



Munich Personal RePEc Archive

Analytic Approximation of Finite-Maturity Timer Option Prices

Li, Minqiang

Bloomberg LP

March 2014

Online at <https://mpra.ub.uni-muenchen.de/54597/>

MPRA Paper No. 54597, posted 19 Mar 2014 15:40 UTC

Analytic Approximation of Finite-Maturity Timer Option Prices

Minqiang Li*, Fabio Mercurio[†]

March 19, 2014

Abstract

We develop an approximation technique for pricing finite-maturity timer options under Heston-like stochastic volatility models. By approximating the distributions of the accumulated variance and the random variance budget exceeding time, we obtain analytic expressions for timer option prices under zero correlation. For nonzero correlation, we use a simple linear combination approximation which matches the asymptotic correlation behavior. Numerical analysis using the Heston model shows that the method is fairly accurate, especially when the volatility of variance is small or the maximum maturity is large.

Keywords: Finite-Maturity Timer Options; Analytic Approximation; Perturbation; Hitting time; Integrated Diffusion

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. E-mail: minqiang.li@gmail.com.

[†]Derivatives Research, Bloomberg LP. E-mail: fmercurio@bloomberg.net.

1 Introduction

In 2007, Société Générale Corporate and Investment Banking introduced a new variance derivative to the market called “timer options” (See Sawyer (2007)). A timer option is similar to a plain-vanilla option, except that the expiry is not deterministic. Rather it is specified as the first time when the accumulated realized variance exceeds a given budget. See Hawkins and Krol (2008), Carr and Lee (2010) and Lee (2012) for some insights on why such a product can be attractive to investors. Interestingly, timer options were studied many years ago by Neuberger (1990) and Bick (1995) when such security did not even exist in the market place. Sawyer (2008) also reported that Société Générale Corporate and Investment Banking started to sell other timer-style options such as “timer out-performance options” and “timer swaps”.

The pricing of timer options is usually done through Monte Carlo simulation. A naive implementation can take tens of hours to get a reasonably accurate estimate of the option price. See C.X. Li (2010, 2013) and Bernard and Cui (2011) for techniques to speed up the simulation. Another approach is to compute the timer option price exactly through multidimensional numerical integration. See Liang, Lemmens and Tempere (2011) where the authors provide such an integration for the Heston model and the 3/2 model. One shortcoming besides being complex and slow is that this method works for a limited set of models. A third approach is to approximate timer option prices through methods such as perturbation. Saunders (2010) considers an asymptotic expansion for stochastic volatility models under fast mean-reversion. Another paper in this category which is very closely related to the current one is Li and Mercurio (2013). The method in Li and Mercurio (2013) is extremely fast and very accurate, with percentage errors in the order of 0.01% for the Heston model for the parameters they used.

However, the method in Li and Mercurio (2013) is designed to work for perpetual timer options. In practice, a timer option often has a contractual maximum maturity. The current paper aims at developing an approximation which works for finite-maturity timer options. Unlike some previous papers which focus on a particular model, we assume a general stochastic volatility framework with a nonzero dividend rate. Our approach is carried out in two steps. We first develop the approximation assuming a zero correlation between the two Brownian motions. We then build upon this result to price timer options under nonzero correlation.

Under zero correlation, we reduce the finite-maturity timer option prices to simple univariate expectations. Those expectations are taken over the accumulated variance and the random variance budget exceeding time. Li and Mercurio (2013) shows that the variance budget hitting times are approximately normal in both the Heston and 3/2 models under small volatility of variance. M. Li (2013) generalizes this result to show that it holds for all Heston-like stochastic volatility models. A duality result between the cumulative distribution functions

of the variance budget exceeding time and accumulated variance then allows one to obtain an approximation for the latter. Utilizing these distribution approximations, the finite-maturity timer option price is expressed as a sum of two components, one coming from exercising the option when the maximum maturity is reached, and another coming from exercising when the variance budget is exceeded.

For nonzero correlation, we take a simple approach by assuming that the finite-maturity timer option price is a linear combination of the corresponding perpetual timer option and plain-vanilla option prices, where the weights are simultaneously replaced by their zero-correlation counterparts. This approximation has an attractive feature that it reduces to the correct price when the maximum maturity is much larger than the target exercise time implied by the variance budget, or vice versa. Therefore, it matches the asymptotic correlation behavior in these two extreme cases. By construction, this approximation is exact when the correlation coefficient is zero. It is also exact when the variance process is deterministic.

