A Comparison of Alternative Non-parametric Estimators of the Short Rate Diffusion Coefficient

Roberto Renò[*], Antonio Roma[†] and Stephen Schaefer[‡]

November 20, 2006

[*]Corresponding Author, Dipartimento di Economia Politica, Università di Siena, Piazza S.Francesco 7, 53100, Siena, Italy. E-mail: reno@unisi.it. Tel: +39 0577232649, fax: +39 0577232661

[†]Dipartimento di Economia Politica, Università di Siena, Piazza S.Francesco 7, 53100, Siena, Italy. E-mail: roma@unisi.it. Tel: +39 0577232692, fax: +39 0577232661

[‡]London Business School, Sussex Place, Regent's Park, London NW1 4SA, UK. E-mail: schaefer@london.edu. Tel: +44 (0)20 7706-6887, fax: +44 (0)20 7724-3317

A Comparison of Alternative Non-parametric Estimators of the Short Rate Diffusion Coefficient

Abstract

In this paper we discuss the estimation of the diffusion coefficient in one-factor models for the short rate via non-parametric methods. We test the estimators proposed by Ait-Sahalia (1996), Stanton (1997) and Bandi and Phillips (2003) on Monte Carlo simulations of the Vasicek and CIR model. We show that the Ait-Sahalia estimator is not applicable for values of the mean reversion coefficient typically displayed by interest rate data, while the Stanton and Bandi-Phillips estimators perform better. Each of the three estimators depends crucially on the choice of the bandwidth parameter. Our analysis shows that the estimators give different results for both the data set analyzed by Ait-Sahalia (1996) and by Stanton (1997). Finally we show that the data sets used by Ait-Sahalia and Stanton are inherently different and, in particular, that very short-term data exhibit characteristics which are inconsistent with a diffusion.

JEL Classification: C14, E43

1 Introduction

Most continuous time term structure models are based on diffusion dynamics for the state variables. In those that employ a single state variable, this is almost always the instantaneous risk-free interest rate. Different specifications for the interest rate process lead to different term structure models. In principle, an arbitrary univariate diffusion process for the instantaneous interest rate may be specified,

$$dr_t = \mu(r_t)dt + \sigma(r_t)dW_t, \tag{1}$$

although in this case the resulting model for the term structure of interest rates may only be solved for numerically. A number of different models for the short rate have been proposed and tested (Chan et al., 1992), which parameterize the drift and diffusion coefficients $\mu(r_t)$, $\sigma(r_t)$ through the specification

$$dr_t = (a + br_t)dt + \sigma r_t^\gamma dz \tag{2}$$

where, for example, $\gamma = 0$ corresponds to the Vasicek (1977) Gaussian model and $\gamma = 1/2$ to the Cox et al. (1985) (CIR) square root model.

Florens-Zmirou (1993) and Ait-Sahalia (1996) pioneered the idea of letting the diffusion coefficient of for the instantaneous interest rate process to be modeled by the data themselves through a non parametric approach. The idea has subsequently been extended to both the drift and diffusion coefficients by Stanton (1997), Jiang and Knight (1997) and, more recently, by Bandi and Phillips (2003). These estimators are based on the non-parametric estimation of the conditional density of the short rate.

The advantage of the non-parametric specification is clearly its flexibility. When applied to actual interest rate data, these non-parametric estimators produce functions $\mu(r)$ and $\sigma(r)$ that appear non-linear in r, and depart from benchmark parametric models such as the CIR square root model. Chapman and Pearson (2000) used Monte Carlo simulations to investigate the properties of the Stanton estimator for the drift and diffusion function, as well as a drift estimator developed using Ait-Sahalia's approach. Adopting the CIR model as the null hypothesis, they conclude that the non-linearity of the drift function $\mu(r)$, estimated through the non-parametric model, may not be a reliable indicator of a truly non-linear underlying model. They conclude however that the Stanton (1997) diffusion estimator is capable of identifying with some precision the square root form of the non-linearity in the volatility parameter. So far, there is no analysis in the literature of the Ait-Sahalia volatility estimator.

In this paper we examine the non-parametric estimators proposed by Ait-Sahalia, Stanton and Bandi-Phillips, with a particular focus on the Ait-Sahalia (henceforth AS) non-parametric diffusion estimator. We show that the AS estimator is problematic. For realistic parameter values, the estimated volatility will be a non-linear and biased function of the interest rate even when the actual volatility is a constant. Non-parametric smoothing does not attenuate these problems. It turns out that the AS estimator provides reasonable estimates only for unrealistically high mean-reversion speed[1]. We also show that using actual (positive) interest rate data, the approach may result in negative estimates of variance. In contrast, the Stanton and Bandi-Phillips estimators appear to be reasonably accurate, although they are slightly biased in the tail of the distribution. Again, the selection of the smoothing parameter affects the results. We also find that the Stanton and Bandi-Phillips estimators provide almost identical estimates of the variance function.

Finally we apply the methods to actual interest rate data. In particular we analyze the dataset used in Ait-Sahalia (1996) and show that its properties do not fit the desired properties for a diffusion.

The rest of the paper is organized as follows. Sections 2, 3, 4 review and discuss the Ait-Sahalia, Stanton and Bandi-Phillips estimators respectively, using of Monte Carlo simulations of univariate processes of the short rate. Section 5 reports estimation results on the 7-day Eurodollar deposit rate and the 3-months T-bill rate, and compares their properties. Section 6 concludes.

2 Analysis of the Ait-Sahalia estimator

2.1 Theory

The Ait-Sahalia estimator (AS) is based on the fact that, if a variable $r(t)$ follows a diffusion described by:

$$dr(t) = \mu(r)dt + \sigma(r)dW(t) \tag{3}$$

in which $W(t)$ is a standard Brownian motion and $\mu(r)$ and $\sigma(r)$ are such that a unique solution of the stochastic differential equation (3) exists, then:

$$\sigma^2(r) = \frac{2\int_{-\infty}^{r} \mu(x)\pi(x)dx}{\pi(r)} \tag{4}$$

where $\pi(r)$ is the unconditional distribution of r under the diffusion (3). Given two of the three functions $\mu(r), \sigma(r), \pi(r)$, equation (4) allows the third to be obtained via integration or differentiation.

[1]By "unrealistically" high we mean in comparison with levels of mean reversion displayed in interest rate data.

Ait-Sahalia (1996) suggests specifying the drift $\mu(r)$ as an affine function of r and then estimating the conditional variance $\sigma^2(r)$. He also suggests replacing $\pi(r)$ in (4) with an estimate derived using a non-parametric approach (Scott, 1992). Suppose we have T observations of r, denoted by $\hat{r}_i, i = 1, \ldots, T$. Then the non-parametric estimator of $\pi(r)$ is given by:

$$\hat{\pi}(r) = \frac{1}{Th_s} \sum_{i=1}^{T} K\left(\frac{r - \hat{r}_i}{h_s}\right) \tag{5}$$

where K is a kernel function that depends on h_s. A typical choice of the kernel, suggested by Ait-Sahalia (1996), is the Gaussian kernel,

$$K(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}}. \tag{6}$$

For the Gaussian kernel the optimal smoothing parameter is given by

$$h_s = h\hat{\sigma}T^{-\frac{1}{5}}, \tag{7}$$

where $\hat{\sigma}^2 = Var[\hat{r}_i]$ and $h = 1.06$ (Scott, 1992, p.131)[2]. It is widely recognized that the choice of the kernel function is much less important than the choice of an appropriate smoothing parameter (Scott (1992), p.133). Although nominally a "non-parametric" approach, the bandwidth h_s (i.e. the scalar h) is in fact a parameter that has to be properly selected.
Under these assumptions, the Ait-Sahalia estimator is given by:

$$\hat{\sigma}^2(r) = \frac{2 \sum_{i=1}^{T} \int_{-\infty}^{r} \mu(s) K\left(\frac{s - \hat{r}_i}{h_s}\right) ds}{\sum_{i=1}^{T} K\left(\frac{r - \hat{r}_i}{h_s}\right)} \tag{9}$$

with $K(x)$ given by (6) and h_s by (7).

2.2 Non-parametric estimates of the density with persistent data

While much is known about the asymptotic behavior of non-parametric estimators, their finite sample properties are largely unknown and their robustness, especially with respect to the choice

[2]The optimal smoothing parameter is chosen with reference to the Mean Integrated Square Error (MISE) of the density estimator, $\hat{f}$, compared to a reference density, f, given by:

$$MISE = \int_t \left(\hat{f}(t) - f(t)\right)^2 dt. \tag{8}$$

Different reference densities imply a different optimal smoothing parameter. See Scott (1992), p. 142 for the conversion ratio of the smoothing parameter under different reference densities.

of the bandwidth parameter h, is suspect. For the AS estimator, we show that estimating the unconditional density of data, that is the denominator in (4), is particularly critical.

To study the small-sample properties of the AS and other estimators we use a Monte Carlo approach and generate sample paths for the interest rate that follow an Ornstein-Uhlenbeck ('Vasicek') process as in the Vasicek (1977) model:

$$dr(t) = k(\alpha - r(t))dt + \sigma dW(t). \tag{10}$$

For this model the diffusion coefficient is, of course, constant and equal to σ^2. The range of possible interest rates is $(-\infty, \infty)$ and the unconditional density function, $\pi(r)$, is a Normal distribution with mean α and variance $\sigma^2/2k$. Although the Vasicek process might be viewed as an unrealistic model of the short-term rate – because, for example, it admits negative values – it is quite adequate to illustrate our first point that, when the degree of mean reversion is (realistically) low, estimating the unconditional density is very difficult.

