Quantitative Finance, Vol. 5, No. 3, June 2005, 237–244

Waiting for returns: using space–time duality to calibrate financial diffusions

MARK KAMSTRA and MOSHE A. MILEVSKY*

Schulich School of Business, York University, North York, Ontario M3J 1P3, Canada

(*Received 17 February 2004; in final form 14 April 2005*)

1. Introduction

Historically, tests of the Geometric Brownian Motion (GBM) model for security prices—and for that matter any diffusion process—have been performed by selecting a fixed interval of time (one day, one week, one month) Δt and then using the increments in logarithmic price $\Delta \ln[P]$ over the predetermined Δt. Under the classical specification of the GBM model, the logarithmic price increments $\Delta \ln[P]$ should be statistically independent from each other and these increments $\Delta \ln[P]$ should be Normally distributed with a mean and variance that is proportional to the time increment Δt. This approach has a long tradition in finance. Research done in the 1950s by Kendall (1953) and Osborne (1959) as well as the work by Fama (1970) all the way through to the contemporary work of Campbell *et al.* (1997) based on Lo and MacKinlay (1988) focuses on a particular time interval Δt.

Thus, for example, Kendall (1953) looked at a time increment of $\Delta t =$ one week on the New York Stock Exchange, and concluded that the logarithmic price increments $\Delta \ln[P]$ have a statistically insignificant serial correlation in addition to being (approximately) normally distributed. In another study, Fama (1970) looked at the 30 Dow Jones Industrial stocks with a $\Delta t =$ one day, and concluded that there is a statistically significant positive serial correlation in logarithmic price increments $\Delta \ln[P]$. Poterba and Summers (1988) found that for a $\Delta t =$ three years, the logarithmic price increments $\Delta \ln[P]$ exhibit a statistically significant negative serial correlation which translates into a long-term mean reversion in prices. Among the many recent studies that document violations of the GBM by looking at the time series properties of returns to various financial instruments are Bakshi *et al.* (2000), Bollerslev *et al.* (1992), Cont (2001), Cont and da Fonseca (2002), and Nelson (1991).

Nevertheless, the broad unifying methodology of this large literature is to select a time interval and then investigate price increments *vis a vis* that time interval. Hence, it is quite common to hear that the GBM-Lognormal model is rejected for hourly data while it is accepted for monthly data but rejected again for yearly data or some combination thereof. In fact, this was the recent conclusion of Levy and Duchin (2004).

In this paper we propose an alternative way of thinking about the appropriate distribution. We investigate the GBM model for fixed $\Delta \ln[P]$ intervals as opposed to fixed Δt intervals. In other words, we start at the beginning of a time series and judiciously select a price increment $\Delta \ln[P] = d$ (for example, 1%) and then measure the amount of time τ_1 it takes the security to move the pre-specified quantity. After the security has moved by $\Delta \ln[P] = d$, we measure the time τ_2 at which the security moves an additional $\Delta \ln[P] = d$ and so on and so forth. The final result is a collection of time increments $(\tau_{i+1} - \tau_i)$ for each pre-specified $\Delta \ln[P]$. We then compare (statistically) the empirical distribution of the $(\tau_{i+1} - \tau_i)$'s to the theoretical distribution they should obey under the GBM model. If, indeed, the price increments are normal, then the $(\tau_{i+1} - \tau_i)$'s—for each particular $\Delta \ln[P]$—should obey the Inverse Gaussian (or Wald) distribution as a result of the Space–Time duality that exists for Brownian motion. We select an entire spectrum of $\Delta \ln[P]$'s (for example, from 1% all the way to 15%) and then extract the appropriate sample of τ_i's (for each $\Delta \ln[P]$) so as to measure goodness of fit and estimate confidence intervals for the implied drift and diffusion coefficients. Our approach should not be confused with, and is very different from, the paradigm of spectral

*Corresponding author. Email: milevsky@yorku.ca

Quantitative Finance
ISSN 1469–7688 print/ISSN 1469–7696 online © 2005 Taylor & Francis
http://www.tandf.co.uk/journals
DOI: 10.1080/14697680500167976

analysis. Spectral analysis attempts to uncover *cycles* in the underlying process by fitting sine and cosine functions to the data. See, for example, Granger and Morgenstern (1963).

To our knowledge, this is the first study of its kind which attempts to verify a particular parametric form and estimate parameters via this duality methodology. This study will also shed light on the persistence of trend as a function of price momentum as well as the velocity of the price process. If mean reversion behaviour exists in the S&P 500, larger price increments and their respective collection of τ_i's will exhibit smaller implied drifts and diffusion coefficients as well as a 'poorer fit' to the Inverse Gaussian distribution. We find the reverse, with larger drift and diffusion coefficients and a better fit to the Inverse Gaussian distribution with larger increments.

In addition—although we do not pursue this directly within the paper—investigating the data via the space/time dual has implications to option pricing, since the optimal exercise policy of an American option revolves around the first passage time (FPT) to a given curve in space/time. The option can therefore be expressed and calibrated to a function of the FPT. We pursue this in a follow-up paper by Kamstra and Milevsky (2005).

The rest of this paper is organized as follows. In order for the paper to be self-contained, section 2 summarizes the theoretical properties of the Inverse Gaussian distribution and demonstrates its relationship to the first passage time of a Brownian Motion. Section 3 discusses the issue of parameter and confidence interval estimation for the Inverse Gaussian distribution *vis a vis* the first passage time distribution. The empirical results using the SP 500 as a test case are tabulated in section 4. In that section we estimate the implied drift and diffusion coefficients by implementing the algorithm developed in the previous sections. Conclusions and directions for further research are offered in section 5.

