



Munich Personal RePEc Archive

Derivatives Pricing on Integrated Diffusion Processes: A General Perturbation Approach

Li, Minqiang

Bloomberg LP.

8 March 2014

Online at <https://mpra.ub.uni-muenchen.de/54595/>
MPRA Paper No. 54595, posted 19 Mar 2014 15:39 UTC

Derivatives Pricing on Integrated Diffusion Processes: A General Perturbation Approach

Minqiang Li^{*†}

March 19, 2014

Abstract

Many derivatives products are directly or indirectly associated with integrated diffusion processes. We develop a general perturbation method to price those derivatives. We show that for any positive diffusion process, the hitting time of its integrated process is approximately normally distributed when the diffusion coefficient is small. This result of approximate normality enables us to reduce many derivative pricing problems to simple expectations. We illustrate the generality and accuracy of this probabilistic approach with several examples in the Heston model, including variance derivatives, European vanilla options, timer forwards, and timer options. Major advantages of the proposed technique include extremely fast computational speed, ease of implementation, and analytic tractability.

Keywords: Integrated diffusion process; Asymptotic expansion; Hitting time; Derivative pricing; Timer options

JEL Classifications: C02; G12; G13

The contents of this article represent the authors' views only and do not represent the opinions of any firm or institution.

^{*}Derivatives Research, Bloomberg LP, 731 Lexington Avenue, New York, NY, 10022. E-mail: minqiang.li@gmail.com.

[†]I thank seminar participants at Bloomberg, Brooklyn College, and Rutgers University for comments and suggestions.

1 Introduction

There are many financial derivatives whose payoff or pricing is related to an integrated diffusion process. Here by an *integrated diffusion process*, we mean a continuous-time stochastic process that is a time integral of a diffusion process. For example, virtually all variance derivative products are associated with the accumulated realized variance, which is often modeled as the time integral of the instantaneous variance for high accumulating frequency. Another example is the continuous-time average price Asian option in which the payoff is a function of the integrated stock process. A third example is interest rate derivatives pricing using short-rate models, in which the integrated short-rate process plays an important role.¹

A common technique for pricing derivatives is through solving the corresponding pricing PDE, either analytically or numerically. Financial derivatives related to integrated diffusion processes pose a challenge for this approach. The reason is that the PDE is usually of high dimension. For example, in pricing variance derivatives, in order to form a Markovian system, one usually has to include simultaneously the instantaneous variance process and the accumulated variance process. Therefore, the pricing PDE also includes both variables. On the other hand, the final payoff of a variance derivative never depends explicitly on the unobservable instantaneous variance. For example, in the case of a volatility swap, to get the fair volatility swap rate today, we just need to compute the expectation of the square root of the accumulated variance at expiry. If we have the explicit probability density of the accumulated variance at expiry, the computation becomes just a simple one-dimensional integration.

The above discussion highlights the potential usefulness of the probabilistic approach based on the risk-neutral expectation since often fewer variables are involved using this approach than the PDE approach. In practice, however, it is not easy to compute the probability densities. Analytical results are only known for a very limited set of models and even in those cases multiple dimensional Fourier inversion is often involved. Therefore, under many circumstances, in order to use the probabilistic approach effectively, it is useful to have the probability density available through analytical approximation means such as perturbation.

The current paper is one step in this direction. The central object of interest in this paper is the random time that the integrated process first exceeds a fixed budget. We study this hitting time directly rather than the integrated process itself for several reasons. First, for a positive diffusion process, once we have the distribution function of the hitting time, by a duality result, we immediately have the distribution function for the integrated process. Second, a technical but important motivation for the current paper is that the PDEs for functionals of the hitting

¹We note that integrated processes are also useful for modeling in biology or economics, where the time integral of quantities such as mortality rate, birth rate, gene mutation, income stream, consumption stream, etc. are often of great interest.

time are sometimes easier to deal with than those for the integrated process itself. Third, in practice, there are derivative securities whose final payoffs are explicit functions of the hitting time. For example, Société Générale Corporate and Investment Banking introduced a new type of variance derivative products called “timer options” in 2007. See Sawyers (2007, 2008). A timer option is similar to a plain-vanilla option, except that it can only be exercised when the accumulated realized variance reaches a given budget. Major banks have since traded timer options. Sawyers (2008) also reports that more complex derivatives with timer features such as timer swaps have been introduced to the over-the-counter derivatives market.

The basic assumption we use is that the diffusion coefficient of the diffusion process is small. We perform an asymptotic perturbation analysis on the moment generating function of the hitting time of the integrated diffusion process. We show that under small diffusion coefficient, the hitting time is approximately normally distributed since its moment generating function has an asymptotic form similar to that of a normal random variable. For many common models including the popular square-root process, the approximate mean and variance can be easily obtained in simple closed form. We also generalize the result to the time integral of *functions* of diffusion processes. The result of approximate normal distribution is very convenient in approximating derivative prices. We give several examples in the paper using the Heston model, including generic variance derivative pricing, plain-vanilla European-style options, timer forwards, and timer options. In all these examples, the final approximated price is either a simple one-dimensional integration or in closed form. Numerical analysis shows that these approximations are fairly accurate when the volatility coefficient is not too large, in addition to being extremely fast and easy to implement.

There have been studies in the literature on integrated diffusion processes. For example, Dufresne (2001) studies the integrated square-root process. Forde and Jacquier (2010) consider the integrated geometric Brownian motion process to price Asian options. The approach we take here is different. While existing literature focuses mostly on exact properties of specific processes, we study the approximate properties of general integrated processes. These two branches of research directions are therefore complementary to each other.

The use of perturbation technique in derivative pricing has a long history and it is difficult to list all the references. The two references which are most closely related to the current paper are Lewis (2000) and Lipton (2001), where the authors consider volatility of volatility expansion for plain-vanilla European-style option prices in the Heston model. One difference is that in this paper we perform an expansion for the moment generating function of the hitting time of any integrated diffusion process and then use it to price many different derivative products.

There are several clear advantages of the proposed technique over alternative numerical

methods such as Monte Carlo or PDE. First, the computation usually takes well below one second or even below one millisecond for computing each derivative price, compared with Monte Carlo or PDE which can take many orders of magnitude more computational time. The benefit in computational time is much greater than it first seems when looking at pricing a single option price. Below I elaborate on this point because it is often overlooked and under-appreciated. Take for example, the computation of credit value adjustment (CVA) required by Basel II and III (see, for example, Gregory (2012)). Roughly speaking, CVA is a time-average of conditional exposure up to a future horizon, weighted by the default time probability density. In a Monte Carlo setup, for each future time grid and each realized intermediate state variable configuration, one needs to run a Monte Carlo to get a price as a function of the state variables. These intermediate future prices are in turn fed into the time-averaging formula. The number of simulations needed in such a nested Monte Carlo is the usual number of simulations needed for a single price, multiplied by the number of scenarios (usually taken to be at least 1000, but can require a lot more if number of factors involved is large), and also by the number of time steps (usually taken to be somewhere from 20 to 100). Furthermore, hedging analysis, stress testing, risk analysis, real-time pricing, and greek computations all exacerbate the Monte Carlo simulation burden. For example, to compute the greek gamma through bumping and repricing, one needs to repeat the Monte Carlo simulations three times. Related, the availability of closed-form approximations allows us to examine the price sensitivity to model parameters more easily. For example, suppose one wants to study the price of a certain derivative as a function of the long-run mean and the mean-reverting strength in the Heston model. In a Monte Carlo setup, we can sample 10 different values of the mean and the mean-reverting strength. If each Monte Carlo takes 1 minute which is easily exceeded in the cases of exotic derivatives such as timer options, then we will need 100 minutes of computational time since each parameter combination requires a separate Monte Carlo simulation. To contrast, it usually takes less than one second to compute 1000 prices with an analytic approximation.

Second, the perturbation technique developed in this paper often gives us closed-form formulas for derivative prices. The closed-form formulas are nice not just for aesthetic reasons. They are very intuitive and one can see clearly the financial meaning of the different parts in the formula. These formulas often satisfy additional attractive properties, such as having the right limits, obeying put-call parity automatically, preserving price positivity or convexity, etc. Often, greeks are also available in closed form. The analysis leading to the explicit formulas also provides additional insights about the derivative pricing problem at hand. For example, in the timer option case, when volatility of variance is small, our analysis shows that the time to exercise the option roughly follows a normal distribution. An investor can then

make heuristic hedging and other decisions based on this observation.

Third, the approximations developed from the perturbation technique are very easy to implement. As we will see from the examples, most of the approximations involve only a one-dimensional numerical integration on the real line. This is to contrast other numerical methods such as the timer option pricing formula in Liang, Lemmens and Tempere (2011) which is a high-dimensional numerical integration in the complex plane involving complicated functions such as modified Bessel functions. These integrals are often very tricky to evaluate numerically due to oscillatory integrand and slow decaying near ends of integration region. Even without these difficulties, multi-dimensional numerical integration is still very expensive computationally.² For example, Li (2013) reports a computing time of about 60 seconds for pricing one perpetual timer option. Another alternative, PDE method, also requires a lot of expertise and care, especially when a high dimension is encountered. To contrast, valuing one perpetual timer option using the proposed perturbation technique takes less than 10^{-4} seconds.

There are limitations with the perturbation technique developed in this paper which a potential user should be aware of. First, only a limited number of models have been solved in this paper, and we have only tested the numerical accuracy in the Heston model for some limited set of parameters.³ The Heston model is used in the testing because it is one of the most popular models for derivatives pricing. Also, while the method applies to generic Heston-like stochastic models, not all models possess simple closed-form formulas. Second, the perturbation technique developed in this paper requires the presence of a small parameter, which might not be available in some real-life applications. Third, other methods such as Monte Carlo might be more versatile in the sense that it is easier to incorporate additional features such as American feature into the pricing engine. However, we believe that these limitations do not diminish the advantages of numerical approximations. It is also our hope that future research will address and overcome some of these limitations.

The rest of the paper is organized as follows. Section 2 develops the approximation for the moment generating function of the hitting time of a general integrated diffusion process. We show that under small diffusion coefficient, the hitting time is approximately normally distributed. Section 3 illustrates the usefulness and accuracy of this probabilistic approach using several examples, namely, generic variance derivatives, European options, timer forwards, and timer options. Section 4 concludes.

²This is not to say that they are not useful. In fact, they are extremely useful because they can provide definite benchmarks to examine the accuracy of approximations or Monte Carlo.

³More testing using different sets of parameters and a different model (3/2-model) has been carried out in Li and Mercurio (2013a, 2013b, 2013c). The approximations are still found to be very accurate. We refer readers to these papers for more details.

2 Approximating the Hitting Time Distribution

2.1 The Setup

We consider a time-homogeneous one-dimensional diffusion process X_t whose dynamics is as follows

$$dX_t = a(X_t) dt + \eta b(X_t) dW_t. \quad (1)$$

Here W_t is a standard Brownian motion. Notice that we have singled out a nonnegative constant η for the diffusion function and will refer to it as the volatility coefficient. The drift and diffusion functions $a(X_t)$ and $b(X_t)$ are assumed to be functions of X_t only. We assume that the state space of X_t is $(0, \infty)$ and that the process does not explode to either zero or infinity in finite time. Despite its simplicity, this specification covers many important models in finance, including the Black-Scholes model for stock price movement, the Vasicek and Cox-Ingersoll-Ross models for the short rate movement, and the Heston model for the instantaneous variance movement, among many others.

