importance for

controlling the slope factor in measuring the impact of prepayment on the forward densities.

We obtain nonparametric estimates of the forward densities based on the observations in the high prepayment periods conditional on different levels of the slope and volatility factors. We then compare these estimates with that obtained from the low prepayment periods, i.e., the rest of the sample. The four slope-volatility cells we consider are the ones with most observations. The first column of Figure 9 provides nonparametric estimates of the forward densities during low and high prepayment periods based on observations within the four slope-volatility cells, while the second column provides the differences between the forward densities during low and high prepayment periods with 95% bootstrapped confidence intervals. The impacts of prepayment activities on the forward densities become very pronounced at 5 and 7 year maturities, which are most relevant for mortgage hedging. For example, when the slope equals 2.2, the forward densities during high prepayment periods are much more negatively skewed than those in low prepayment periods, and these differences are statistically significant based on the 95% bootstrapped confidence bands. This is consistent with the notion that investors in MBS can hedge their potential losses from prepayment by buying OTM floors.

Our results confirm the findings of Duarte (2006) and show that prepayment affects the entire forward density and consequently the pricing of interest rate options across moneyness. Given that our findings hold even after controlling for the slope and volatility factors, they suggest that part of the USV factors could be driven by activities in the mortgage market.

4.3 State-Price Densities Implicit in Cap Prices

In addition to the forward densities, we also estimate the physical densities of LIBOR rates at different maturities. Combining the forward and physical densities, we provide nonparametric estimates of the SPDs over different horizons. Figure 10 provides nonparametric estimates of the physical densities of LIBOR rates at maturities of 2, 3, 4, and 5 years at the three levels of the slope and volatility factors.¹⁴ The 95% confidence intervals are calculated based on the asymptotic distribution of the kernel density estimator. The most important result from Figure 10 is that the physical densities are not as negatively skewed and widely dispersed as the forward densities. This suggests that the high dispersion and negative skewness of the forward densities are caused by the SPDs rather than the physical densities.

We still see clear dependence of the physical densities on the slope and volatility factors. Different

¹⁴We only report results up to five years because we do not have enough data on log-returns of LIBOR rates for longer maturities.

from the forward densities, the physical densities become more compact and closer to being normally distributed when the term structure is steeper than average. The dispersion of the physical densities is the largest when the slope is at the average level. This result could be due to mean reversion in interest rates. Also different from the forward densities, when the spot volatility is above the average, the physical densities become more compact and symmetric, which suggests a mean reverting volatility process. The 2- and 3-year LIBOR rates are more widely dispersed than the 4- and 5-year LIBOR rates. A comparison between the forward and physical densities shows that the slope and volatility factors are more persistent under the forward measures.

Figure 11 provides nonparametric estimates of the SPDs projected onto LIBOR spot rates at four different maturities for the three different levels of the slope and volatility factors.¹⁵ The most important findings from Figure 11 is that the SPDs exhibit a pronounced U-shape as a function of future LIBOR rates, especially at the 4- and 5-year maturities. This result suggests that investors attach high values to payoffs received when future LIBOR rates are either extremely high or low. This is consistent with the notion that low interest rates tend to be associated with economic slowdowns or even recessions, while high interest rates tend to be associated with high inflations. Investors with large bond portfolio holdings can hedge their potential losses due to rising interest rates using OTM caps. On the other hand, investors with large holdings in mortgage-backed securities can hedge their potential losses due to prepayments resulted from declining interest rates using OTM floors.¹⁶ Or in general, investors can hedge potential losses due to high inflation or recession using OTM caps or floors, respectively. These effects could explain the high state prices at both low and high levels of interest rates. The shape of the SPDs estimated from interest rate options differ significantly from that estimated using index options, which is typically a declining function of the level of the equity market.

The SPDs at the 4- and 5-year maturities have much more pronounced U-shape than those at the 2- and 3-year maturities. This is mainly driven by the fact that the physical densities at the 4- and 5-year maturities are more compact than those at the 2- and 3-year maturities, while the forward densities at all maturities have similar shapes. By the law of large numbers and mean-reversion in interest rates, fluctuations in interest rates should be canceled out over longer horizons. However, if interest rates indeed have gone up or down a lot over a longer horizon, it means that the rates

¹⁵The estimator of the SPD in (33), $\hat{H}_k(u|Z)$, depends on the discount factor $D_{k+1}(t)$. For ease of comparison of SPD estimates across maturities, we report that ration between $\hat{H}_k(u|Z)$ and $D_{k+1}(t)$ in Figure 11, i.e., $\frac{\hat{P}_{u|Z}^{Q^{k+1}}(u|Z)}{\hat{P}_{u|Z}^F(u|Z)}$.

¹⁶See Duarte (2006) for excellent discussion on the impacts of mortgage prepayments on interest rate volatility.

probably have consistently gone up or down over the time period, respectively. This further implies that extremely high or low interest rates correspond to really bad states of the economy, which lead to the high prices of risks in those states.

The SPDs at the 2- and 3-year maturities depend more significantly on the level and volatility factors. For example, at the 3-year maturity, when the slope of the term structure is very steep, the left arm of the U-shape is much more pronounced. Steep yield curves are typically observed in recessions when the Fed tends to lower short-term interest rates to spur growth. However, conditioning on a steep yield curve, low realizations of future interest rates mean that the Fed has not been successful in stimulating the economy, which has probably run into deeper recession. Therefore, those states with low interest rates are really bad states, which lead to high state prices. On the other hand, we find that the right arm of the U-shape becomes much more pronounced when the term structure is flat. With a flat yield curve, the Fed is probably raising short-term interest rate to slow the growth of the economy and inflation. Conditioning on that, if the rate increases a lot, then it means that the Fed has not been successful and the economy is probably suffering from high inflation, which again leads to high state prices.

The above analysis reveals some interesting features of the physical and forward densities of LIBOR rates as well as the SPDs.¹⁷ They show that both the physical and forward densities depend significantly on the slope and volatility factors of LIBOR rates. The SPDs show that interest rate options allow us to study investor preferences from a different perspective than index options. Given that each market only contains a subset of information about the pricing kernel, our analysis shows that we should explore the implications of asset pricing models in different markets.

5. Conclusion

In this paper, we extract the rich information on term structure dynamics contained in the prices of interest rate caps using nonparametric methods. Methodologically, we extend the constrained locally polynomial approach of Aït-Sahalia and Duarte (2003) to a multivariate setting and (for the first time) estimate the forward densities of LIBOR rates and the SPDs conditional on the slope and volatility factors of LIBOR rates. The multivariate constrained locally polynomial approach has excellent finite sample performances and guarantees that the nonparametric estimates satisfy

¹⁷One should be careful in making inferences based on the asymptotic distribution of the physical densities because interest rates are highly persistent. However, we believe that this fact is unlikely to affect our results on the SPDs because the U-shape is very pronounced and depends only on the fact that the physical densities are more compact than the forward densities.

necessary theoretical restrictions. Empirically, we find that the forward densities of LIBOR rates deviate significantly from the log-normal distribution and are strongly negatively skewed. Both the forward densities and the SPDs depend significantly on the volatility of LIBOR rates, and there is a significant impact of mortgage prepayment activities on the forward densities. The SPDs exhibit a pronounced U-shape as a function of future LIBOR rates, suggesting that the state prices are high at both extremely low and high interest rates, which tend to be associated with periods of economic recessions and high inflations, respectively. Our results highlight the importance of unspanned stochastic volatility and especially refinancing activities in the mortgage markets for term structure modeling.