We replicate 1,000 paths of the Vasicek diffusion model (10) for a range of values for k, the mean reversion coefficient, and N, the number of observations. Here and throughout the paper we simulate daily observations, that is $\Delta t = 1/252$, where the number of days per year is assumed to be 252. Moreover, we draw the starting point for each sample path from the (true) unconditional distribution. In each case we set $\alpha = 8.3\%$ and σ such that the unconditional standard deviation is 3%. For each sample path we estimate the mean and the standard deviation of $r(t)$ and, as $N \to +\infty$, these should converge to α and $\sigma/\sqrt{2k}$ respectively, i.e. the mean and standard deviation of the unconditional distribution. Table 1 reports the average values of the mean and standard deviation of $r(t)$ for the replications, along with confidence bands and shows that the rate of convergence of the sample estimates can be very slow when mean reversion is low. In particular, the standard deviation of the unconditional distribution is significantly underestimated in small samples. For example, the generated value of the standard deviation is 3%, but with $N = 10,000$ observations, the estimated standard deviation has a mean of 1.73% and a 95th percentile of 2.76%. The bias is even larger at $N = 5,000$, and still appreciable at $N = 50,000$, even with $k = 0.5$. On the other hand, with high mean reversion ($k = 5$) the estimates are unbiased. With $N = 10,000$ observations, for example, and $k = 5$, the 5^{th} and 95^{th} percentiles of the estimate of the standard deviation are 2.73% and 3.23%. Moreover, with a small sample size, even the mean is poorly estimated. In the case of the mean, unlike the standard deviation, the estimate is unbiased but the dispersion of the estimates can be very large. Obviously, since densities are constrained to be integrated to one, if the unconditional standard deviation is systematically underestimated, the estimated non-parametric density will have a peak at the estimated mean (which will be unlikely to be the true mean) and will have thinner tails than the true unconditional distribution. The problem of estimating kernels with persistent data has been also studied by Pritsker (1998), who finds out that "inference based on the asymptotic distribution of kernel density estimators can be very misleading".

These results show how estimates of the unconditional density of $r(t)$ vary with the length of the data series and the degree of mean reversion. Where, in the range we have considered, do actual data fit? First, even with daily data, $5,000$ observations represent around 20 years of data, $10,000$ observations 40 years and $50,000$ observations an entirely unrealistic 200 years. The sample size in Stanton paper is $N = 7,795$, and to our knowledge the largest ever used is $N = 8,522$ in Johannes (2004). AS has $N = 5,505$ observations. Thus AS's sample size is approximately equal to our "small" sample and Stanton's lies between our "small" and "medium" sized samples. Both samples, however, are "large" in an economic sense, however, in that both rely on models having stable parameters for a period of 20-30 years. This is particularly worrisome in the case of the mean parameter, α, and many authors from Brennan and Schwartz onwards have advocated a two (or more) factor process for the short rate in which the mean itself changes stochastically (Brennan and Schwartz, 1980; Balduzzi et al., 1996; Dai and Singleton, 2000). This misspecification is a further potential problem for non-parametric estimation of a single factor specification.

The second parameter that strongly affects the reliability of small sample estimates of the unconditional density of $r(t)$ is the degree of mean reversion k, and Table 2 reports some estimates of the mean reversion parameter, k, from four studies in the literature. Three of these report estimates in the range $0.1 - 0.2$; the exception is the study by Ait-Sahalia (1996) who reports a value of 0.978. AS uses a daily series of 7-day Eurodollar deposit rates. These data have some features, discussed below, that cast serious doubt on the reliability of this high estimate of k

In summary, our conclusion from this part of the analysis is that the likely value of k is in the region of 0.1-0.2 and that, in this case, even with sample sizes of 5,000 and 10,000, reliable estimation of the unconditional density is almost impossible.

2.3 Performance of the estimator

In this section we carry out a Monte-Carlo analysis of the AS estimator (12) when the data are generated by the Vasicek model. Initially we assume that the drift is known, that is, in computing the estimator of $\sigma^2(r)$ we use the values of k and α used to generate the simulated samples. An advantage of using a constant conditional variance is that in this case, with a Gaussian kernel, we can easily find an analytic expression for the estimator of $\sigma^2(r)$ in (9). Indeed, in the case of the Vasicek model (10), it is straightforward to show that the Ait-Sahalia estimator has the form:[3]

$$\hat{\sigma}^2(r) = 2kh_s^2 + \frac{h_s\sqrt{2\pi}\sum_{t=1}^{T} k(\alpha - \hat{r}_t)\left[1 + Erf\left(\frac{r-\hat{r}_t}{\sqrt{2}h_s}\right)\right]}{\sum_{t=1}^{T} e^{-\frac{(r-\hat{r}_t)^2}{2h_s^2}}}. \tag{12}$$

From equation (7), as the number of observations increases ($T \rightarrow \infty$) and if the second moment of the unconditional distribution is finite, then $h_s \rightarrow 0$ and the first term on the right hand side of (12), a constant, goes asymptotically to zero. The asymptotic estimator is then given by:

$$\sigma^2(r) = \lim_{T\rightarrow\infty} \frac{h_s\sqrt{2\pi}\sum_{t=1}^{T} k(\alpha - \hat{r}_t)\left[1 + Erf\left(\frac{r-\hat{r}_t}{\sqrt{2}h_s}\right)\right]}{\sum_{t=1}^{T} e^{-\frac{(r-\hat{r}_t)^2}{2h_s^2}}} \tag{13}$$

However, in any finite sample, the conditional variance estimator (12) displays not only sampling error but, in the case of the Vasicek model when the numerator may become negative, it is not even guaranteed to be positive. The two panels of Figure 1 show the behavior of the numerator as a function of $r(t)$ for $h = 4$, conditioning on $r = 0.08$ and $r = 0.18$ respectively, using the FGLS parameter values estimated by Ait-Sahalia (1996). As the estimated conditional variance will be approximately proportional to the expected value of the numerator under the unconditional distribution of the interest rate $r(t)$, it is clear that the estimated conditional variance may become negative.

Figures 2,3,4 show, respectively, the true density (the denominator in 4), the numerator in (4) and the true variance $\sigma^2(r)$ together with their estimates computed from a single simulated series of size N. We give results for values of N between 5,000, i.e., approximately twenty years of daily data, and 5,000,000, i.e., approximately 20,000 years of daily data, again by using a single simulation in order to study the convergence of the estimator as $N \rightarrow \infty$.

In the previous subsection we showed that, with low mean reversion,s the variance of the unconditional distribution is systematically underestimated and the mean is poorly estimated. These deficiencies are clearly illustrated in Figure 2. For low and average mean reversion we obtain reliable estimates of the density only for $N = 50,000$ or larger. Again, especially with low or average mean reversion, errors in estimating the marginal density can have a dramatic

[3] The function $Erf(x)$ is defined as:

$$Erf(x) = \frac{2}{\sqrt{\pi}}\int_0^x e^{-t^2}dt \tag{11}$$

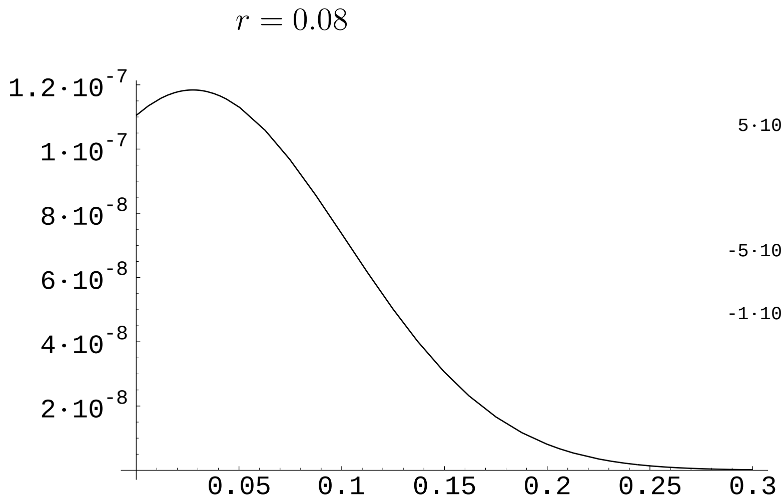

Figure 1: Numerator with Ait-Sahalia parameters conditioned on $r = 8\%$ (left) and $r = 18\%$ (right).

effect on the estimate of the numerator, as illustrated in Figure 3. This, in turn, affects the variance estimate which is very poor for small sample sizes (Figure 4).

The particular case for $k = 0.05,\ N = 5,000$, at the bottom-left corner of Figures 2,3,4 is worth noting. Here, in the center of the distribution, the density is overestimated and the numerator underestimated. These two effects combine to produce substantial underestimation of the variance. Asymptotically, the estimates are quite good but, for low mean reversion ($k = 0.05$), only for values of N larger than around $500,000$. In contrast, for high mean reversion, which unfortunately is not a feature of the data, we obtain reasonable estimates of the density and numerator even in small samples.

2.4 Small sample properties

We now turn to compute the small sample properties of the estimator. The charts in Figure 5 were computed using $1,000$ replications of samples paths with $N = 6,000$. The charts in the first column of Figure 5 show the mean value of $\hat{\sigma}$, together with the 5% and 95% confidence limits, for values for k of 5.0 (first row), 0.5 (second row) and 0.05 (third row). In the first column, the estimator is computed using the *true* values of k and α. For the charts in the second column we use the true value of k and an estimated value of α. For the third column, estimated values of both parameters are used.