We use ξ_t to denote the time-integrated process of X_t , defined as follows

$$\xi_t = \xi + \int_0^t X_u du. \quad (2)$$

Here $\xi_0 = \xi$ is the value of the integrated process at time 0. We assume that $\xi \geq 0$.⁴ Notice that ξ_t is nonnegative and increasing in t . As motivated by the reasons listed in the introduction, the central object of interest is the hitting time τ^B for the process ξ_t to hit a certain level B where $B \geq \xi$. That is,

$$\tau^B \equiv \inf \{t \geq 0 : \xi_t = B\} = \inf \left\{ t \geq 0 : \int_0^t X_u du = B - \xi \right\}. \quad (3)$$

We are interested in the distribution of the random time τ^B . Therefore, we consider its moment generating function:

$$M_{\tau^B}(\lambda) \equiv \mathbb{E}_0[e^{\lambda \tau^B} \mid \xi_0 = \xi, X_0 = X]. \quad (4)$$

We assume that the process X_t is such that $M_{\tau^B}(\lambda)$ is well-defined and exists for a continuous range of real values of λ .

⁴The accumulation can start from sometime in the past, so that $\xi > 0$. For example, this can correspond to a timer option initiated in the past, so that it has accumulated some nonzero amount of realized variance at time 0. The variable ξ is needed in the PDEs since they involve partial derivatives with respect to ξ , but after solving the PDEs, they are usually set to value 0 by thinking of B as the remaining variance budget rather than the original budget in the contract. The situation is similar to the Black-Scholes PDE where the variable t is used in the PDE. But after the Black-Scholes formula is obtained as a function of t , one usually assumes today is time 0 and set $t = 0$, and reinterprets T as the remaining maturity $\tau = T - t$.

Since (X_t, ξ_t) is jointly Markovian which is sufficient to determine whether B is exceeded or not, $M_{\tau^B}(\lambda)$ is a function of ξ and X only, and not a function of the current time t . This situation is similar to that of perpetual American options in the Black-Scholes model. Therefore, we let

$$\Pi(\xi, X; \lambda) \equiv M_{\tau^B}(\lambda). \quad (5)$$

For notational ease, we will often omit the parameter λ in $\Pi(\xi, X; \lambda)$ and just write $\Pi(\xi, X)$. We can interpret $\Pi(\xi, X)$ as the price of a zero-coupon timer bond which pays 1 dollar when the budget B is exceeded with the risk-free rate being constant and equal to $-\lambda$.

By the Feynman-Kac theorem applied to the random exit time τ^B , $\Pi(\xi, X)$ satisfies the following partial differential equation

$$X\Pi_\xi + a(X)\Pi_X + \frac{1}{2}\eta^2 b^2(X)\Pi_{XX} + \lambda\Pi = 0, \quad (6)$$

with the boundary condition

$$\Pi(B, X) = 1. \quad (7)$$

In the PDE above, we have used subscripts to denote partial derivatives. For general $a(X)$ and $b(X)$ functions, the above PDE is difficult to solve exactly. Therefore, below we take a perturbation approach.

2.2 The Approximation

We approximate the moment generating function of τ^B under the assumption that the volatility coefficient η is small. Small η expansion is considered for plain-vanilla options in Lewis (2000) and Lipton (2001). In Li and Mercurio (2013a), it is shown that for the Heston model and the 3/2 model, τ^B is approximately normally distributed for small η in the sense that the asymptotic expansion of $M_{\tau^B}(\lambda)$ is exactly in the form of the moment generating function of a normal random variable. Below we show that this is true for any time-homogeneous one-dimensional diffusion process.

We make an important remark that by “small η ” here, we do not require that η is smaller than 1. Rather, we mean that the effect of η should be small. This could be measured by the long-run variance of the process X_t , for example. If $b(X_t)$ is small, then it is possible that η is much larger than 1 even though the effect of η is still very small. It is the effect of η on the variability of the process X_t that matters.

It is useful for developing perturbation series purpose to have a constantly zero boundary condition. Therefore, we define the function $p(\xi, X)$ by

$$\Pi(\xi, X) \equiv e^{p(\xi, X)}. \quad (8)$$

Because $\Pi(\xi, X) > 0$, this is well-defined. The quantity $p(\xi, X)$ is actually the cumulant generating function of the random variable τ^B conditioning on that the current process value is X and the current integrated process value is ξ . From Equation (6), $p(\xi, X)$ satisfies the nonlinear partial differential equation below

$$Xp_\xi + a(X)p_X + \lambda + \frac{1}{2}\eta^2 b^2(X) [p_{XX} + (p_X)^2] = 0, \quad (9)$$

with the boundary condition

$$p(B, X) = 0. \quad (10)$$

The above partial differential equation is exact. We can solve it approximately by asymptotically expanding in η . It is clear that when $\eta = 0$, we get an ordinary first-order differential equation which can be solved exactly by method of characteristics. This gives us the zeroth-order expansion $p_0(\xi, X)$ in η^2 . The first-order expansion in η^2 (second-order in η) can then be obtained by replacing $[p_{XX} + (p_X)^2]$ with $[p_{0,XX} + (p_{0,X})^2]$. This second-order expansion of $p(\xi, X)$ in turn gives a second-order expansion for $\Pi(\xi, X)$. It turns out that because of the special structure of the PDE for $p(\xi, X)$, the first-order expansion in η^2 for $p(\xi, X)$ is actually quadratic in λ , as Proposition 1 below states. This is very interesting because it means that in a certain sense, τ^B is approximately normally distributed for small η . This is true regardless of the functional forms of $a(X)$ and $b(X)$. A detailed proof is given in Appendix.

Proposition 1. *Assume that $\mathbb{E}_0 e^{\lambda\tau^B}$, $\mathbb{E}_0[e^{\lambda\tau^B}\tau_B]$ and $\mathbb{E}_0[e^{\lambda\tau^B}\tau_B^2]$ are finite for some range of λ values in \mathbb{R} containing 0. The moment generating function of τ^B has the following asymptotic expansion form*

$$M_{\tau^B}(\lambda) \equiv \mathbb{E}_0 e^{\lambda\tau^B} = e^{\lambda(T_0 + \eta^2 H_0) + \lambda^2 \eta^2 H_1} + o(\eta^2), \quad (11)$$

where T_0 , H_0 and H_1 are not functions of λ or η .⁵ Furthermore, $T_0 \geq 0$ and $H_1 \geq 0$ with equality if and only if $B = \xi$. Therefore, for $B > \xi$, in the sense above, τ^B is approximately normally distributed with mean μ and variance σ^2 , where μ and σ^2 are given by

$$\mu = \mu(B) = T_0 + \eta^2 H_0, \quad (12)$$

$$\sigma^2 = \sigma^2(B) = 2\eta^2 H_1. \quad (13)$$

⁵Explicit expressions for evaluating T_0 , H_0 and H_1 are given in Equations (110), (114) and (115) in Appendix. However, in practice, it is usually easier to directly solve the PDE perturbatively than to use these integral equations. Once the PDE is solved perturbatively to first order in η^2 , we can use Proposition 1 to read off T_0 , H_0 and H_1 because by Proposition 1 they are multiplied by λ , $\lambda\eta^2$ and $\lambda^2\eta^2$, respectively. Also, in professional software such as Mathematica, one does not need to perform the characteristic transformation to the original PDE oneself, as we have done in the proof of Proposition 1 in Appendix. For many models, the software can solve the perturbed PDEs in the original variables directly and all one needs to do is to simplify the results by defining variable combinations that appear multiple times in the results. Most of the time, these definitions correspond to the characteristic transformations. However, Equations (110), (114) and (115) in Appendix can be useful numerically when a model has complicated drift and diffusion functions.

We make a few remarks below. First, it is useful to take a look at Proposition 1 for the degenerate case $\eta = 0$. In this case, the process X_t evolves deterministically according to $dX_t/dt = a(X_t)$. The time τ^B to hit the budget B becomes exactly $T_0 = T_0(\xi, X)$, which satisfies the first-order PDE below:

$$XT_{0,\xi} + a(X)T_{0,X} + 1 = 0, \quad (14)$$

with the boundary condition

$$T_0(B, X) = 0. \quad (15)$$

Therefore, the moment generating function $M_{\tau^B}(\lambda)$ degenerates to that of a Dirac delta function at T_0 , that is, $e^{\lambda T_0}$. This provides a sanity check for Proposition 1.

The approximate mean $\mu(B)$ and variance $\sigma^2(B)$ are functions of B , ξ and X as well as other model parameters. For notational simplicity, we only emphasize their dependence on B in Proposition 1. They actually depend on ξ and B through the difference $B - \xi$, but later on we will always assume without loss of generality that $\xi = 0$. The quantities μ and σ^2 are also approximations in the following sense. If we let $m_1(\xi, X)$ and $m_2(\xi, X)$ be the true first and second moments of τ^B when the initial integrated process value is ξ and the initial state variable value is X , then it is easy to see that they satisfy the following PDEs: (see, for example, Chapter 15 of Karlin and Taylor (1991) for a derivation)

$$Xm_{1,\xi} + a(X)m_{1,X} + \frac{1}{2}\eta^2b^2(X)m_{1,XX} + 1 = 0, \quad (16)$$

$$Xm_{2,\xi} + a(X)m_{2,X} + \frac{1}{2}\eta^2b^2(X)m_{2,XX} + 2m_1 = 0. \quad (17)$$

By formally differentiating Equation (6) with respect to λ once and twice and setting $\lambda = 0$, we can easily see that μ satisfies Equation (16) asymptotically to order η^2 . Similarly $\mu^2 + \sigma^2$ satisfies Equation (17) asymptotically.⁶

In many actual applications, the final payoff of the derivative is a function of ξ_T instead of τ^B . Therefore, it is useful to have an approximation for the distribution function of ξ_T . Let $F_Z(\cdot)$ denote the distribution function of a random variable Z . For simplicity we assume without loss of generality that $\xi = 0$. Then, for $T > 0$, we can approximate the cumulative distribution function $F_{\xi_T}(x)$ as:

$$F_{\xi_T}(x) \equiv \mathbb{P}(\xi_T < x) = 1 - F_{\tau^x}(T) \approx N\left(\frac{\mu(x) - T}{\sigma(x)}\right), \quad (18)$$

⁶In fact, all raw moments and central moments of τ^B can be approximated to second order in η in the sense of both matching the actual expectation and satisfying their respective PDEs asymptotically.

where $N(\cdot)$ is the cumulative normal distribution function. The first equality in the above statement is easily seen by noticing the following duality between τ^x and ξ_T :

$$\{\tau^x > T\} = \{\xi_T < x\}. \quad (19)$$

Therefore, we have

$$\mathbb{P}(\xi_T < x) = \mathbb{P}(\tau^x > T). \quad (20)$$

The last approximate equality in Equation (18) is due to Proposition 1. Although it is difficult to show analytically, for reasonable parameter values we have tried using the Heston model, we always observe numerically that $F_{\xi_T}(0^+) = 0$, $F_{\xi_T}(+\infty) = 1$, and that $F_{\xi_T}(x)$ is monotonically increasing in x . For small η , the above approximation is very good and captures some important features of the simulated distribution of ξ_T .