REFERENCES

- Ahn, Dong-Hyun, Robert F. Dittmar, A. Ronald Gallant, 2002, Quadratic term structure models: Theory and evidence, *Review of Financial Studies* 15, 243-288.
- Aït-Sahalia, Yacine, 1996a, Nonparametric pricing of interest rate derivative securities, *Econometrica* 64, 527-560.
- Aït-Sahalia, Yacine, 1996b, Testing continuous-time models of the spot interest rate, *Review of Financial Studies* 9, 385-426.
- Aït-Sahalia, Yacine, and Andrew Lo, 1998, Nonparametric estimation of state-price densities implicit in financial asset prices, *Journal of Finance* 53, 499-547.
- Aït-Sahalia, Yacine, and Andrew Lo, 2000, Nonparametric risk management and implied risk aversion, *Journal of Econometrics* 94, 9-51.
- Aït-Sahalia, Yacine, and Jefferson Duarte, 2003, Nonparametric option pricing under shape restrictions, *Journal of Econometrics* 116, 9-47.
- Andersen, Torben G., and Luca Benzoni, 2006, Do bonds span volatility risk in the U.S. Treasury market? A specification test of affine term structure models, Working paper, Northwestern University.
- Andersen, Torben G., and Jesper Lund, 1997, Estimating continuous time stochastic volatility models of the short-term interest rate, *Journal of Econometrics* 77, 343-378.
- Ball, Clifford A., and Walter Torous, 1999, The stochastic volatility of short-term interest rates: Some international evidence, *Journal of Finance* 54, 2339-2359.
- Beber, Alessandro, and Michael Brandt, 2006, The effects of macroeconomic news on beliefs and preferences: Evidence from the options market, *Journal of Monetary Economics* forthcoming.
- Black, Fischer, 1976, The pricing of commodity contracts, *Journal of Financial Economics* 3, 167-179.
- Boes, Mark-Jan, Feike C. Drost, and Bas J.M. Werker, 2005, Nonparametric risk-neutral return and volatility distributions, Working paper, Tilburg University.

- Brace, Alan, Darius Gatarek, and Marek Musiela, 1997, The market model of interest rate dynamics, *Mathematical Finance* 7, 127-155.
- Breeden, Douglas, and Robert Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *Journal of Business* 51, 621-651.
- Brenner, Robin J., Richard H. Harjes, and Kenneth F. Kroner, 1996, Another look at alternative models of short-term interest rate, *Journal of Financial and Quantitative Analysis* 31, 85-107.
- Chen, Ren-raw, and Louis Scott, 2001, Stochastic volatility and jumps in interest rates: An empirical analysis, Working paper, Rutgers University.
- Collin-Dufresne, Pierre, and Robert S. Goldstein, 2002, Do bonds span the fixed income markets? Theory and evidence for unspanned stochastic volatility, *Journal of Finance* 57, 1685-1729.
- Dai, Qiang, and Kenneth J. Singleton, 2003, Term structure dynamics in theory and reality, *Review of Financial Studies* 16, 631-678.
- Deuskar, Prachi, Anurag Gupta, and Marti G. Subrahmanyam, 2003, Liquidity effects and volatility smiles in interest rate option markets, Working paper, New York University.
- Duarte, Jefferson, 2006, The causal effect of mortgage refinancing on interest rate volatility: Empirical evidence and theoretical implications, *Review of Financial Studies*, forthcoming.
- Duffie, Darrell, 2001, *Dynamic Asset Pricing Theory*, Princeton University Press, Princeton, New Jersey.
- Duffie, Darrell, Jun Pan, and Kenneth J. Singleton, 2000, Transform analysis and asset pricing for affine jump-diffusions, *Econometrica* 68, 1343-1376.
- Fan, Jianqing, and I. Gijbels, 1996, *Local Polynomial Modelling and its Applications*, Chapman and Hall, London.
- Fan, Rong, Anurag Gupta, and Peter Ritchken, 2003, Hedging in the possible presence of unspanned stochastic volatility: Evidence from swaption markets, *Journal of Finance* 58, 2219-2248.
- Gupta, Anurag, and Marti Subrahmanyam, 2005, Pricing and hedging interest rate options: Evidence from cap-floor markets, *Journal of Banking and Finance* 29, 701-733.

- Han, Bing, 2002, Stochastic volatilities and correlations of bond yields, Working paper, Ohio State University.
- Heath, David, Robert Jarrow, and Andrew Morton, 1992, Bond pricing and the term structure of interest rates: A new methodology, *Econometrica* 60, 77-105.
- Heidari, Massoud, and Liuren Wu, 2003, Are interest rate derivatives spanned by the term structure of interest rates? *Journal of Fixed Income* 13, 75-86.
- Jackwerth, Jens Carsten, 2000, Recovering risk aversion from option prices and realized returns, *Review of Financial Studies* 13, 433-451.
- Jarrow, Robert, Haitao Li, and Feng Zhao, 2007, Interest rate caps “smile” too! But can the LIBOR market models capture the smile?, *Journal of Finance* 62, 345-382.
- Jarrow Robert, and Andrew Rudd, 1982, Approximate option valuation for arbitrary stochastic processes, *Journal of Financial Economics* 10, 347-369.
- Leippold, Markus, and Liuren Wu, 2002, Asset pricing under the quadratic class, *Journal of Financial and Quantitative Analysis* 37, 271-295.
- Li, Haitao, and Feng Zhao, 2006, Unspanned stochastic volatility: Evidence from hedging interest rate derivatives, *Journal of Finance* 61, 341-378.
- Longstaff, Francis, Pedro Santa-Clara, and Eduardo Schwartz, 2001, The relative valuation of caps and swaptions: Theory and empirical evidence, *Journal of Finance* 56, 2067–2109.
- Litterman, Robert, and Jose Scheinkman, 1991, Common factors affecting bond returns, *Journal of Fixed Income* 1, 62–74.
- Miltersen, Kristian, Klaus Sandmann, and Dieter Sondermann, 1997, Closed-form solutions for term structure derivatives with lognormal interest rates, *Journal of Finance* 52, 409-430.
- Piazzesi, Monika, 2003, Affine term structure models, *Handbook of Financial Econometrics*, forthcoming.
- Rosenberg Joshua, and Robert Engle, 2002, Empirical pricing kernels, *Journal of Financial Economics* 64, 341-372.

Trolle Anders, and Eduardo Schwartz, 2007, A general stochastic volatility model for the pricing of interest rate derivatives, *Review of Financial Studies* forthcoming.

Appendix. Proof of Proposition 1.

For a fixed (x_1, x_2) , we first augment the sample to the size of the original observations by including the unfiltered observations that are not in $D(x_2; \bar{h})$. Next we define

$$\begin{aligned} k_{l,i} &= K_{l,h}(x_{il} - x_l), l = 1, 2, \\ k_{i,j} &= (x_{i1} - x_1)^2 k_{1,i} k_{1,j} k_{2,i} k_{2,j}, \\ M_{i,j} &= (m_i - m_j) / (x_{i1} - x_1). \end{aligned}$$

The estimator of the first-order derivative equals

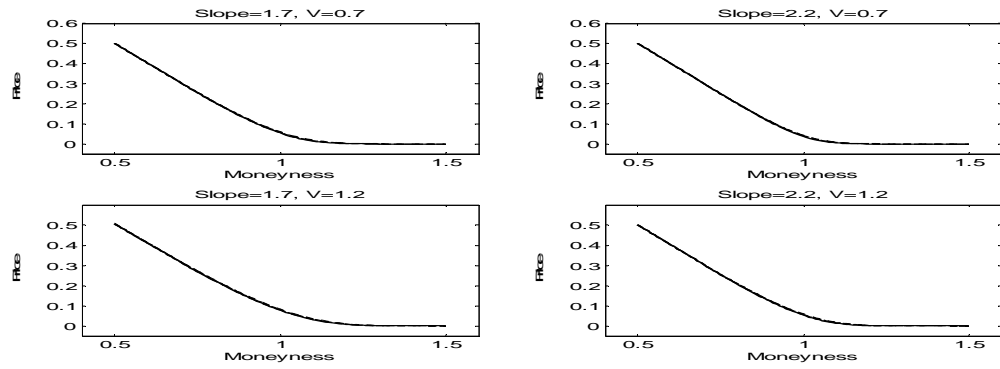
$$\frac{\partial \hat{m}(x_1, x_2)}{\partial x_1} = \frac{\sum_{i=1}^{n-1} \sum_{j=i+1}^n M_{i,j} k_{i,j}}{\sum_{k=1}^{n-1} \sum_{l=k+1}^n k_{k,l}}.$$

Since $k_{2,i,j} = 0$ outside the set $D(x_2; \bar{h})$, we can reduce the sum on this set only, i.e.,