The proposed method for pricing finite-maturity timer options has several attractive properties. First, the method is extremely fast. Most of the computational cost is a one-dimensional numerical integration. Second, numerical study shows that the approximate timer option prices are fairly accurate. The approximation is extremely accurate when the volatility of variance is small or when the maximum expiry is not tightly binding. Third, the approximate prices reduce to the correct exact formulas under many limiting cases. Finally, the method in principle works for any Heston-like stochastic volatility model. For any such model, the task is to compute the approximate mean and variance of the random variance budget exceeding time. Once these two quantities are computed, it becomes mechanical to just plug them into our approximation formula to compute the timer option price.

There are however some limitations in the current paper. First, although the method works for any Heston-like model, in practice it is attractive for the variance process to be simple so that the approximate mean and variance of the hitting time can be computed in closed form. Second, our approximation is derived under the assumption of a small diffusion coefficient, which might not be available in some real-life applications. Finally, in this paper, only the Heston model with a limited set of parameters has been tested for accuracy.¹

The rest of the paper is organized as follows. Section 2 describes the model setup and develops the approximation for finite-maturity timer option prices, first under the assumption of a zero correlation and then for general nonzero correlations. Section 3 performs a numerical study on the accuracy of the approximation. Section 4 concludes.

¹Some sensitivity analysis has been carried out in the numerical analysis section by varying individual parameters. See this section for the parameter ranges we have tested. We also remark that Li and Mercurio (2013) have tested the accuracy of the hitting time approximation in the case of perpetual timer options for Heston model with different parameters, and for the 3/2-model. The results are found to be fairly accurate.

2 Pricing Timer Options

2.1 The Model Setup

We consider a general time-homogenous stochastic volatility framework. Under the risk-neutral measure, the stock process S_u and instantaneous variance process V_u evolve as

$$dS_u = (r - \delta)S_u du + \sqrt{V_u}S_u dW_u^S, \quad (1)$$

$$dV_u = a(V_u) du + \eta b(V_u) dW_u^V. \quad (2)$$

Here we assume that the interest rate r and dividend rate δ are both constants, and the two Brownian motions W_u^S and W_u^V are correlated with constant coefficient ρ . The functions $a(V)$ and $b(V)$ are assumed to be such that the process V_u has a domain $(0, \infty)$, and does not explode to either 0 or ∞ . We have isolated out the volatility of variance coefficient η in the diffusion function. Similar setup for pricing timer options is also used in Li (2010), Bernard and Cui (2011), as well as Liang, Lemmens and Tempere (2011). However, unlike in these papers, we assume a nonzero dividend rate δ which is more relevant in practice. This framework incorporates the Heston model (Heston 1993) and the 3/2-model as special cases.

Now define the accumulated variance process to be

$$\xi_u = \xi + \int_0^u V_s ds. \quad (3)$$

Here $\xi_0 = \xi$ is the accumulated variance at time 0. Let τ^B be the random time remaining for a pre-specified variance budget $B > 0$ to be exceeded. That is,

$$\tau^B \equiv \inf \{u > 0 : \xi_u = B\} = \inf \left\{ u > 0 : \int_0^u V_s ds = B - \xi \right\}. \quad (4)$$

A finite-maturity timer call option has a mandated maximum maturity T . The option matures when the variance budget is exhausted, but no later than the prescribed maximum maturity T . That is, the finite-maturity timer option is exercisable at the random time τ , where

$$\tau = \tau^B \wedge T \equiv \min(\tau^B, T). \quad (5)$$

When the option is exercised at random time τ , the buyer of the option receives a payoff $(S_\tau - K)^+$ where K is the strike price specified in the contract. Similarly, a timer put option pays $(K - S_\tau)^+$ when exercised at time τ . A perpetual timer option is similar to a finite-maturity timer option except that there is no contractual maximum maturity. So a perpetual timer option is exercisable at exactly the random time τ^B .

In practice, the realized variance up to day N is computed using the daily closing prices as follows:

$$\sum_{i=1}^N \left[\log \left(\frac{S_{t_i}}{S_{t_{i-1}}} \right) \right]^2. \quad (6)$$

This daily accumulated variance is used in the timer option contract to determine the random exercise time. The approximation of using ξ_u instead of the above discrete sum is standard in the variance derivative pricing literature. The discretization bias induced is usually very small with daily accumulation. The variance budget B is usually quoted according to a target expiration T_{target} as $B = \sigma_{\text{imp}}^2 T_{\text{target}}$, where σ_{imp} is the implied volatility for a European option on the underlying with maturity T_{target} and the same strike K . On contract initiation for a finite-maturity timer option, the maximum maturity T is usually much larger than T_{target} . In the term sheets we have seen, it can be as large as six times the target expiration.