2.3 Examples

The proof in Appendix shows a procedure to compute $\mu(B)$ and $\sigma(B)$ needed in Proposition 1. We first compute the characteristic coordinates, and then T_0 , H_0 and H_1 needed for $\mu(B)$ and $\sigma(B)$ are simply given by integrals. For models with simple $a(X)$ and $b(X)$ functions, the integrals can be performed analytically. We give a few examples below. Readers interested in the details of the calculations are referred to Li and Mercurio (2013a). For simplicity, we assume that currently $\xi = 0$ so that the quantity B in formulas below should be interpreted as the remaining budget $B - \xi$.

Example 1: (Square root process where $dV = \kappa(\theta - V)dt + \eta\sqrt{V} dW$)

Here the state variable X_t is V_t . We will interpret V_t as the instantaneous variance as in the Heston (1993) model. It is worth mentioning that this square-root process is also frequently used to model short rates, as in Cox, Ingersoll and Ross (1985). For notational ease, we will assume that the current time is 0. We use V_0 to denote the current instantaneous variance, which is more natural than V . By solving the PDE for $p(\xi, V)$ to second order in η , the final T_0 , H_0 and H_1 are given by⁷

$$T_0 = \frac{1}{\kappa} \log R, \quad (21)$$

⁷We can also use the integral equations in the Appendix. However, it's much more straightforward to solve the PDE perturbatively and then eyeball the results to get T_0 , H_0 and H_1 . We can do it because T_0 is multiplied by λ , H_0 is multiplied by $\lambda\eta^2$ and H_1 is multiplied by $\lambda^2\eta^2$ by Proposition 1. In fact, our implementation to get the results for the Heston model only takes a few lines of code in Mathematica. The first line solves the zero-order PDE. The built-in function `DSolve` is perfect for this purpose. The second line solves the first-order PDE in η^2 , again using `DSolve`. The rest a few lines are pure algebraic simplification. It takes only a few seconds to run the Mathematica code, but takes considerably much more time to type the results up and double-check that no human mistakes have been made!

and

$$H_0 = \frac{(R-1)[2R^2z^2 + R(2-5z-2z^2) - 2 - z]}{4\kappa^2R^2(1+z)^3\theta} + \frac{3z \log R}{2\kappa^2(1+z)^3\theta}, \quad (22)$$

$$H_1 = \frac{(R-1)(1+2R^2z + R(2z-3))}{4\kappa^3R^2(1+z)^2\theta} - \frac{(2z-1) \log R}{2\kappa^3(1+z)^2\theta}, \quad (23)$$

with

$$R = e^{z-z_0+\kappa\frac{B}{\theta}}, \quad (24)$$

$$z_0 \equiv \frac{V_0 - \theta}{\theta}, \quad (25)$$

$$z \equiv W\left(z_0 e^{z_0} \cdot e^{-\kappa\frac{B}{\theta}}\right), \quad (26)$$

where $W(\cdot)$ is Lambert's product-log function defined implicitly as $x = W(x)e^{W(x)}$. See Corless et al (1996) for information on product-log function. We refer readers to Li and Mercurio (2013a) for more details on the derivation of the above formulas.

Since $R \geq 1$, T_0 is nonnegative. Notice also that T_0 is the implicit solution of

$$\theta T_0 + (V_0 - \theta) \frac{1 - e^{-\kappa T_0}}{\kappa} = B. \quad (27)$$

This is exactly the deterministic time when $\xi_{T_0} = B$ for $\eta = 0$. The easiest way to see that $H_1 \geq 0$ is through numerical plotting. A three-dimensional plotting is possible because the denominators of both H_0 and H_1 are positive, and the numerators of both H_0 and H_1 can be written as a function of the two variables z_0 and $\kappa B/\theta$.

Figure 1 plots the probability density function and cumulative distribution function of ξ_T in the Heston model using our approximation as well as the histogram and empirical cumulative distribution function from Monte Carlo simulation. Here the process ξ_t is the accumulated realize variance process. Parameters used here are: $V_0 = 0.087$, $\kappa = 2$, $\theta = 0.09$, $T = 1.5$ years, and $\eta = 0.250$. As we see, the approximation is fairly good when compared to the histogram from simulation. Both graphs show almost zero mass in the regions $\xi_T < 0.06$ and $\xi_T > 0.4$, and both are left skewed. The theoretical expectation of ξ_T is given by

$$\mathbb{E}_0 \xi_T = \theta T + (V_0 - \theta) \frac{1 - e^{-\kappa T}}{\kappa}. \quad (28)$$

Notice that it does not depend on η . With the given parameters, the expectation is about 0.1336. Numerical integration shows that the approximate density gives the expectation of ξ_T as 0.1338 with a percentage error of around 0.1%. The two cumulative distribution functions are also very similar. The approximate cumulative distribution function has the desired property of being strictly increasing.

Example 2: (3/2 model where $dV = \kappa V(\theta - V) dt + \eta V^{\frac{3}{2}} dW$)

Here the state variable X_t is V_t . See Ahn and Gao (1999) for some univariate analysis on the 3/2 model. This model has been used to model both short rate and instantaneous variance. In this model, T_0 , H_0 and H_1 are given by

$$T_0 = \frac{1}{\kappa\theta} \log \left(\frac{V_0 + \theta(e^{\kappa B} - 1)}{V_0} \right), \quad (29)$$

$$H_0 = \frac{4V_0 [1 + (\log R - 1)R] + \theta[-3 + (4 - 4\log R)R + (2\log R - 1)R^2]}{4\kappa^2 [V_0 + \theta(R - 1)]^2}, \quad (30)$$

$$H_1 = \frac{4R - [3 - 2\log R]R^2 - 1}{4\kappa^3 [V_0 + \theta(R - 1)]^2}, \quad (31)$$

with $R = e^{\kappa B}$. Since $R \geq 1$, it is very easy to check that $T_0 \geq 0$ and $H_1 \geq 0$.

Example 3: (Geometric Brownian Motion where $dS = (r - \delta)Sdt + \eta S dW$)

Here the state variable is S . For simplicity, we assume $r \neq \delta$. The degenerate case $r = \delta$ is simpler and can be solved similarly. We solve the PDE of $p(\xi, S)$ to second order in η . The functions T_0 , H_0 and H_1 can then be read off as

$$T_0 = \frac{1}{\hat{r}} \log R, \quad (32)$$

$$H_0 = \frac{B\hat{r}(B\hat{r} + 2S) - 2(B\hat{r} + S)^2 \log R}{4\hat{r}^2(B\hat{r} + S)^2}, \quad (33)$$

$$H_1 = \frac{2(B\hat{r} + S)^2 \log R - B\hat{r}(3B\hat{r} + 2S)}{4\hat{r}^3(B\hat{r} + S)^2}, \quad (34)$$

where $R = (S + B\hat{r})/S$ and $\hat{r} = r - \delta$. It is easy to verify that $T_0 > 0$ and $H_1 > 0$ if $B > 0$.

2.4 A Generalization

We now discuss an interesting generalization of the previous setup.⁸ The result in this subsection is not used in the numerical study section, and readers interested in the applications of the previous results can skip this subsection completely.

It turns out that the hitting time of the following integrated diffusion process is also approximately normally distributed (to be shown below):

$$\int_0^t f(X_u) du, \quad (35)$$

where X_t is any diffusion process not necessarily living on $(0, \infty)$, and $f(\cdot)$ is any second-order differentiable function. The only requirement is that $f(\cdot)$ is positive and X_t has a small

⁸We thank an anonymous referee for suggesting this fruitful extension.

parameter η in its diffusion function. Such a setup has been considered in Cui (2013). Notice that in this case f might not be one-to-one, and the filtration generated by $f(X_u)$ might be strictly smaller than that of X_u . One example in point is the Schöbel-Zhu (1999) stochastic volatility model where instead of modeling the instantaneous variance V_t , one models the “signed volatility” v_t (so here $X_t \equiv v_t$ is the state variable):

$$dS_t = rS_t dt + v_t S_t dW_t^S, \quad (36)$$

$$dv_t = \kappa(\theta - v_t)dt + \eta dW_t^v. \quad (37)$$

The “signed volatility” v_t follows an Ornstein-Uhlenbeck process and has a state space $(-\infty, \infty)$. The instantaneous variance is given by $V_t = v_t^2$. The joint state variables are (S_t, v_t) . Notice that v_t contains strictly more information than V_t since V_t does not contain the information about the sign of v_t . In this case, the hitting time of the integrated variance $\int_0^t v_u^2 du$ (rather than $\int_0^t v_u du$) is of interest in finance.

Following Cui (2013), we write $f(x) = m^2(x)$ for some function m to emphasize that f is positive. The integrated process we are interested in is

$$\xi_t = \xi + \int_0^t m^2(X_u) du. \quad (38)$$

Notice that $d\xi_t = m^2(X_t)dt$. The hitting time τ^B is the first time a budget B is reached by the ξ_t process:

$$\tau^B \equiv \inf \{t \geq 0 : \xi_t = B\} = \inf \left\{ t \geq 0 : \int_0^t m^2(X_u) du = B - \xi \right\}. \quad (39)$$

Since (X_t, ξ_t) is a Markovian system sufficient to determine whether B is exceeded or not, the moment generating function of τ^B is again a function of the current states ξ and X . By Feynman-Kac theorem Π satisfies the slightly more general PDE below:

$$m^2(X)\Pi_\xi + a(X)\Pi_X + \frac{1}{2}\eta^2 b^2(X)\Pi_{XX} + \lambda\Pi = 0, \quad (40)$$

with the boundary condition

$$\Pi(B, X) = 1. \quad (41)$$

When $m(x) = \sqrt{x}$, the PDE above reduces to the one in the case we have considered previously. The same method of characteristics we have used in the Proof of Proposition 1 applies with little modification.⁹ The upshot is that τ^B is still approximately normally distributed. All

⁹Specifically, we need to modify z to be

$$z = \Phi \left(B - \xi + \int_{X^*}^X \frac{m^2(u)}{a(u)} du \right). \quad (42)$$

three equations in Proposition 1 are still valid. The exact formulas for the mean and variance can be obtained, similar to what we have done in Appendix. However, as we have remarked previously in the footnote, in practice it is much easier to directly perturb the PDE in Equation (40) and read off the functions T_0 , H_0 and H_1 from the results.

In some sense, we have already seen one example in this generalized setup. The 3/2-model studied in Example 2 is the reciprocal of the Heston model in Example 1, as can be verified using Ito's lemma. If we let X_t follow the Heston model, then our generalized result immediately says that the hitting time of $\int_0^t 1/X_u du$ is also approximately normally distributed.