Some of the results are striking. For $k = 5.0$ the performance of the estimator, even with N as "low" as $6,000$, is quite good and, interestingly, the performance actually improves when estimated values of k and α are used in place of their true values. The improvement is particularly

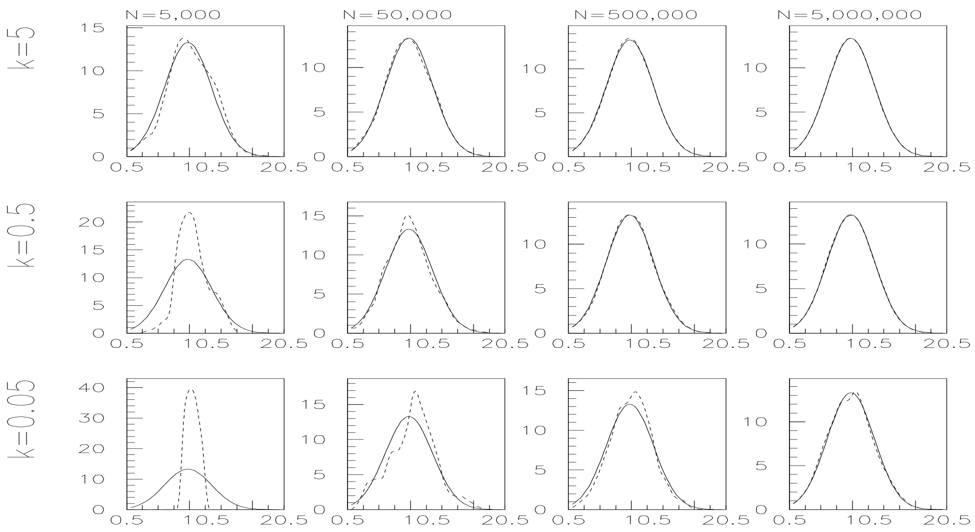

Figure 2: Non-parametric estimates of the density (dashed line) on a single simulation of the Vasicek model (10) of length N, for different values of N, as displayed, with $\alpha = 8.3\%$, σ is such that the unconditional standard deviation $\sqrt{\sigma^2/2k}$ is 3%, $h = 1.06$ and different levels of mean reversion:$k = 5$ (top row), $k = 0.5$ (central row) and $k = 0.05$ (bottom row). For comparison, the true density is shown (solid line).

marked when the true value of α is replaced by its estimate and the reason for this is that, by using the estimated value the mean of the unconditional distribution and the sample mean coincide and the numerator and denominator of the estimator become more "consistent". For $k = 0.5$ and 0.05, however, the performance of the estimator is quite poor even when estimated values of k and α are used. In particular, for large values of r, the mean is increasing in r even though the true value of r in this case is constant.

2.5 The choice of the bandwidth parameter h

A critical issue in any non-parametric study is the choice of the bandwidth parameter h and the literature provides no widely accepted recipes for this choice. In this section, we study the dependence of the Ait-Sahalia estimator on h.

Figures 6 and 7 show the results on simulations of the Vasicek models with the same parameters as in Figure 5, but with $h = 3$ and $h = 5$ respectively. There are two main reasons for considering a higher value of h. First, it will lead to a smoother estimate of the density in the tails of the distribution where fewer observations are available. Second, we compensate for the

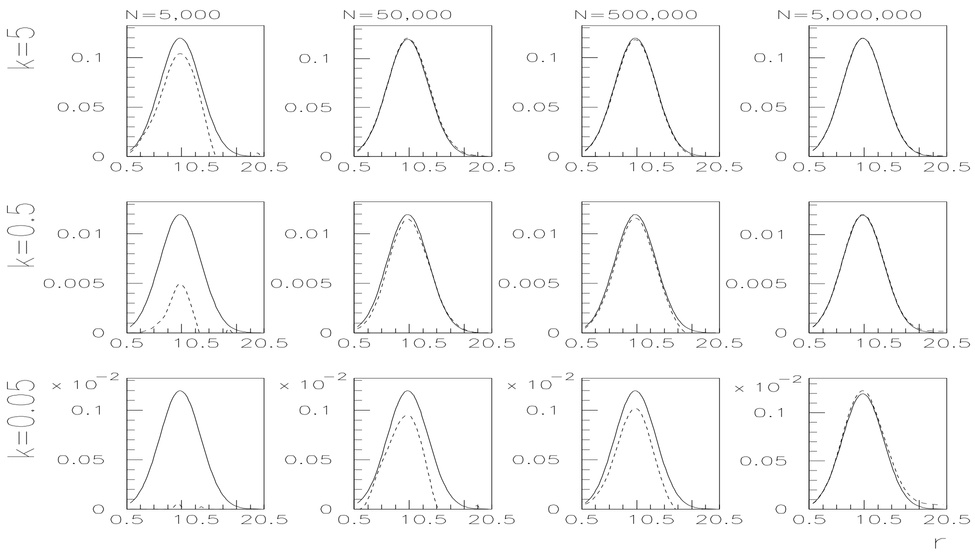

Figure 3: Non-parametric estimates of the numerator in the Ait-Sahalia estimator (dashed line), on a single simulation of the Vasicek model (10) of length N, for different values of N, as displayed, with $\alpha = 8.3\%$, σ is such that the unconditional standard deviation $\sqrt{\sigma^2/2k}$ is 3%, $h = 1.06$ and different levels of mean reversion:$k = 5$ (top row), $k = 0.5$ (central row) and $k = 0.05$ (bottom row). For comparison, the true numerator is shown (solid line).

underestimation of $\hat{\sigma}$ in equation (7) by choosing a high value of h.

Figures 6 and 7 show that the estimator is highly sensitive on the value on h in small samples ($N = 6,000$). In particular, for high and moderate mean reversion the higher values of h lead to an upward bias. Note that in the expression (12) h_s is a scale factor for the estimator of the variance in small sample. As expected, by increasing h we reduce the explosive behavior for large values of r, since the estimate of the density in the tail is smoother. If we repeat the experiment with larger sample size, the estimator converges slowly to the generated variance.

2.6 Performance when short rate follows the CIR Process

It is possible that our results so far are specific to the Ornstein-Uhlenbeck case when the variance of the short rate is not a function of its level. In this section we therefore repeat the analysis just described but when the true model describing the evolution of the short rate is the Cox et al. (1985) model (CIR):

$$dr(t) = k(\alpha - r(t))dt + \sigma\sqrt{r(t)}dW(t). \tag{14}$$

As before we compute 1,000 sample paths for the short rate, each of length 6,000, and, in Figure 8, show the mean value of $\hat{\sigma}$, together with the 5% and 95% confidence limits. The parameter values are given in the caption. The results in the three columns are for values of the bandwidth parameter h of 1.06, 3.0 and 5.0.

Once again, results are not encouraging. For $h = 1.06$ the estimator works reasonably only when there is a low level of persistence in the data, i.e. for high k. The results for $k = 0.5$ and 0.05 are, as in the previous case, poor.

In summary, the small sample performance of the Ait-Sahalia estimator is poor except in the case when the short rate displays a very high degree of mean reversion. In this case it is possible to obtain a reasonable estimate of the marginal density. However, in the empirically relevant case when the degree of mean reversion is low, the estimator does not produce reliable estimates of the variance and can easily suggest the presence of a non-linear relation between the variance of interest rates and their level when, in fact, no such relation exists.

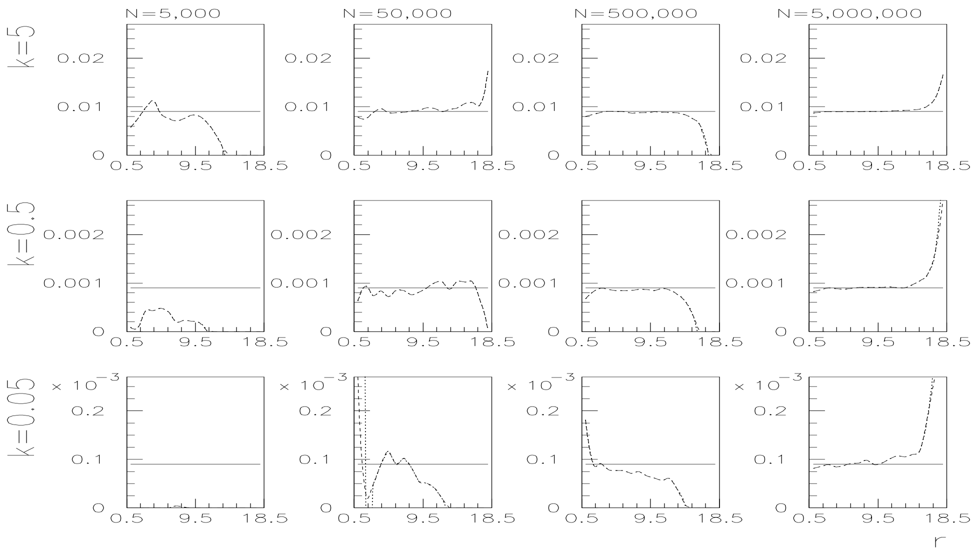

Figure 4: Ait-Sahalia estimator of the variance (dashed line), on a single simulation of the Vasicek model (10) of length N, for different values of N, as displayed, with $\alpha = 8.3\%$, σ is such that the unconditional standard deviation $\sqrt{\sigma^2/2k}$ is 3%, $h = 1.06$ and different levels of mean reversion:$k = 5$ (top row), $k = 0.5$ (central row) and $k = 0.05$ (bottom row). For comparison, the true variance is shown (solid line).