2. The first passage time distribution

Let P_t denote the price of a security or index at time $t \geq 0$. The standard (a.k.a. Black–Scholes) assumption in finance is to assume that the dynamics of P_t obey the following stochastic differential equation:

$$d(\ln[P_t]) = \nu\, dt + \sigma\, dB_t. \tag{1}$$

The parameter ν is often called the geometric mean return or growth rate so that $MED[P_1] = P_0\, e^{\nu}$ and the expected value is $E[P_1] = P_0\, e^{\nu+0.5\sigma^2}$. Either way, we let $X_t = \ln[P_t]$, which simplifies the main diffusion process to

$$dX_t = \nu\, dt + \sigma\, dB_t. \tag{2}$$

Thus, the logarithm of security (or index) prices obeys a non-standard Brownian motion with drift. Let us now start at some point in time denoted by zero, such that $X_0 = x_0$. Furthermore, choose an increment denoted by d. Let

$$\tau_1 = \inf\{s; X_s \geq x_0 + d\}. \tag{3}$$

Likewise, let

$$\tau_2 = \inf\{s; X_s \geq X_{\tau_1} + d\}. \tag{4}$$

Further,

$$\tau_3 = \inf\{s; X_s \geq X_{\tau_2} + d\}. \tag{5}$$

Finally,

$$\tau_i = \inf\{s; X_s \geq X_{\tau_{i-1}} + d\}. \tag{6}$$

Thus, $(\tau_{i+1} - \tau_i)$ is the sequence of first passage times of the stochastic process X_t to the barriers demarcated by increments of d. It corresponds to the random amount of time it takes the stochastic process P_t to move by $e^d - 1 = D$ percent. In can be shown (see Seshadri (1993) or Wasan (1969)) that the probability density function of the time increments is Inverse Gaussian distributed. The probability density function (pdf) of the Inverse Gaussian (IG) random variable is a two-parameter (β, λ) function that can be expressed as follows:

$$g(t \mid \beta, \lambda) = \frac{\sqrt{\lambda}}{\sqrt{2\pi t^3}} \exp\left(-\frac{\lambda(t-\beta)^2}{2\beta^2 t}\right), \quad t > 0. \tag{7}$$

The pdf is defined for $\beta > 0$ and $\lambda > 0$. The mean (expected value) of the Inverse Gaussian random variable is β, while the variance is β^3/λ. The cumulative distribution function (c.d.f.), which we denote by $G(T \mid \beta, \lambda)$, of the Inverse Gaussian random variable cannot be expressed in closed form; however, it can be expressed as a function of the c.d.f. of the standard normal random variable $\Phi[x]$ in the following elegant way (see Chhikara and Folks (1989) for details):

$$G(T \mid \beta, \lambda) = \int_0^T g(t \mid \beta, \lambda)\, dt = \Phi\left[\sqrt{\frac{\lambda}{T}}\left(\frac{T}{\beta} - 1\right)\right] + e^{(2\lambda/\beta)}\Phi\left[-\sqrt{\frac{\lambda}{T}}\left(1 + \frac{T}{\beta}\right)\right]. \tag{8}$$

For the first passage time, the parameters will be $\beta = d/\nu$ and $\lambda = d^2/\sigma^2$. Thus, the expected amount of time it will take the stochastic process P_t to move D percent is d/ν, and the variance in the amount of time will be $d\sigma^2/\nu^3$.

3. Parameter estimation

The Maximum Likelihood Estimate for the value of β is

$$\hat{\beta} = \frac{1}{n}\sum_{i=1}^{n} \tau_i. \tag{9}$$

It is also an unbiased estimate for the value of β. The UMVUE for λ is

$$\hat{\lambda} = \frac{n-1}{\sum_{i=1}^{n}[(1/\tau_i) - (1/\hat{\beta})]}. \tag{10}$$

See Wasan (1969) for a derivation of the confidence intervals for β, λ.

A $(1-\alpha)$ percent confidence interval for the value of λ is

$$\left(\frac{\hat{\lambda}}{n-1}\cdot\chi^2_{\alpha/2} \le \lambda \le \frac{\hat{\lambda}}{n-1}\cdot\chi^2_{1-\alpha/2}\right), \tag{11}$$

where $\chi^2_{\alpha/2}$ denotes the value from the chi square distribution with $n-1$ degrees of freedom. Since $\lambda = d^2/\sigma^2$, we can obtain a $(1-\alpha)$ percent confidence interval for the value of σ:

$$\left(\frac{d}{\sqrt{[\hat{\lambda}/(n-1)]\cdot\chi^2_{1-\alpha/2}}} \le \sigma \le \frac{d}{\sqrt{[\hat{\lambda}/(n-1)]\cdot\chi^2_{\alpha/2}}}\right). \tag{12}$$

Likewise, a $(1-\alpha)$ percent confidence interval for β is (where $t_{1-\alpha/2}$ denotes the value from the student t distribution, with n degrees of freedom)

$$\left(\hat{\beta}\left[1+\sqrt{\frac{\hat{\beta}}{n\hat{\lambda}}}\cdot t_{1-\alpha/2}\right]^{-1} \le \beta \le \hat{\beta}\left[1-\sqrt{\frac{\hat{\beta}}{n\hat{\lambda}}}\cdot t_{1-\alpha/2}\right]^{-1}\right), \tag{13}$$

provided that $(\hat{\beta}/n\hat{\lambda})^{1/2}\cdot t_{1-\alpha/2} < 1$. Otherwise, the confidence interval is

$$\left(\hat{\beta}\left[1+\sqrt{\frac{\hat{\beta}}{n\hat{\lambda}}}\cdot t_{1-\alpha/2}\right]^{-1} \le \beta \le \infty\right). \tag{14}$$