The Schöbel-Zhu (1999) model, unfortunately, does not give simple formulas when solving it using our perturbation approach.¹⁰ The reason is that in this model we cannot obtain the zeroth-order solution T_0 in an explicit form. The instantaneous variance when $\eta = 0$ is given by

$$V_t = v_t^2 = (\theta + (v_0 - \theta)e^{-\kappa t})^2. \quad (44)$$

The deterministic time T_0 to reach a variance budget B is the solution of

$$\int_0^{T_0} (\theta + (v_0 - \theta)e^{-\kappa t})^2 dt = B - \xi. \quad (45)$$

Unfortunately, it's not possible to write T_0 as an explicit function of ξ and v_0 easily unless $\theta = 0$, or $\kappa = 0$, or $v_0 = \theta$. All special cases are of considerably less interest in practice. Since the derivatives of T_0 need to be fed into the first-order perturbed PDE for Π , the lack of an explicit expression for T_0 prevents an explicit formula for Π .

There are, nonetheless, some interesting new models that are explicitly solvable in this generalized setup, and we discuss two of them below. The first model is similar to the Schöbel-Zhu (1999) model, except that the volatility process follows a 3/2-process:

$$dS_t = rS_t dt + v_t S_t dW_t^S, \quad (46)$$

$$dv_t = \kappa v_t (\theta - v_t) dt + \eta v_t^{3/2} dW_t^v. \quad (47)$$

Equation (106) then changes to

$$m^2(X)z_\xi + a(X)z_X = 0. \quad (43)$$

Modifications to other equations are straightforward.

¹⁰We emphasize that this by no means implies that our perturbation technique is no longer valid. The procedure still works. The distribution of τ^B is still approximately normal. It is just that the generalized versions of Equations (110), (114) and (115) in Appendix cannot be integrated out to give *explicit* formulas. Numerical methods can still be used to perform these integrals. The numerically obtained $\mu(B)$ and $\sigma(B)$ can still be used to price timer options, for example.

We are interested in the distribution of τ^B , which is the first time for the process $\int_0^t V_u du$ to exceed the remaining budget $B - \xi$. Following the PDE convention, we let $v_0 = v$. The PDE for its moment generating function $\Pi(\xi, v)$ is given by

$$v^2 \Pi_\xi + \kappa v (\theta - v) \Pi_v + \frac{1}{2} \eta^2 v^3 \Pi_{vv} + \lambda \Pi = 0, \quad (48)$$

with the boundary condition $\Pi(B, v) = 1$. This equation can be solved perturbatively to first order in η^2 . The function T_0 can be read off from the zeroth-order solution. It is given by (for simplicity, we let $\xi = 0$ in the following)

$$T_0 = \frac{1}{\kappa \theta} \log \left(\frac{z_0(1+z)}{z(1+z_0)} \right), \quad (49)$$

where

$$z_0 = \frac{v - \theta}{\theta}, \quad (50)$$

$$z = W \left(z_0 e^{z_0} \cdot e^{-\kappa \frac{B}{\theta}} \right). \quad (51)$$

Here $W(\cdot)$ is the product-log function. The functions H_0 and H_1 can be read off from the first-order solution in η^2 by Proposition 1 and are given by

$$H_0 = \frac{(1 + z(14 + z)) \log(z_0/z)}{2\kappa^2(1+z)^4\theta} + \frac{(z - z_0) \left[-3z_0 + z(5 + z(4 + z) + 24z_0 + z(17 + 4z)z_0 - 2(8 + z)z_0^2 - 2z_0^2) \right]}{4\kappa^2(1+z)^4z_0^2\theta}, \quad (52)$$

and

$$H_1 = \frac{(2 - z(2 + z)) \log(z_0/z)}{\kappa^3(1+z)^4\theta^2} + \frac{8z(2 + z)z_0 - 12(1 + z)z_0^2 + 8z_0^3 + z_0^4 - z^2(2 + z)^2}{4\kappa^3(1+z)^4z_0^2\theta^2}. \quad (53)$$

It can be quickly checked that when $B = 0$, $H_0 = H_1 = 0$, as we would expect from Proposition 1.

In the 3/2 model we above, since v_t is always positive, the function $m^2(x) = x^2$ is one-to-one. Therefore, another approach we could have taken is to use Ito's lemma to write down the dynamics of V_t and use the result in Proposition 1.¹¹ This is a valid alternative approach.

¹¹The dynamics of V_t is given by

$$dV_t = V_t \left(2\kappa\theta - (2\kappa - \eta^2) \sqrt{V_t} \right) dt + 2\eta V_t^{5/4} dW_t^v. \quad (54)$$

Notice that both the drift and diffusion functions involve η . Also, the drift is not mean-reverting unless $2\kappa > \eta^2$. This condition is not needed for the volatility process v_t to be mean-reverting. However, in practice, such a condition might be attractive to have to make the variance process also mean-reverting.

The current approach of using v_t as the state variable is slightly simpler because both the drift and diffusion functions of V_t involve the parameter η , and therefore the perturbation using V_t as the state variable is a little bit more involved. In general, however, m^2 might not be one-to-one, and the filtration generated by X_t might be strictly larger than $m^2(X_t)$. In these cases, we will have to use the generalization above and Equation (40) since the filtration generated by $m^2(X_t)$ might not be sufficient to determine τ^B . Below we consider such a model.

This second model we consider is a “signed” stochastic volatility in the same spirit of Stein and Stein (1991) and Schöbel and Zhu (1999). The “signed” volatility v_t is modeled as

$$dv_t = \kappa v_t(\theta^2 - v_t^2) dt + \eta dW_t^v, \quad (55)$$

where $\kappa > 0$, $\theta > 0$ and η are constant parameters. This process has a state space $(-\infty, \infty)$. We see that the drift function is positive for very negative v_t , and is negative for very positive v_t . Therefore, the process is globally mean-reverting. However, the drift function has three zeros at $-\theta$, 0 , and θ . It has a tendency to return to either $-\theta$ or θ when $|v_t| > \theta$, and has a tendency to further drift away from 0 when $|v_t| < \theta$. The stationary density of v_t is bimodal with the two modes at $-\theta$ and θ . Similar processes have been used to study the bifurcation behavior in nonlinear dynamics.¹² Let $v_0 = v$. If $\eta = 0$, the process is always nonnegative or nonpositive and is given by the solution:

$$v_t = v \sqrt{\frac{\theta^2}{v^2 + (\theta^2 - v^2)e^{-2\kappa\theta^2 t}}}. \quad (56)$$

As t goes to infinity, it either goes to $-\theta$ or θ depending on the sign of v . We are still interested in the hitting time of integrated variance process ξ_t where $d\xi_t = v_t^2 dt$. The function T_0 can be obtained by solving the following ODE:

$$v^2 T_{0,\xi} + \kappa v(\theta^2 - v^2) T_{0,v} + 1 = 0, \quad (57)$$

with the boundary condition $T_0(B, v) = 0$. Or it can be solved by inverting the relation

$$\int_0^{T_0} v_u^2 du = B - \xi. \quad (58)$$

In any case, the result is given by (for simplicity, we let $\xi = 0$ so B is the remaining budget)

$$T_0 = \frac{1}{2\kappa\theta^2} \log R, \quad (59)$$

¹²Readers with a physics background will recognize that the drift is the negative of the gradient of double-well potential $U(r) = \kappa r^4/r - \kappa\theta^2 r^2/2$. Double-well potentials have deep connections to bifurcation phenomena such as superconductivity, Higgs mechanism and spontaneous symmetry breaking, and have found applications to protein folding in biology.

where

$$R \equiv \frac{v^2 + (e^{2\kappa B} - 1)\theta^2}{v^2}. \quad (60)$$

Notice that $R \geq 1$ so $T_0 \geq 0$. The functions H_0 and H_1 are given by

$$H_0 = \frac{(R-1)\left((1-R)v^2 + (R-2)\theta^2\right) + \left((R-1)v^2 + \theta^2\right)R \log R}{4\kappa^2 R^2 v^2 \theta^6}, \quad (61)$$

$$H_1 = \frac{(R-1)\left(2(1+R)\theta^2 - (1+5R)v^2\right) + 2R\left((2+R)v^2 - 2\theta^2\right) \log R}{8\kappa^3 R^2 v^2 \theta^8}. \quad (62)$$

As a quick sanity check, we see that $H_0 = H_1 = 0$ when $B = 0$. The approximate mean and variance of τ^B are then given by the last two equations in Proposition 1.

3 Applications to Derivatives Pricing

In what follows, we consider several examples to illustrate the potential usefulness of the distributional approximation developed in the previous section. We consider a general stochastic volatility framework. The types of derivative securities we consider include: a variance derivative whose payoff is a function of the realized variance, plain-vanilla European option, a perpetual timer forward contract, and a perpetual timer option. We reduce the prices of these derivatives to either a simple one-dimensional numerical integration or a closed-form expression. Numerical study using the Heston model is also carried out to demonstrate the accuracy of the approximations.

3.1 Variance Derivatives Pricing

Variance derivatives are derivatives written on the realized accumulated variance, which is usually computed as

$$\sum_{i=1}^N \left[\log \left(\frac{S_{t_i}}{S_{t_{i-1}}} \right) \right]^2, \quad (63)$$

where N is the number of days until maturity T , and S_{t_j} 's are the stock prices on day j . Notice that the summation goes from 1 to N since we set $T_N = T$. One standard approach is to model the underlying and instantaneous variance as following a standard stochastic volatility model in the risk-neutral measure \mathbb{Q} :

$$dS_u = (r - \delta)S_u du + \sqrt{V_u}S_u dW_u^S, \quad (64)$$

$$dV_u = a(V_u) du + \eta b(V_u) dW_u^V, \quad (65)$$

where r is the constant risk-free interest rate, δ is the constant dividend yield, and W_u^S and W_u^V are two standard Brownian motions with a constant correlation ρ . The daily accumulated variance is usually modeled using its continuous counterpart ξ_t , which is defined as

$$\xi_t = \int_0^t V_u \, du. \quad (66)$$

We are interested in computing quantities of the following form

$$G = \mathbb{E}_0[g(\xi_T)]. \quad (67)$$

For example, for $g(x) = x/T$ or $g(x) = \sqrt{x/T}$, G is the annualized variance swap rate or the volatility swap rate. Without loss of generality, we assume that $g(0) = 0$. Otherwise, we can always define $\tilde{g}(x) = g(x) - g(0)$ and $G = g(0) + \mathbb{E}_0[\tilde{g}(\xi_T)]$. We also assume that g satisfies suitable integrability and differentiability conditions.

If we assume that η is small, we can approximate G using Proposition 1. After integration by parts, the result is given as a one-dimensional numerical integration

$$G = \int_0^\infty (1 - F_{\xi_T}(B))g'(B) \, dB \approx \int_0^\infty N\left(\frac{T - \mu(B)}{\sigma(B)}\right)g'(B) \, dB. \quad (68)$$

Given a stochastic volatility model with drift and diffusion functions $a(V)$ and $b(V)$, the task is then to compute the approximate mean $\mu(B)$ and approximate standard deviation $\sigma(B)$ of the hitting time in Proposition 1. For the Heston model and the 3/2 model, these expressions can be computed easily from T_0 , H_0 and H_1 given in Examples 1 and 2.

