

Maximum penalized quasi-likelihood estimation of the diffusion function

JEFF HAMRICK[†], YIFEI HUANG[‡], CONSTANTINOS KARDARAS^{*‡}
and MURAD S. TAQQU[‡]

[†]Department of Mathematics and Computer Science, Rhodes College,
2000 N. Parkway, Memphis, TN 38112, USA

[‡]Department of Mathematics and Statistics, Boston University,
111 Cummington Street, Boston, MA 02215, USA

(Received 14 August 2010; in final form 29 March 2011)

We develop a *maximum penalized quasi-likelihood estimator* for estimating in a non-parametric way the diffusion function of a diffusion process, as an alternative to more traditional kernel-based estimators. After developing a numerical scheme for computing the maximizer of the penalized maximum quasi-likelihood function, we study the asymptotic properties of our estimator by way of simulation. Under the assumption that overnight London Interbank Offered Rates (LIBOR), the USD/EUR, USD/GBP, JPY/USD, and EUR/USD nominal exchange rates, and the 1-month, 3-month Treasury bill yields, and 30-year Treasury bond yields are generated by diffusion processes, we use our numerical scheme to estimate the diffusion function.

Keywords: Penalized quasi-likelihood; Diffusions; Nonparametric estimation

JEL Classification: C14

1. Introduction

One of the key achievements in the field of financial engineering is the representation of the price of contingent claims as the expectation of their discounted future payoffs under the so-called *risk-neutral probability*, i.e. a probability measure under which discounted (by the exponential of the integrated short rate) traded asset prices are martingales. In complete markets, this representation permits the computation of the unique arbitrage-free price, as well as the hedging strategy needed to remove all the risk associated with issuing (or writing) the contingent claim.

From a practical viewpoint, there is a major issue with choosing the form of the risk-neutral probability measure, as it depends on the model specification. In other words, the stochastic movement of asset prices must be modeled in a way that is consistent with observed data collected from the market. As equivalent changes of probability leave the quadratic variation of a process intact, market data provide a possible way to pin down the volatility

component; then, the drift rate under the risk-neutral measure is simply set equal to the applicable short rate.

The most elementary continuous-time model for asset prices is probably geometric Brownian motion, for which the log-price dynamics are characterized by both constant drift rate and constant volatility. This particular choice of parameters is not consistent with various asset prices such as stock prices, exchange rates, or interest rates. Therefore, more elaborate models have to be utilized. One direct generalization of geometric Brownian motion is local volatility models, where the diffusion coefficient is a function of the underlying asset's level.

In this paper, we consider the class of diffusion models with coefficient functions that are allowed to depend on the asset level. We aim to develop a non-parametric estimation procedure for the diffusion function based on a maximum penalized quasi-likelihood method, following the work of Eggermont and LaRiccia (2001, 2009).

In section 2, we consider a one-dimensional diffusion model for the movement in the price of a financial asset, and briefly review the history of attempts to estimate the diffusion function σ . In section 3, we develop a quasi-likelihood function for a diffusion process and then add a penalization term (also known as a *regularization* term) to

*Corresponding author. Email: kardaras@bu.edu

obtain a *penalized quasi-likelihood function*. After establishing the existence of a maximizer θ_* of the penalized quasi-likelihood function, we use techniques from the calculus of variations to justify a property of θ_* that, in turn, permits us to introduce a numerical scheme for calculating θ_* along discrete, non-uniform design points. In section 4, we study two simulated diffusion processes and note that the mean integrated squared error of our estimator converges to zero at rates that seem to be comparable to the rates of convergence achieved by kernel-based estimators. Then, in section 5, we use software to estimate the diffusion function for overnight London Interbank Offered Rates (LIBOR), 1-month Treasury bill yields and 30-year Treasury bond yields, and the USD/EUR, USD/GBP, and JPY/USD exchange rates.

2. Discussion of the problem

2.1. The market model

We consider a diffusion model for the price movement of a financial asset. In particular, we study a one-dimensional diffusion $(Y_t)_{t \in [0, T]}$ with $Y_0 = y$ and dynamics given by

$$dY_t = b(Y_t)dt + \sigma(Y_t)dW_t, \quad t \in [0, T]. \quad (2.1)$$

Here, W is a standard Brownian motion, $b: \mathbb{R} \mapsto \mathbb{R}$ and $\sigma: \mathbb{R} \mapsto \mathbb{R}_{++} \equiv (0, \infty)$ are Borel-measurable functions, and $T > 0$ is a fixed time horizon. With $C([0, T]; \mathbb{R})$ denoting the canonical path-space of continuous functions equipped with the Borel sigma field, the dynamics described in (2.1) are valid under the probability $\mathbb{P}_y^{(b, \sigma)}$ on $C([0, T]; \mathbb{R})$, which is such that $\mathbb{P}_y^{(b, \sigma)}[Y_0 = y] = 1$.

In order for the problem to be well-posed, we assume that the stochastic differential equation (2.1) has a weak solution that is unique in the sense of probability law. For conditions that guarantee the existence of a weak solution to (2.1), see Karatzas and Shreve (1991, theorem 5.4, p. 332). Of course, uniqueness of a weak solution is a necessary condition for the well-posedness of statistical estimation problems involving diffusions.

Remark 1: The assumption that the functions b and σ have domain \mathbb{R} is made because in the analysis we shall use a conditional Gaussian approximation for the transition densities of the diffusion Y . In practice, the diffusion can live on any sub-interval of \mathbb{R} , such as $(0, \infty)$; this situation will create only theoretical obstacles. Ultimately, it is usually the case that a scale transformation of Y will result in a related diffusion with full support on \mathbb{R} ; for example, a log-transformation of a diffusion supported on $(0, \infty)$ results in a new diffusion supported on \mathbb{R} .

2.2. Non-parametric estimation of the diffusion function

We consider the problem of estimating in a non-parametric way the coefficient σ that appears in equation (2.1) under the assumption that $\mathbb{P}_y^{(b, \sigma)}$ is the historical (or statistical, or real-world) probability law. The estimation

of σ is of crucial importance, since it stays fixed under the equivalent changes of probability measure that are necessary for pricing and hedging contingent claims. More specifically, recall that when we pass from the historical probability law to the risk-neutral probability law, σ remains unaltered while b changes.

In the case in which we *continuously* observe data over the time interval $[0, T]$, i.e. when the whole path $(Y_t)_{t \in [0, T]}$ is observed, perfect estimation of σ is possible, at least in the window of observations $[m_T, M_T]$, where we define

$$m_T := \min_{t \in [0, T]} Y_t \quad \text{and} \quad M_T := \max_{t \in [0, T]} Y_t. \quad (2.2)$$

Indeed, the quadratic variation of Y under $\mathbb{P}^{(b, \sigma)}$ is $\langle Y, Y \rangle = \int_0^t \sigma^2(Y_s)ds$ and hence $\sigma^2(Y_t) = \partial \langle Y, Y \rangle_t / \partial t$ for $t \in [0, T]$. In this case of continuous observations, given perfect estimation of σ , Kutoyants (2004) provides substantial results regarding estimation of the drift b in both parametric and non-parametric settings.