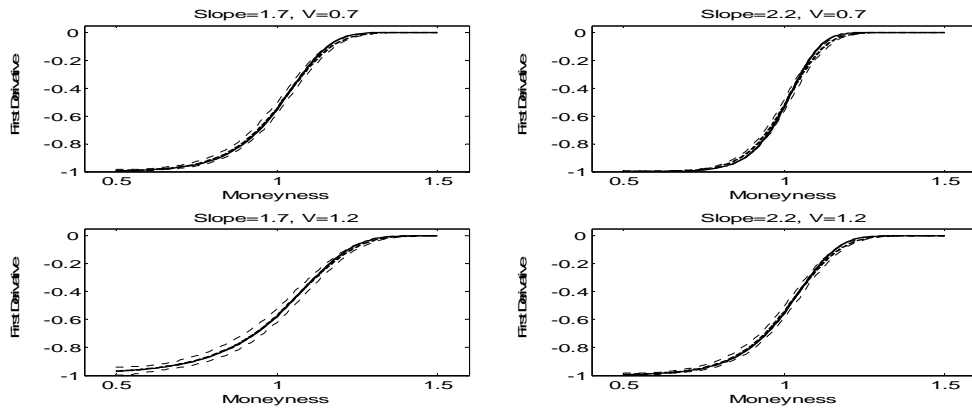
$$\begin{aligned} \frac{\partial \hat{m}(x_1, x_2)}{\partial x_1} &= \frac{\sum_{i=1}^{d-1} \sum_{j=i+1}^d M_{i,j} k_{i,j}}{\sum_{k=1}^{d-1} \sum_{l=k+1}^d k_{k,l}}, \\ \frac{\partial \hat{m}^2(x_1, x_2)}{\partial x_1^2} &= \frac{\left(\sum_{i=1}^{d-1} \sum_{j=i+1}^d M_{i,j} k'_{i,j} \right) \left(\sum_{k=1}^{d-1} \sum_{l=k+1}^d k_{k,l} \right) - \left(\sum_{i=1}^{d-1} \sum_{j=i+1}^d M_{i,j} k_{i,j} \right) \left(\sum_{k=1}^{d-1} \sum_{l=k+1}^d k'_{k,l} \right)}{\left(\sum_{k=1}^{d-1} \sum_{l=k+1}^d k_{k,l} \right)^2} \end{aligned}$$

Ait-Sahalia and Duarte (2003) show $-1 \leq \frac{\partial \hat{m}(x_1, x_2)}{\partial x_1} \leq 0$ and $\frac{\partial \hat{m}^2(x_1, x_2)}{\partial x_1^2} \geq 0$ for the univariate case. It can be easily shown from the definition that the same result applies here.

1.A. Nonparametric Estimate of the Caplet Price Function



1.B. Nonparametric Estimate of the First Derivative of the Caplet Price Function



1.C. Nonparametric Estimate of the Forward Density

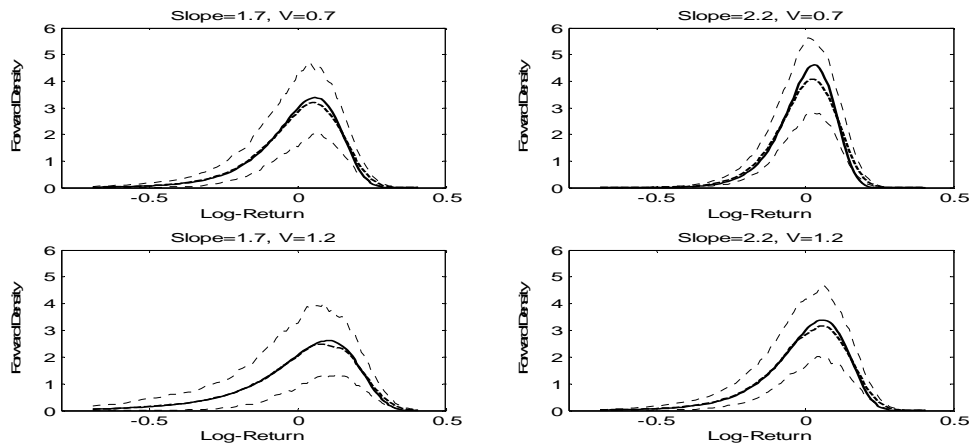
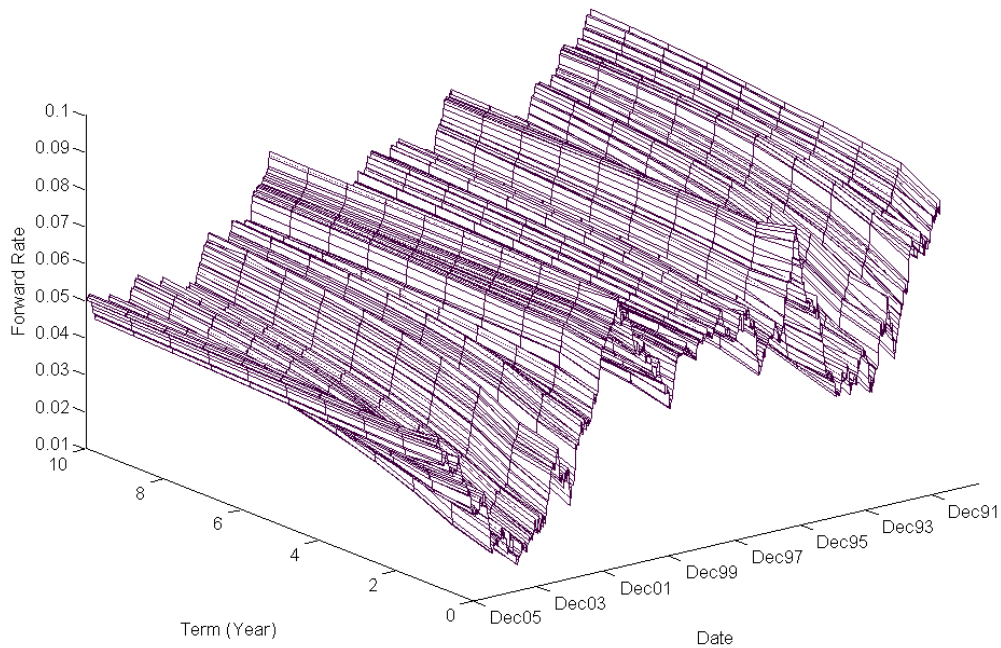



Figure 1. Simulation evidence on the finite sample performance of the multivariate constrained locally polynomial estimator. We assume that caplet prices are determined by the Black (1976) formula, in which the log of the volatility is a linear function of moneyness, slope, and volatility of LIBOR rates with observation errors. We generate 500 random samples of the volatility, each with 2,000 observations, by producing random draws of moneyness, slope, and volatility from three independent uniform distributions with 20, 10, and 10 grid points, respectively. Based on the simulated caplet prices, we provide nonparametric estimates of the caplet price function, the first derivative of the price function, and the forward density conditional on the slope and volatility factors using the multivariate constrained locally polynomial estimator. For each of the functions that we are trying to estimate, we report the true function (solid line), the average estimate over the 500 random samples (thick dashed line), and the 95% confidence band over the 500 estimates (thin dashed line).

2.A. The Term Structure of LIBOR Forward Rates



2.B. Time Series of the Slope Factor



2.C. Distribution of the Slope Factor






Figure 2. The term structure of LIBOR forward rates and the slope factor of the term structure. Based on daily three-month LIBOR rates with maturities of three, six, and twelve months, as well as daily two-, three-, four-, five-, seven- and ten-year swap rates between August 13, 1990 and December 8, 2005 obtained from Datastream, we obtain daily term structure of LIBOR forward rates as shown in Figure 2.A. We define the slope factor as the difference between the 10- and 2-year LIBOR forward rates and provide its time series plot and distribution in Figure 2.B and 2.C, respectively. For convenience, we report the slope factor in percentage terms in Figure 2.C.


3.A. Average Black Implied Volatilities




3.B. Average ATM Black Implied Volatilities




3.C. Time Series of ATM Black Implied Volatilities



3.D. Principle Components of Filtered Spot Volatilities



3.E. Time Series of the Volatility Factor



3.F. Distribution of the Volatility Factor






Figure 3. The Black implied volatilities of caplets and the spot volatility factor. Figure 3.A plots the average Black-implied volatilities of caplets across moneyness and maturity. Figure 3.B plots the average implied volatilities of ATM caplets. Figure 3.C plots the time series of ATM implied volatilities of 1-, 3-, and 5-year caplets. Figure 3.D reports the first three principal components of EGARCH-filtered spot volatilities of changes of LIBOR rates. Figure 3.E contains the time series of the volatility factor, while Figure 3.F reports the distribution of the volatility factor.