Below we examine the accuracy of the above approximation for the volatility swap rate in the Heston model. That is, $g(x) = \sqrt{x/T}$. To remove the singularity at $B = 0$ in $g'(B)$, it is useful to perform a change of variable $y = \sqrt{B}$ in Equation (68) to get

$$G_{\text{Approx}} = \frac{1}{\sqrt{T}} \int_0^\infty N\left(\frac{T - \mu(y^2)}{\sigma(y^2)}\right) \, dy. \quad (69)$$

The above approximation is valid for any $a(V)$ and $b(V)$ functions. In our numerical analysis, we consider the popular Heston model. In this model, the volatility swap rate can be computed using the known moment generating function of ξ_T which provides us a benchmark to check for accuracy. Specifically, we have

$$G_{\text{NI}} = \frac{1}{\sqrt{T}} \mathbb{E}_0[\sqrt{\xi_T}] = \frac{1}{2\sqrt{\pi}\sqrt{T}} \int_0^\infty \frac{1 - M_{\xi_T}(-\lambda)}{\lambda^{3/2}} \, d\lambda, \quad (70)$$

where the subscript NI stands for numerical inversion, and $M_{\xi_T}(-\lambda)$ is the moment generating function $M_{\xi_T}(-\lambda) \equiv \mathbb{E}[e^{-\lambda\xi_T}]$. We call this method numerical inversion because it involves numerically inverting from the Laplace transform of ξ_T to get the half-integer moment. Notice

that the usual Laplace inverse transformation will involve contour integration in the complex plane. Equation (70) is very nice because it is an inversion on the real line. This is a well-known result in mathematics and has connections to fractional calculus, see Wolfe (1975). In finance, this inversion trick has been popularized by Schürger (2002) and subsequently by Gatheral (2006), among others. In the Heston model, M_{ξ_T} is given explicitly by (see Cox, Ingersoll and Ross (1985))

$$M_{\xi_T}(-\lambda) = e^{\phi - \lambda \psi V_0}, \quad (71)$$

with

$$\phi = \phi(\lambda) = \frac{\kappa\theta}{\eta^2}(\kappa + \gamma)T - \frac{2\kappa\theta}{\eta^2} \log \left(1 + \frac{(\gamma + \kappa)(e^{\gamma T} - 1)}{2\gamma} \right), \quad (72)$$

$$\psi = \psi(\lambda) = \frac{2(1 - e^{-\gamma T})}{(\gamma + \kappa)(1 - e^{-\gamma T}) + 2\gamma e^{-\gamma T}}, \quad (73)$$

and

$$\gamma = \gamma(\lambda) = \sqrt{\kappa^2 + 2\lambda\eta^2}. \quad (74)$$

While theoretically very pleasing, in the actual implementation, some care is needed for the numerical integration in Equation (70) because of the singularity at $\lambda = 0$ and the slow decay rate as λ goes to infinity. This slow decay becomes very challenging in practice when V_0 or ψ is small because a very wide integration region needs to be used to approximate the positive half real line. In these cases, one needs to either come up with some transformation which regularizes the integral, or dramatically increase the number of evaluation points used in the integration which has a negative impact on the computational speed. In contrast, the numerical integration in Equation (69) is more well-behaved. The integrand is bounded from 0 to 1. Also, although the integration region is from 0 to ∞ , in practice only a small region needs to be integrated over because the cumulative distribution function is very close to 0 for relatively large y . Of course, Equation (70) is exact theoretically, while Equation (69) is an approximation that is accurate when η is small. We note also that somewhat ironically a naïve implementation of Equation (70) will break down for very small η since ϕ in Equation (72) can evaluate to infinity for zero η . In this case, one needs to treat the η dependence of ϕ very carefully in Equation (72) for small η . Alternatively, Equation (69) can act as a drop-in replacement for Equation (70) when η is small.

Table 1 shows the volatility swap rates for both our approximation (**Approx**) and the benchmark (**NI**) from numerical inversion. Three numerical integration routines in MATLAB are employed to cross-verify the results from numerical integration: **quad** which uses the adaptive Simpson quadrature method, **quadl** which uses the adaptive Lobatto quadrature

method, and `quadgk` which uses the adaptive Gauss-Kronrod quadrature method. To the accuracy reported in Table 1, all three routines give identical results. We vary the maturity T , the current instantaneous variance V_0 , and the volatility coefficient η . The mean-reversion strength κ is fixed at 2.0, and the long-run variance θ is fixed at 0.09. Feller's condition requires $\eta < 0.6$. Therefore, we consider the three η values in the table. All volatility swap rates reported are in percentage terms. It is interesting to notice that the volatility swap rate in the Heston model with the given parameters is a decreasing function of the volatility coefficient η .¹³ As we see, the approximation is very accurate in general, and especially when the volatility coefficient η is small or when the maturity T is large.

3.2 European Option Pricing

We now consider the pricing of European-style options under general stochastic volatility models specified in Equations (64) and (65). We consider the special case of $\rho = 0$. The current approximation technique does not readily generalize itself to nonzero ρ . By the mixing technique in Hull and White (1987), the price of a European call option with strike K and maturity T can be computed as

$$C_{\text{vanilla}} = \int_0^\infty C^{\text{BS}}(S_0, K, r, \delta, T, x) dF_{\xi_T}(x), \quad (75)$$

where C^{BS} is the Black-Scholes price given by

$$C^{\text{BS}}(S_0, K, r, \delta, T, x) = S_0 e^{-\delta T} N(d_1(x)) - K e^{-rT} N(d_2(x)), \quad (76)$$

with

$$d_1(x) = \frac{\log(S_0 e^{(r-\delta)T}/K)}{\sqrt{x}} + \frac{\sqrt{x}}{2}, \quad (77)$$

$$d_2(x) = \frac{\log(S_0 e^{(r-\delta)T}/K)}{\sqrt{x}} - \frac{\sqrt{x}}{2}. \quad (78)$$

By integration by parts, we can rewrite C_{vanilla} as

$$C_{\text{vanilla}} = \left(S_0 e^{-\delta T} - K e^{-rT} \right)^+ + S_0 e^{-\delta T} \int_0^\infty (1 - F_{\xi_T}(y^2)) n(d_1(y^2)) dy, \quad (79)$$

where $n(\cdot)$ is the standard normal probability density function. The above formula is exact and expresses C_{vanilla} as the sum of two components. The first component is the value of the option if the price process is deterministic and grows to the forward price. The second component is a strictly positive adjustment due to the fact that the stock process is stochastic.

¹³This is to be expected. The variance swap rate in the Heston model does not depend on η for fixed κ and θ . The fact that volatility swap rate is decreasing in η for fixed κ and θ is due to convexity of the square-root function.

By using the approximation for the cumulative distribution function $F_{\xi_T}(\cdot)$ in the last section, we can approximate the European call option price as

$$C_{\text{vanilla}} \approx \left(S_0 e^{-\delta T} - K e^{-rT}\right)^+ + S_0 e^{-\delta T} \int_0^\infty N\left(\frac{T - \mu(y^2)}{\sigma(y^2)}\right) n(d_1(y^2)) dy. \quad (80)$$

The above approximation works for any Heston-like stochastic volatility model as long as we can compute the functions $\mu(\cdot)$ and $\sigma(\cdot)$. In the case of the Heston model, it offers an alternative to the volatility of volatility expansion in Lewis (2000) and Lipton (2001). It is interesting to notice that the expansion there was developed using complex Fourier inversion while we have worked strictly in the original real space. Also, while it is possible for the price expansion in Lewis (2000) and Lipton (2001) to be negative, the approximate price above is always positive. The integrand is also very well-behaved and decays very fast. A shortcoming is that the current approximation only works for zero correlation.

The complex Fourier inversion is theoretically exact, but its implementation has a number of pitfalls and is far from trivial, especially when maturity is very short or very large, or when the option is far-away from the money. See for example, Carr and Madan (1999), and Kahl and Jäckel (2005). One advantage of Equation (80) is that it is a very simple integral in the real space. The integrand is bounded in value, and in practice only a limited region needs to be integrated over due to the fact that the integrand decreases exponentially for large y . Our implementation shows that it is extremely fast with a computational time well below one second. Of course, the downside is that it is an approximation, accurate when η is small.

We use the Heston model to test the accuracy of the approximation for the European call price in Equation (80). The parameters used are: $\kappa = 2.0$, $\theta = 0.09$, $S_0 = 100$, $r = 0.03$, and $\delta = 0$. We vary the option maturity T , the instantaneous variance V_0 , the strike price K , and the volatility coefficient η . The results are reported in Table 2. The exact prices (FI) are computed from numerical Fourier inversion using the known characteristic function for the log stock price under the Heston model. See for example, Lewis (2000). Again, three numerical integration routines `quad`, `quadl` and `quadgk` in MATLAB are used to cross-verify the numerical integration results. As we see from Table 2, the approximation is very accurate for all parameter combinations of (T, V_0, K, η) we have considered.

3.3 Perpetual Timer Forward Pricing

Here we illustrate the usage of the perturbation result obtained in previous section with a perpetual timer forward. This is a contract to exchange one share of the underlying with K units of cash at a random future date τ^B at which the daily accumulated realized variance first exceeds a predefined variance budget B .

We use the same general stochastic volatility framework in the last subsection. By risk-neutral pricing, the price Π of the timer forward contract is given by

$$\Pi = \Pi(S_0, \xi, V) = \mathbb{E}_0[e^{-r\tau^B}(S_{\tau^B} - K)]. \quad (81)$$

Notice that because of the perpetual nature, Π does not depend on the calendar time t .

In Appendix, we show that we can simplify the above expectation to be the following

$$\Pi = \Pi(S_0, \xi, V) = S_0 \tilde{\mathbb{E}}_0[e^{-\delta\tau^B}] - K \mathbb{E}_0[e^{-r\tau^B}], \quad (82)$$

where $\tilde{\mathbb{E}}_0$ is taken under the measure $\tilde{\mathbb{Q}}$ in which the instantaneous variance process follows

$$dV = \tilde{a}(V)dt + \eta b(V)d\tilde{W}. \quad (83)$$

Here the modified drift function is given by

$$\tilde{a}(V) \equiv a(V) + \rho\eta\sqrt{V}b(V), \quad (84)$$

and \tilde{W} is a Brownian motion under measure $\tilde{\mathbb{Q}}$.