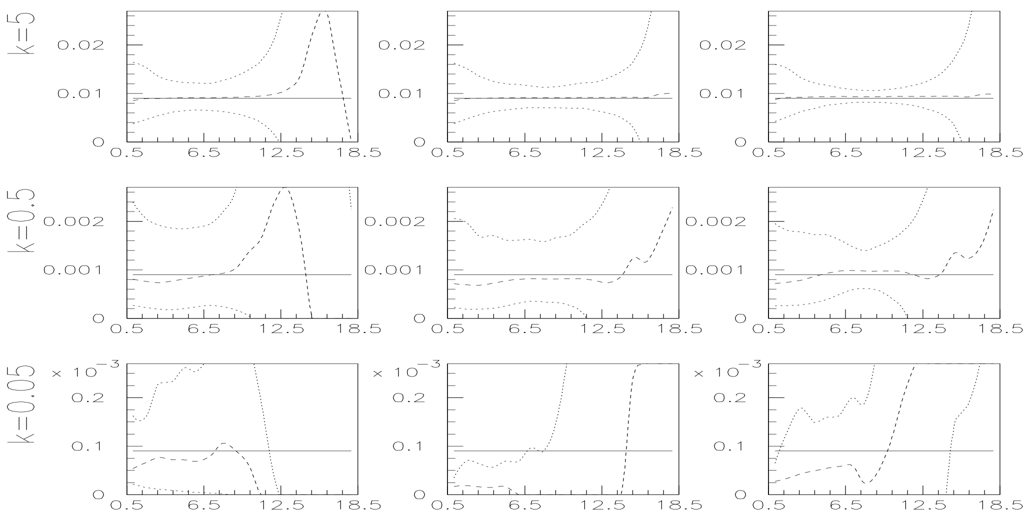

Figure 5: Average Ait-Sahalia estimator computed on 1,000 simulation of the Vasicek model (10), with $\alpha = 8.3\%$, σ such that the unconditional standard deviation $\sqrt{\sigma^2/2k}$ is 3%, $\mathbf{h} = \mathbf{1.06}$, sample size $N = 6,000$ and different levels of mean reversion: $k = 5$ (top row), $k = 0.5$ (central row) and $k = 0.05$ (bottom row). In the first column, we use the true values of α and k when estimating $\sigma^2(r)$; in the second column, we use the true value of k and estimate α from data; in the third column, we estimate both k and α from data.The average estimate is the dashed line, while the true value is the solid line. The dotted line show 95% and 5% confidence limits.

3 Analysis of the Stanton Estimator

The estimator proposed by Stanton (1997) differs from the Ait-Sahalia estimator and, in particular, does not require that the drift is known. Stanton (1997) shows that, given discretely sampled data, the drift and diffusion coefficients in (3) may be represented as:

$$\mu(r_t) = \frac{1}{\Delta} E_t[r_{t+\Delta} - r_t] + O(\Delta) \tag{15}$$

$$\sigma(r_t) = \sqrt{\frac{1}{\Delta} E_t[(r_{t+\Delta} - r_t)^2]} + O(\Delta) \tag{16}$$

Disregarding higher order terms, the drift and diffusion coefficients for a given level of r may be approximated, as the conditional expectations of (15) and (16). These expected values may be computed using the non-parametric estimate of the conditional density, which is obtained via the kernel estimator, as previously described:

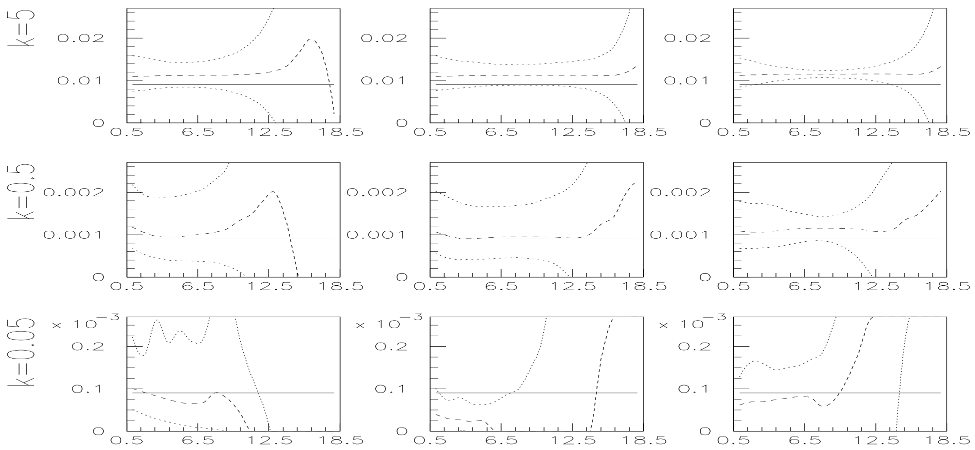

Figure 6: As in Figure 5, but with $h = 3$.

$$\hat{\mu}(r) = E[r_{t+1} - r_t | r_t = r] = \frac{\sum_{t=1}^{T-1} (\hat{r}_{t+1} - \hat{r}_t) K\left(\frac{r - \hat{r}_t}{h_s}\right)}{\sum_{t=1}^{T} K\left(\frac{r - \hat{r}_t}{h_s}\right)} \tag{17}$$

$$\hat{\sigma}^2(r) = E[(r_{t+1} - r_t)_t^2 | r_t = r] = \frac{\sum_{t=1}^{T-1} (\hat{r}_{t+1} - \hat{r}_t)^2 K\left(\frac{r - \hat{r}_t}{h_s}\right)}{\sum_{t=1}^{T} K\left(\frac{r - \hat{r}_t}{h_s}\right)} \tag{18}$$

The variance estimator suggested by Stanton (18) has the attractive feature that the same non-parametric estimator appears in both the numerator and the denominator and so errors arising from this source may offset at least to some extent. Even though both the numerator and the denominator (which is the same as that for Figure 2) are not well approximated in small samples, the bias is, in this case, in the same direction and so cancels out in the ratio. As a consequence the variance is better estimated than in the case of the AS estimator.

Our results confirm that this is indeed the case. Figure 9 shows the estimate of the numerator in (18) and of the estimated variance from a single simulation of the Vasicek model (10) with the same parameters as used previously. As before, the single path simulation is used to test the properties of the estimator as the sample size goes to infinity.

In the Ait-Sahalia estimator the kernel estimate of the marginal density also appears but only in the denominator and so, in this case, there is no scope for errors to offset as in the Stanton

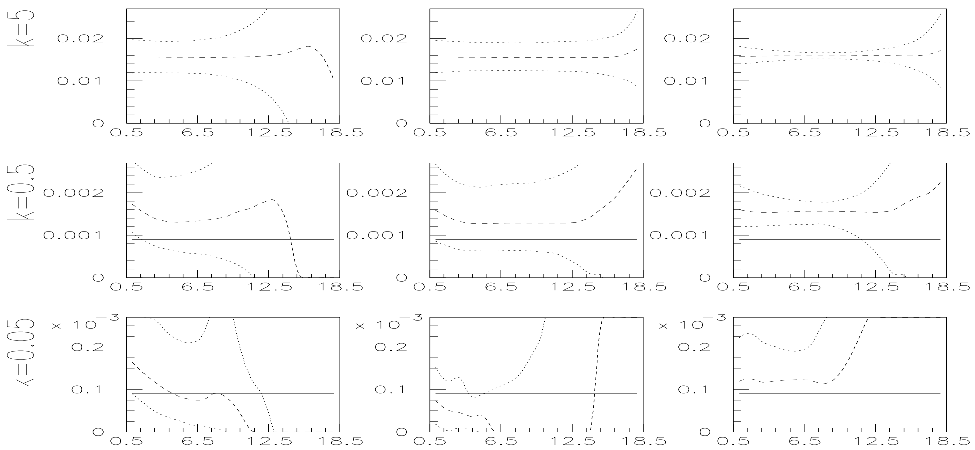

Figure 7: As in Figure 5, but with $h = 5$.

estimator.

Even in the case of the Stanton estimator, however, h is still a crucial parameter that needs to be fine-tuned. Figure 10 shows that, in the case of the Vasicek model, the performance of the estimator is improved substantially by increasing h. This causes the width of the confidence intervals to shrink without introducing a bias. Things are slightly different when the Stanton method is used to estimate variance from data generated by the CIR model (14). The results are shown in figure 11 which shows that, while the Stanton estimator performs much better than the Ait-Sahalia estimator, for larger values of h (3.0 and 5.0) small biases arise for both low and high values of r. [4]

4 Analysis of the Bandi-Phillips estimator

The estimator proposed in Bandi and Phillips (2003) is the following:

$$\hat{\sigma}^2(r) = \frac{\sum_{i=1}^{T} K\left(\frac{r-\hat{r}_i}{h_s}\right)\left(\frac{1}{m_i}\sum_{j=0}^{m_i}\left[\hat{r}_{t_{i,j}+1} - \hat{r}_{t_{i,j}}\right]^2\right)}{\sum_{i=1}^{T} K\left(\frac{r-\hat{r}_i}{h_s}\right)} \tag{19}$$

where $t_{i,j}$ is a subset of indices such that

$$t_{i,0} = \inf\left\{t \geq 0 : |\hat{r}_t - \hat{r}_i| \leq \varepsilon_s\right\},$$

[4] Stanton (1997) uses $h = 4$.