Of course, the case of continuous observations is only a theoretical idealization. In practice, we are usually presented with discrete observations $Y_{t_0}, Y_{t_1}, \dots, Y_{t_n}$, where $0 = t_0 < \dots < t_n = T$. Typically, researchers study two distinct problems for the case in which we only have discrete observations at our disposal. In the case of *low sampling frequency*, the sampling interval remains fixed, and the behavior of estimators is studied as T tends to infinity. (Bandi and Phillips (2003) refer to this as the *long-time asymptotic* approach—see, for example, Gobet *et al.* (2004) for the non-parametric approach in this case.) We shall consider the problem of estimating the diffusion function σ when we have *high sampling frequency*, that is the mesh of the partition, as defined by $\max_{i=1, \dots, n} |t_i - t_{i-1}|$, tends to zero, but the time horizon T remains fixed. See Bandi and Phillips (2003) for more on non-parametric estimation of the drift and the diffusion function in the case of high-frequency data, which is referred to as the *infill asymptotic* approach. The infill asymptotic approach has recently been combined with observation of the integrated diffusion process to estimate the drift and diffusion functions (Comte *et al.* 2007).

The earliest research on the problem of non-parametric estimation of the diffusion function $\sigma: \mathbb{R} \mapsto (0, \infty)$ seems to have been undertaken by Florens-Zmirou (1993). Jacod (2000) utilized the kernel-smoothing estimation of σ in a manner motivated by kernel estimation of the density functions, and a rate of convergence of the estimator of order $n^{-m/(2m+1)}$ was obtained for the case in which σ is m times continuously differentiable. A similar approach was undertaken by Banon (1978), Jiang and Knight (1997) and Soulier (1998). Hoffmann (1999a) developed an L_p -loss estimator for σ that, under suitable conditions, has a minimax rate of convergence that is also of order $n^{-m/(2m+1)}$. The same author constructs estimators for the coefficients of a diffusion process based on adaptive wavelet thresholding (with respect to an unknown degree of smoothness in the coefficient functions) (Hoffmann 1999b). Wavelet methods were also used by Genon-Catalot *et al.* (1992) to estimate the drift and diffusion functions and, again, the rate of

convergence was found to match the classical rate of convergence for kernel-based estimators. A kernel smoothing approach is used by Fan *et al.* (2007) to define spatial and temporal estimators for σ . This technique was later adapted by Hamrick and Taqqu (2009) to test diffusions for stationarity.

In this paper, we propose instead to use a method based on the maximum penalized quasi-likelihood. We provide estimators of σ that empirically behave at least as well as the other estimators developed in the literature up to the present time. We note additionally the connection between the method that we will propose and kernel estimators via reproducing kernel Hilbert spaces, as is presented, for example, by Eggermont and LaRiccia (2009, pp. 20–26).

3. Estimation using maximum penalized quasi-likelihood

3.1. Zero drift

The whole analysis below will be carried out assuming that $b \equiv 0$. Analysis for other values for the drift coefficient could be similarly carried out, albeit in a more complicated way. However, this would not eventually serve any purpose in the case of high-frequency data. In fact, any analysis regarding consistency and rates of convergence that applies to the case of zero drift applies also immediately to the case of non-zero drift. The reason is that, in view of Girsanov’s theorem, under very mild integrability assumptions on b/σ , the probabilities $\mathbb{P}^{(b,\sigma)}$ and $\mathbb{P}^{(0,\sigma)}$ are equivalent. In fact, in section 4 we shall present a simulation study that explores the efficiency of our proposed diffusion coefficient estimator in the case of a diffusion with non-zero drift.

3.2. Quasi-likelihood

Under the assumption that $b \equiv 0$ and that our observations are of high frequency, for each $i = 1, \dots, n$, the conditional law of Y_{t_i} given $(Y_{t_0}, Y_{t_1}, \dots, Y_{t_{i-1}})$ is approximately† normal with mean $Y_{t_{i-1}}$ and standard deviation $\sigma(Y_{t_{i-1}})\sqrt{t_i - t_{i-1}}$. Use of this approximation permits us to construct the *weighted quasi-log-likelihood* of the sample, which is defined by

$$\begin{aligned} \text{qll}(\sigma \mid Y_{t_0}, \dots, Y_{t_n}) \\ = \frac{1}{n} \sum_{i=1}^n \left\{ -\log(\sigma(Y_{t_{i-1}})) - \frac{1}{2} \left(\frac{Y_{t_i} - Y_{t_{i-1}}}{\sigma(Y_{t_{i-1}})\sqrt{t_i - t_{i-1}}} \right)^2 \right\} \end{aligned} \tag{3.1}$$

(the symbol ‘qll’ for the above function stands for ‘quasi-log-likelihood’).

Of course, any estimator $\hat{\sigma}$ that satisfies $\hat{\sigma}(Y_{t_{i-1}}) = |R_{t_{i-1}}|$ at the values in the observation set $\mathfrak{D} := \{Y_{t_0}, Y_{t_1}, \dots, Y_{t_{n-1}}\}$, where

$$R_{t_{i-1}} := \frac{Y_{t_i} - Y_{t_{i-1}}}{\sqrt{t_i - t_{i-1}}}, \quad i = 1, \dots, n, \tag{3.2}$$

would correspond to a maximum likelihood estimator with perfect fit to the data. (In particular, such an estimator is not unique.) As is typical in infinite-dimensional estimation problems (another example of which is non-parametric density estimation), naive interpolations of the points $\{(Y_{t_{i-1}}, |R_{t_{i-1}}|)\}_{i=1, \dots, n}$ result in estimators that oscillate wildly and are nonsensical. In order to effectively resolve this issue, one needs to impose some condition on the estimators. Non-parametric estimation procedures frequently call for the function being estimated to possess some degree of differentiability, which restores the well-posedness of the optimization problem. In this paper, we undertake such an approach, penalizing lack of smoothness in estimates of σ via a *maximum penalized quasi-likelihood* method.

3.3. Maximum penalized quasi-likelihood estimator

Since $\sigma > 0$, the transformation

$$\theta := -\log(\sigma)$$

is well-defined. The weighted quasi-likelihood is now

$$\text{qll}(\theta \mid Y_{t_0}, \dots, Y_{t_n}) = \frac{1}{n} \sum_{i=1}^n \left\{ \theta(Y_{t_{i-1}}) - \frac{1}{2} R_{t_{i-1}}^2 e^{2\theta(Y_{t_{i-1}})} \right\}. \tag{3.3}$$

For the numerical schemes that we will be developing, it is useful to rewrite the above expression in terms of the order statistics of the points in \mathfrak{D} . For $j = 1, \dots, n$, let p_j denote the name of the rank j point in \mathfrak{D} , with the smallest number having rank one. Set‡

$$y_j := Y_{p_j} \quad \text{and} \quad r_j := R_{p_j}. \tag{3.4}$$

For future reference, we set $y_0 := -\infty$ and $y_{n+1} := +\infty$; note that these two points do *not* belong in our observations set \mathfrak{D} . With this new notation, we rewrite the weighted quasi-likelihood as

$$\text{qll}(\theta) = \frac{1}{n} \sum_{j=1}^n \left\{ \theta(y_j) - \frac{1}{2} r_j^2 e^{2\theta(y_j)} \right\},$$

where we drop the dependence of qll on the sample. As mentioned before, we incorporate a term that will penalize estimates for lack of smoothness to produce the penalized quasi-log-likelihood

$$\text{pqll}(\theta; m, \lambda) = \text{qll}(\theta) - \frac{\lambda}{2} \int_{\mathbb{R}} |\theta^{(m)}(z)|^2 dz, \tag{3.5}$$

†Obviously, the quality of the approximation will improve as the mesh tends to zero.