4.A. 1-Year



4.B. 3-Year



4.C. 5-Year



4.D. 7-Year


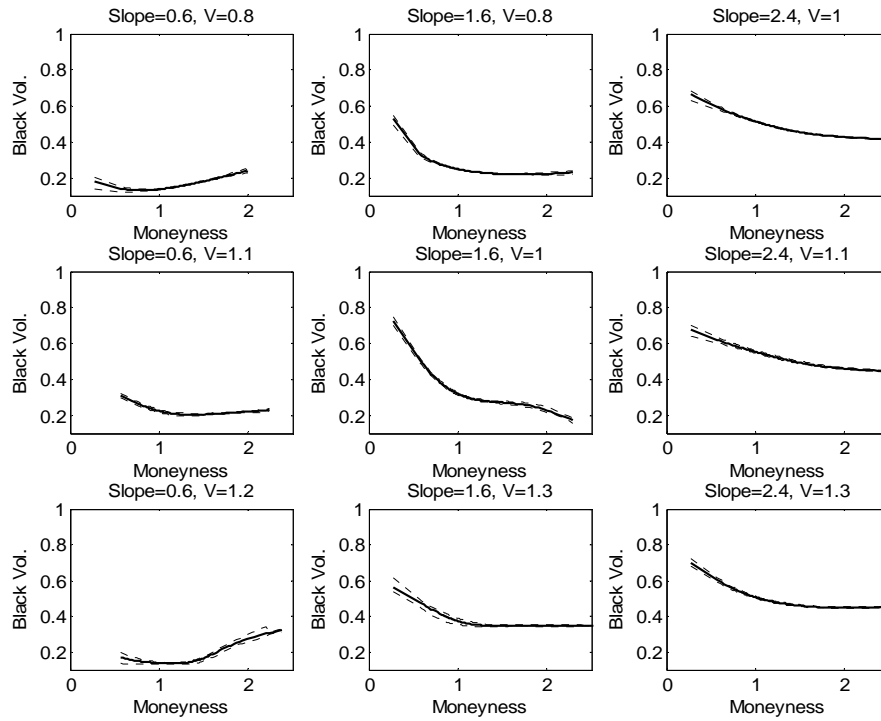


Figure 4. Nonparametric estimates of the Black implied volatilities of caplets as a function of moneyness and the slope factor. The slope factor is defined as the difference between the 10- and 2-year three-month LIBOR forward rates. The slope ranges from 0% (flat forward curve) to 3% (steep forward curve).

5.A. 1-Year



5.B. 3-Year

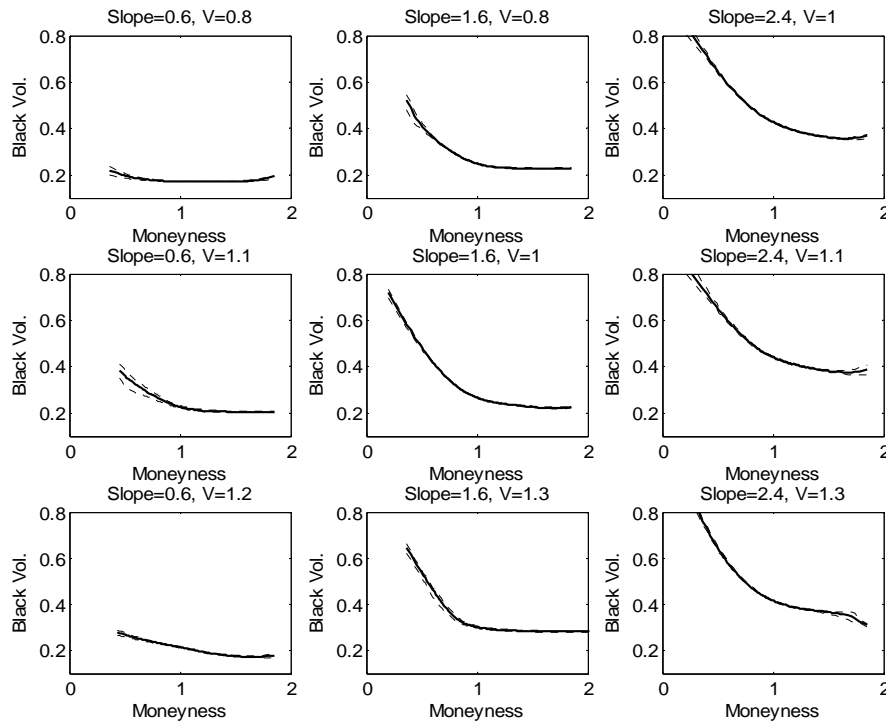
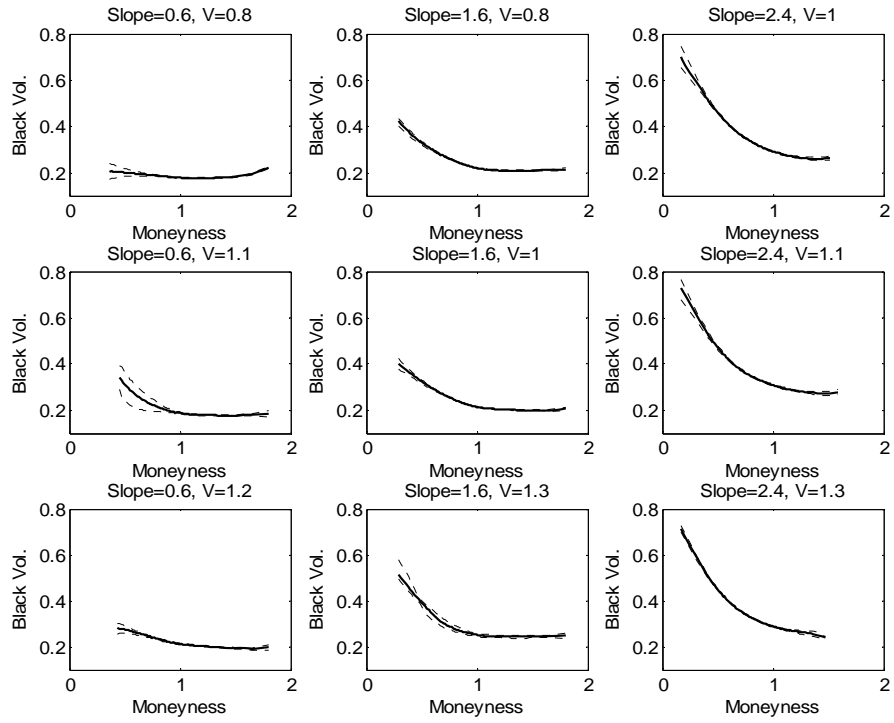


Figure 5. Nonparametric estimates of the Black implied volatilities of caplets as a function of moneyness at three levels of the slope and volatility factors. The slope factor is defined as the difference between the 10- and 2-year three-month LIBOR forward rates. The volatility factor is defined as the first principal component of EGARCH-filtered spot volatilities and has been normalized to a mean that equals one. The three levels of the slope factor correspond to flat, average, and steep forward curves, while the three levels of the volatility factor corresponds to low, medium, and high levels of volatility.

5.C. 5-Year



5.D. 7-Year

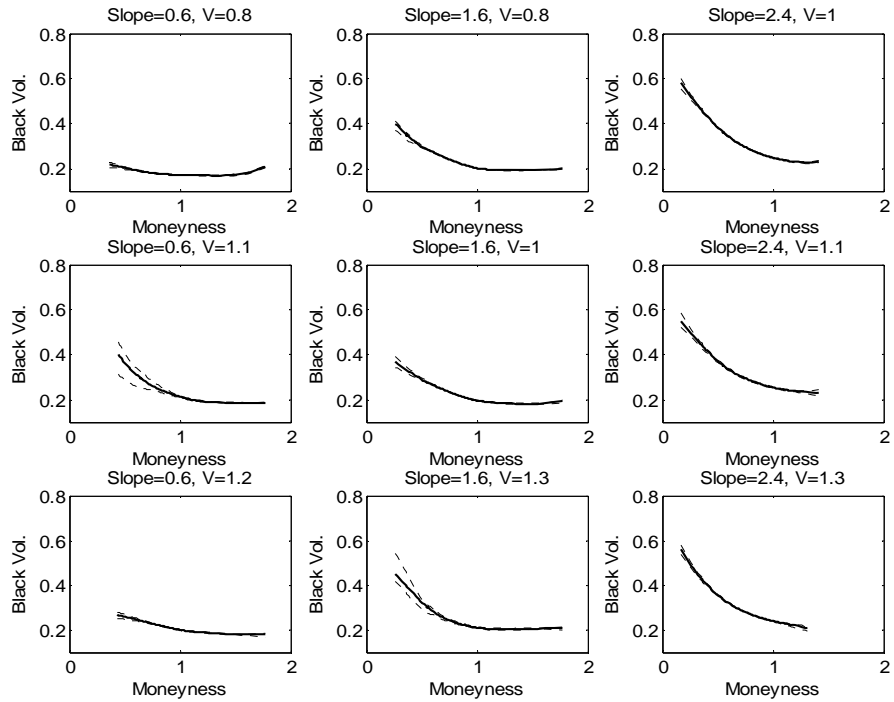




Figure 5. Nonparametric estimates of the Black implied volatilities of caplets as a function of moneyness at three levels of the slope and volatility factors (continued).