Assume that the current time is 0. For notational simplicity, we will assume that $\xi = 0$ so that B should be interpreted as the remaining variance budget $B - \xi$. Let the current stock price be S_0 and the current instantaneous variance be V_0 . Let C_{fin} denote the timer call option price with a remaining maximum maturity $T > 0$ and a remaining variance budget $B > 0$. Similarly for P_{fin} . By risk-neutral pricing, we have

$$C_{\text{fin}} = \mathbb{E}_0 [e^{-r\tau}(S_\tau - K)^+]. \quad (7)$$

For the put, we have $P_{\text{fin}} = \mathbb{E}_0 [e^{-r\tau}(K - S_\tau)^+]$. The perpetual option prices are given by

$$C_{\text{perp}} = \mathbb{E}_0 [e^{-r\tau^B}(S_{\tau^B} - K)^+], \quad P_{\text{perp}} = \mathbb{E}_0 [e^{-r\tau^B}(K - S_{\tau^B})^+]. \quad (8)$$

If one has the joint distribution of (τ, S_τ) or (τ^B, S_{τ^B}) , the pricing of timer option is a simple two-dimensional integration. In practice, such a joint distribution is often not readily available. This makes the pricing of timer options a difficult problem.

Another approach is through solving the partial differential equations the timer option prices satisfy. For the finite-maturity timer option price, a simple application of the general Feynman-Kac theorem using the exit time τ gives the following equation

$$C_t + VC_\xi + a(V)C_V + \frac{1}{2}\eta^2 b^2(V)C_{VV} + (r - \delta)SC_S + \frac{1}{2}VS^2C_{SS} + \rho\eta\sqrt{V}b(V)SC_{SV} - rC = 0, \quad (9)$$

with the boundary conditions $C(T, S, \xi, V) = (S - K)^+$, and $C(t, S, B, V) = (S - K)^+$. Here for simplicity, we have written the timer call price as $C(t, S, \xi, V)$ and dropped the dependence on other parameters. The equation for the timer put is identical but with boundary conditions replaced by $(K - S)^+$. A similar application of the general Feynman-Kac theorem using the exit time τ^B gives the partial differential equation for the perpetual timer call:

$$VC_\xi + a(V)C_V + \frac{1}{2}\eta^2 b^2(V)C_{VV} + (r - \delta)SC_S + \frac{1}{2}VS^2C_{SS} + \rho\eta\sqrt{V}b(V)SC_{SV} - rC = 0, \quad (10)$$

with the boundary condition $C(S, B, V) = (S - K)^+$.

2.2 Approximating the Distributions of ξ_T and τ^B

Li and Mercurio (2013) develops an approximation for pricing perpetual timer options by perturbing the partial differential equation it satisfies under the assumption that η is small. The approximation is found to be extremely fast and accurate. Therefore, in this paper, we focus on approximating the finite-maturity timer option prices. We use the same assumption that η is small as in Li and Mercurio (2013). It turns out to be difficult to extend the method in Li and Mercurio (2013) directly to finite-maturity timer options because the finite-maturity timer option prices satisfy a higher-dimensional partial differential equation. Therefore, we take a different and more probabilistic approach below.

We start by approximating the probability densities of τ^B and ξ_T under the assumption that η is small. We then use these densities approximations to approximate the timer option prices. Small η expansion is also considered for plain-vanilla options in Lewis (2000) and Lipton (2001). We make an important remark that here by small η , we do not mean that η has to be smaller than 1. Rather by the loose term “small η ”, we mean that the effect of η is small. We will comment more on this point later.

In Li and Mercurio (2013), it is shown that for the Heston model and the 3/2 model, τ^B is approximately normally distributed for small η . M. Li (2013) shows that this is true for the integrated process of any time-homogeneous diffusion process. The result is summarized below. Readers interested should refer to M. Li (2013) for a detailed proof.