‡We shall be assuming that there are no ties in the ranking. For an example of how we handle ties in the rankings, see section 5. Note however that, in theory, ties occur with probability zero.

where $\lambda > 0$ is a penalization factor, $m \in \mathbb{N}$, and $\theta^{(m)}$ is the derivative of order m of θ . We shall call θ_* a maximum penalized quasi-likelihood estimator if it is a solution to the problem

$$\theta_* := \arg \max_{\theta \in \mathbb{W}^{(m,2)}} \text{pqll}(\theta; m, \lambda),$$

where the maximization is over the space

$$\mathbb{W}^{(m,2)} := \left\{ f: \mathbb{R} \rightarrow \mathbb{R} \mid f^{(m-1)} \text{ exists, is absolutely continuous, and } \int_{\mathbb{R}} |f^{(m)}(z)|^2 dz < \infty \right\}.$$

(Note that $\mathbb{W}^{(m,2)}$ as defined above is not the usual Sobolev space, which would be a particular subset of $L^2(\mathbb{R})$ —functions in $\mathbb{W}^{(m,2)}$ might fail to be square-integrable.)

It can be shown that the maximizer θ_* of functionals like $\mathbb{W}^{(m,2)} \ni \theta \mapsto \text{pqll}(\theta; m, \lambda)$ is a natural spline of order $2m - 1$ with knots $\mathfrak{D} = \{y_1, \dots, y_n\}$ (Eubank 1988).[†] It follows that θ_* is at least $2m - 1$ times differentiable, as well as piecewise polynomial of order $2m - 1$ on all intervals $\{(y_k, y_{k+1})\}_{k=0, \dots, n}$. (Recall that we are using the conventions $y_0 = -\infty$ and $y_{n+1} = +\infty$.) In particular, $\theta_*^{(2m-1)}$ is piecewise constant on the intervals $\{(y_k, y_{k+1})\}_{k=0, \dots, n}$. The fact that θ_* is a natural spline means that all derivatives of order $m, \dots, 2m - 1$ vanish outside $[y_1, y_n]$. Since θ_* is at least $2m - 2$ times continuously differentiable, this implies that $\theta_*^{(i)}(y_1) = 0 = \theta_*^{(i)}(y_n)$ for $i = m, \dots, 2m - 2$; furthermore, $\theta_*^{(2m-1)}(y_1-) = 0 = \theta_*^{(2m-1)}(y_n+)$, where $\theta_*^{(2m-1)}(y-)$ and $\theta_*^{(2m-1)}(y+)$ will denote the left-hand and right-hand limit, respectively, of $\theta_*^{(2m-1)}$ at $y \in \mathbb{R}$.

Remark 2: By definition, the estimators θ_* have a number of derivatives equal to zero at the points y_1 and y_n . Of course, the ‘true’ $\theta = -\log(\sigma)$ is not expected to display such behavior near the endpoints. The method of penalization over-smooths the estimator close to the extreme observations y_1 and y_n —this situation is unavoidable, since there are very few observations near these points. It would be an interesting topic for future research to investigate in a rigorous way the severity of this effect close to the extreme observations.

In all that follows, we shall be freely using the fact that θ_* has the above special structure. As $\theta_*^{(2m-1)}$ is not uniquely defined at the points $\mathfrak{D} = \{y_1, \dots, y_n\}$, we agree to pick θ_* such that $\theta_*^{(2m-1)}$ is right-continuous, i.e. we enforce $\theta_*^{(2m-1)}(y_i) = \theta_*^{(2m-1)}(y_i+)$ to hold for all $i = 1, \dots, n$.

Lemma 3.1: Let θ_* be the maximizer of equation (3.5). Then, for any $\delta \in \mathbb{W}^{(m,2)}$,

$$\frac{1}{n} \sum_{j=1}^n \{\delta(y_j)(1 - r_j^2 e^{2\theta_*(y_j)})\} = (-1)^{m-1} \lambda \int_{\mathbb{R}} \theta_*^{(2m-1)}(z) \delta'(z) dz. \tag{3.6}$$

Proof: To see this fact, assume that θ_* is the maximizer of equation (3.5). For any $\delta \in \mathbb{W}^{(m,2)}$,

$$0 = \frac{\partial}{\partial \epsilon} \text{pqll}(\theta_* + \epsilon \delta; m, \lambda) \Big|_{\epsilon=0} = \frac{\partial}{\partial \epsilon} \left(\text{qll}(\theta_* + \epsilon \delta) - \frac{\lambda}{2} \int_{\mathbb{R}} |\theta_*^{(m)}(z) + \epsilon \delta^{(m)}(z)|^2 dz \right) \Big|_{\epsilon=0}.$$

By (3.3), we obtain

$$\frac{\partial}{\partial \epsilon} \text{qll}(\theta_* + \epsilon \delta) \Big|_{\epsilon=0} = \frac{1}{n} \sum_{j=1}^n \{\delta(y_j) - \delta(y_j) r_j^2 e^{2\theta_*(y_j)}\}.$$

Furthermore, it is straightforward that

$$\frac{\partial}{\partial \epsilon} \left(\frac{\lambda}{2} \int_{\mathbb{R}} |\theta_*^{(m)}(z) + \epsilon \delta^{(m)}(z)|^2 dz \right) \Big|_{\epsilon=0} = \lambda \int_{\mathbb{R}} \theta_*^{(m)}(z) \delta^{(m)}(z) dz.$$

Therefore, the first-order conditions for optimality become

$$\frac{1}{n} \sum_{j=1}^n \{\delta(y_j) - \delta(y_j) r_j^2 e^{2\theta_*(y_j)}\} = \lambda \int_{\mathbb{R}} \theta_*^{(m)}(z) \delta^{(m)}(z) dz.$$

Since $\theta_*^{(m)}$ vanishes outside $[y_1, y_n]$, integration by parts implies that

$$\lambda \int_{\mathbb{R}} \theta_*^{(m)}(z) \delta^{(m)}(z) dz = -\lambda \int_{\mathbb{R}} \theta_*^{(m+1)}(z) \delta^{(m-1)}(z) dz.$$

Now, repeatedly integrating by parts and using the fact that $\theta_*^{(i)}$ vanishes outside $[y_1, y_n]$ for $i = m, \dots, 2m - 1$ yields the result. \square

By considering appropriate functions δ and substituting them into equation (3.6), we obtain the following result, which will be the basis for our algorithm to compute θ_* .