Now, since $\beta = d/\nu$, by inverting the confidence interval for β we can obtain a C.I. for ν,

$$\left(d\left[1-\sqrt{\frac{\hat{\beta}}{n\hat{\lambda}}}\cdot t_{1-\alpha/2}\right]\hat{\beta}^{-1} \le \nu \le d\left[1+\sqrt{\frac{\hat{\beta}}{n\hat{\lambda}}}\cdot t_{1-\alpha/2}\right]\hat{\beta}^{-1}\right), \tag{15}$$

provided that $(\hat{\beta}/n\hat{\lambda})^{1/2}\cdot t_{1-\alpha/2} < 1$. Otherwise, the confidence interval for ν is

$$\left(0 \le \nu \le d\left[1+\sqrt{\frac{\hat{\beta}}{n\hat{\lambda}}}\cdot t_{1-\alpha/2}\right]\hat{\beta}^{-1}\right). \tag{16}$$

In general, small data sets tend to result in one-sided confidence intervals.

4. Empirical results

Using the principles set out in the previous section we can now derive point estimates and confidence intervals for the values of ν, σ—the expected growth rate and volatility of returns—as implied from β, λ from the first passage time data. We used the daily closing prices on the S&P 500 cash index, for a period of time spanning January 1952 to December 2003, resulting in 13 109 data points for the stochastic process P_t. We then computed the amount of time it takes the S&P 500 to move a prespecified percentage D. Under the Null Hypothesis that P_t obeys a geometric Brownian motion, the collection of these *time* increments should obey an Inverse Gaussian distribution.

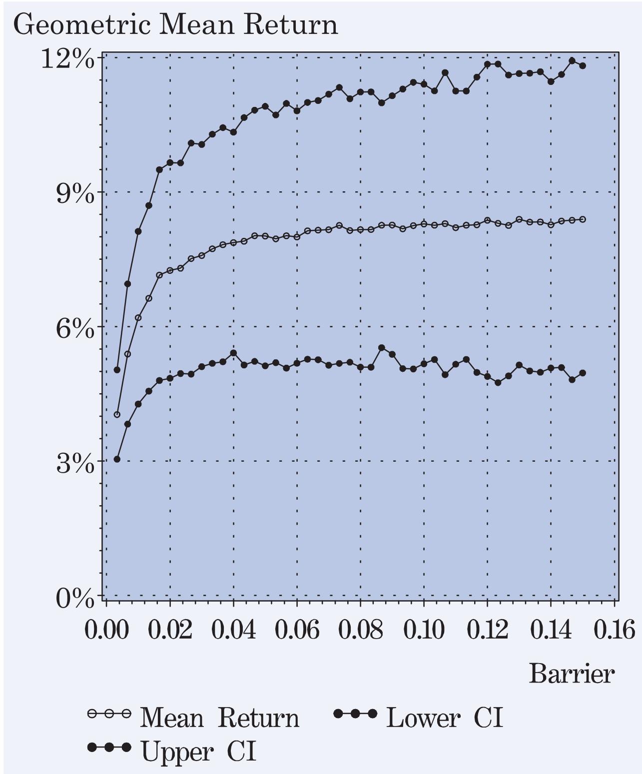

Figure 1. S&P 500 1952–2004 return mean estimate, with confidence interval.

Figure 1 shows a graphical representation of the confidence interval for ν as a function of the percent increment. The line indicated with circles is the point estimate of the expected return, while the solid dotted lines indicate the plus or minus two standard deviation confidence about the mean. As one can see, larger increments in space imply a larger range for the ν of the diffusion process and somewhat larger point estimates. Figure 2 shows a graphical representation of the confidence interval for σ as a function of the percent increment, again with the line of circles representing the point estimate, now of the volatility, and the solid dotted lines the confidence interval about that point estimate. In this case we obtain a more dramatic result with larger increments in space, implying a much larger value for the σ.

Recall that in theory—under the constant parameter GBM assumption—both graphs should be flat to within statistical variations and the size of the data set. It is important to note, however, that the kink in this graph may be due in part to sample-size truncation issues involved with using daily returns. After all, if we are searching for 1% moves and are only looking at daily numbers there is a (strong) chance that the S&P 500 moved up by more than 1% during the course of the day, and then reversed itself to close at a less-than-1% change. The cumulative effect of this truncation is that we (erroneously) conclude the market did not increase by 1%—when it did—and thus the underlying drift is not as high. This has far reaching implications beyond just intra-day moves. For example, the market might take 3.2 trading days to increase by 1%, but in our data set it will be recorded as (much longer) four trading days,