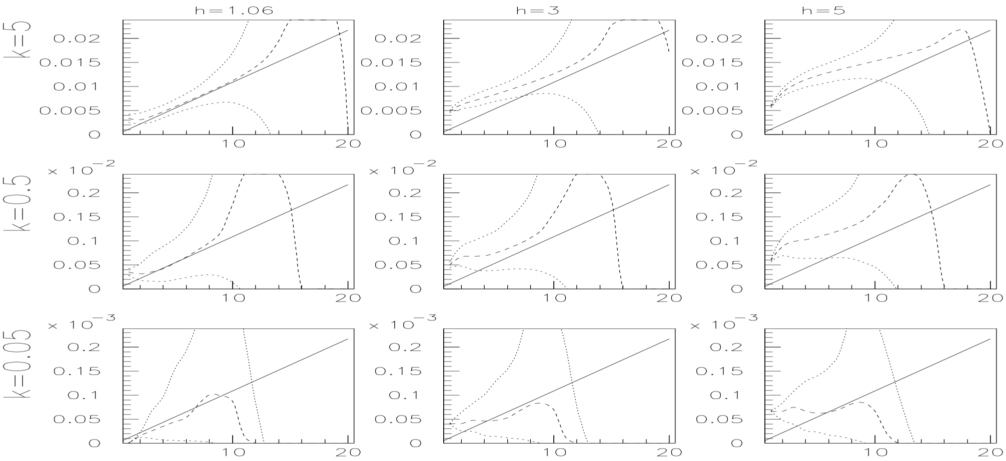

Figure 8: Average Ait-Sahalia estimator computed on 1,000 simulation of the CIR model (14), with $\alpha = 8.3\%$, σ is such that the unconditional standard deviation $\sqrt{\sigma^2/2k}$ is 3%, sample size $N = 6,000$ and different levels of mean reversion: $k = 5$ (top row), $k = 0.5$ (central row) and $k = 0.05$ (bottom row), and bandwidth parameter: $h = 1.06$ (first column), $h = 3$ (second column) and $h = 5$ (third column). The average estimate is the dashed line, while the true value is the solid line. The dotted line show 95% and 5% confidence limits.

and

$$t_{i,j+1} = \inf\left\{t \geq t_{i,j} + 1 : |\hat{r}_t - \hat{r}_i| \leq \varepsilon_s\right\},$$

m_i is the number of times that $|\hat{r}_t - \hat{r}_i| \leq \varepsilon_s$ and ε_s is a parameter to be selected[5]. Looking at expressions (18),(19), we can see that the difference between Stanton and Bandi-Phillips estimators is that, while the Stanton estimator weights the observation r_t by the quadratic variation at time t, the Bandi-Phillips estimator weights the observation r_t with the average quadratic variation of all observation which are "close" to r.

As for the previous estimators, we implement the Bandi-Phillips method on simulated paths of the Vasicek and CIR models. We use $\varepsilon_s = 1.5\%$ as suggested in Bandi (2002). Figures 12 and 13 show the results for the Vasicek and CIR models respectively. The performance of the Bandi-Phillips and Stanton estimators is almost identical with the errors from the former just slightly lower.

[5] See Bandi and Phillips (2003) for details.

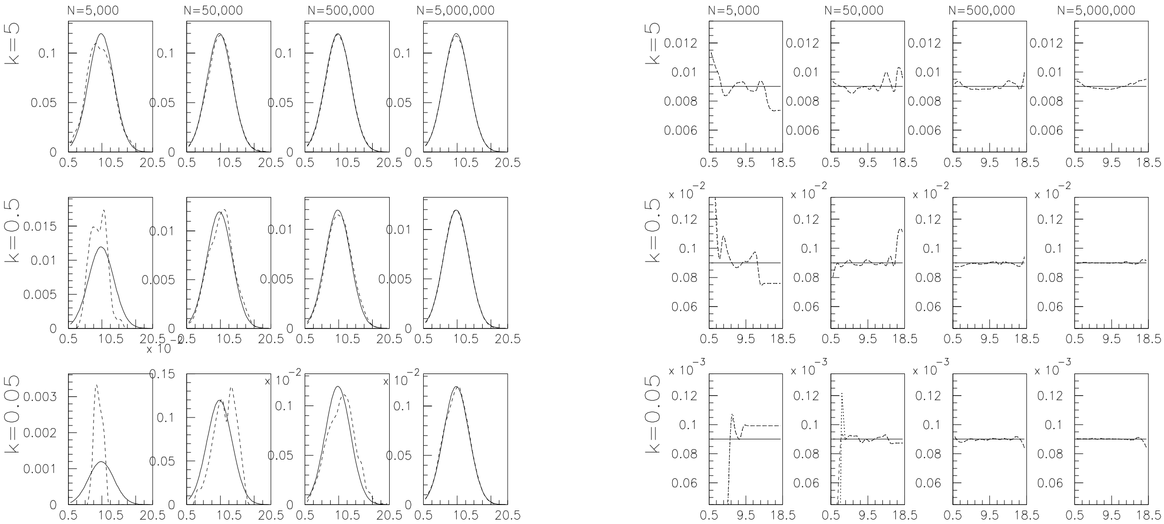

Figure 9: Non-parametric estimates of the numerator (left) and variance (right) for the Stanton estimator (dashed line), on a single simulation of the Vasicek model (10) of length N, for different values of N, as displayed, with $\alpha = 8.3\%$, σ is such that the unconditional standard deviation is 3%, $h = 1.06$ and different levels of mean reversion:$k = 5$ (top row), $k = 0.5$ (central row) and $k = 0.05$ (bottom row). For comparison, the true numerator is shown (solid line).

5 Data analysis

The analysis in the previous Section, using simulated interest rate data, casts serious doubt on the performance of non-parametric estimators on In this Section, we take a careful look of the data used in Ait-Sahalia (1996) and Stanton (1997).

The data used in Ait-Sahalia (1996) consist of the seven-day Eurodollar deposit rate from June 1, 1973 to February 25, 1995, a total of 5505 daily observations. Stanton (1997)'s data span an even longer period, from January 1965 to July 1995, and consist of daily yields on the three-month U.S. Treasury Bill, a total of 7975 observations. The time series of these two datasets are shown in Figure 14 and, even from visual inspection, it is clear that the two data sets are very different, and the Ait-Sahalia data has many more "spikes" which are typical of very short-term interest-rates. Indeed, Duffee (1996) suggests that rates on instruments with less than three months to maturity should not be used for these purposes since they have too much idiosyncratic variation.

To implement the Ait-Sahalia estimator, the drift function has to be assumed. We follow Ait-Sahalia (1996) and choose a linear specification $\mu(x) = k(\alpha - x)$. The parameters k, α must be

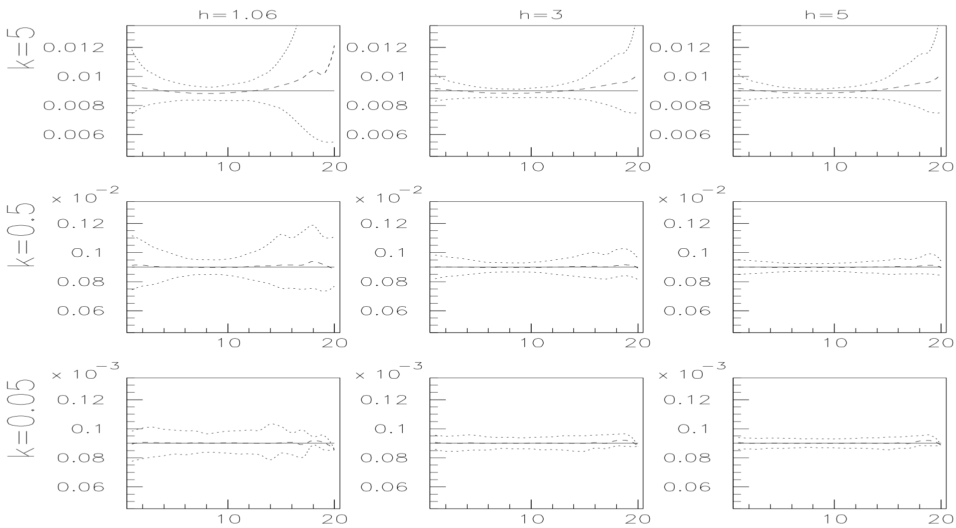

Figure 10: Average Stanton estimator computed on 1,000 simulation of the Vasicek model (10), with $\alpha = 8.3\%$, σ is such that the unconditional standard deviation is 3%, sample size $N = 6,000$ and different levels of mean reversion: $k = 5$ (top row), $k = 0.5$ (central row) and $k = 0.05$ (bottom row), and bandwidth parameter: $h = 1.06$ (first column), $h = 3$ (second column) and $h = 5$ (third column). The average estimate is the dashed line, while the true value is the solid line. The dotted line show 95% and 5% confidence limits.

estimated from the data and may be obtained in a number of ways. In Table 3 we report the parameter estimates obtained from Ait-Sahalia data using a number of alternative econometric methods. In the Table, OLS denotes Ordinary Least Squares regression, GLS denotes the use of a Cochrane-Orcutt correction for autocorrelation of the residuals and FGLS* denotes the method used by Ait-Sahalia. This is a two-stage procedure where the non-parametric variance function is used to weight observations.

We have estimated the conditional variance function on the Ait-Sahalia dataset using the Stanton and Bandi-Phillips methods and also using the Ait-Sahalia method with a linear drift specification. The two panels on the left of Figure 15 show the estimated variance function under the three methods and with two different bandwidths, $h = 1.06$ (the Gaussian optimal bandwidth) and $h = 4$ (that used by Stanton). While the Stanton and Bandi-Phillips estimators give fairly similar results, the Ait-Sahalia estimator gives a larger variance for low interest rates and a smaller variance for high interest rates than the other estimators. All three estimators give smoother estimates for larger values of h.