Proposition 3.2: The maximum penalized quasi-likelihood estimator θ_* is such that

$$\theta_*^{(2m-1)}(y) = \frac{(-1)^m}{n\lambda} \left(k - \sum_{j=1}^k r_j^2 e^{2\theta_*(y_j)} \right) \tag{3.7}$$

holds for all $y \in (y_k, y_{k+1})$, where $k = 1, \dots, n$. (Recall that $y_{n+1} = +\infty$, by convention.)

Proof: Fix $y \in (y_k, y_{k+1})$ for some $k \in \{1, \dots, n\}$ and consider the function $\delta_y = \mathbb{1}_{(-\infty, y)}$. In that case, $-\delta'_y$ is a Dirac mass at y . Of course, equation (3.6) cannot be applied directly to a non-differentiable function like δ_y . To circumvent this issue, we use an approximating procedure. Consider the sequence of functions $\mathbb{R} \ni z \mapsto \delta_{y,N}(z) = \Phi(N(y - z))$ for $N \in \mathbb{N}$, where Φ is the cumulative distribution function of the standard normal distribution. Note that $\delta_{y,N} \in \mathbb{W}^{(m,2)}$ for $m \geq 1$ and that $(\delta_{y,N})_{N \in \mathbb{N}}$ converges pointwise to δ_y at all points except for y . By equation (3.6) and typical arguments utilizing the distributional convergence

[†]The solution θ_* is an example of an M -type estimator (see, for example, Cox (1983)). For more information about the general problem of fitting splines with restricted sets of values using penalty functions, see Utreras (1981).

of $\mathbb{R} \ni z \mapsto -(\partial/\partial z)\delta_{y,N}(z)$ to the Dirac mass at y , the equation

$$\begin{aligned} & \sum_{j=1}^k \{\delta_{y,N}(y_j)(1 - r_j^2 e^{2\theta_*(y_j)})\} \\ &= (-1)^{m-1} n\lambda \int_{\mathbb{R}} \theta_*^{(2m-1)}(z) \left\{ \frac{\partial}{\partial z} \delta_{y,N}(z) \right\} dz \end{aligned}$$

is valid for all $N \in \mathbb{N}$. It also leads to equation (3.7) upon sending N to infinity as soon as one notices that $\lim_{N \rightarrow \infty} \sum_{j=1}^k \delta_{y,N}(y_j) = k$ for $y \in (y_k, y_{k+1})$ and that $\theta_*^{(2m-1)}$ is constant on (y_k, y_{k+1}) . Recalling that we are considering the right-continuous version of $\theta_*^{(2m-1)}$, we conclude that equation (3.7) is true for all $y \in [y_k, y_{k+1})$. \square

3.4. Computing the estimator

Assume that values for the penalization parameters λ and m are pre-specified. We will discuss possible ways of choosing λ and m in subsection 3.5. We know $Y_{t_0}, Y_{t_1}, \dots, Y_{t_n}$ and hence the pairs $(y_1, r_1), (y_2, r_2), \dots, (y_n, r_n)$ given by equation (3.4). We want to determine $\theta_*^{(i)}(y_j)$ for $i=0, \dots, 2m-1$ and $j=1, \dots, n$. These values of θ_* and its higher-order derivatives along the observation set $\mathfrak{D} = \{y_1, \dots, y_n\}$ can be used to fit a spline, which then generates an estimate of θ (and, therefore, of $\sigma = \exp(-\theta)$ as well) on the interval $[y_1, y_n]$. The basic idea of the algorithm is to use equation (3.7) to obtain $\theta^{(2m-1)}$ at a particular y_j and then to ‘work downwards’ to the lower-order derivatives.

We introduce further notation and definitions to support the description of the iterative procedure induced by equation (3.7). For $a = (a^{(0)}, \dots, a^{(m-1)})$, define $\Theta(\cdot; a)$ to be the spline of order $2m-1$ with knots \mathfrak{D} , with the properties that $\Theta^{(i)}(y_1; a) = a^{(i)}$ for $i=0, \dots, m-1$, $\Theta^{(i)}$ vanishes on $(-\infty, y_1)$ for $i=m, \dots, 2m-1$, and

$$\Theta^{(2m-1)}(y; a) = \frac{(-1)^m}{n\lambda} \left(k - \sum_{j=1}^k r_j^2 e^{2\Theta(y_j; a)} \right) \quad (3.8)$$

for all $y \in [y_k, y_{k+1})$, $k=1, \dots, n$, where again we are considering the right-continuous version of $\Theta^{(2m-1)}(\cdot; a)$ and set $y_{n+1} = \infty$ by convention. Note that the above properties characterize the spline $\Theta(\cdot; a)$ entirely in terms of a recursive procedure; we shall discuss how to compute all the values of $\Theta(\cdot; a)$ given $a = (a^{(0)}, \dots, a^{(m-1)}) \in \mathbb{R}^m$ in the sequel. For the time being, and as a warm-up for the algorithm that will be presented below, note that equation (3.8) implies that

$$\begin{aligned} \Theta^{(2m-1)}(y_k; a) &= \Theta^{(2m-1)}(y_{k-1}; a) + \frac{(-1)^m}{n\lambda} (1 - r_k^2 e^{2\Theta(y_k; a)}), \\ & \quad k = 1, \dots, n, \end{aligned} \quad (3.9)$$

upon agreeing that $\Theta^{(2m-1)}(y_0; a) = 0$ as a matter of convention.

In view of proposition 3.2, and since the estimator θ_* must be a natural spline, it will hold that $\theta_*(\cdot) = \Theta(\cdot; a_*)$ where $a_* \in \mathbb{R}^m$ is such that $\Theta^{(i)}(y_n; a_*) = 0$ for

$i=m, \dots, 2m-1$. Therefore, our goal is to obtain the root of the nonlinear equation $F(a) = 0$, where the mapping $F: \mathbb{R}^m \rightarrow \mathbb{R}^m$ is defined by

$$F(a) := (\Theta^{(i)}(y_n; a))_{i=m, \dots, 2m-1}.$$

In order to do this, an efficient way of computing $F(a)$ for a given $a = (a^{(0)}, \dots, a^{(m-1)}) \in \mathbb{R}^m$ is required. The following pseudocode illustrates the computation of $F(a)$ in the case $m=2$, which is the value we mostly use for the numerical computations. It is straightforward to adapt the code for any value of $m \in \mathbb{N}$.

Input $a = (a^{(0)}, a^{(1)}) = (\Theta^{(0)}(y_1; a), \Theta^{(1)}(y_1; a))$.

Set $\Theta^{(2)}(y_1; a) = 0, \Theta^{(3)}(y_0; a) = 0$.