6.A. 1-Year



6.B. 3-Year



6.C. 5-Year



6.D. 7-Year


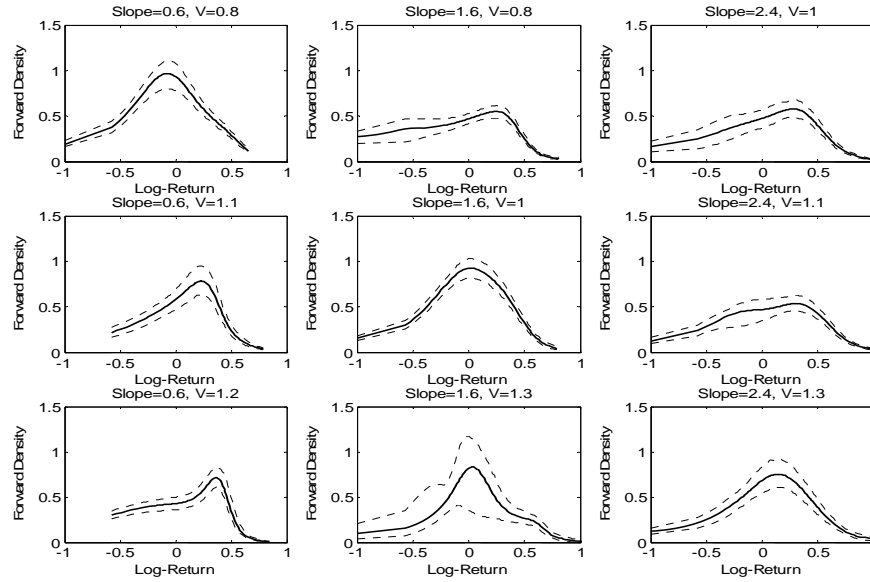


Figure 6. Nonparametric estimates of the LIBOR forward densities at different levels of the slope factor. The slope factor is defined as the difference between the 10- and 2-year three-month LIBOR forward rates. The slope ranges from 0% (flat forward curve) to 3% (steep forward curve).

7.A. 1-Year



7.B. 3-Year

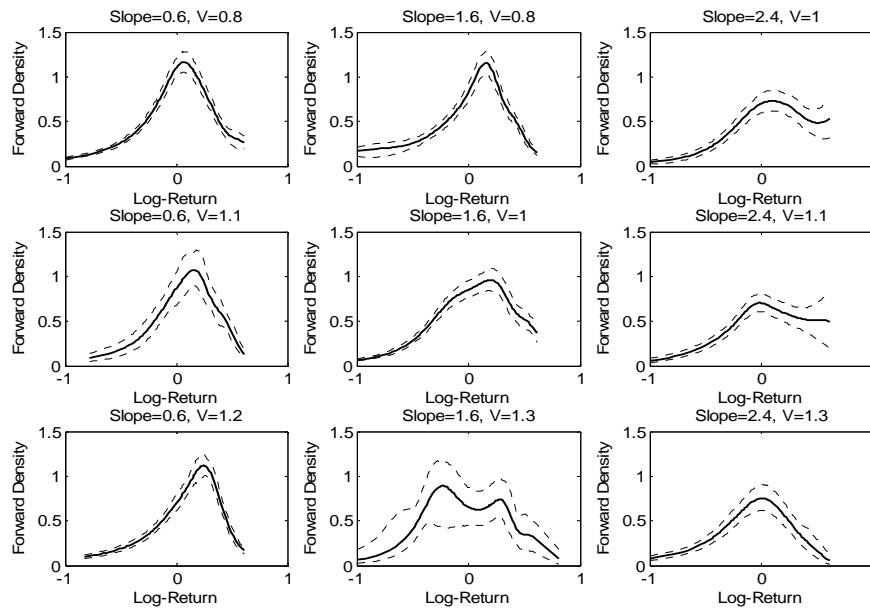
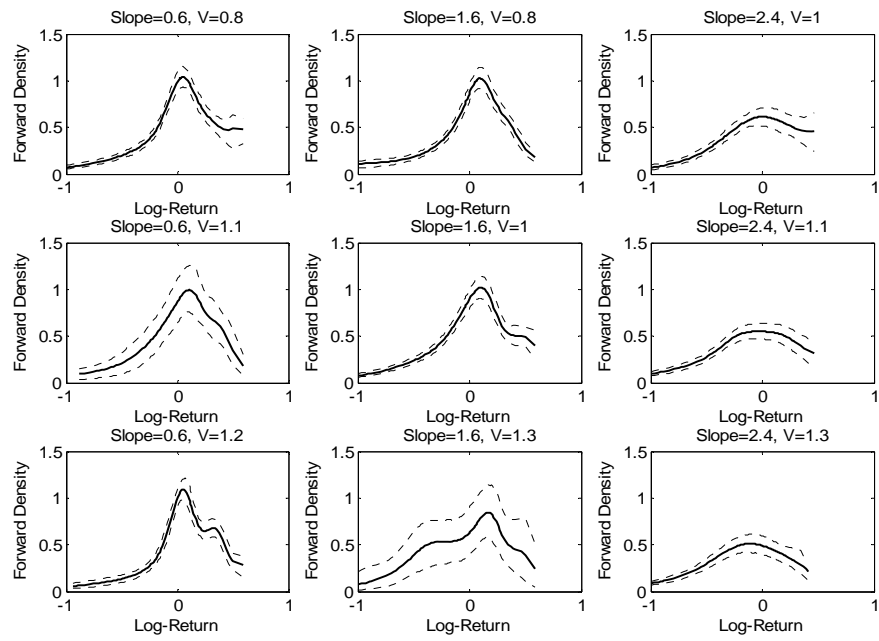


Figure 7. Nonparametric estimates of the LIBOR forward densities at three levels of the slope and volatility factors. The slope factor is defined as the difference between the 10- and 2-year three-month LIBOR forward rates. The volatility factor is defined as the first principal component of EGARCH-filtered spot volatilities and has been normalized to a mean that equals one. The three levels of the slope factor correspond to flat, average, and steep forward curves, while the three levels of the volatility factor corresponds to low, medium, and high levels of volatility.

7.C. 5-Year



7.D. 7-Year

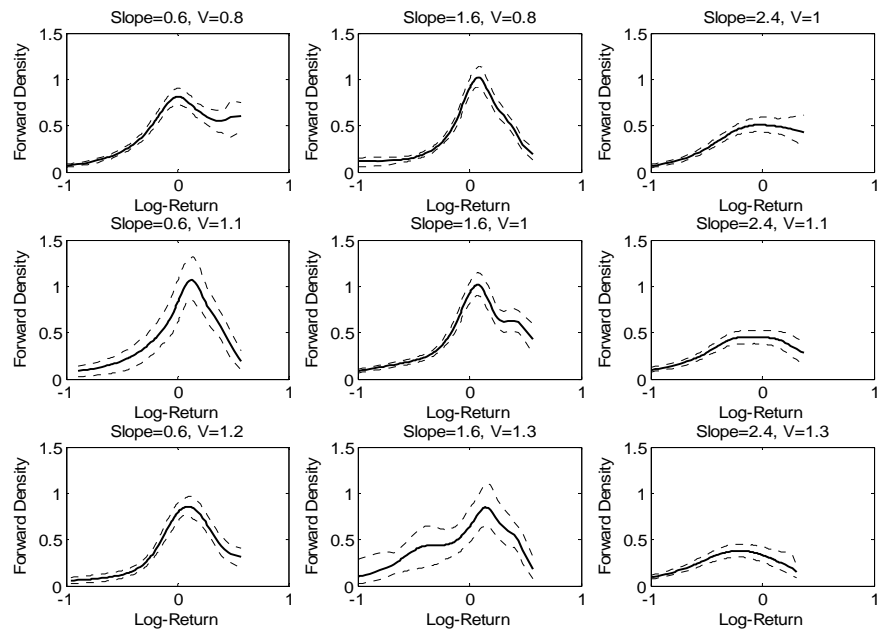


Figure 7. Nonparametric estimates of the LIBOR forward densities at three levels of the slope and volatility factors (continued).