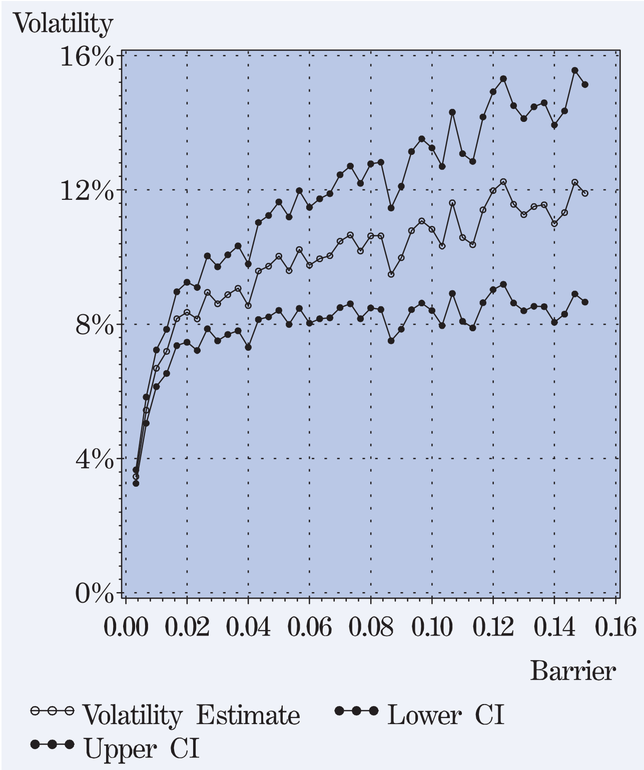

Figure 2. S&P 500 1952–2004 return volatility estimate, with confidence interval.

which creates an artificial downward bias on the implied ν and σ. Of course, as we increase the size of D, the extent to which this occurs is much less, since it is highly unlikely that we missed a 10% move in the S&P 500 because we only examined daily closing prices.

Table 1 displays the point estimates for the parameters $\beta, \lambda, \nu, \sigma$ and associated standard estimates, together with the number of data points that were observed. Thus, for example, there were only 27 movements of 15% between 1952 and 2004. This small number may limit the inferences we can draw from this data set from 15% moves. Table 2 displays the 95% confidence intervals for the above-mentioned parameter values. Recall that under the *constant parameter GBM* assumption the estimated values for β, λ should only depend on the (logarithmic) space increment d, via the relationship $\beta = d/\nu$ and $\lambda = d^2/\sigma^2$. Thus, if the ν, σ for the return generating process are truly constant, then, for example, the β value estimated for $d = 2\%$ increments should be twice the β value estimated at $d = 1\%$ increments. As tables 1 and 2 indicate this is not the case and the parameter estimates are not scaling by d and d^2. Once again, this is an indication that the underlying generating process is likely not GBM with constant parameters, though this result may also be due in part to sample-size truncation issues involved with using daily returns, as discussed above.

Table 3 displays the results from performing a Kolmogorov–Smirinov (KS) test for goodness-of-fit of the crossing time intervals to an Inverse Gaussian distribution. It is interesting to note that, within any given increment d above the 1% case, the data does not fail a KS test for goodness-of-fit to an Inverse Gaussian distribution. And, while some of this might be due to the low power of the KS test, a casual examination of the data shows that the plots of the CDF of the data versus the Inverse Gaussian distribution (figures 3–9) reveals a good match between the empirical and theoretical distribution, in particular where the data is most dense, up to the 70th percentile or so of the cumulative.

5. Extension to non-lognormal returns

As mentioned earlier, the First Passage Time (FPT) distribution of the logarithmic prices X_t to a level D will satisfy an Inverse Gaussian (I.G.) distribution *if and only if* the logarithmic prices themselves are Normally distributed. Indeed, when the process X_t is something other than a non-standard Brownian motion, i.e. when e^{X_t} is no longer a geometric Brownian motion, the collection of time increments τ_i will not be I.G. and, although it is beyond the scope of this paper to derive and present FPT distributions for all possible parameterization of X_t, in this section we briefly describe how one could go about deriving a related probability for a general process and thus use the space–time duality method for investigating more general diffusions.

In order to adhere to common notation and terminology in the continuous-time finance literature, assume the price process itself obeys the following one-dimensional diffusion:

$$\mathrm{d}Y_t = \nu(Y_t, t)Y_t\,\mathrm{d}t + \xi(Y_t, t)Y_t\,\mathrm{d}B_t, \qquad Y_0 = y. \tag{17}$$

This representation covers our earlier geometric Brownian motion—when $\nu(Y_t, t) = \nu + 0.5\sigma^2$ and $\xi(Y_t, t) = \sigma$ are constants—as well as more general mean reverting and time-dependent cases. In this case, the probability $H(y, t)$ that Y_t hits or breaches a level denoted by D during a time period denoted by t, satisfies a so-called Kolmogorov partial differential equation (PDE), denoted by

$$\frac{\partial H(y,t)}{\partial t} + \nu(y,t)y\frac{\partial H(y,t)}{\partial y} + \frac{1}{2}\xi^2(y,t)y^2\frac{\partial H^2(y,t)}{\partial y^2} = 0, \tag{18}$$

where the current position of the process (x, s) is implicit in the notation, with a terminal condition $H(D, s) = 1$, if $D > y$ and zero otherwise as well as a boundary condition $H(D, t) = 1$ if $y \le D$. See the book by Oksendal (2004) for more details. Thus, for example, under a particular parameterization of equation (17), we can solve for the probability of observing a $D = 1\%$ move within a $s = 1$-day period. We can then compare the *theoretical probability* dictated by equation (18) against the *observed frequency* of 1% moves in one day. And although this is not exactly the FPT density, we can employ standard goodness-of-fit methods to test whether in fact the original (dual) diffusion Y_t satisfies the postulated process in question.

Table 1. S&P 500 annualized percentage returns 1952/01/01–2003/12/31.