For $k = 1, \dots, n-1$, **set**:

$$\begin{aligned} \Theta^{(3)}(y_k; a) &= \Theta^{(3)}(y_{k-1}; a) + (n\lambda)^{-1} (1 - r_k^2 e^{2\Theta^{(0)}(y_k; a)}), \\ \Theta^{(2)}(y_{k+1}; a) &= \Theta^{(2)}(y_k; a) + \Theta^{(3)}(y_k; a)(y_{k+1} - y_k), \\ \Theta^{(1)}(y_{k+1}; a) &= \Theta^{(1)}(y_k; a) + \Theta^{(2)}(y_k; a)(y_{k+1} - y_k) \\ & \quad + (1/2)\Theta^{(3)}(y_k; a)(y_{k+1} - y_k)^2, \\ \Theta^{(0)}(y_{k+1}; a) &= \Theta^{(0)}(y_k; a) + \Theta^{(1)}(y_k; a)(y_{k+1} - y_k) \\ & \quad + (1/2)\Theta^{(2)}(y_k; a)(y_{k+1} - y_k)^2 \\ & \quad + (1/6)\Theta^{(3)}(y_k; a)(y_{k+1} - y_k)^3. \end{aligned}$$

Next k

Set $\Theta^{(3)}(y_n; a) = \Theta^{(3)}(y_{n-1}; a) + (n\lambda)^{-1} (1 - r_n^2 e^{2\Theta^{(0)}(y_n; a)})$.

Return $F(a) = (\Theta^{(2)}(y_n; a), \Theta^{(3)}(y_n; a))$.

Note that the philosophy of our algorithm is a variant of the so-called ‘shooting method’ (Asher and Petzold 1988, p. 177). When close to the root a_* , Newton’s method can be utilized to get a better ‘aim’, which will result in fast convergence of this iterative scheme. (Newton’s method has to be used with care—if far away from the root, the method is not likely to work and the numerical scheme will fail to converge. For further comments on this problem, see subsection 3.6.) Newton’s method requires computation of the partial derivatives $(\partial/\partial a^{(i)})F$, which can be easily computed along with $F(a)$. Indeed, note that $\varphi_i(\cdot; a) := (\partial/\partial a^{(i)})\Theta^{(j)}(\cdot; a)$ for $i=0, \dots, m-1$ and $j=0, \dots, 2m-1$ satisfy

$$\varphi_i^{(j)}(y_1; a) = \frac{\partial \Theta^{(j)}(y_1)}{\partial a^{(i)}} = \frac{\partial a^{(j)}}{\partial a^{(i)}} = \begin{cases} 1, & \text{if } i = j, \\ 0, & \text{if } i \neq j. \end{cases}$$

Furthermore, upon differentiating (3.8), we obtain

$$\begin{aligned} \varphi_i^{(2m-1)}(y; a) &= \frac{\partial \Theta^{(2m-1)}(y)}{\partial a^{(i)}} \\ &= \frac{(-1)^{m+1}}{n\lambda} \sum_{j=1}^k r_j^2 e^{2\Theta^{(0)}(y_j; a)} \varphi_i^{(0)}(y_j; a) \end{aligned}$$

for all $y \in [y_k, y_{k+1})$, where $k=1, \dots, n$. Therefore, computing these derivatives can be performed for little extra cost in the same iterative procedure used to compute $\Theta(y_j; a)$ for $j=1, \dots, n$. Indeed, one has to simply differentiate the iterative equations with respect to a and obtain iterative equations for the derivatives. As soon as the information about the m partial first-order derivatives of




Figure 1. We generate a sample path with $\Delta t = 1/2^{17}$, no drift, and $\sigma \equiv 3$ and we also create two reductions of this sample path with $\Delta t = 1/2^{16}$ and $\Delta t = 1/2^{13}$. Along with the true σ , the resulting σ_* are shown when $\lambda = 20(\Delta t)^{2/3}$ and $m = 1$.

F at a particular $a \in \mathbb{R}^m$ is obtained, we can implement Newton's method for finding the root of the equation $F(a) = 0$. Newton's method proceeds by successive approximation. Consider an initial guess $a_1 = (a_1^{(i)})_{i=0, \dots, m-1}$ and inputs $\beta > 0$ and $\epsilon > 0$, where β controls the step size in Newton's method and ϵ determines when Newton's method is terminated.† Then, for $k = 1, \dots$, while $|F(a_k)| \geq \epsilon$ (where $|\cdot|$ denotes the usual Euclidean norm), one sets $a_{k+1} = a_k - \beta \Phi^{-1}(a_k) F(a_k)$. In the previous formula, Φ is the $m \times m$ matrix with entry $\varphi_i^{(j)}(y_n; a)$ in the i th row and the j th column for $i = 0, \dots, m-1$ and $j = 0, \dots, m-1$. (Note that we are numbering rows and columns from 0 to $m-1$, in order to be consistent with our notation.)

Remark 3: There is accompanying software that implements the above algorithm, together with instructions regarding its use, available upon request from the authors.

3.5. Choosing the penalization parameters

In the literature on non-parametric regression using the penalized likelihood method (also called the regularization method), the most typical choice for the order of differentiation to penalize is $m = 2$ (see, for example, Cox and O'Sullivan (1990) or Green and Silverman (1994)). The choice $m = 2$ gives rise to estimators that are natural cubic splines and are visually very attractive. The use of $m = 1$ results in estimated functions that are quite 'wiggly' (see, for example, figure 1). Use of $m \geq 3$ is computationally involved; for this reason, we refrain from such practice.

The choice of the penalization coefficient λ is more subtle. One can use cross-validation techniques (see, for example, Fan and Yao (2005) or Green and Silverman (1994)). However, it is often the case that simple visual inspection of the graphs is sufficient.

Theoretical results from the theory of non-parametric regression using the penalized likelihood method provide a hint to understanding how the penalization coefficient λ should decrease with increasing sample size n when the degree of smoothness (as given by m) of the target function is fixed. More precisely, for fixed m , it is conjectured that the choice $\lambda_n \sim n^{-2m/(2m+1)}$ as $n \rightarrow \infty$ will result in convergence of the estimator to the true value of the order $n^{-m/(2m+1)}$. For theoretical background, see Cox and O'Sullivan (1990). In the next section, we shall provide empirical results on rates of convergence for $m = 1$ and $m = 2$ that seem to support this conjecture.

3.6. Practical remarks regarding convergence of our algorithm

If one uses the proposed penalization $\lambda_n \sim n^{-2m/(2m+1)}$ for a given $m \in \mathbb{N}$ as $n \rightarrow \infty$, the numerical scheme suggested in subsection 3.4 tends to be unstable for large sample sizes. In order to get a feeling for the reason, note that $n\lambda_n \sim n^{1/(2m+1)}$ converges to infinity when $n \rightarrow \infty$, albeit slower than n . Since one has to repeat the iteration given by (3.9) n times to obtain $F(a)$, it is reasonable to suspect (and it actually happens in practice) that for choices of $a \in \mathbb{R}^m$ that are far away from a_* the vector $\Phi^{-1}(a)F(a)$ (recall that Φ is the $m \times m$ matrix of first-order partial derivatives of F) consists of entries with huge magnitude. This situation results in failure of convergence of the algorithm, since the updating step in Newton's method takes one farther away from the sought-after root. In order to overcome this difficulty, it is often desirable to first consider a subset of the sample, effectively thinning the observations. Once the data set has been thinned enough in order for the method to converge, one can use the root derived from the thinned data set as initial seed for a denser subset of the data. Continuing this way, one finally computes the estimator of σ using the whole sample. Note that we use this method in order to obtain the estimators in the following two sections.

4. Simulation study and empirical rates of convergence

In this section, we shall conduct an empirical investigation of the rate at which our estimators converge to the true volatility function.