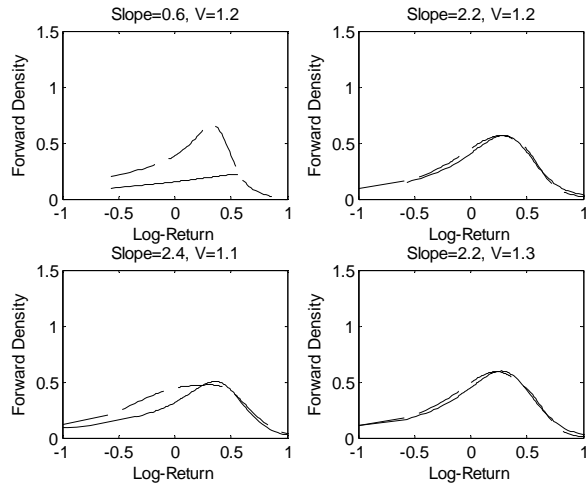
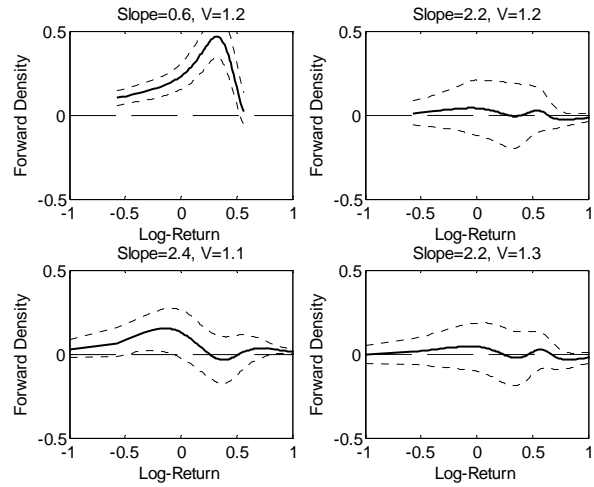


Figure 8. Mortgage Bankers Association of America (MBAA) weekly refinancing index. This figure reports the log of the number of weekly mortgage loan applications from Mortgage Bankers Association of America (MBAA). The shaded areas denote high prepayment periods, which are defined as weeks with changes in loan applications that are in the top 20 percentile of all weekly changes and the subsequent three weeks after each spike.

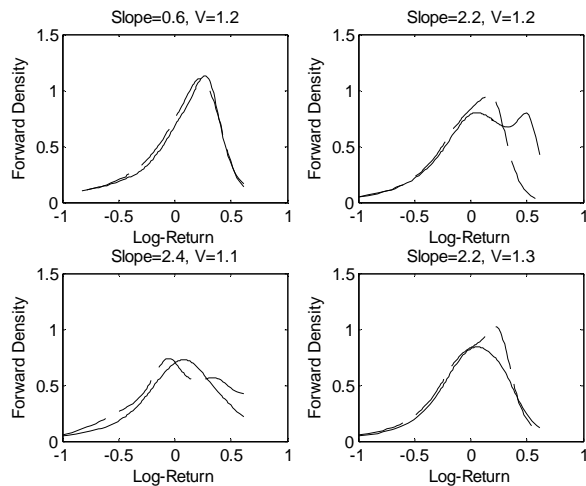
9.A. 1-Year, Low (dash) and High (solid)



9.B. 1-Year, Low (dash) - High (solid)



9.C. 3-Year, Low (Dash) and High (Solid)



9.D. 3-Year, Low (Dash) - High (Solid)




Figure 9. Impacts of mortgage prepayment activities on the LIBOR forward densities. The first column provides nonparametric estimates of LIBOR forward densities during low and high prepayment periods based on observations within the four slope-volatility cells, while the second column provides the differences between the nonparametric forward densities during the low and high prepayment periods with 95% bootstrapped confidence intervals. High prepayment periods are defined as weeks with changes in loan applications that are in the top 20 percentile of all weekly changes and the subsequent three weeks after each spike. Low prepayment periods represent the rest of the sample. The slope factor is defined as the difference between the 10- and 2-year three-month LIBOR forward rates. The volatility factor is defined as the first principal component of EGARCH-filtered spot volatilities and has been normalized to a mean that equals one.

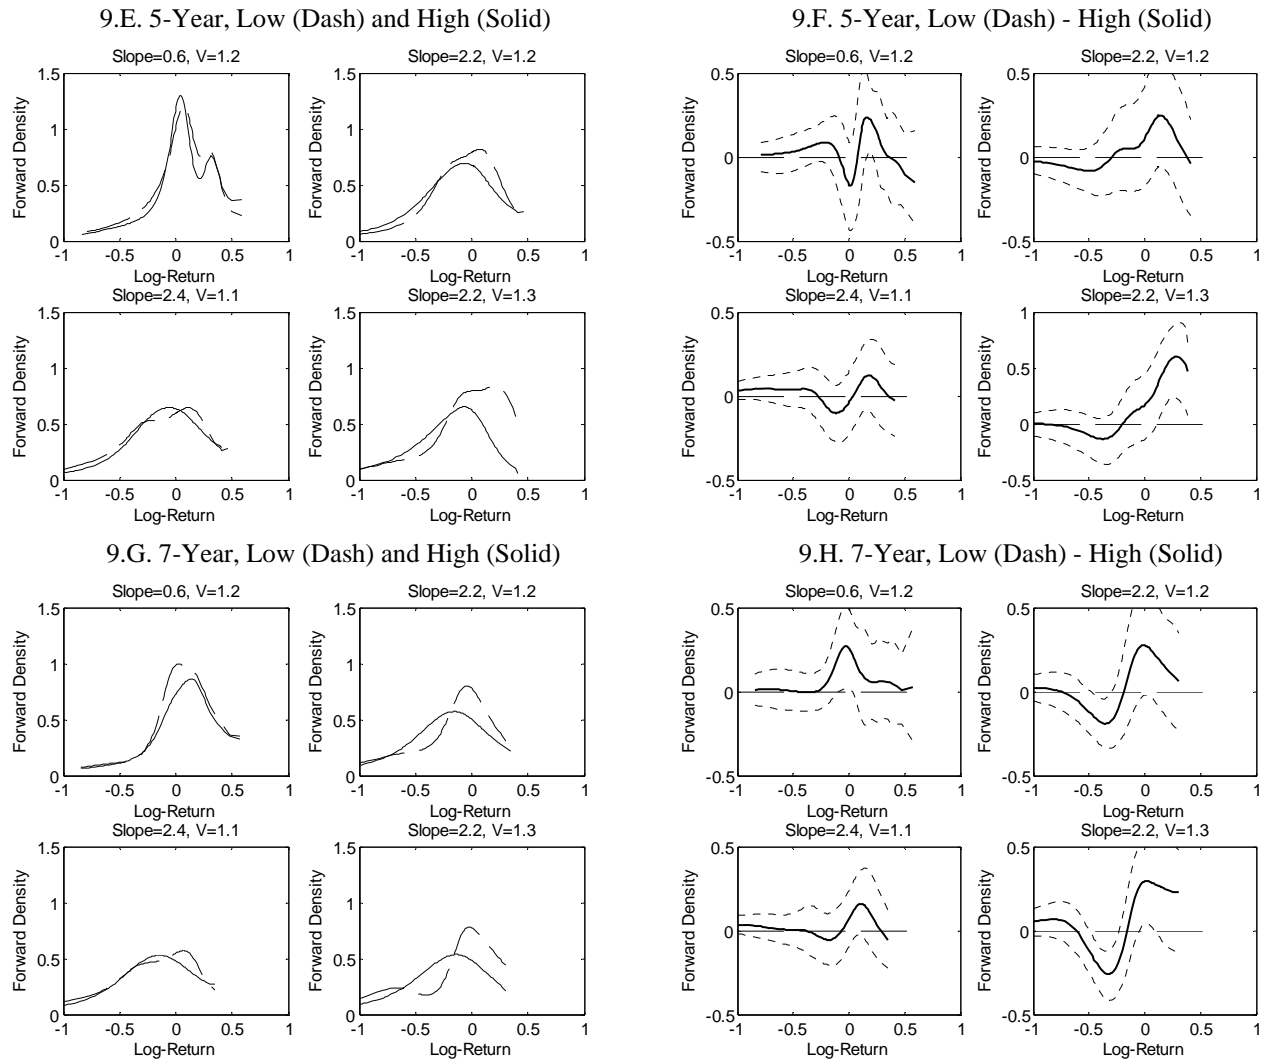
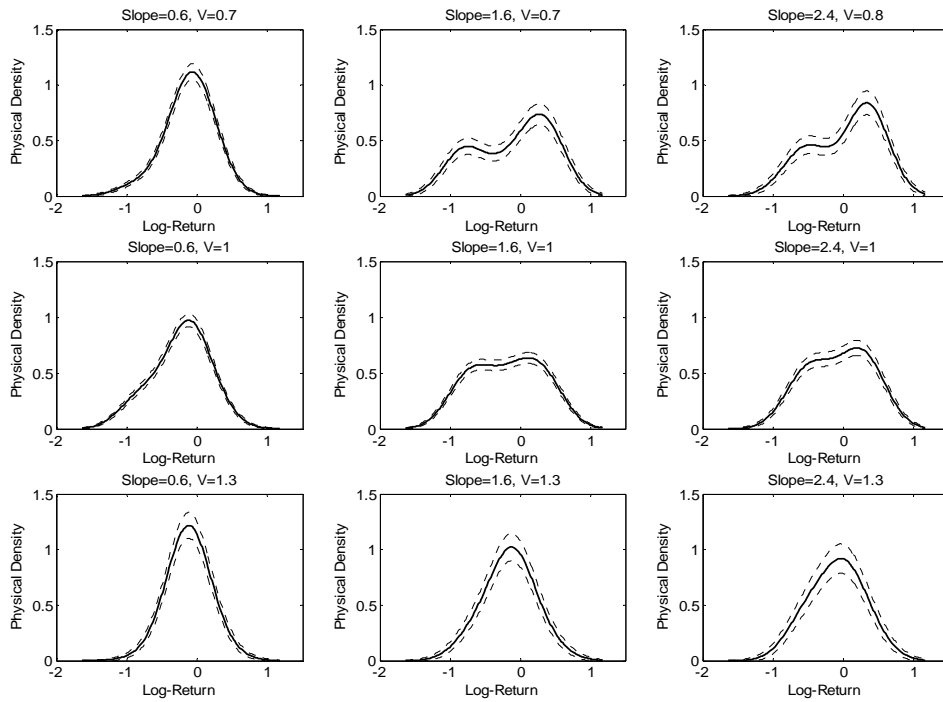