Proposition 1. *Assume that η is small in $dV_u = a(V_u) du + \eta b(V_u) dW_u^V$. Assume further 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 any $\lambda \in \mathbb{R}$. 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 deterministic functions of B and V_0 , and not functions of λ or η . Furthermore, we have $T_0 \geq 0$ and $H_1 \geq 0$ with equality if and only if $B = 0$. Therefore, for $B > 0$, τ^B is approximately normally distributed with mean $\mu(B)$ and variance $\sigma^2(B)$, where $\mu(B)$ and $\sigma^2(B)$ are given by

$$\mu(B) = T_0 + \eta^2 H_0, \quad \sigma^2(B) = 2\eta^2 H_1. \quad (12)$$

For any $T > 0$, the cumulative distribution function $F_{\xi_T}(x) \equiv \mathbb{P}(\xi_T < x)$ of the accumulated variance ξ_T can be approximated as:

$$F_{\xi_T}(x) \approx \tilde{F}_{\xi_T}(x) \equiv N\left(\frac{\mu(x) - T}{\sigma(x)}\right). \quad (13)$$

Equation (13) follows rigorously from the small- η perturbation in Appendix of M. Li (2013), and there is no additional distribution approximation here. Also, while the distribution of τ^B is approximately normal, the distribution of ξ_T obtained from duality is not normal even though it has a form of cumulative normal distribution function. In fact, in the Heston model, numerically we find that the approximate distribution of ξ_T above looks like a shifted noncentral-chi-squared distribution, and matches the true distribution very well.

For small η , the above approximation is reasonably good and captures some important features of the simulated distribution of ξ_T . Notice that in the limit $\eta = 0$ so that the instantaneous variance process is deterministic, M_{τ^B} reduces to that of a Dirac-delta function at T_0 , as it should be.

The approximate mean $\mu(B)$ and variance $\sigma^2(B)$ are functions of V_0 and other parameters. For notational simplicity, we only emphasize their dependence on B above. M. Li (2013) shows a procedure to compute $\mu(B)$ and $\sigma(B)$. 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 involving the characteristic coordinates. For many models, the integrals can be done explicitly. We give two examples below.² Readers interested in the details of the calculation are referred to Li and Mercurio (2013) or M. Li (2013). For simplicity, we assume that currently $\xi = 0$ so the quantity B in formulas below should be interpreted as the remaining variance budget $B - \xi$.

Example 1: (Heston model where $dV = \kappa(\theta - V)dt + \eta\sqrt{V} dW^V$)

We assume Feller condition is satisfied. The functions T_0 , H_0 and H_1 are given by

$$T_0 = \frac{1}{\kappa} \log R, \quad (14)$$

$$H_0 = \frac{(R-1)[2R^2z^2 + R(2-5z-2z^2) - 2 - z]}{4\kappa^2R^2(1+z)^3\theta} + \frac{3z \log R}{2\kappa^2(1+z)^3\theta}, \quad (15)$$

$$H_1 = \frac{(R-1)(1+2R^2z + R(2z-3))}{4\kappa^3R^2(1+z)^2\theta} - \frac{(2z-1) \log R}{2\kappa^3(1+z)^2\theta}, \quad (16)$$

with

$$R = e^{z-z_0+\kappa\frac{B}{\theta}}, \quad z_0 \equiv \frac{V_0 - \theta}{\theta}, \quad z \equiv W\left(z_0 e^{z_0} \cdot e^{-\kappa\frac{B}{\theta}}\right), \quad (17)$$

where $W(\cdot)$ is Lambert's product log function defined implicitly as $x = W(x)e^{W(x)}$.

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

$$\theta T_0 + (V_0 - \theta) \frac{1 - e^{-\kappa T_0}}{\kappa} = B. \quad (18)$$

²We emphasize that Proposition 1 is valid for any drift and diffusion functions. However, for complicated models, we might not be able to compute $\mu(B)$ and $\sigma(B)$ in closed form. In these cases, $\mu(B)$ and $\sigma(B)$ can still be computed numerically from their integral representations, as shown in Appendix of M. Li (2013).

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 both positive, and the numerators of both H_0 and H_1 can be written as a function of the two variables z_0 and $\kappa B/\theta$.

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

See Ahn and Gao (1999) for some univariate analysis on the 3/2 model. In this model, T_0 , H_0 and H_1 are given by

$$T_0 = \frac{1}{\kappa\theta} \log \left(\frac{V_0 + \theta(e^{\kappa B} - 1)}{V_0} \right), \quad (19)$$

$$H_0 = \frac{4V_0 [1 + (\log R - 1)R] + \theta [-3 + (4 - 4 \log R)R + (2 \log R - 1)R^2]}{4\kappa^2 [V_0 + \theta(R - 1)]^2}, \quad (20)$$

$$H_1 = \frac{4R - [3 - 2 \log R] R^2 - 1}{4\kappa^3 [V_0 + \theta(R - 1)]^2}, \quad (21)$$