We have also attempted to apply all three methods using Stanton's (1997) dataset. However,

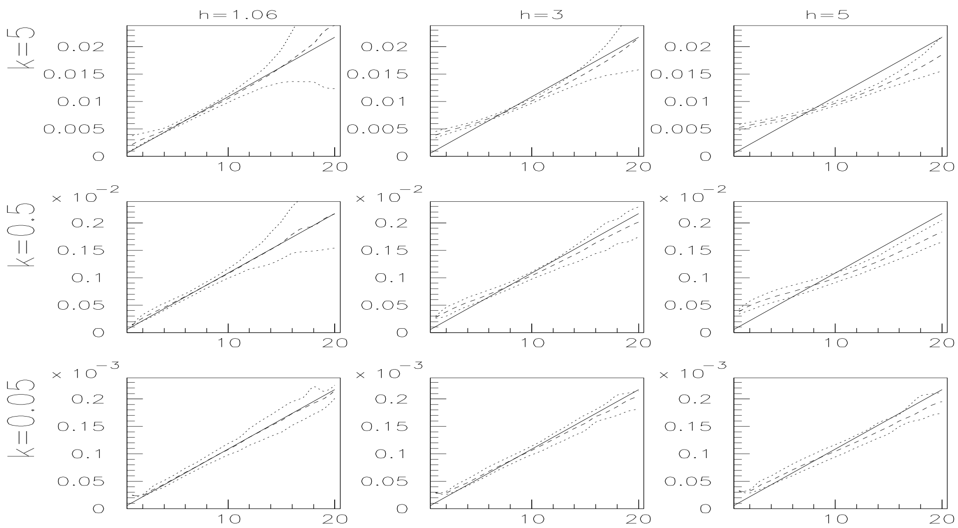

Figure 11: Average Stanton estimator computed on 1,000 simulation of the CIR model (14), with $\alpha = 8.3\%$, σ is such that the unconditional standard deviation is 3%, sample size $N = 6,000$ and different levels of mean reversion: $k = 5$ (top row), $k = 0.5$ (central row) and $k = 0.05$ (bottom row), and bandwidth parameter: $h = 1.06$ (first column), $h = 3$ (second column) and $h = 5$ (third column). The average estimate is the dashed line, while the true value is the solid line. The dotted line show 95% and 5% confidence limits.

when the mean reversion parameters, which are required for the Ait-Sahalia method, were estimated on this dataset, it became apparent that the data are close to non-stationary. The OLS autoregressive coefficient was 0.998998 which translates into a k of $-log(0.998996) * 255 = 0.25615$. The absence of the spikes that are present in the 7-day rate used by Ait-Sahalia is almost certainly responsible for the lower estimated intensity of the mean reversion in this case. The estimated parameters were are reported in Table 4.

The two panels on the right of Figure 15 show the estimates for the Stanton (1997) dataset. With the low level of estimated mean reversion, it was impossible, using the Ait-Sahalia method, to obtain meaningful conditional variance estimates with $h = 1.06$ and for rates larger than 15%. This is consistent with the results of our Monte Carlo simulations which show instability in the Ait-Sahalia estimator when mean reversion is low. With larger values of h, all three estimators appear more stable. It is worth noting that, for some values of r, even though they span more or less the same period, the two data sets give rise to estimates of variance that differ by an order of magnitude. It is possible that this simply reflects the difference in the maturity of the two rates but, in our view, it is more likely to result from the presence of large

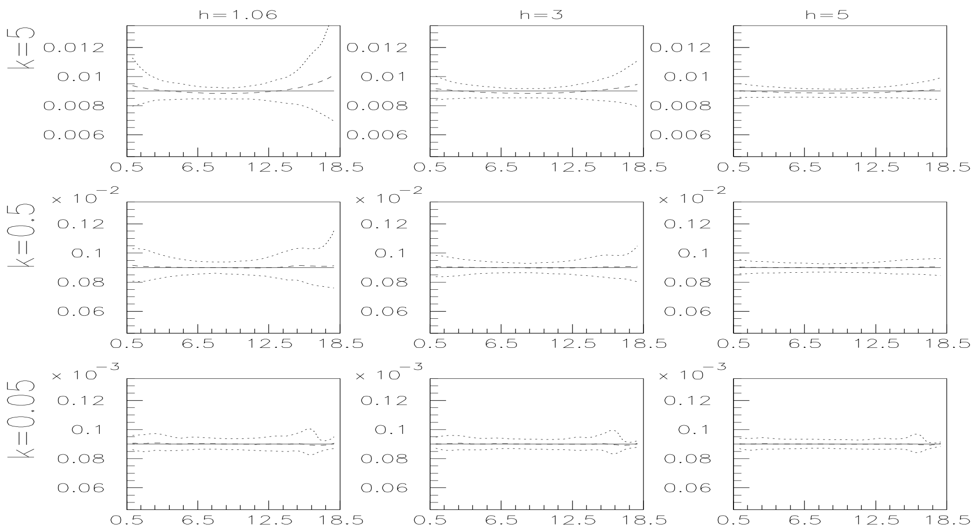

Figure 12: Average Bandi-Phillips estimator computed on 1,000 simulation of the Vasicek model (10), with $\alpha = 8.3\%$, σ is such that the unconditional standard deviation is 3%, sample size $N = 6,000$ and different levels of mean reversion: $k = 5$ (top row), $k = 0.5$ (central row) and $k = 0.05$ (bottom row), and bandwidth parameter: $h = 1.06$ (first column), $h = 3$ (second column) and $h = 5$ (third column). The average estimate is the dashed line, while the true value is the solid line. The dotted line show 95% and 5% confidence limits.

spikes in the Ait-Sahalia data and their relative absence in the Stanton data.

It is clear that, both from our simulations and from application to actual data, the conditional variance estimated by, on one hand, the Stanton and Bandi-Phillips methods and, on the other, the Ait-Sahalia method, differ substantially when the number of observations are available for the estimation is low and, in particular, for high values of the interest rate (say, above 12%). For these purposes, twenty years of daily data is a "low" number of observations. We also note that, as pointed out above, using the Ait-Sahalia technique, the estimated conditional variance may be negative. This occurs with the Stanton dataset and a bandwidth parameter of $h = 1.06$ when the interest rate on which the estimate is conditioned is large.

5.1 Further analysis of the Ait-Sahalia (1996) dataset

When estimating diffusion models for interest rates, a short term market interest rate must be used as an empirical proxy for the instantaneous short rate required by the theory. From a theoretical point of view, there is no strong basis for choosing between the 3-months rate or the

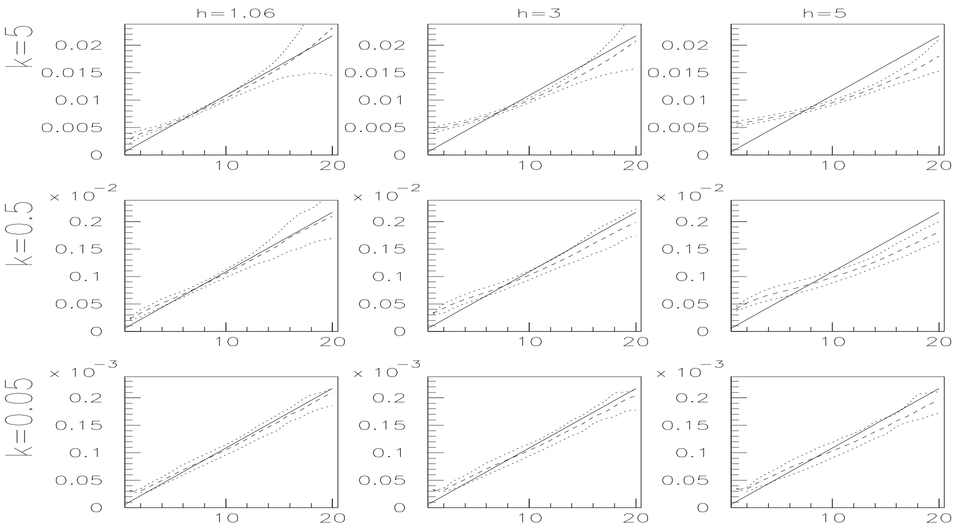

Figure 13: Average Bandi-Phillips estimator computed on 1,000 simulation of the CIR model (14), with $\alpha = 8.3\%$, σ is such that the unconditional standard deviation is 3%, sample size $N = 6,000$ and different levels of mean reversion: $k = 5$ (top row), $k = 0.5$ (central row) and $k = 0.05$ (bottom row), and bandwidth parameter: $h = 1.06$ (first column), $h = 3$ (second column) and $h = 5$ (third column). The average estimate is the dashed line, while the true value is the solid line. The dotted line show 95% and 5% confidence limits.

7-day rate: indeed the 7-day rate might be regarded as preferable, since its maturity is closer to that of the theoretical ideal. Empirically, however, there could be substantial differences. In particular, it is well known that interest rates for very short maturities display liquidity and calendar effects which may be inconsistent with the diffusion model, see e.g. Duffee (1996); Hamilton (1996).

We show that there are indeed substantial differences in the statistical properties of the Ait-Sahalia (1996) and Stanton (1997) time series. To gauge the difference, we restrict ourselves to the time-span of the Ait-Sahalia paper, that is 5,505 daily observations from January, 1973 to February, 1995. We study the time series of the 3-months rate and the 7-day rate; the latter coincides with the data set in Ait-Sahalia (1996), the first is a subset of the data set in Stanton (1997).