Barrier (%)	n	ν	(std)	σ	(std)	β	(std)	λ	(std)
1	299	6.2	(0.96)	6.69	(0.27)	0.161	(0.025)	0.022	(0.002)
2	175	7.26	(1.2)	8.35	(0.45)	0.276	(0.046)	0.057	(0.006)
3	122	7.58	(1.24)	8.61	(0.55)	0.396	(0.065)	0.121	(0.016)
4	95	7.87	(1.23)	8.55	(0.62)	0.508	(0.079)	0.219	(0.032)
5	77	8.02	(1.45)	10.03	(0.81)	0.624	(0.113)	0.249	(0.04)
6	64	8	(1.41)	9.76	(0.86)	0.751	(0.132)	0.378	(0.067)
7	56	8.16	(1.51)	10.48	(0.99)	0.858	(0.159)	0.447	(0.084)
8	49	8.16	(1.53)	10.63	(1.07)	0.98	(0.185)	0.565	(0.114)
9	44	8.26	(1.44)	9.98	(1.06)	1.089	(0.19)	0.813	(0.173)
10	40	8.29	(1.56)	10.83	(1.21)	1.206	(0.227)	0.853	(0.191)
11	36	8.21	(1.52)	10.58	(1.25)	1.342	(0.25)	1.08	(0.255)
12	33	8.37	(1.74)	11.98	(1.47)	1.433	(0.298)	1.002	(0.247)
13	31	8.39	(1.63)	11.26	(1.43)	1.55	(0.301)	1.332	(0.338)
14	28	8.27	(1.6)	11	(1.47)	1.694	(0.328)	1.621	(0.433)
15	27	8.39	(1.71)	11.9	(1.62)	1.787	(0.365)	1.59	(0.433)

Using daily returns from S&P 500 for the period 1950 to 2004, the table displays point estimates for the parameters β, λ, ν, σ. Note that the β, λ parameters are estimated directly from the data, while the ν, σ are 'solved' by the analytic relationship between the Inverse Gaussian and Normal distribution.

Table 2. S&P 500 95% confidence interval 1952/01/01–2003/12/31.

Barrier (%)	n	ν	σ	β	λ
1	299	(4.32 8.08)	(6.16 7.22)	(0.11 0.21)	(0.02 0.03)
2	175	(4.91 9.61)	(7.47 9.23)	(0.19 0.37)	(0.05 0.07)
3	122	(5.15 10.01)	(7.53 9.69)	(0.27 0.52)	(0.09 0.15)
4	95	(5.46 10.28)	(7.33 9.77)	(0.35 0.66)	(0.16 0.28)
5	77	(5.18 10.86)	(8.44 11.62)	(0.4 0.85)	(0.17 0.33)
6	64	(5.24 10.76)	(8.07 11.45)	(0.49 1.01)	(0.25 0.51)
7	56	(5.2 11.12)	(8.54 12.42)	(0.55 1.17)	(0.28 0.61)
8	49	(5.16 11.16)	(8.53 12.73)	(0.62 1.34)	(0.34 0.79)
9	44	(5.44 11.08)	(7.9 12.06)	(0.72 1.46)	(0.47 1.15)
10	40	(5.23 11.35)	(8.46 13.2)	(0.76 1.65)	(0.48 1.23)
11	36	(5.23 11.19)	(8.13 13.03)	(0.85 1.83)	(0.58 1.58)
12	33	(4.96 11.78)	(9.1 14.86)	(0.85 2.02)	(0.52 1.49)
13	31	(5.2 11.58)	(8.46 14.06)	(0.96 2.14)	(0.67 1.99)
14	28	(5.13 11.41)	(8.12 13.88)	(1.05 2.34)	(0.77 2.47)
15	27	(5.04 11.74)	(8.72 15.08)	(1.07 2.5)	(0.74 2.44)

The table displays the 95% confidence interval for the estimated parameters β, λ, ν, σ.

In some cases, equation (18) can be solved analytically, as we implicitly did earlier in the paper. Of course, under the most general cases for $\nu(y, t)$ and $\xi(y, t)$, one must resort to numerical methods. Nevertheless, it is possible to obtain the hitting/crossing probabilities for processes other than simple Brownian motions which opens the door for an alternative method of calibrating and testing the return generating process for investment returns.

6. Conclusions

We have proposed an alternative method for calibrating financial diffusions. We choose a specific increment in price space, say a 1% return barrier, and examine the amount of time it takes the stochastic process to move the predetermined increment. This is instead of focusing on a particular increment in time—such as an hour, day or month—as do most conventional estimation procedures. This methodology benefits from its ability to capture changes in distribution that depend on the price (space) increment in question. We also believe this approach better fits the perspective and needs of investors who are interested in how long they will have to wait in order to achieve pre-specified target returns. Our empirical results re-enforce previous results obtained in the literature that the stochastic price process for S&P 500 equity returns does not conform to the standard geometric Brownian motion (GBM) model as evidenced by the fact that our implied growth and volatility rates are not constant. Interestingly, we do find a reasonably good fit of the GBM to first passage time data for any given fixed barrier, but these parameters are unstable across different return barriers. In other words, if we only had access to historical data for how long it took the S&P 500 to grow x%—as opposed to the daily or monthly

Table 3. S&P 500 1952/01/01–2003/12/31.

Barrier (%)	Data points	K.S. value	R/NR at 10% sig.
1	299	1.298	Reject
2	175	0.783	Do not reject
3	122	0.684	Do not reject
4	95	0.784	Do not reject
5	77	0.611	Do not reject
6	64	0.874	Do not reject
7	56	0.609	Do not reject
8	49	0.511	Do not reject
9	44	0.633	Do not reject
10	40	0.549	Do not reject
11	36	0.629	Do not reject
12	33	0.671	Do not reject
13	31	0.576	Do not reject
14	28	0.675	Do not reject
15	27	0.570	Do not reject

The table displays results from a Kolmogorov–Smirinov (K.S.) goodness-of-fit test of the data—for each level of D—against an Inverse Gaussian (I.G.) distribution. The Null Hypothesis for our K.S. test is that the data was generated from an I.G. distribution with β, λ parameters specified in table 2. The Null Hypothesis is rejected if the test statistic is 'too large', which means that the distance between the empirical CDF and candidate CDF are 'too far' from each other. At the 5% significance level the critical value of the K.S. statistic is approximately *1.358* and at the 10% significance the critical value is approximately *1.223*.