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

We make an important remark on the value of η in the 3/2-model here. Because the function $b(V) = V^{3/2}$ is typically very small since the instantaneous variance V is usually much smaller than 1, η will usually be large when compared with 1. However, the effect of η can still be small. That is, a Heston model and a 3/2-model with very different values of η can have similar level of variability in the instantaneous variance process. Take the following Heston model parameters for example: $\kappa_H = 2$, $\theta_H = 0.09$, and $\eta_H = 0.375$ (we have added a subscript H to indicate that these are Heston parameters, similarly for the 3/2 parameters). When calibrated to the same set of market data, the parameters of Heston model and 3/2 model should have roughly the same mean-reverting strength near the mode of the invariant distribution, roughly the same long-run mean level, and roughly the same long-run variance. By the results in Ahn and Gao (1999) and Cox, Ingersoll and Ross (1985), these three heuristic criteria give three equations:

$$\kappa_{\frac{3}{2}} \theta_{\frac{3}{2}} = \kappa_H, \quad \frac{2\kappa_{\frac{3}{2}} \theta_{\frac{3}{2}}}{2\kappa_{\frac{3}{2}} + \eta_{\frac{3}{2}}^2} = \theta_H, \quad \frac{2\kappa_{\frac{3}{2}} \theta_{\frac{3}{2}}^2 \eta_{\frac{3}{2}}^2}{(2\kappa_{\frac{3}{2}} + \eta_{\frac{3}{2}}^2)^2} = \frac{\eta_H^2 \theta_H}{2\kappa_H}. \quad (22)$$

Solving these three equations, we get approximately that $\theta_{\frac{3}{2}} = 0.125$, $\kappa_{\frac{3}{2}} = 15.980$, and $\eta_{\frac{3}{2}} = 3.533$. Notice here $\eta_{\frac{3}{2}}$ is much larger than 1, but still its effect is roughly comparable to an η of 0.375 in the Heston model. Intuitively, it is clear that the accuracy of our approximation depends on the level of fluctuations in the instantaneous variance (a good measure is the long-run variance), rather than on the nominal values of η 's in different models.³ It is conceivable

³Li and Mercurio (2013) have studied the accuracy of the approximation in Proposition 1 in the context of

that our method should yield qualitatively similar degrees of accuracy when applied to different instantaneous variance models that behave very similarly on the macro scale, that is, having roughly the same mean-reverting strength on average, roughly the same long-run mean, and roughly the same long-run variance. This should be the case if two stochastic volatility models are calibrated properly using the same option price data from the market.

2.3 Approximating Timer Option Prices for Zero ρ

We consider the simpler problem in this subsection where we assume $\rho = 0$. The starting point of our approach is to write the timer option price as a mixing of Black-Scholes prices. Similar mixing idea was used to price European-style plain-vanilla options in Hull and White (1987), Romano and Touzi (1997), and Lewis (2000), among many others. Lipton (2001) also contains a discussion on this approach.

Let us focus on finite-maturity timer call options first. The pricing of timer puts in our approach requires little additional effort and will be discussed later. Notice that we can always decompose the timer call option price into two components as follows:

$$C_{\text{fin}} = C_{\text{fin}}^B + C_{\text{fin}}^T, \quad (23)$$

with

$$C_{\text{fin}}^B = \mathbb{E}_0 [e^{-r\tau} (S_\tau - K)^+ \cdot 1_{\tau=\tau^B}], \quad (24)$$

$$C_{\text{fin}}^T = \mathbb{E}_0 [e^{-r\tau} (S_\tau - K)^+ \cdot 1_{\tau=T}]. \quad (25)$$

We will use $N(\cdot)$ to denote the cumulative normal distribution function and $n(\cdot)$ to denote the standard normal probability density function. The notation $n(\cdot; m, \Sigma^2)$ will be used to denote the normal probability density function with mean m and variance Σ^2 .

The following proposition shows that we can also express each of the two components of the timer call option price as a mixing of Black-Scholes prices.

Proposition 2. *Under zero correlation, the two components C_{fin}^B and C_{fin}^T are given by a mixing of the Black-Scholes prices. More specifically, we have*

$$C_{\text{fin}}^B = \mathbb{E}_0 [C^{\text{BS}}(S_0, K, r, \delta, \tau^B, B) \cdot 1_{\tau^B < T}], \quad (26)$$

$$C_{\text{fin}}^T = \mathbb{E}_0 [C^{\text{BS}}(S_0, K, r, \delta, T, \xi_T) \cdot 1_{\xi_T < B}], \quad (27)$$

perpetual timer options. For the 3/2 model, a value of 8.56 for η is used. The resulting approximate prices are still found to be very accurate. The percentage errors however are slightly larger than the corresponding ones for the Heston model. This is because with $\eta = 8.56$, the effective long-run variance in the 3/2 model is larger than that in the Heston model. Had we used an η of about 3.533, the accuracy of the two models would be very similar. This is strong indication that the methodology in this paper should work equally well for the 3/2-model.