Equation (82) is exact and allows us to use the asymptotic expansion in Proposition 1. This gives the following approximation for Π :

$$\Pi \approx S_0 e^{-\delta\tilde{\mu}(B) + \frac{1}{2}\delta^2\tilde{\sigma}^2(B)} - K e^{-r\mu(B) + \frac{1}{2}r^2\sigma^2(B)}, \quad (85)$$

where $\mu(B)$ and $\sigma^2(B)$ are the approximate mean and variance of τ^B under measure \mathbb{Q} , and $\tilde{\mu}(B)$ and $\tilde{\sigma}^2(B)$ are the approximate mean and variance of τ^B under measure $\tilde{\mathbb{Q}}$. All four quantities can be computed using the method of characteristics illustrated in Appendix. For practical purposes, it is often easier to assume that $\tilde{a}(V)$ has no explicit η dependence. That is, we absorb the η dependence into other model parameters in $\tilde{a}(V)$. In many cases, $\tilde{a}(V)$ turns out to be formally identical to $a(V)$ so no additional effort is needed to compute $\tilde{\mu}(B)$ and $\tilde{\sigma}^2(B)$. This absorbing approach will produce a generalized asymptotic expansion series instead of the usual power series asymptotic expansion. It gives slightly different formulas for $\tilde{\mu}(B)$ and $\tilde{\sigma}^2(B)$, but to second order in η the results are equivalent. Notice that by Proposition 1, Π given above satisfies the pricing PDE to second order in η :

$$V\Pi_\xi + a(V)\Pi_V + \frac{1}{2}\eta^2 b^2(V)\Pi_{VV} + (r - \delta)S\Pi_S + \rho\eta\sqrt{V}b(V)\Pi_{SV} + \frac{1}{2}VS^2\Pi_{SS} - r\Pi = o(\eta^2), \quad (86)$$

with the boundary condition $\Pi(S, B, V) = S - K$. Therefore, for η not too large, the approximation in Equation (85) should be very accurate. Compared with the complicated PDE

above, the advantage of Equation (85) is obvious. While a typical implementation of a three-dimensional PDE can take multiple seconds¹⁴, the approximation in Equation (85) takes less than 10^{-4} seconds given the access of a fast algorithm for the product-log function needed in the Heston model. Fast algorithms for the product-log function are readily available in many software products. For example, in Mathematica, a single evaluation of the product-log function `ProductLog` takes about the same time as a cumulative normal distribution function.

From the above approximation, we can approximate the fair delivery price K_* as

$$K_* = S_0 e^{-\delta \tilde{\mu}(B) + \frac{1}{2} \delta^2 \tilde{\sigma}^2(B)} e^{r\mu(B) - \frac{1}{2} r^2 \sigma^2(B)} + o(\eta^2). \quad (87)$$

This is the delivery price that makes the forward contract to have value $\Pi = 0$ today.

For the Heston model and the 3/2 model, the computation of $\tilde{\mu}(B)$ and $\tilde{\sigma}^2(B)$ requires no extra effort once we have computed $\mu(B)$ and $\sigma^2(B)$, which are given in Examples 1 and 2 in the last section. This is because $\tilde{a}(V)$ takes exactly the same parametric form as $a(V)$ once we absorb the η dependence. For example, in the Heston model, we have

$$\tilde{a}(V) = \kappa(\theta - V) + \rho\eta V = \tilde{\kappa}(\tilde{\theta} - V), \quad (88)$$

where $\tilde{\kappa} = \kappa - \rho\eta$, and $\tilde{\theta} = \kappa\theta/\tilde{\kappa}$. The adjustment for the 3/2 model turns out to be identical. This is not a coincidence since the instantaneous variance process in the 3/2 model is the reciprocal of that in the Heston model. Since $\tilde{a}(V)$ is formally identical to $a(V)$ except with different parameters, $\tilde{\mu}(B)$ and $\tilde{\sigma}(B)$ are also formally identical to $\mu(B)$ and $\sigma(B)$. All we need to do to get $\tilde{\mu}(B)$ and $\tilde{\sigma}(B)$ is to replace κ with $\tilde{\kappa}$, and θ with $\tilde{\theta}$ in Examples 1 and 2. By absorbing a linear η term into $\tilde{a}(V)$, $\tilde{\mu}(B)$ and $\tilde{\sigma}(B)$ are now in the form of generalized asymptotic series.

We again carry out a numerical study for the accuracy of the timer forward fair delivery price using the Heston model. For simplicity we assume $\delta = 0$. The reason we choose this special case is that in this case we do not need to solve the three-dimensional PDE in Equation (86). Instead, the exact fair delivery price is computed by numerically solving the PDE that the quantity $\Upsilon(\xi, V) \equiv \mathbb{E}_0[e^{-r\tau^B}]$ satisfies:

$$V\Upsilon_\xi + a(V)\Upsilon_V + \frac{1}{2}\eta^2 b^2(V)\Upsilon_{VV} - r\Upsilon = 0, \quad (89)$$

with the boundary condition $\Upsilon(B, V) = 1$. The fair delivery price is given by

$$K_* = \frac{S_0}{\Upsilon}. \quad (90)$$

¹⁴Of course, this number is only a rough estimate. It depends heavily on the accuracy goal and on how sophisticated the implementation is, for example, whether GPU is used, or whether special structure of a PDE is exploited, etc.

This is easy to show from Equation (81) since in this zero-dividend case $\mathbb{E}_0[e^{-r\tau^B} S_{\tau^B}] = S_0$. To avoid implementation errors, we employ the standard routine `NDSolve` in Mathematica, which uses many different methods to numerically solve PDEs, including method of lines, implicit backward differentiation, Runge-Kutta methods, etc. We use a precision goal of 10^{-8} . Parameters such as step size are automatically chosen by the internal algorithms of `NDSolve`.

The results are reported in Table 3. We vary r , V_0 and η as in the table. Other parameters we use are $S_0 = 100$, $\kappa = 2.0$, $\theta = 0.09$, and $B = 0.09$. For each combination of parameter values (r, V_0, η) , we report both the exact fair delivery price (PDE) and the approximated one (`Approx`) in Equation (87). As we see in Table 3, the approximation is very accurate for all combinations of parameter values (r, V_0, η) .

3.4 Perpetual Timer Option Pricing

The perturbation developed in the last section can also be used to price more complicated derivatives. Here we demonstrate this by considering the pricing of perpetual timer call options in the Heston model with $\rho = 0$. We consider this special case since this is a nice application of the approximation we have developed.¹⁵

A perpetual timer call option pays $(S_{\tau^B} - K)^+$ and is only exercisable at time τ^B , where B is a contractual variance budget. Timer options have been studied in Bick (1995), Bernard and Cui (2011), Liang, Lemmens, and Tempere (2011), Li (2013), and Li and Mercurio (2013a, 2013b).

By risk-neutral pricing, the price of a timer call is given by

$$C_{\text{perp}} = \mathbb{E}_0 \left[e^{-r\tau^B} (S_{\tau^B} - K)^+ \right]. \quad (91)$$

By Ito's lemma, we have

$$\log S_t = \log S_0 + (r - \delta)t - \frac{1}{2} \int_0^t V_u du + \int_0^t \sqrt{V_u} dW_u^S. \quad (92)$$

Let \mathcal{F}^V be the filtration generated by the stochastic process V_u . That is, $\mathcal{F}^V = \sigma(V_u : u \geq 0)$. Notice τ^B is measurable in \mathcal{F}^V . Conditioning on \mathcal{F}^V , $\log S_{\tau^B}$ is normally distributed with mean $\log S_0 + (r - \delta)\tau^B - \frac{1}{2}B$ and variance B . Integrating out the randomness due to W_u^S in Equation (91) then gives

$$C_{\text{perp}} = \mathbb{E}_0 \left[\mathbb{E}_0[e^{-r\tau^B} (S_{\tau^B} - K)^+ | \mathcal{F}^V] \right] = \mathbb{E}_0 \left[C^{\text{BS}}(S_0, K, r, \delta, \tau^B, B) \right], \quad (93)$$

¹⁵Readers interested in extensions are referred to Li and Mercurio (2013a, 2013b), where the authors have developed approximations for both perpetual and finite-maturity timer options with general values of ρ . The approximations are found to be very accurate through extensive numerical study.

where C^{BS} is the Black-Scholes price defined in Equation (76).

So far the analysis is exact. We can now approximate C_{perp} by using the approximate distribution for τ^B in Proposition 1 to get

$$C_{\text{perp}} \approx \int_{-\infty}^{\infty} C^{\text{BS}}(S_0, K, r, \delta, u, B) n(u; \mu(B), \sigma^2(B)) du. \quad (94)$$

This integral can be evaluated in closed form to get (see, for example, Appendix of Li, Deng and Zhou (2008) for evaluating the integral)

$$C_{\text{perp}} \approx S_0 \cdot \mathbf{J}(a_+, b, -\delta, \mu(B), \sigma^2(B)) - K \cdot \mathbf{J}(a_-, b, -r, \mu(B), \sigma^2(B)), \quad (95)$$

where

$$a_{\pm} = \frac{\log(S_0/K)}{\sqrt{B}} \pm \frac{\sqrt{B}}{2}, \quad (96)$$

$$b = \frac{r - \delta}{\sqrt{B}}, \quad (97)$$

and the function \mathbf{J} is given by

$$\mathbf{J}(a, b, s, m, \Sigma) = e^{ms + \Sigma s^2/2} \cdot N\left(\frac{a + b(m + s\Sigma)}{\sqrt{1 + b^2\Sigma}}\right). \quad (98)$$

Interestingly, the timer call price approximation in Equation (95) resembles the Black-Scholes formula for plain-vanilla European-style options. It also has many attractive properties. First, when $r = \delta = 0$, the formula reduces to known exact result for timer call price. Second, when $\eta = 0$ so that the exercise time is deterministic, the formula reduces to the Black-Scholes formula for plain-vanilla options. Third, the Black-Scholes form makes it easier to compute the Greeks of the timer call option due to the following symmetry:

$$S_0 \frac{\partial}{\partial S_0} \mathbf{J}(a_+, b, -\delta, \mu(B), \sigma^2(B)) = K \frac{\partial}{\partial S_0} \mathbf{J}(a_-, b, -r, \mu(B), \sigma^2(B)). \quad (99)$$

This can be verified by tedious but straightforward calculation. For example, by using the above symmetry, the Delta is given by the following simple expression

$$\Delta_{C_{\text{perp}}} \approx \mathbf{J}(a_+, b, -\delta, \mu(B), \sigma^2(B)). \quad (100)$$

Notice that the right-hand-side is always positive. Therefore, the approximated timer call price is always positive and strictly increasing in S_0 , as it should be. Similarly, the Gamma can also be easily computed and seen to be positive. Finally, we emphasize that our formula is in closed form and therefore extremely fast to evaluate. When implemented in both Mathematica and MATLAB, our program takes less than 10^{-4} seconds to evaluate one timer option price.

We test the accuracy of the timer call price approximation in Equation (95) using the Heston model. The results are reported in Table 4. The fixed parameters are: $r = 0.03$, $\delta = 0$, $S_0 = 100$, $\kappa = 2.0$, $\theta = 0.09$, and $B = 0.09$. We vary the instantaneous variance V_0 , the strike K , and the volatility coefficient η . We use Monte Carlo prices (MC) as benchmarks. The standard errors from Monte Carlo simulations are all in the order of 10^{-3} or smaller. We use 4 million sample paths with a time step of every two hours. Only the instantaneous variance process needs to be simulated since the Brownian motion W_u^S can be integrated out. See Bernard and Cui (2013) for more details and for a more sophisticated implementation. The integrated variance ξ_T is used as a control variate for variance reduction, whose mean is given analytically in the Heston model. Our implementation takes about 30 minutes for each simulation with the desired accuracy goal. As we see, for all parameter combinations (V_0, K, η) , the approximation is extremely accurate. Unlike some approximation method such as moment matching, the approximation in Equation (95) is accurate for all values of K . In fact, when $\delta = 0$, the approximation goes to the correct limit 0 when K goes to infinity, and goes to the correct limit S_0 when K goes to 0.