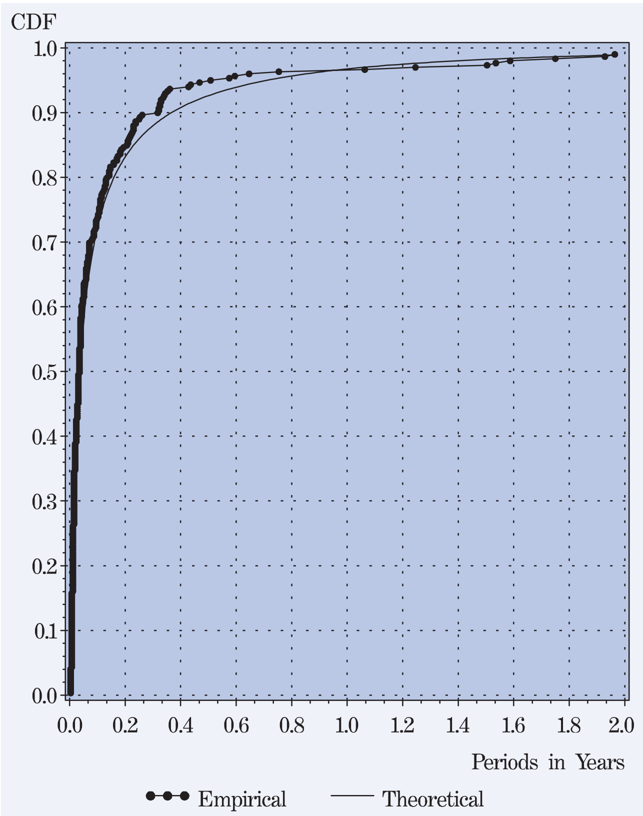

Figure 3. Barrier = 1%.

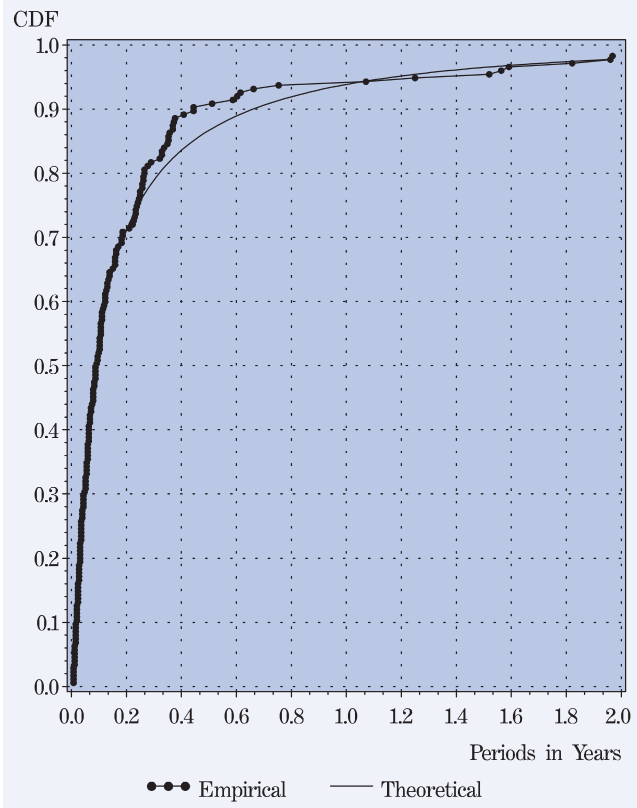

Figure 4. Barrier = 2%.

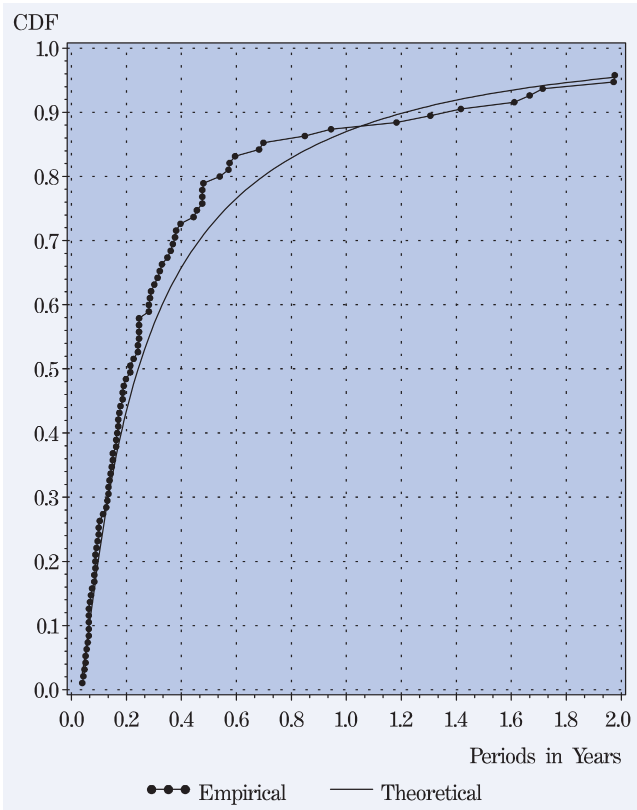

Figure 5. Barrier = 3%.

returns—we could not reject the Null Hypothesis that equity returns are lognormally distributed. It is only when we compare the implied parameters across price increments that the GBM model fails. And, although some of this instability may come from the coarseness of our data, measured daily, it is unlikely to be solely due to this truncation time issue since this effect persists at larger increments as well.