As a simple benchmark case when $m = 1$, we consider the case of a driftless Brownian motion with $\sigma \equiv 3$. We generate a path of the Brownian motion using *exact* simulation, and produce data points with a step size of $\Delta t = 2^{-17}$ over the unit interval. Continuing, we remove every other realization of the sample path to arrive at a 'reduction' of the original sample path with step sizes of $\Delta t = 2^{-16}$. We then remove every other observation three more times to create yet another reduction with step size $\Delta t = 2^{-13}$. The resulting estimates σ_* of the true volatility function $\sigma \equiv 3$ are shown in figure 1 using $\lambda = 20(\Delta t)^{2/3}$.

†In most of the examples in this paper, we took $\epsilon = 10^{-10}$ and $\beta = 0.1$, the latter to prevent overstepping in Newton's method.

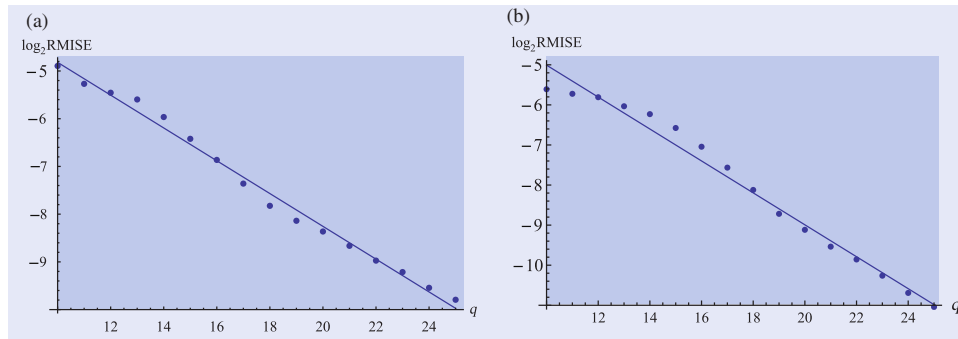


Figure 2. We generate a sample path with $\Delta t = 1/2^{25}$ from the diffusion whose dynamics are given by equation (4.1). Then, we create 15 reductions of this sample path with $\Delta t = 1/2^{10}, \dots, 1/2^{25}$. (a) In the case $m=1$, we plot $\log_2 \text{RMISE}(\sigma_*^{(1,q)}, \sigma)$ against q for $q = 10, \dots, 25$. The least-squares line is $-1.383 - 0.343q$, which is consistent with results from kernel-based estimation schemes, which suggests that the slope should be $-1/3 \approx -0.333$. (b) In the case $m=2$, we again plot $\log_2 \text{RMISE}(\sigma_*^{(2,q)}, \sigma)$ against q for $q = 10, \dots, 25$. The least-squares line is $-1.024 - 0.398q$, which is consistent with our conjecture which suggests that the slope should be $-2/5 = -0.4$.

(This choice is consistent with the discussion in subsection 3.5.) Notice that convergence to the constant volatility function $\sigma \equiv 3$ seems to be quite fast.

To study the empirical rates of convergence more precisely, and to illustrate that our method works well even for diffusions with non-zero drift, we present a more involved example. Consider the diffusion Y with $Y_0 = 1/2$ and dynamics

$$dY_t = -Y_t^2(1 - Y_t)dt + Y_t(1 - Y_t)dW_t, \quad t \in [0, 1]. \quad (4.1)$$

The diffusion Y is $(0, 1)$ -valued; in fact, a straightforward use of Itô's formula shows that

$$Y_t = \frac{\exp(W_t - t/2)}{1 + \exp(W_t - t/2)}, \quad t \in [0, 1]. \quad (4.2)$$

Note that Y does *not* have zero drift and that $\sigma(y) = y(1 - y)$. Furthermore, equation (4.2) gives a way to simulate Y without any discretization error, since one only needs to simulate W , which can be done exactly. We simulate a sample path with a step size 2^{-25} and $T = 1$, and also create reduced versions of this sample path with step sizes Δt of $2^{-10}, 2^{-11}, \dots, 2^{-25}$. For the case $m=1$, we use $\lambda = 30(\Delta t)^{2/3}$, while for the case $m=2$ we use $\lambda = 20(\Delta t)^{4/5}$. For both $m=1$ and $m=2$ and $q = 10, \dots, 25$ (where the step size δ is 2^{-q}), we produce penalized maximum quasi-likelihood estimators $\sigma_*^{(m,q)}$. We then compute the root mean-integrated squared error

$$\text{RMISE}(\sigma_*^{(m,q)}, \sigma) = \sqrt{\frac{T}{n} \sum_{j=1}^n (\sigma_*^{(m,q)}(y_j) - \sigma(y_j))^2},$$

where $n = 2^q T$ is the sample size, as a proxy to the quantity

$$\begin{aligned} & \sqrt{\int_0^T (\sigma_*^{(m,q)}(Y_t) - \sigma(Y_t))^2 dt} \\ &= \sqrt{\int_{m_T}^{M_T} \left(\frac{\sigma_*^{(m,q)}(y)}{\sigma(y)} - 1 \right)^2 L_T^Y(y) dy}, \end{aligned}$$

where $L_T^Y(y)$ is the semimartingale local time of Y at level $y \in \mathbb{R}$ accumulated up to time T . We then plot $\log_2(\text{RMISE}(\sigma_*^{(m,q)}, \sigma))$ against q and execute least-squares linear fits to determine the rate of convergence as we 'fill in' the sample path (see figure 2). The regression line in the case $m=1$ was $-1.383 - 0.343q$. In the case $m=2$, the regression line is $-1.024 - 0.398q$. In both cases, the slope of the line is roughly consistent with convergence results for kernel-based estimation schemes, which suggests that the slope should be $-m/(2m+1)$, giving $-1/3$ for $m=1$ and $-2/5$ for $m=2$. We, therefore, conjecture (as already mentioned) that the rate of convergence is of order $n^{-m/(2m+1)}$ when one chooses m as penalization differentiation order and uses $\lambda_n \sim n^{-2m/(2m+1)}$ as $n \rightarrow \infty$. Note that this conjectured rate is optimal, as demonstrated by Hoffmann (1999b).

5. Application to exchange rates and interest rates

5.1. Exchange rates

The three exchange rates were taken from the Federal Reserve Economic Data (FRED) database, which is maintained by the St. Louis branch of the Federal Reserve Bank. (The FRED database can be accessed at <http://research.stlouisfed.org/fred2>. Registering for a username and password is required as of May 31, 2010.) For purposes of comparability, we study all three exchange rates from January 4, 1999 through May 21, 2010 (the Euro debuted at the beginning of 1999). The FRED database contains noon buying rates in New York City for cable transfers payable in foreign currencies. It includes data only for days on which financial markets are open.