4 Conclusion

Many derivatives products are directly or indirectly associated with integrated diffusion processes. The hitting time of such an integrated diffusion process plays an very important role in pricing those derivative products. Through perturbation technique, we show that for any diffusion processes, this hitting time is approximately normally distributed when the diffusion coefficient is small. This distributional approximation of the hitting times enables us to reduce many pricing problems to simpler one-dimensional expectations. We illustrate the generality and accuracy of this probabilistic approach using several examples.

The approach has several limitations which we acknowledge here. It cannot handle time-varying drift and diffusion functions right now. It also cannot handle jumps. It also requires that the diffusion coefficient be small which might not be satisfied in some practical applications. Also, although the method in principle works for any drift and diffusion functions, to render it effective in practice we still needs these functions to be simple to obtain closed-form formulas. Also, only Heston model with a limited set of parameters has been tested in this paper. To implement our approximation in product-quality code, extensive testing needs to be carried out beforehand, especially on the greeks.

There are several future research directions one can follow. First, it is useful to work out the approximation for more diffusion processes and find more interesting applications. Second, we have focused on small diffusion coefficient in this paper which could be a limitation in actual

applications. It would therefore be useful to derive other approximations under different assumptions such as strong or weak mean reversions in the drift. For specific models where some parameters can be assumed to be small, one might be able to develop approximations using the perturbation technique we have developed.

Appendix

Proof of Proposition 1:

Let $X^* > 0$ be an arbitrary integration limit. Define two dimensionless variables z_0 and z as follows

$$z_0 = \frac{X - X^*}{X^*}, \quad (101)$$

$$z = \Phi \left(B - \xi + \int_{X^*}^X \frac{u}{a(u)} du \right). \quad (102)$$

We remark that in a concrete model, it might be easier to just use $z_0 = X - X^*$ without the scaling or simply $z_0 = X$. The function Φ is most conveniently chosen to be

$$\Phi(x) = \frac{f^{-1}(x) - X^*}{X^*}, \quad (103)$$

where $f^{-1}(x)$ is the inverse of the function

$$f(X) = \int_{X^*}^X \frac{u}{a(u)} du. \quad (104)$$

With this choice of Φ , it is easy to see that $z = z_0$ when $\xi = B$. This allows us to write the solution of the ordinary differential equation in simple integral form, as we will see below.

The purpose of the above definitions is to change the coordinates from (ξ, V) to the characteristic coordinate (z_0, z) . Let us express p in terms of the characteristic coordinates z and z_0 :

$$p(\xi, X) = \tilde{p}(z, z_0). \quad (105)$$

It is easy to check that

$$X z_\xi + a(X) z_X = 0. \quad (106)$$

By utilizing the above equation and chain differential rule, to zeroth-order in η , the PDE for $p(\xi, X)$ simplifies to an ODE for $\tilde{p}(z, z_0)$ due to cancelation of the \tilde{p}_z terms:

$$\tilde{a}(z_0) \tilde{p}_{z_0} + \lambda = 0, \quad (107)$$

where $\tilde{a}(z_0) \equiv a(X)$, and with the boundary condition

$$\tilde{p}(z_0, z_0) = 0. \quad (108)$$

Therefore, the zeroth-order approximation is given by a simple integral

$$p_0(\xi, X) = \tilde{p}(z, z_0) = \int_{z_0}^z \frac{\lambda}{\tilde{a}(u)} du \equiv \lambda T_0(\xi, X). \quad (109)$$

When $\eta = 0$, the process X_t is deterministic, and the quantity $T_0(\xi, X)$ is just the deterministic time to exceed the budget B . Therefore, we have $T_0(\xi, X) \geq 0$ with $T_0(\xi, X) = 0$ if and only if $\xi = B$. Notice that the solution of T_0 is given by

$$T_0(\xi, X) = \int_{z_0}^z \frac{1}{\tilde{a}(u)} du. \quad (110)$$

To second order in η , the second-order PDE for $p(\xi, X)$ can be approximated by a first-order ODE for $\tilde{p}(z, z_0)$ as follows

$$\tilde{a}(z_0)\tilde{p}_{z_0} + \lambda + \frac{1}{2}\eta^2\tilde{b}^2(z_0) [\tilde{p}_{0,XX} + (\tilde{p}_{0,X})^2] = 0. \quad (111)$$

Here $\tilde{b}(z_0) \equiv b(X)$, and $\tilde{p}_{0,X}$ denotes the partial derivative of $p_0(\xi, X)$ with respect to X , but expressed in terms of the characteristic coordinates z and z_0 . Similarly for $\tilde{p}_{0,XX}$. The solution of $p(\xi, X)$ to second order in η is then given by a simple integral

$$p(\xi, X) = \tilde{p}(z, z_0) = \int_{z_0}^z \frac{\lambda + \frac{1}{2}\eta^2\tilde{b}^2(u) [\lambda\tilde{T}_{0,XX} + \lambda^2(\tilde{T}_{0,X})^2]}{\tilde{a}(u)} du + o(\eta^2) \quad (112)$$

$$\equiv \lambda[T_0 + \eta^2 H_0] + \lambda^2\eta^2 H_1 + o(\eta^2). \quad (113)$$

By comparing the last two equalities, the functions H_0 and H_1 are given by the following integrals

$$H_0 = \int_{z_0}^z \frac{\tilde{b}^2(u)\tilde{T}_{0,XX}}{2\tilde{a}(u)} du, \quad (114)$$

$$H_1 = \int_{z_0}^z \frac{\tilde{b}^2(u)(\tilde{T}_{0,X})^2}{2\tilde{a}(u)} du. \quad (115)$$

This proves that $\mathbb{E}_0[e^{\lambda\tau^B}]$ possesses an asymptotic expansion as follows

$$\mathbb{E}_0[e^{\lambda\tau^B}] = e^{\lambda[T_0 + \eta^2 H_0] + \lambda^2\eta^2 H_1} + o(\eta^2). \quad (116)$$

Since by assumption $\mathbb{E}_0[e^{\lambda\tau^B}\tau_B]$ and $\mathbb{E}_0[e^{\lambda\tau^B}\tau_B^2]$ are finite for a range of real λ values containing 0, differentiating the above expansion series once and twice and setting $\lambda = 0$ gives the asymptotic expansion series for the first two moments of τ^B as follows

$$\mathbb{E}_0[\tau^B] = T_0 + \eta^2 H_0 + o(\eta^2), \quad (117)$$

$$\mathbb{E}_0[(\tau^B)^2] = (T_0 + \eta^2 H_0)^2 + 2\eta^2 H_1 + o(\eta^2). \quad (118)$$

Therefore, the variance of τ^B has the following asymptotic expansion series

$$\text{Var}(\tau^B) = 2\eta^2 H_1 + o(\eta^2). \quad (119)$$

Dividing both sides by η^2 and taking the limit of $\eta \rightarrow 0^+$, we see that $H_1 \geq 0$. Furthermore, $H_1 = 0$ if and only if $\text{Var}(\tau^B) = 0$. \square

Proof of Equation (82):

This is essentially a change of measure technique. See Li and Mercurio (2013a). However, because random time is involved, it is more elementary to prove the equation from the PDE perspective. To show Equation (82), let

$$\Phi(S, \xi, V) = S\Psi(\xi, V) \equiv \mathbb{E}_0[e^{-r\tau^B} S_{\tau^B}]. \quad (120)$$

Notice that Ψ does not depend on S because of the homogeneity of the payoff $e^{-r\tau^B} S_{\tau^B}$ and the homogeneity of the stock process. By the Feynman-Kac theorem applied using the exit time τ^B , $\Phi(S, \xi, V)$ satisfies the following PDE

$$V\Phi_\xi + a(V)\Phi_V + \frac{1}{2}\eta^2 b^2(V)\Phi_{VV} + (r - \delta)S\Phi_S + \rho\eta\sqrt{V}b(V)\Phi_{SV} + \frac{1}{2}VS^2\Phi_{SS} - r\Phi = 0, \quad (121)$$

with the boundary condition $\Phi(S, B, V) = S$, and where we use subscripts to denote partial derivatives. Since $\Phi(S, \xi, V) = S\Psi(\xi, V)$, we can easily show that $\Psi(\xi, V)$ satisfies the following PDE

$$V\Psi_\xi + \tilde{a}(V)\Psi_V + \frac{1}{2}\eta^2 b^2(V)\Psi_{VV} - \delta\Psi = 0, \quad (122)$$

with the boundary condition $\Psi(B, V) = 1$. A reverse application of the Feynman-Kac theorem using again the exit time τ^B now shows that

$$\Psi(\xi, V) = \tilde{\mathbb{E}}_0[e^{-\delta\tau^B}]. \quad (123)$$

\square

References

- Ahn, D-H., B. Gao. 1999. A parametric nonlinear model of term structure dynamics. *Review of Financial Studies* 12(4), 721–762.
- Bernard, C., Z. Cui. 2011. Pricing timer options. *Journal of Computational Finance* 15(1), 69–104.
- Bick, A. 1995. Quadratic-variation-based dynamic strategies. *Management Science* 41(4), 722–732.
- Carr, P., and D.B. Madan. Option valuation using the fast Fourier transform. *Journal of Computational Finance* 2(4), 61–73.
- Corless, R.M., G. H. Gonnet, D. E. G. Hare, D. J. Jeffrey, D. E. Knuth. 1996. On the LambertW function. *Advances in Computational Mathematics* 5(1), 329–359.
- Cox, J.C., J.E. Ingersoll, and S.A. Ross. 1985. A theory of the term structure of interest rates. *Econometrica* 53(2), 385–408.
- Cui, Z. 2013. Martingale property and pricing for time-homogeneous diffusion models in finance. Ph.D. Dissertation, University of Waterloo.
- Dufresne, D. 2001. The integrated square-root process. Working paper, University of Montreal.
- Forde, M., and A. Jacquier. 2010. Robust approximations for pricing Asian options and volatility swaps under stochastic volatility. *Applied Mathematical Finance* 17(3), 241–259.
- Gatheral, J. 2006. *The Volatility Surface: A Practitioner’s Guide*. John Wiley & Sons.
- Gregory, J. 2012. *Counterparty Credit Risk and Credit Value Adjustment: A Continuing Challenge for Global Financial Markets*. John Wiley & Sons.
- Heston, S.L. 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6(2), 327–343.
- Hull, J., and A. White. 1987. The pricing of options on assets with stochastic volatilities. *Journal of Finance* 42(2), 281–300.
- Kahl, C., and P. Jäckel. 2005. Not-so-complex logarithms in the Heston model. *Wilmott Magazine* 19(9), 94–103.
- Karlin, S., and H.M. Taylor. 1991. *A Second Course in Stochastic Processes*. Academic Press.
- Lewis, A.L. 2000. *Option Valuation Under Stochastic Volatility*. Finance Press, Newport Beach, California.
- Li, C.X. 2013. Bessel processes, stochastic volatility, and timer options. *Mathematical Finance*, forthcoming. DOI:10.1111/mafi.12041