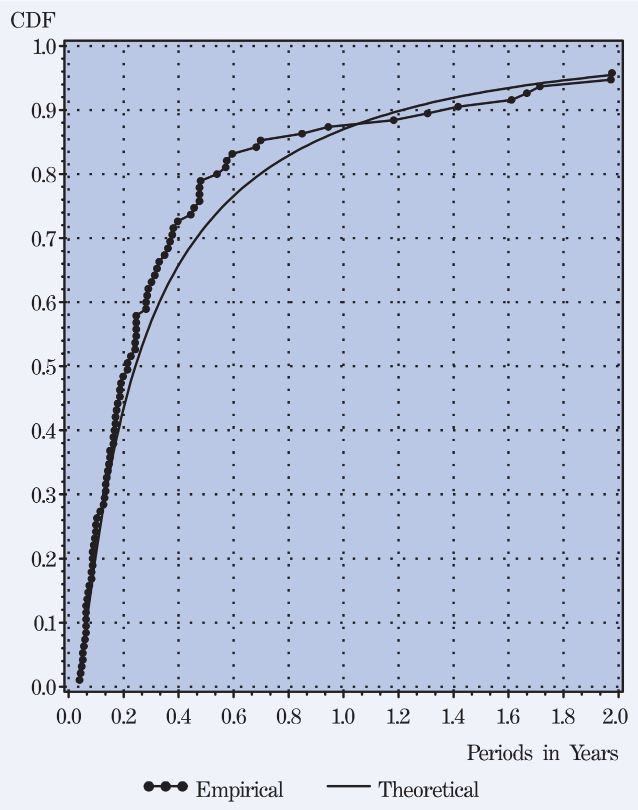

Figure 6. Barrier = 4%.

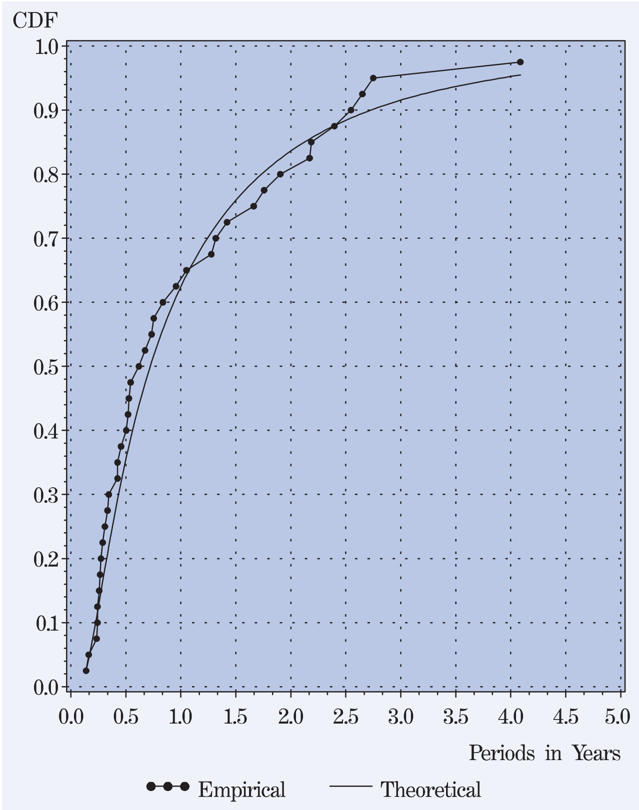

Figure 8. Barrier = 10%.

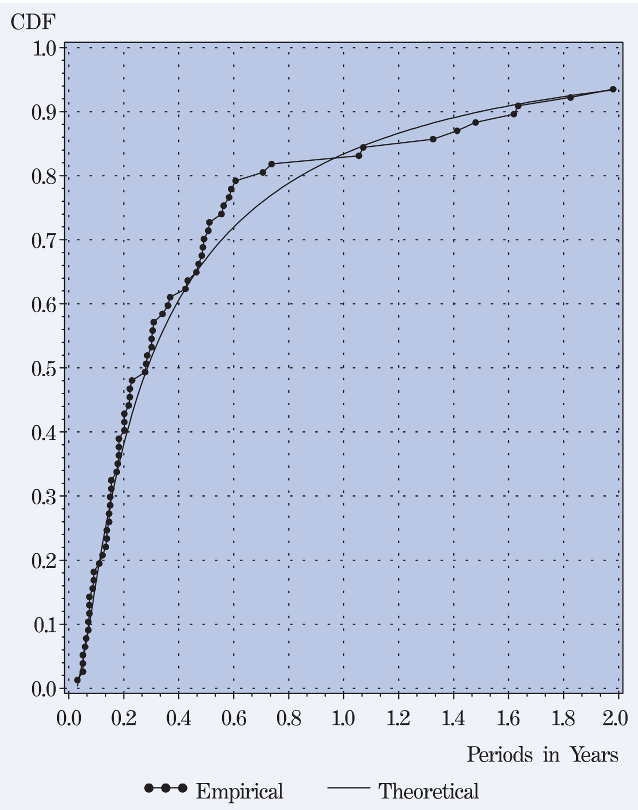

Figure 7. Barrier = 5%.