Figure 9. Impacts of mortgage prepayment activities on the LIBOR forward densities (continued).

10.A. 2-Year



10.B. 3-Year

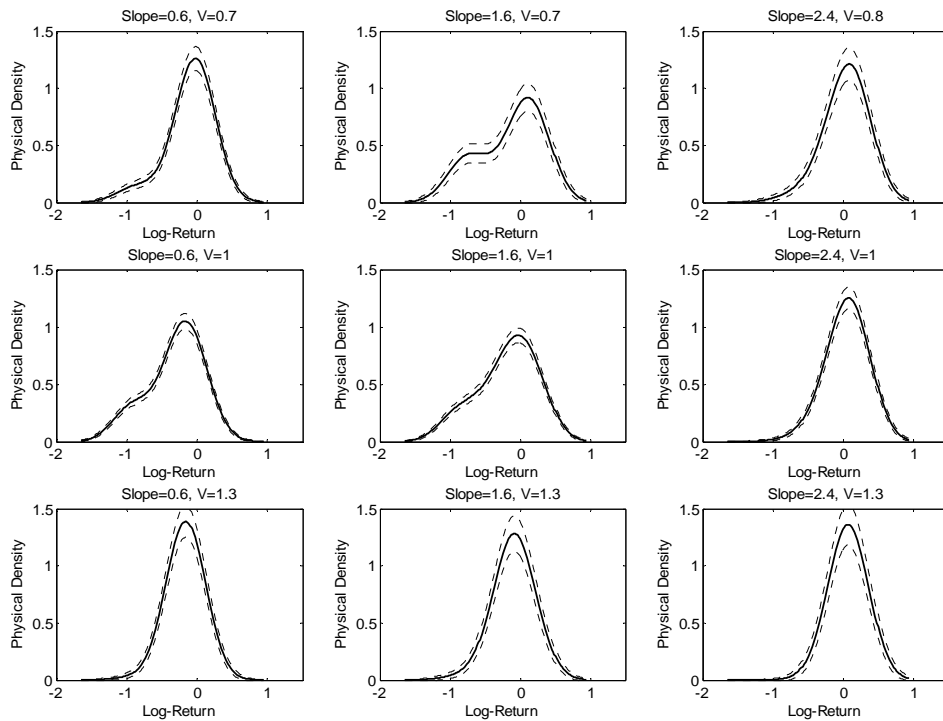
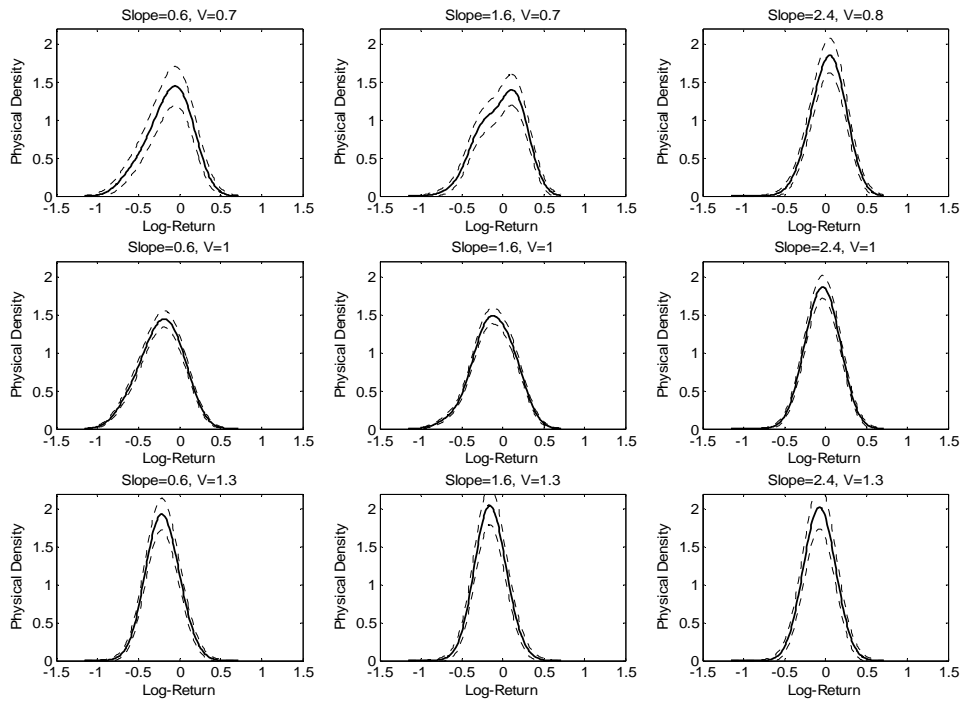


Figure 10. Nonparametric estimates of the physical densities of LIBOR rates at three levels of the slope and volatility factors. The slope factor is defined as the difference between the 10- and 2-year three-month LIBOR forward rates. The volatility factor is defined as the first principal component of EGARCH-filtered spot volatilities and has been normalized to a mean that equals one. The three levels of the slope factor correspond to flat, average, and steep forward curves, while the three levels of the volatility factor corresponds to low, medium, and high levels of volatility. The dotted lines are the 95% confidence interval.