Although exchange rate data at weekly and monthly reporting frequencies is also available through the FRED database, our non-parametric estimation technique performs best when the frequency of the data is relatively high. However, we did not attempt to use intra-day exchange rate data, even though such data is increasingly




Figure 3. Under the assumption that the USD/EUR, USD/GBP, JPY/USD, and EUR/USD exchange rates for the period from January 4, 1999 to May 21, 2010 are governed by diffusion processes, these figures show estimates σ_* of the diffusion function σ .

available for free on the Internet (see, for example, the weblink <http://www.forexrate.co.uk/forexhistoricaldata.php>). Although bid–ask spreads on exchange rates are generally quite low, much of the apparent volatility observed in financial asset prices observed with sufficiently high frequency (for example, one-minute time intervals) is due to buyers buying at the ask price and sellers selling at the bid price. This ‘togglng’ between bid and ask prices during very high frequency trading biases volatility estimation upwards, rendering rules like the ‘square root of time multiplied by the volatility’ ineffectual (Jorion 2009).

We first treat the data by computing the R_{t_i} mentioned in equation (3.2). In particular, we are careful to compute $\sqrt{t_i - t_{i-1}}$ precisely and in calendar time, rather than by assuming that a calendar year has approximately 250 trading days and then dividing $Y_{t_i} - Y_{t_{i-1}}$ by $\sqrt{1/250}$. This adjustment is small but potentially important, particularly when using high-frequency data. For example, the standard deviation of the returns on the USD/EUR exchange rate over weekdays (or holidays) is 0.6457733%. Over weekends, the standard deviation is 0.675802%. With 624 weekend/holiday trading days and 2866 weekday returns, the variance ratio test yields a test statistic of 1.0953 and a rejection of the null hypothesis of equality of variances at the 10% (although not the 5%) level of significance. In general, we will appropriately adjust for the length of the underlying time intervals, since there appears to be a slight but meaningful tendency for

information to be released over the weekend that causes Friday-to-Monday returns to have somewhat higher dispersion than regular weekday returns.

We also pre-process the data in one additional way to make it suitable for the numerical scheme articulated in section 3. In this section, we required that sorted values of the asset prices (denoted y_j) be distinct from one another. In almost all diffusion models, the theoretical probability that two values Y_{t_i} and Y_{t_j} are equal is zero when $i \neq j$. In practical situations in financial markets, in which asset prices are only recorded to a finite number of decimal places, ties are possible and sorting the raw asset price data becomes an ill-defined task. To manage this problem, we tentatively compute the values of $R_{t_i}^2$, take their mean, and add to the raw asset price data a Gaussian random number with mean zero and variance equal to a very small constant times the mean of the $R_{t_i}^2$. We then recompute $|R_{t_i}|$. By slightly perturbing the raw data, all ties are randomly broken and, in practice, the movement in the time series is very negligible. (One could also perform the perturbation only to the tied data—in practice, the two approaches lead to almost identical output.)

In figure 3(a), we plot the estimates of the USD/EUR exchange rate volatility for two different values of λ for the period from January 4, 1999 to May 21, 2010. Consistent with the discussion in section 4, we choose λ to be $0.02n^{-4/5}$ and $0.05n^{-4/5}$, where n is equal to the number of returns in our time series—2866 days, to be precise.

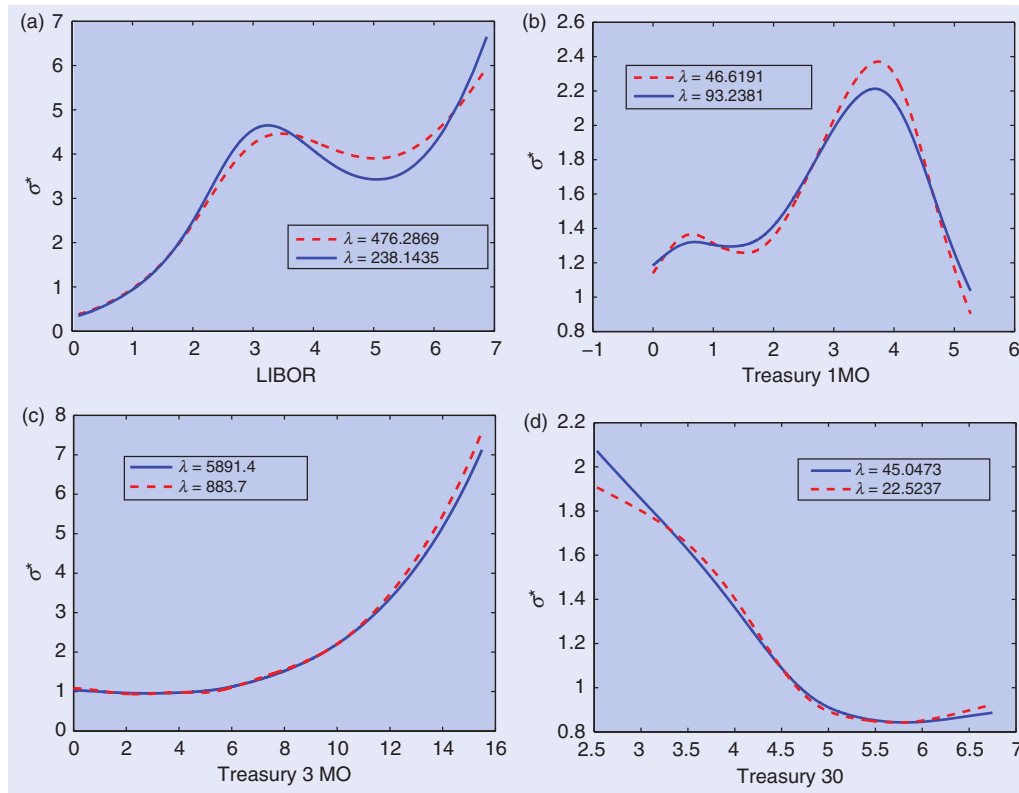


Figure 4. Under the assumption that overnight LIBOR, 1-month U.S. Treasury bill yields, 3-month U.S. Treasury bill yields, and 30-year U.S. Treasury bond yields are governed by diffusion processes, these figures show estimates σ_* of the diffusion function σ . The periods of time for which the estimates σ_* were derived are, respectively, January 2, 1999 through May 26, 2010 for LIBOR, July 31, 2001 through May 31, 2010 for 1-month U.S. Treasury bill yields, January 4, 1982 through May 21, 2010 for 3-month U.S. Treasury bill yields, and January 4, 1999 through May 21, 2010 for 30-year U.S. Treasury bond yields.

The two estimates of the volatility in figure 3(a) seem to indicate that using a linear function for σ in equation (2.1) is not appropriate. Instead, the volatility of the USD/EUR exchange rate seems to increase at a fairly rapid pace around values of 1.30 USD/EUR.

In figure 3(b), we similarly show the estimates σ_* for the USD/GDP exchange rate. The USD/GDP exchange rate volatility varies between 16.5% and 19% over $[y_1, y_n]$ from equation (2.2). Note that while the estimated diffusion function for the USD/EUR exchange rate appears to be monotonically increasing, the estimated USD/GBP exchange rate is not. Generally, diffusion models of financial phenomena assume that the diffusion function is either constant (like the Vasicek model) or that the diffusion function is a constant multiplied by some increasing function (geometric Brownian motion or the Cox–Ingersoll–Ross process). We note that it would be interesting, and potentially the focus of future work, to use our estimated diffusion functions in hypothesis testing—for example, in this setting one could test whether the assumption that the USD/GDP exchange rate volatility is constant can be rejected at a high level of statistical significance.