If we rank the largest absolute daily changes, we observe that almost all fall during the monetary experiment period, that is between 1979 and 1982. Table 5 shows the 20 largest interest rate changes, and the date in which they occurred. A number of differences between the two data set are immediately apparent. First, the largest changes are much larger for the 7-day data

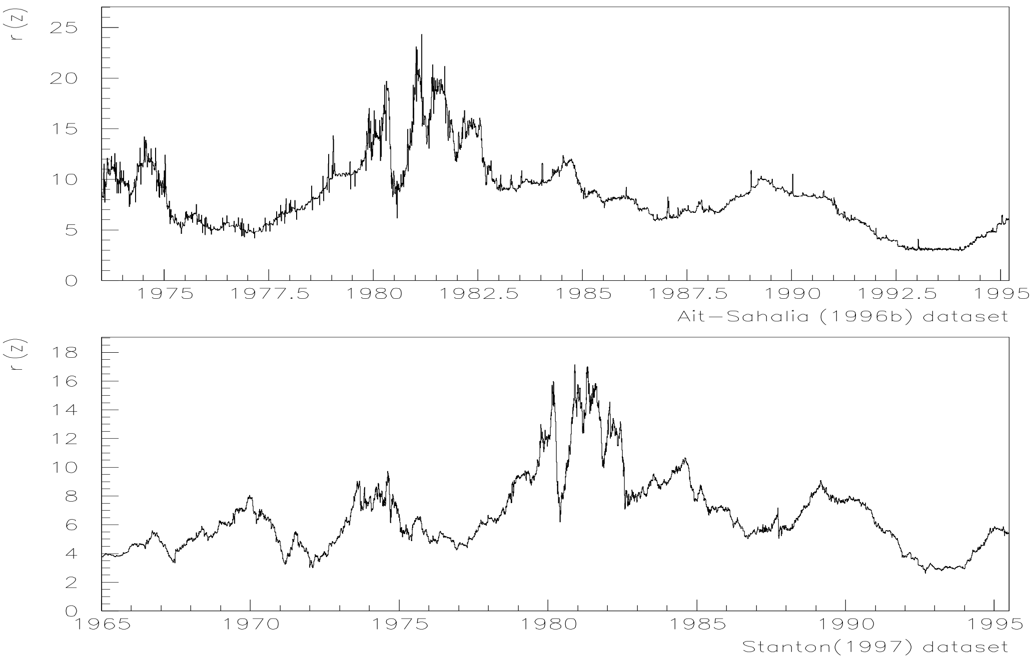

Figure 14: The time series of data used in Ait-Sahalia (1996) (top) and Stanton (1997) (bottom).

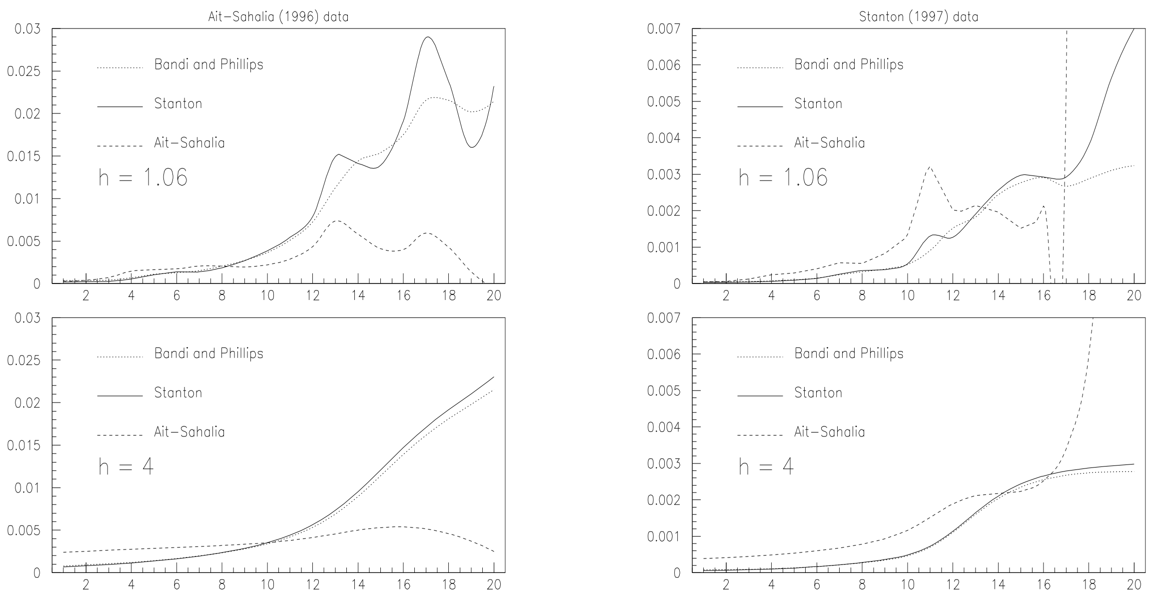

Figure 15: Estimates on data

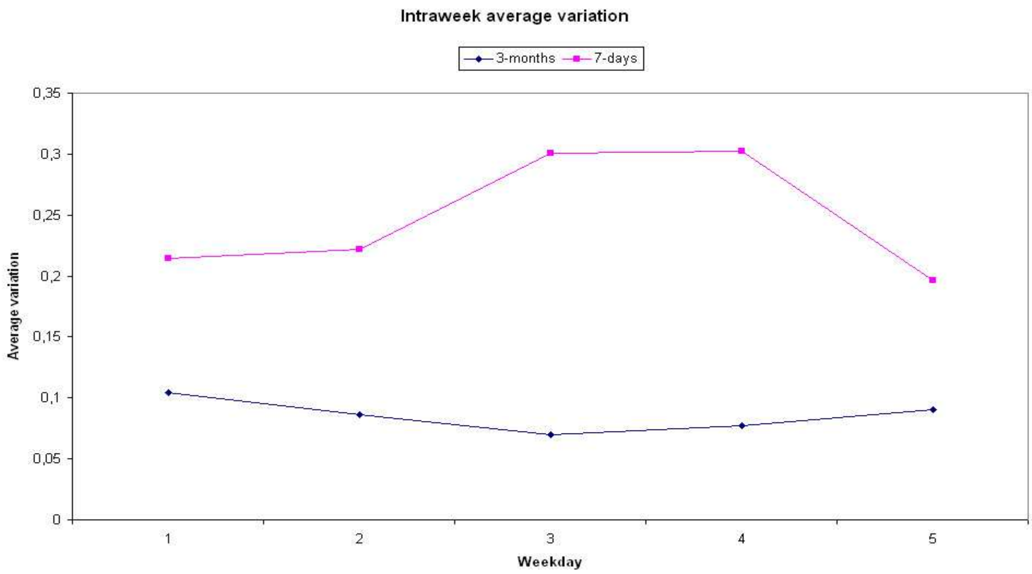

Figure 16: Average absolute rate change in different weekdays. In this Figure, day 1 is Monday, day 5 is Friday.

set. Second, the two time series display the largest changes for different days; the intersection between the dates in the two columns of Table 5 is almost null: there is only one day in common (27 March 1980) and on this day the changes have opposite signs! This suggests that the two interest rate time series are inherently different with regard to large changes, even if their correlation is very high. Finally, from Table 5 it is clear that the largest changes in the 3-months rate occur mainly on Mondays (weekday 2 - 13 out of the 20 largest changes). However, for the 7-day rate the largest changes occur on Wednesdays and Thursdays.

Figure 16 plots the average absolute change in different weekdays for the entire dataset. This figure shows clearly the difference between the two data sets in terms of day-of-the-week effects. The 3-months rate displays the familiar U-shape, see e.g. French and Roll (1986), with higher volatility on Mondays and Fridays, and with small changes within the week. The 7-day rate instead displays an unusual inverse U-shape, with large changes between Wednesdays and Thursdays and the other days. This effect is most likely due to the fact that the 7-day Eurodollar futures is used to manage liquidity by banks, and the EOM (End-Of-Maintenance) day falls exactly on Wednesday. Thus, volatility in the 7-day rate is mostly due to liquidity effects rather than diffusive variations of the short rate. Further, eight of the 20 largest changes are between a Tuesday and a Wednesday or a Wednesday and a Thursday and are immediately reversed the following day, see e.g. February, 3^{rd}, 1981 and July, 1^{st}, 1980 in Table 5.

The volatility of the 7-day rate also displays a monthly seasonal effect. From Figure 17, which

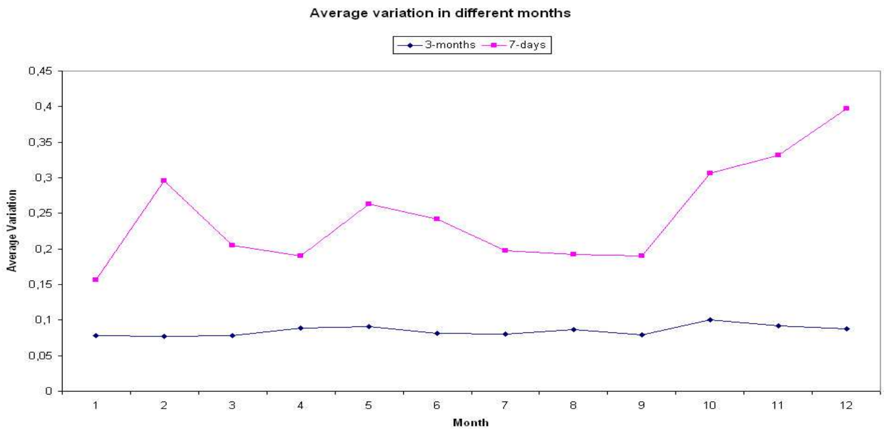

Figure 17: Average absolute daily change along the year in different months.

shows the average absolute daily change by month, we see that, while the volatility of the 3-months rate varies very little by month, the 7-day rate is markedly higher in last quarter and in February than in the remainder of the year.

This simple analysis leads us to the conclusion that the 7-day rate is strongly influenced by calendar and liquidity effects, while the the 3-months rate is not.

Even more importantly – and indeed paradoxically – the spikes present in the Ait-Sahalia data set (which are associated with calendar effects) are responsible for the high level of mean reversion estimated from these data and which is, in turn, necessary for the Ait-Sahalia estimator to be computed. Thus, use of the 7-day Eurodollar deposit rate facilitates estimation of the diffusion coefficient using the Ait-Sahalia estimator, while at the same time data themselves appear inconsistent with a continuous diffusion.

6 Conclusions

Our conclusions regarding the reliability of non-parametric estimators of the diffusion coefficient for the short rate are not encouraging. With the number of observations that is likely to be available, each of the three methods we have investigated – Ait-Sahalia, Stanton and Bandi-Phillips – may produce non-linearities in the relation between conditional variance and the level of rates that are spurious. Applying the methods to data in which the conditional variance is in fact constant, we find that all three methods may, for a single time series of 5,000 observations (i.e., twenty years of daily data), produce a pattern of conditional volatility that is strongly non-linear. The Ait-Sahalia method is particularly prone to this problem but the problem also arises with the other two methods.