10.C. 4-Year



10.D. 5-Year

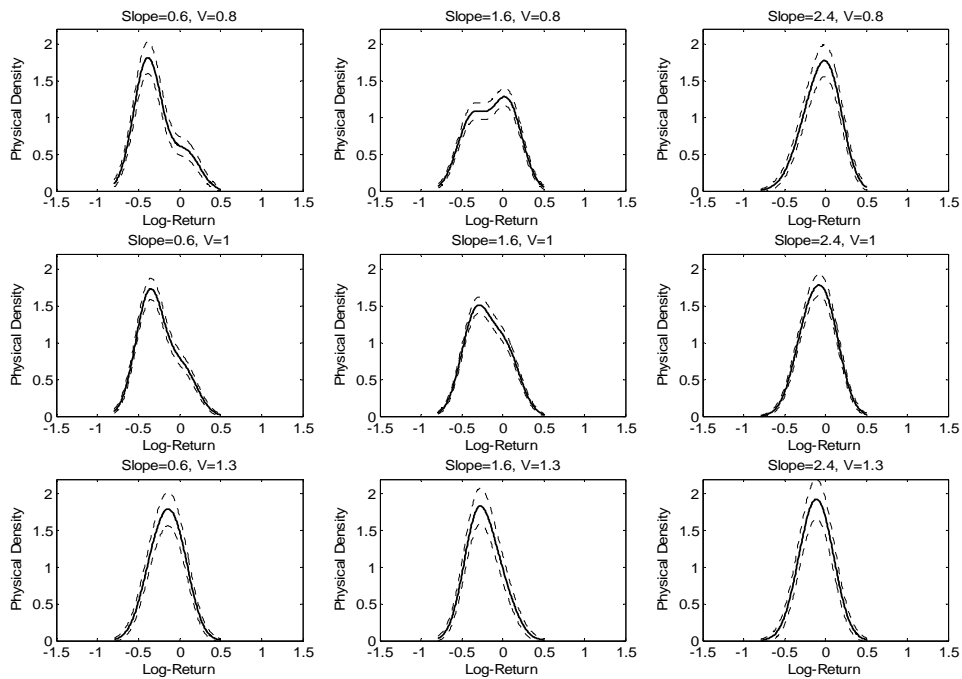
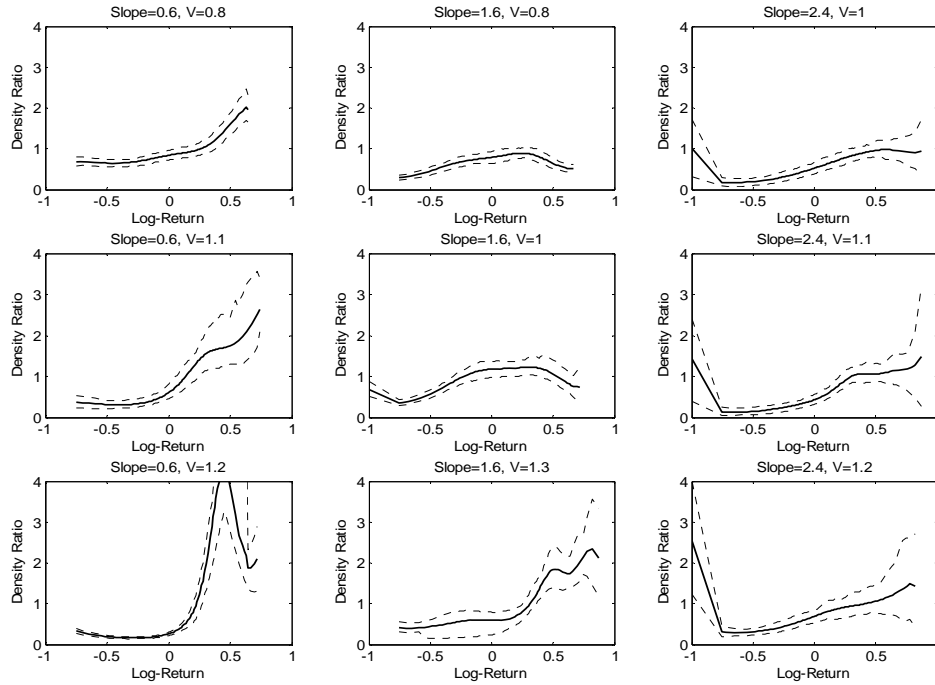


Figure 10. Nonparametric estimates of the physical densities of LIBOR rates at three levels of the slope and volatility factors (continued).

11.A. 2-Year



11.B. 3-Year

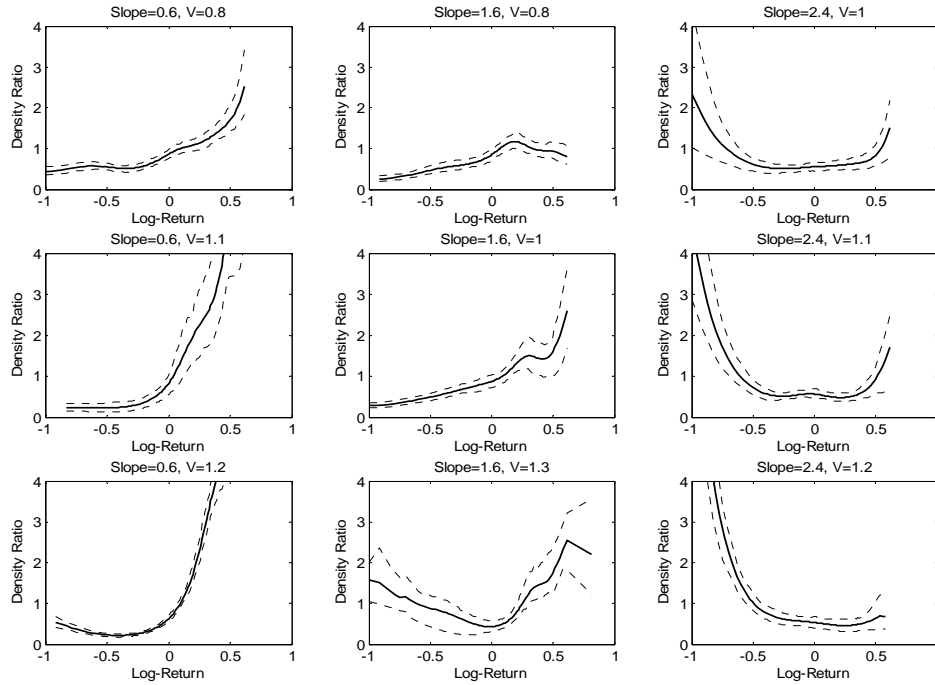



Figure 11. Nonparametric estimates of the state price densities projected onto LIBOR spot rates at three different levels of the slope and volatility factors. The slope factor is defined as the difference between the 10- and 2-year three-month LIBOR forward rates. The volatility factor is defined as the first principal component of EGARCH-filtered spot volatilities and has been normalized to a mean that equals one. The three levels of the slope factor correspond to flat, average, and steep forward curves, while the three levels of the volatility factor corresponds to low, medium, and high levels of volatility. The dotted lines are the 95% confidence interval.

11.C. 4-Year



11.D. 5-Year

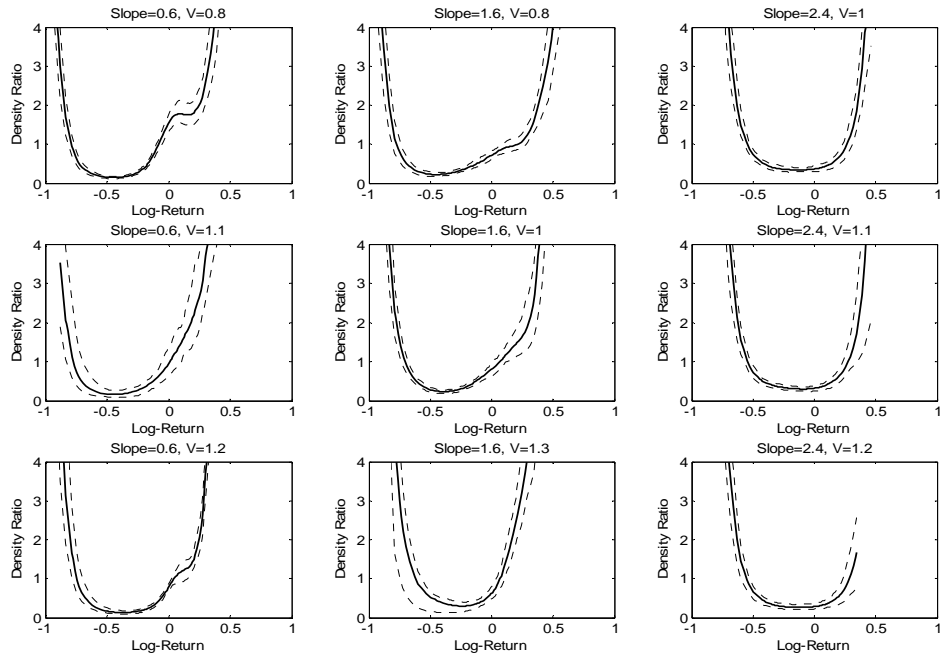


Figure 11. Nonparametric estimates of the state price densities projected onto LIBOR spot rates at three different levels of the slope and volatility factors (continued).