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

$$C^{\text{BS}}(S_0, K, r, \delta, t, I) = S_0 e^{-\delta t} N(d_1) - K e^{-rt} N(d_2), \quad (28)$$

with

$$d_1 = d_1(t, I) = \frac{\log(S_0 e^{(r-\delta)t}/K)}{\sqrt{I}} + \frac{\sqrt{I}}{2}, \quad (29)$$

$$d_2 = d_2(t, I) = \frac{\log(S_0 e^{(r-\delta)t}/K)}{\sqrt{I}} - \frac{\sqrt{I}}{2}. \quad (30)$$

Proof: By Ito's lemma, we have

$$\log S_t = \log S_0 + (r - \delta)t - \frac{1}{2} \int_0^t V_u du + \int_0^t \sqrt{V_u} dW_u^S. \quad (31)$$

Let \mathcal{F}_t^V be the filtration generated by the process V_u from time 0 to time t . Conditioning on \mathcal{F}_t^V , $\log S_t$ is normally distributed with mean $\log S_0 + (r - \delta)t - \frac{1}{2}\xi_t$ and variance ξ_t .

Consider the C_{fin}^T term first. By law of iterated expectation, we have

$$C_{\text{fin}}^T = \mathbb{E}_0 [\mathbb{E}_0 [e^{-rT} (S_T - K)^+ | \mathcal{F}_T^V] \cdot 1_{\tau^B > T}] \quad (32)$$

$$= \mathbb{E}_0 [C^{\text{BS}}(S_0, K, r, \delta, T, \xi_T) \cdot 1_{\xi_T < B}]. \quad (33)$$

The proof of the C_{fin}^B term is similar by noticing that conditioning on $\mathcal{F}_{\tau^B}^V$, $\log S_{\tau^B}$ is normally distributed with mean $\log S_0 + (r - \delta)\tau^B - B/2$ and variance B . \square

By integration by parts, we can also express C_{fin}^B and C_{fin}^T as follows:

$$C_{\text{fin}}^B = C^{\text{BS}}(S_0, K, r, \delta, T, B) F_{\tau^B}(T) - \Delta C_{\text{fin}}^B, \quad (34)$$

$$C_{\text{fin}}^T = C^{\text{BS}}(S_0, K, r, \delta, T, B) F_{\xi_T}(B) - \Delta C_{\text{fin}}^T, \quad (35)$$

with

$$\Delta C_{\text{fin}}^B = \int_0^T F_{\tau^B}(x) C_x^{\text{BS}}(S_0, K, r, \delta, x, B) dx, \quad (36)$$

$$\Delta C_{\text{fin}}^T = \int_0^B F_{\xi_T}(y) C_y^{\text{BS}}(S_0, K, r, \delta, T, y) dy. \quad (37)$$

Here $F_{\tau^B}(x)$ is the cumulative distribution function of τ^B , and $C_x^{\text{BS}}(S_0, K, r, \delta, x, B)$ and $C_y^{\text{BS}}(S_0, K, r, \delta, T, y)$ are partial derivatives given by

$$C_x^{\text{BS}}(S_0, K, r, \delta, x, B) = -\delta S_0 e^{-\delta x} N(d_1(x, B)) + r K e^{-rx} N(d_2(x, B)), \quad (38)$$

$$C_y^{\text{BS}}(S_0, K, r, \delta, T, y) = S_0 e^{-\delta T} \frac{n(d_1(T, y))}{2\sqrt{y}}. \quad (39)$$

Since by definition $F_{\tau^B}(T) + F_{\xi_T}(B) = 1$, by summing up Equation (34) and (35), an equivalent way of writing the finite-maturity timer option price is

$$C_{\text{fin}} = C^{\text{BS}}(S_0, K, r, \delta, T, B) - \Delta C_{\text{fin}}^B - \Delta C_{\text{fin}}^T. \quad (40)$$

This equation is exact with no approximation used. It is easier to understand this equation assuming $\delta = 0$. In this case, $C^{\text{BS}}(S_0, K, r, \delta, T, B)$ is an upper bound for C_{fin} , and the last two terms correct for the fact that the realized variance of a timer option is not necessarily the upper bound B (the vega correction), and the exercise time is not necessarily the upper bound T (the theta correction).