When the short rate is generated by a CIR process we find that, with a low level of mean reversion both the Stanton and Bandi-Phillips methods are biased in the tails. This problem arises because, when the data are persistent, it is effectively impossible to obtain reliable estimates of the density non-parametrically with sample sizes that are at all realistic. We also find that, for all three “non-parametric” methods, the choice of bandwidth remains an important parameter that needs to be determined.

Finally, although they overlap substantially in time, the datasets used by Ait-Sahalia and Stanton have entirely different statistical properties and give rise to quite different estimates of the conditional variance of interest rates. This is due to calendar and liquidity effects which affect the 7-day Eurodollar deposit time series.

References

Ait-Sahalia, Y. (1996). Nonparametric pricing of interest rate derivative securities. *Econometrica 64*(3), 527–560.

Andersen, T. and J. Lund (1997). Estimating continuous-time stochastic volatility models of the short-term interest rate. *Journal of Econometrics 77*, 343–377.

Balduzzi, P., S. Das, S. Foresi, and R. Sundaram (1996). A simple approach to three factor affine term structure models. *Journal of Fixed Income 6*, 43–53.

Bandi, F. (2002). Short-term interest rate dynamics: a spatial approach. *Journal of Financial Economics 65*, 73–110.

Bandi, F. and P. Phillips (2003). Fully nonparametric estimation of scalar diffusion models. *Econometrica 71*(1), 241–283.

Brennan, M. J. and E. S. Schwartz (1980). Analysing convertible bonds. *Journal of Financial and Quantitative Analysis 15*, 907–929.

Chan, K., A. Karolyi, F. Longstaff, and A. Sanders (1992). An empirical comparison of alternative models of the short-term interest rate. *Journal of Finance 47*(3), 1209–1227.

Chapman, D. and N. Pearson (2000). Is the short rate drift actually nonlinear? *Journal of Finance 55*(1), 355–388.

Cox, J., J. Ingersoll, and S. Ross (1985). A theory of the term structure of interest rates. *Econometrica 53*, 385–406.

Dai, Q. and K. Singleton (2000). Specification analysis of affine term-structure models. *Journal of Finance 55*(5), 1943–1978.

Duffee, G. (1996). Idiosyncratic variation of Treasury Bill yield spread. *Journal of Finance 51*, 527–552.

Durham, G. (2003). Likelihood-based specification analysis of continuous-time models of the short-term interest rate. *Journal of Financial Economics 70*(2), 463–487.

Florens-Zmirou, D. (1993). On estimating the diffusion coefficient from discrete observations. *Journal of Applied Probability 30*, 790–804.

French, K. and R. Roll (1986). Stock return variances: the arrival of information and the reaction of traders. *Journal of Financial Economics 17*, 5–26.

Hamilton, J. (1996). The daily market for federal funds. *Journal of Political Economy 104*, 26–56.

Jiang, G. and J. Knight (1997). A nonparametric approach to the estimation of diffusion processes, with an application to a short-term interest rate model. *Econometric Theory 13*, 615–645.

Johannes, M. (2004). The statistical and economic role of jumps in continuous-time interest rate models. *Journal of Finance 59*, 227–260.

Pritsker, M. (1998). Nonparametric density estimation and tests of continuous time interest rate models. *Review of Financial Studies 11*(3), 449–487.

Scott, D. (1992). *Multivariate Density Estimation.* John Wiley & Sons.

Stanton, R. (1997). A nonparametric model of term structure dynamics and the market price of interest rate risk. *Journal of Finance 52*, 1973–2002.

Vasicek, O. (1977). An equilibrium characterization of the term structure. *Journal of Financial Economics 5*, 177–188.

Table 1: We simulate $1,000$ replications of sample paths of N observations of the Vasicek model (10), with $\alpha = 8.3\%$, σ is such that the unconditional standard deviation is 3%, and for different values of k $(5, 0.5, 0.05)$ and N (5000, 10000 , 50000). $r(0)$ is sampled from the unconditional distribution. The Table shows the average estimated mean and standard deviation on the simulated paths, together with the 95% and 5% confidence limits of the distribution of the estimated values. The true value of standard deviation is 3%, the true value of mean is 8.3%. The Table shows that the unconditional variance is underestimated for low values of k.

N=5,000						
	average mean (%)	5% lower bound	95% upper bound	average s.d. (%)	5% lower bound	95% upper bound
$k = 0.05$	8.35	5.68	11.24	1.35	0.75	2.23
$k = 0.50$	8.31	6.43	10.37	2.59	1.80	3.50
$k = 5.00$	8.30	7.65	8.99	2.96	2.61	3.32

N=10,000						
	average mean (%)	5% lower bound	95% upper bound	average s.d. (%)	5% lower bound	95% upper bound
$k = 0.05$	8.31	5.41	11.32	1.73	1.04	2.76
$k = 0.50$	8.30	6.74	9.85	2.77	2.12	3.51
$k = 5.00$	8.30	7.79	8.80	2.98	2.73	3.23

N=50,000						
	average mean (%)	5% lower bound	95% upper bound	average s.d. (%)	5% lower bound	95% upper bound
$k = 0.05$	8.34	6.31	10.28	2.58	1.84	3.49
$k = 0.50$	8.31	7.60	8.99	2.95	2.61	3.28
$k = 5.00$	8.30	8.08	8.52	3.00	2.89	3.11

Table 2: Estimate of the mean reversion parameter k in the literature.

Paper	Estimate	Model	Method
Ait-Sahalia (1996)	0.978	$dr = (\alpha + kr)dt + \sigma(r)dW(t)$	FGLS
Chan et al. (1992)	0.1779	$dr = (\alpha + kr)dt + \sigma r^{\gamma} dW(t)$	GMM
Andersen and Lund (1997)	0.173	stochastic volatility	EMM
Durham (2003)	0.1875	$dr = (\alpha + kr)dt + \sqrt{\beta_1 + \beta_2 r} dW(t)$	Max. Lik.
"	0.1049	$dr = (\alpha + kr)dt + \beta_1 r^{\beta_2} dW(t)$	Max. Lik.
"	0.1056	$dr = (\alpha + kr)dt + \sqrt{\beta_1 + \beta_2 r + \beta_3 r^{\beta_4}} dW(t)$	Max. Lik.

Table 3: Ait-Sahalia (1996) dataset

	OLS	GLS	FGLS*
α	0.083082	0.082536	0.084387
k	1.6088	0.94014	0.97788

Table 4: Stanton (1997) dataset

	OLS	GLS
α	0.068423	0.068009
k	0.25615	0.32935

Table 5: Ranks the twenty largest interest rate daily changes, in absolute value, of the 3-months T-bill rate and of the 7-days Eurodollar deposit rate. In the column "weekday", day 2 is Monday, day 6 is Friday.

3-months			7-days		
Day	Change	Weekday	Day	Change	Weekday
4-May-81	1,34%	2	3-Feb-81	4,37%	3
19-Dec-80	-1,27%	6	18-Dec-80	4,31%	5
1-Feb-82	1,16%	2	4-Feb-81	-4,05%	4
5-Jan-81	-1,13%	2	1-Nov-79	-3,93%	5
9-Oct-79	1,12%	3	18-Nov-80	3,68%	3
20-Jul-81	1,07%	2	1-Jul-80	-3,67%	3
6-Apr-81	1,03%	2	2-Jul-80	3,55%	4
3-Mar-80	0,99%	2	27-Mar-80	3,49%	5
5-Dec-80	0,99%	6	23-Dec-74	3,43%	2
22-Feb-82	-0,94%	2	27-Dec-78	-3,42%	4
7-Jan-81	0,92%	4	7-Feb-80	-3,42%	5
16-Sep-74	-0,86%	2	19-Nov-80	-3,30%	4
28-Apr-80	-0,81%	2	17-Dec-80	-3,29%	4
15-Jun-81	-0,81%	2	14-Nov-79	-2,98%	4
2-Aug-82	-0,81%	2	8-Apr-81	-2,92%	4
1-Apr-80	0,80%	3	21-May-81	2,78%	5
27-Mar-80	-0,80%	5	27-Aug-81	-2,78%	5
17-Nov-80	0,80%	2	25-Nov-80	2,66%	3
24-Mar-80	0,78%	2	18-Dec-74	-2,60%	4
15-Nov-73	-0,77%	5	18-Mar-80	-2,60%	3

NON-TECHNICAL SUMMARY

Volatility estimation, here named as the *diffusion coefficient* estimation, is a crucial topic in the financial econometrics literature for all quoted assets including bonds. Recently, many statistical methods have been introduced in which volatility is modeled as a function of the state variable (here, the short term interest rate), and this function is estimated without using a specific model. This methodology is called non-parametric. In this paper, we study the performance of non-parametric methods in estimating the volatility of the short term interest rate. In particular we focus on three approaches, introduced by Ait-Sahalia (1996); Stanton (1997); Bandi and Phillips (2003). Some of our results are technical in nature. We show that all the considered estimators depend crucially on the degree of mean reversion, the length of the data set under study and the bandwidth parameter used for non-parametric estimation. We also show that, for estimating the short rate, it is advisable not to use interest rate instruments with too short maturities.

Most effectively, we show that for realistic sample sizes and typical interest rate time series, the Ait-Sahalia estimator cannot estimate the interest rate variance, while the Stanton and Bandi-Phillips estimator, which provide almost identical numerical results, deliver more reliable estimates. In summary, we show that non-parametric methods offer great flexibility, but they have to be implemented carefully to get reliable results.