Finally, in figure 3(c) and (d), we show the estimates σ_* for the JPY/USD and EUR/USD exchange rates, respectively. Note that the estimated diffusion function for the JPY/USD exchange rate is also not monotonic and that it

peaks when the exchange rate is between 95 and 100 Japanese yen per U.S. dollar. The estimated diffusion function associated with the EUR/USD exchange rate looks fairly similar to the function in figure 3(a).

5.2. LIBOR rates, Treasury bill yields, and Treasury bond yields

We also obtain constant-maturity 1-month U.S. Treasury bill yields, constant-maturity 3-month U.S. Treasury bill yields, and constant-maturity 30-year Treasury bond yields from the Federal Reserve Economic Data (FRED) database. (The time series designations in the FRED database for the constant-maturity 1-month U.S. Treasury bill yields, the constant-maturity 3-month U.S. Treasury bill yields, and 30-year U.S. Treasury bond yields are DGS1MO, DGS3MO, and GS30, respectively.) Again, the yields are recorded each day at 12:00 noon Eastern Standard Time. The FRED database contains yields recorded at the end of every week and the end of every month, but we chose to work with yields recorded at the end of each day on which financial markets were open. Additionally, we computed the R_{t-1} in equation (3.2) correctly, fully accounting for long weekends and market holidays. We slightly randomly perturbed the raw data to break ties, as discussed in the previous section. Unlike the Treasury bill and bond yield data, we obtain

overnight London Interbank Offered Rates (LIBOR) from Thompson Reuters Datastream. (For more information about Datastream, see <http://thomsonreuters.com/products>.) The LIBOR rate is the rate that large banks use to borrow and lend from one another on the overnight market.

The periods of time for which the estimates σ_* were derived were January 2, 1999 through May 26, 2010 for LIBOR, July 31, 2001 through May 31, 2010 for 1-month U.S. Treasury bill yields, January 4, 1982 through May 21, 2010 for 3-month U.S. Treasury bill yields, and January 4, 1999 through May 21, 2010 for 30-year U.S. Treasury bond yields.

In figure 4(a), we plot the estimates of London Interbank Offered Rate (LIBOR) volatility for two different values of λ . Our procedure generates an estimate of the volatility function, σ_* , that is clearly neither linear nor a constant multiple of the square-root function (as in the case of the Cox–Ingersoll–Ross process). Indeed, the volatility function is not even monotonic. Although similarly non-monotonic, the estimate σ_* for one-month U.S. Treasury bill yields is low when yields are both relatively low and relatively high. The volatility of constant-maturity one-month U.S. Treasury bill yields seems to be highest when yields are around 3.75%. The estimates of σ for 3-month U.S. Treasury bill yields and 30-year Treasury bond yields are remarkably different. The volatility function associated with 3-month U.S. Treasury bill yields appears to be a monotonic increasing function of the underlying yield. In contrast, the volatility function associated with 30-year U.S. Treasury bond yields is a monotonically decreasing function of the underlying yield. No function appears to adhere closely to any of the standard choices of σ in the financial literature.

Acknowledgements

The third and fourth authors gratefully acknowledge partial support by the National Science Foundation, grant numbers DMS-0908461 and DMS-1007616, respectively.

References

Asher, U. and Petzold, L., *Computer Methods for Ordinary Differential Equations and Differential Algebraic Equations*, 1988 (SIAM: Philadelphia).

- Bandi, F. and Phillips, P., Fully nonparametric estimation of scalar diffusion models. *Econometrica*, 2003, **71**(1), 241–283.
- Banon, G., Nonparametric identification for diffusion processes. *SIAM J. Control Optimiz.*, 1978, **16**(3), 380–395.
- Comte, F., Genon-Catalot, V. and Rozenholc, Y., Penalized nonparametric mean square estimation of the coefficients of diffusion processes. *Bernoulli*, 2007, **13**(2), 514–543.
- Cox, D., Asymptotics for m -type smoothing splines. *Ann. Statist.*, 1983, **11**(2), 530–551.
- Cox, D. and O’Sullivan, F., Asymptotic analysis of penalized likelihood and related estimators. *Ann. Statist.*, 1990, **18**(4), 1676–1695.
- Eggermont, P. and LaRiccia, V., *Maximum Penalized Likelihood Estimation*, Vol. I, 2001 (Springer: New York).
- Eggermont, P. and LaRiccia, V., *Maximum Penalized Likelihood Estimation*, Vol. II, 2009 (Springer: New York).
- Eubank, R., *Spline Smoothing and Nonparametric Regression*, 1988 (Marcel Dekker: New York).
- Fan, J., Fan, Y. and Lv, J., Aggregation of nonparametric estimators for volatility matrix. *J. Financial Econometr.*, 2007, **5**, 321–357.
- Fan, J. and Yao, Q., *Nonlinear Time Series: Nonparametric and Parametric Methods*, 2005 (Springer: New York).
- Florens-Zmirou, D., On estimating the diffusion coefficient from discrete observations. *J. Appl. Probab.*, 1993, **30**(4), 790–804.
- Genon-Catalot, V., Larédo, C. and Picard, D., Nonparametric estimation of the diffusion coefficient by wavelet methods. *Scand. J. Statist.*, 1992, **19**, 319–335.
- Gobet, E., Hoffmann, M. and Reiss, M., Nonparametric estimation of scalar diffusions based on low frequency data. *Ann. Statist.*, 2004, **32**(5), 2223–2253.
- Green, P. and Silverman, B., *Nonparametric Regression and Generalized Linear Models*, 1994 (Chapman & Hall: London).
- Hamrick, J. and Taqqu, M., Testing diffusion processes for non-stationarity. *Math. Meth. Oper. Res.*, 2009, **69**(3), 509–551.
- Hoffmann, M., l_p Estimation of the diffusion coefficient. *Bernoulli*, 1999a, **5**(3), 447–481.
- Hoffmann, M., Adaptive estimation in diffusion processes. *Stochast. Process. Appl.*, 1999b, **79**(1), 135–163.
- Jacod, J., Non-parametric kernel estimation of the coefficient of a diffusion. *Scand. J. Statist.*, 2000, **27**(1), 83–96.
- Jiang, G. and Knight, J., A nonparametric approach to the estimation of diffusion processes, with an application to a short-term interest rate model. *Econometr. Theory*, 1997, **13**(5), 615–645.
- Jorion, P., *Financial Risk Manager Handbook*, 2009 (Wiley: Hoboken, NJ).
- Karatzas, I. and Shreve, S., *Brownian Motion and Stochastic Calculus*, 1991 (Springer: New York).
- Kutoyants, Y., *Statistical Inference for Ergodic Diffusion Processes*, 2004 (Springer: London).
- Soulier, P., Nonparametric estimation of the diffusion coefficient of a diffusion process. *Stochast. Anal. Appl.*, 1998, **16**, 185–200.
- Utreras, F., On computing robust splines and applications. *SIAM J. Sci. Statist. Comput.*, 1981, **2**, 153–163.