Although our actual implementation uses $C_{\text{fin}} = C_{\text{fin}}^B + C_{\text{fin}}^T$, the above equation is very important in understanding the performance of the final approximation. The distribution approximation in the last subsection was based on an asymptotic expansion of the moment generating function of τ^B . Because it is not a direct approximation of the distribution function itself, the quantiles $F_{\tau^B}(T)$ and $F_{\xi_T}(B)$ are not necessarily approximated very accurately. However, the above representation shows that C_{fin} is not directly a function of those quantiles. Because ΔC_{fin}^B and ΔC_{fin}^T are usually much smaller than $C^{\text{BS}}(S_0, K, r, \delta, T, B)$ and integrals of the quantiles are involved in the two correction terms rather than particular quantiles, a mildly accurate approximation in the quantiles might give fairly accurate approximation for C_{fin} . As we will see in the numerical analysis section, this is indeed the case.

Proposition 2 is exact with no approximation made. Notice that C_{fin}^B is given by a univariate expectation over τ^B , while C_{fin}^T is given by a univariate expectation over ξ_T . Therefore, if we want to compute the timer option price through Monte Carlo simulation, Proposition 2 allows us to just simulate the one-dimensional process V_u instead of the joint process (S_u, V_u) . Later on, we will compute the benchmark timer option prices in our numerical analysis using this approach. We remark that although Proposition 2 reduces the computational time of Monte Carlo simulation significantly, it can still be computationally very costly.

Proposition 2 is very useful when the distributions of τ^B or ξ_T are known in closed form or can be computed relatively easily. Unfortunately, the distributions of τ^B or ξ_T are usually not easy to obtain or computationally costly. Therefore, it is useful to develop a fast and accurate approximation. Proposition 1 now becomes handy since it gives us the explicit approximate densities need in Proposition 2. The result is given in the next proposition. The C_{fin}^B term is given explicitly in terms of bivariate normal distribution functions while the C_{fin}^T term is given as a one-dimensional integration involving the cumulative distribution function $F_{\xi_T}(x)$.

Proposition 3. *Assume that $\rho = 0$. Then C_{fin}^B can be approximated as follows:*

$$C_{\text{fin}}^B \approx S_0 \cdot \mathbf{I}(a_+, b, -\delta, \mu(B), \sigma(B), T) - K \cdot \mathbf{I}(a_-, b, -r, \mu(B), \sigma(B), T), \quad (41)$$

where $\mu(B)$ and $\sigma^2(B)$ are given in Proposition 1, and

$$a_{\pm} = \frac{\log(S_0/K)}{\sqrt{B}} \pm \frac{\sqrt{B}}{2}, \quad b = \frac{r - \delta}{\sqrt{B}}. \quad (42)$$

The function \mathbf{I} is given by

$$\mathbf{I}(a, b, s, m, \Sigma, \beta) = e^{ms + \Sigma^2 s^2 / 2} \cdot N_2 \left(\frac{\beta - (m + s\Sigma^2)}{\Sigma}, \frac{a + b(m + s\Sigma^2)}{\sqrt{1 + b^2\Sigma^2}}; -\frac{b\Sigma}{\sqrt{1 + b^2\Sigma^2}} \right), \quad (43)$$

where $N_2(z_1, z_2; \rho)$ is the cumulative bivariate normal distribution function with correlation coefficient ρ . The C_{fin}^T term can be approximated as

$$C_{\text{fin}}^T \approx C^{\text{BS}}(S_0, K, r, \delta, T, B) \cdot \tilde{F}_{\xi_T}(B) - S_0 e^{-\delta T} \cdot \int_0^{\sqrt{B}} n(d_1(T, y^2)) \tilde{F}_{\xi_T}(y^2) dy, \quad (44)$$

where $\tilde{F}_{\xi_T}(\cdot)$ is the approximate cumulative distribution function of the accumulated variance ξ_T given in equation (13).

Proof: Consider the C_{fin}^B term first. From Proposition 1, τ^B is approximately normally distributed with mean $\mu(B)$ and variance $\sigma^2(B)$. Therefore, by Proposition 2,

$$\begin{aligned} C_{\text{fin}}^B &\approx \int_{-\infty}^T S_0 e^{-\delta x} N \left(\frac{\log(S_0/K) + (r - \delta)x}{\sqrt{B}} + \frac{1}{2}\sqrt{B} \right) n(x; \mu(B), \sigma^2(B)) dx \\ &\quad - \int_{-\infty}^T K e^{-rx} N \left(\frac{\log(S_0/K) + (r - \delta)x}{\sqrt{B}} - \frac{1}{2}\sqrt{B} \right) n(x; \mu(B), \sigma^2(B)) dx. \end{aligned} \quad (45)$$

The above integrals can be performed explicitly by using the following identity (see, for example, Appendix B of Toft (1996) for a proof)

$$\int_{-\infty}^{\beta} N(a + bx) e^{sx} n(x; m, \Sigma^2) dx = \mathbf{I}(a, b, s, m, \Sigma, \beta), \quad (46)$$

where the expression of \mathbf{I} is given in Proposition 3.