- Li, M., S.-J. Deng, J. Zhou. 2008. Closed-form approximations for spread option prices and greeks. *Journal of Derivatives* 16(4), 58-80.
- Li, M., F. Mercurio. 2013a. Closed-form approximation of timer option prices under general stochastic volatility models. Working paper.
- Li, M., F. Mercurio. 2013b. Analytic approximation of finite-maturity timer option prices. Working paper.
- Li, M., F. Mercurio. 2013c. Time for a timer. *Risk* (December 2013), 88–93.
- Liang, L.Z.J., D. Lemmens, and J. Tempere. 2011. Path integral approach to the pricing of timer options with the Duru-Kleinert time transformation. *Physical Review E* (83).
- Lipton, A. 2001. *Mathematical Methods for Foreign Exchange: a Financial Engineers Approach*. World Scientific.
- Sawyer, N. 2007. SG CIB launches timer options. *Risk* 20(7), 6.
- Sawyer, N. 2008. Equity derivative house of the year: Société Générale. *Risk* 21(1).
- Schürger, K. 2002. Laplace transforms and suprema of stochastic processes. In: *Advances in Finance and Stochastics*, Springer. 285–294
- Schöbel, R., J. Zhu. 1999. Stochastic volatility with an Ornstein–Uhlenbeck process: An extension. *Review of Finance* 3(1), 23–46.
- Stein, E.M., and J.E. Stein. 1991. Stock price distributions with stochastic volatility. *Review of Financial Studies* 4, 727–752.
- Wolfe, S.J. 1975. On moments of probability distribution functions. In: *Fractional Calculus and Its Applications* (Ed: B. Ross). Lecture Notes in Mathematics 457, 306–316.

Table 1: **Volatility Swap Rates in the Heston Model**

We report the volatility swap rates for both our approximation (**Approx**) and the benchmark (**NI**) from numerical inversion. We vary T and V_0 . The other parameters used are: $\kappa = 2.0$ and $\theta = 0.09$. All volatility swap rates reported are in percentage terms.

T	V_0	$\eta = 0.125$		$\eta = 0.250$		$\eta = 0.375$	
		Approx	NI	Approx	NI	Approx	NI
0.5	0.04	24.10	24.10	23.95	23.93	23.73	23.65
	0.09	29.95	29.94	29.82	29.78	29.67	29.52
	0.16	36.59	36.59	36.52	36.45	36.44	36.21
1.0	0.04	26.09	26.07	25.90	25.90	25.62	25.61
	0.09	29.94	29.92	29.76	29.76	29.47	29.47
	0.16	34.62	34.60	34.45	34.45	34.16	34.17
1.5	0.04	27.83	27.83	27.67	27.68	27.41	27.44
	0.09	29.95	29.95	29.79	29.80	29.53	29.56
	0.16	32.69	32.69	32.53	32.54	32.27	32.31

Table 2: **European Call Price in the Heston Model**

We report call prices from both our approximation (**Approx**) and the benchmark (**FI**) from Fourier inversion. We vary T , V_0 , K and η . The other parameters used are: $\kappa = 2.0$, $\theta = 0.09$, $S_0 = 100$, $r = 0.03$, and $\delta = 0$.

T	V_0	K	$\eta = 0.125$		$\eta = 0.250$		$\eta = 0.375$	
			Approx	FI	Approx	FI	Approx	FI
0.5	0.04	80	21.76	21.76	21.77	21.77	21.82	21.80
		100	7.51	7.51	7.45	7.46	7.40	7.38
		120	1.61	1.61	1.60	1.60	1.63	1.60
	0.09	80	22.42	22.42	22.44	22.43	22.48	22.44
		100	9.13	9.13	9.10	9.09	9.06	9.02
		120	2.81	2.81	2.81	2.80	2.82	2.78
	0.16	80	23.41	23.41	23.42	23.40	23.45	23.40
		100	10.98	10.98	10.96	10.94	10.93	10.87
		120	4.38	4.38	4.38	4.36	4.39	4.33
1.0	0.04	80	24.39	24.39	24.39	24.39	24.40	24.39
		100	11.77	11.77	11.70	11.70	11.59	11.59
		120	4.86	4.86	4.82	4.82	4.76	4.75
	0.09	80	25.28	25.28	25.27	25.27	25.25	25.25
		100	13.26	13.26	13.19	13.19	13.08	13.08
		120	6.27	6.27	6.22	6.22	6.14	6.14
	0.16	80	26.47	26.47	26.45	26.45	26.41	26.41
		100	15.07	15.07	15.00	15.00	14.88	14.89
		120	8.05	8.05	7.99	8.00	7.90	7.91
1.5	0.04	80	26.97	26.97	26.95	26.95	26.92	26.92
		100	15.24	15.24	15.16	15.16	15.03	15.04
		120	8.01	8.01	7.94	7.94	7.83	7.84
	0.09	80	27.83	27.83	27.80	27.81	27.76	27.77
		100	16.53	16.53	16.45	16.45	16.31	16.33
		120	9.33	9.33	9.26	9.26	9.13	9.15
	0.16	80	28.98	28.98	28.95	28.95	28.88	28.90
		100	18.16	18.16	18.08	18.08	17.94	17.96
		120	11.01	11.01	10.94	10.94	10.80	10.83

Table 3: **Timer Forward Fair Delivery Price in the Heston Model**

We report the timer forward prices for both our approximation (Approx) and the benchmark (PDE) from numerical PDE solutions. We vary r , V_0 and η . The other parameters used are: $\kappa = 2.0$, $\theta = 0.09$, $S_0 = 100$, and $B = 0.09$.

r	V_0	$\eta = 0.125$		$\eta = 0.250$		$\eta = 0.375$	
		Approx	PDE	Approx	PDE	Approx	PDE
0.03	0.04	103.89	103.89	104.03	104.03	104.27	104.24
	0.09	103.09	103.09	103.24	103.23	103.49	103.44
	0.16	102.18	102.19	102.31	102.31	102.52	102.50
0.06	0.04	107.92	107.94	108.21	108.21	108.71	108.63
	0.09	106.28	106.28	106.59	106.56	107.09	106.99
	0.16	104.41	104.42	104.67	104.66	105.10	105.06

Table 4: **Timer Call Price in the Heston Model**

We report the timer option prices for both our approximation (**Approx**) and the benchmark (**MC**) from Monte Carlo simulations. We vary V_0 , K , and η . The other parameters used are: $\kappa = 2.0$, $\theta = 0.09$, $S_0 = 100$, $r = 0.03$, $\delta = 0$, and $B = 0.09$. Standard errors from Monte Carlo simulations are all in the order of 10^{-3} or smaller.

V_0	K	$\eta = 0.125$		$\eta = 0.250$		$\eta = 0.375$	
		Approx	MC	Approx	MC	Approx	MC
0.04	80	25.76	25.76	25.85	25.84	25.99	25.96
	100	13.67	13.66	13.73	13.73	13.85	13.83
	120	6.54	6.54	6.58	6.58	6.66	6.65
0.09	80	25.31	25.31	25.40	25.39	25.54	25.51
	100	13.31	13.31	13.38	13.37	13.49	13.47
	120	6.31	6.31	6.35	6.35	6.43	6.41
0.16	80	24.79	24.79	24.87	24.86	24.99	24.98
	100	12.90	12.89	12.95	12.95	13.05	13.05
	120	6.04	6.04	6.08	6.08	6.14	6.14

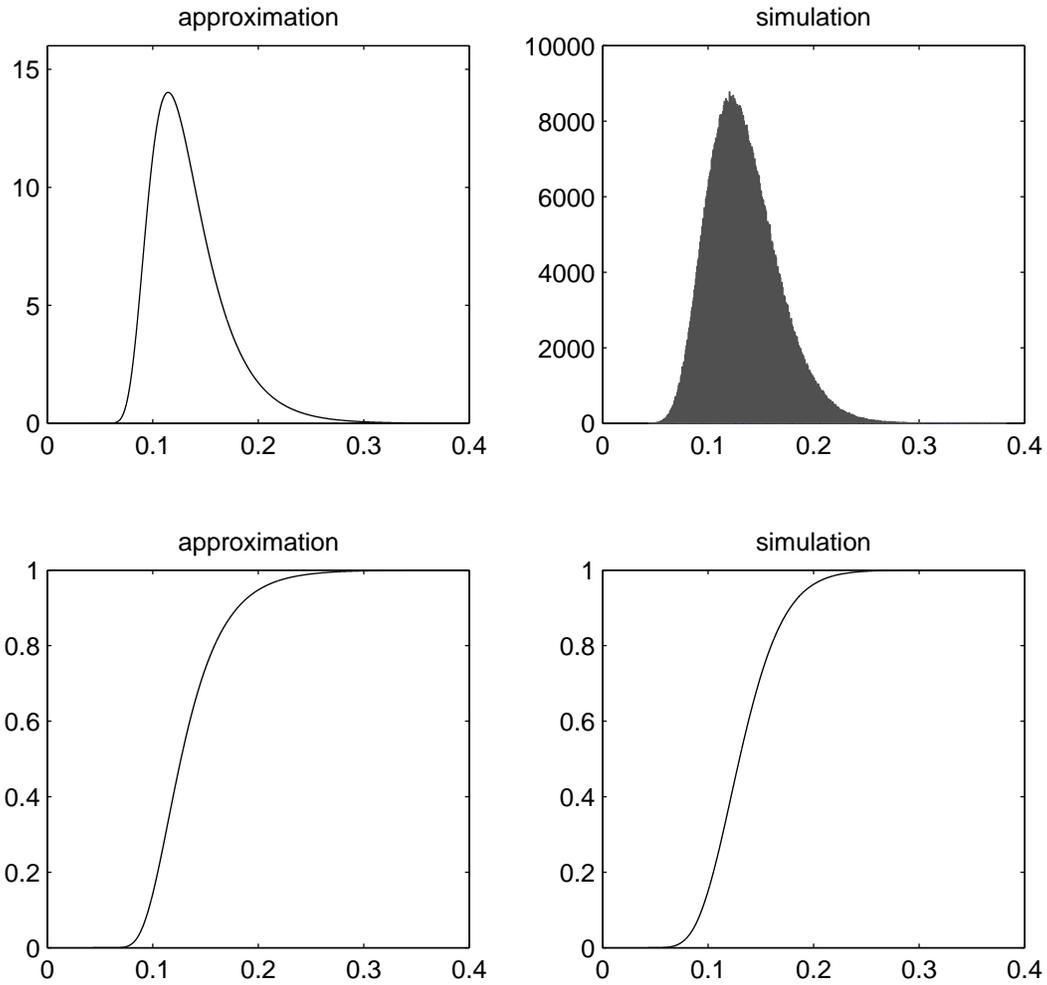


Figure 1: **Approximated and simulated distributions of the accumulated variance ξ_T .** The left two plots are the probability density and cumulative distribution function from the approximation. The right two plots are the histogram and empirical distribution function from Monte Carlo simulation.