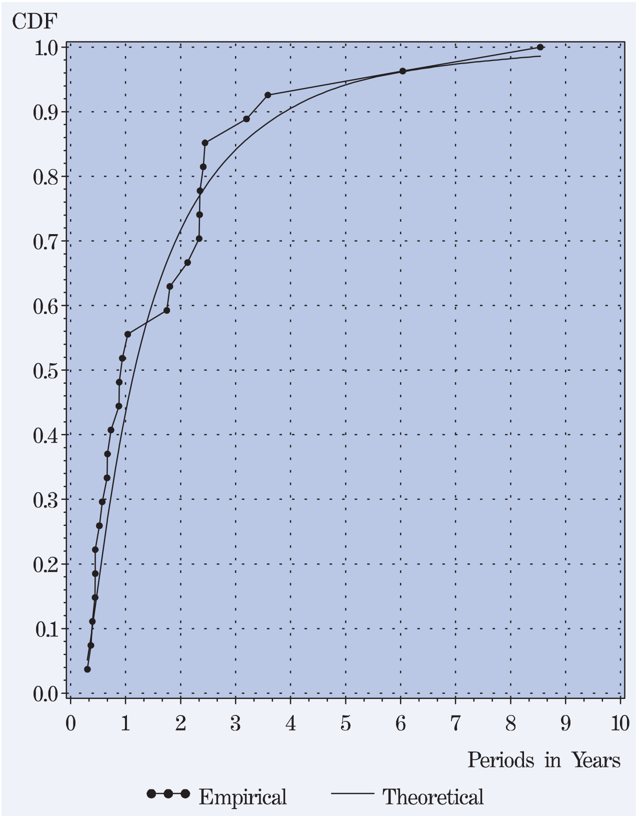

Figure 9. Barrier = 15%.

This study also sheds light on mean reversion in returns. If mean reversion behaviour exists in the S&P 500, larger price increments and their respective collection of first passage times should exhibit smaller implied drifts and diffusion coefficients as well as a 'poorer fit' to the Inverse Gaussian distribution. We find the reverse, with larger drift and diffusion coefficients and a better fit to the Inverse Gaussian distribution with larger increments.

Further research entails calibrating and testing first passage times for alternative prices processes—such as currencies, commodities and interest rates—at higher frequency and in particular on individual stocks. The same principle of space time duality can be used to derive the distribution of first passage times for other stochastic processes. Along these lines, the authors (Kamstra and Milevsky 2005) are currently working on classifying the FPT distribution for processes such as stochastic volatility and mean reverting diffusions with applications to American option pricing. Indeed, even if returns are generated by infinite variance stable distributions, as originally argued by Mandelbrot (1963), then the finite variance first passage times could be analysed instead of the actual returns.

Acknowledgements

The authors would like to thank Eliezer Prisman, Tom Salisbury, Georges Monette, Helen Massam, Tim Krehbiel, Peter Carr and participants at the York University S.S.B. seminar series in Finance and the participants at the Eastern Finance Association's annual meeting as well as the editor and anonymous referee for helpful comments and discussions.

References

Bakshi, G., Cao, C. and Chen, Z., Do call prices and the underlying stock always move in the same direction? *Review of Financial Studies*, 2000, **13**, 549–584.

Bollerslev, T., Chou, R.Y. and Kroner, K.F., ARCH modeling in finance: a review of the theory and empirical evidence. *Journal of Econometrics*, 1992, **52**, 5–59.

Campbell, J.Y., Lo, A.W. and MacKinlay, A.C., *The Econometrics of Financial Markets*, 1997 (Princeton University Press: Princeton, NJ).

Chhikara, R.S. and Folks, J.L., *The Inverse Gaussian Distribution, Theory, Methodology and Applications*, 1989 (Marcel Dekker: New York).

Cont, R., Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 2001, **1**, 223–236.

Cont, R. and da Fonseca, J., Dynamics of implied volatility surfaces. *Quantitative Finance*, 2002, **2**, 45–60.

Fama, E.F., Efficient capital markets: a review of theory and empirical work. *The Journal of Finance*, 1970, **25**, 383–417.

Granger, C.W.J. and Morgenstern, O., Spectral analysis of New York stock market prices. *Kyklos*, 1963, **16**, 1–27.

Kamstra, M. and Milevsky, M.A., Waiting to exercise: first passage times and American option pricing. Work in progress, 2005.

Kendall, M.G., The analysis of economic time-series, Part I: prices. *Journal of the Royal Statistical Society*, 1953, **96**, 11–25.

Levy, H. and Duchin, R., Asset return distributions and the investment horizon. *Journal of Portfolio Management*, 2004, **30**, 47–62.

Lo, A.W. and MacKinlay, A.C., Stock market prices do not follow random walks: evidence from a simple specification test. *The Review of Financial Studies*, 1988, **1**, 41–46.

Mandelbrot, B.B., The variation of certain speculative prices. *The Journal of Business*, 1963, **36**, 394–419.

Nelson, D.B., Conditional heteroskedasticity in asset returns: a new approach. *Econometrica*, 1991, **59**, 347–370.

Oksendal, B., *Stochastic Differential Equations*, 5th ed., 2004 (Springer: Berlin).

Osborne, M.F.M., Brownian motion in the stock market. *Operations Research*, 1959, **7**, 145–173.

Poterba, J.M. and Summers, L.H., Mean reversion in stock market prices: evidence and implications. *The Journal of Financial Economics*, 1988, **22**, 27–60.

Seshadri, V., *The Inverse Gaussian Distribution: a Case Study in Exponential Families*, 1993 (Clarendon Press: Oxford).

Wasan, M.T., First passage time distribution of Brownian motion with positive drift. Queen's Papers in Pure and Applied Mathematics, No. 19, Queen's University, Kingston, Ontario, 1969.