For the C_{fin}^T term, let $f_{\xi_T}(x)$ denote the true density of ξ_T . Then, by Proposition 2,

$$C_{\text{fin}}^T = \int_0^B C^{\text{BS}}(S_0, K, r, \delta, T, x) f_{\xi_T}(x) dx. \quad (47)$$

Integration by parts, we get

$$C_{\text{fin}}^T = C^{\text{BS}}(S_0, K, r, \delta, T, B) \cdot F_{\xi_T}(B) - \int_0^B S_0 e^{-\delta T} n(d_1(T, x)) \frac{1}{2\sqrt{x}} F_{\xi_T}(x) dx. \quad (48)$$

A change of variable $y = \sqrt{x}$ in addition to substituting $F_{\xi_T}(x)$ with $\tilde{F}_{\xi_T}(x)$ gives the expression in the Proposition. This change of variable is helpful in performing the integral numerically as it avoids the singularity at $x = 0$. \square

So far we have focused on pricing timer call options. Our approximation method can also be applied to price timer put options. Notice that Proposition 2 requires little modification. We give the results below. Proof is very similar to that of timer calls and omitted.

Proposition 4. *Assume that $\rho = 0$. The price of a timer put can be decomposed into two components as follows:*

$$P_{\text{fin}} = P_{\text{fin}}^B + P_{\text{fin}}^T, \quad (49)$$

where

$$P_{\text{fin}}^B = \mathbb{E}_0 [e^{-r\tau} (K - S_\tau)^+ \cdot 1_{\tau=\tau^B}] = \mathbb{E}_0 [P^{\text{BS}}(S_0, K, r, \delta, \tau^B, B) \cdot 1_{\tau^B < T}], \quad (50)$$

$$P_{\text{fin}}^T = \mathbb{E}_0 [e^{-r\tau} (K - S_\tau)^+ \cdot 1_{\tau=T}] = \mathbb{E}_0 [P^{\text{BS}}(S_0, K, r, \delta, T, \xi_T) \cdot 1_{\xi_T < B}]. \quad (51)$$

Here P^{BS} is the usual Black-Scholes formula for plain-vanilla put options.

The P_{fin}^B component can be approximated as

$$P_{\text{fin}}^B \approx K \cdot \mathbf{I}(-a_-, -b, -r, \mu(B), \sigma(B), T) - S_0 \cdot \mathbf{I}(-a_+, -b, -\delta, \mu(B), \sigma(B), T). \quad (52)$$

The P_{fin}^T component can be approximated as

$$P_{\text{fin}}^T \approx P^{\text{BS}}(S_0, K, r, \delta, T, B) \cdot \tilde{F}_{\xi_T}(B) - S_0 e^{-\delta T} \cdot \int_0^{\sqrt{B}} n(d_1(T, y^2)) \tilde{F}_{\xi_T}(y^2) dy. \quad (53)$$

Notice that the second and computationally more expensive term in P_{fin}^T is exactly the same as the second term in C_{fin}^T . Therefore, if we need to simultaneously compute the timer call and timer put prices, it only needs to be computed once. Also, due to the symmetry in the Black-Scholes formula, an equivalent alternative expression is

$$P_{\text{fin}}^T \approx P^{\text{BS}}(S_0, K, r, \delta, T, B) \cdot \tilde{F}_{\xi_T}(B) - K e^{-rT} \cdot \int_0^{\sqrt{B}} n(d_2(T, y^2)) \tilde{F}_{\xi_T}(y^2) dy. \quad (54)$$

The proposed formulas for timer calls and puts have the attractive property that they reduce to the correct limits under various special parameter values. First, when $K = 0$ and $\delta = 0$, we have $C_{\text{fin}} = S_0$ and $P_{\text{fin}} = 0$ as they should be. To see this is the correct limit, notice that $e^{-rt} S_t$ is a martingale in this case assuming suitable integration conditions are satisfied. Similarly, when $S_0 = 0$ and $r = 0$, we have the correct limit $C_{\text{fin}} = 0$ and $P_{\text{fin}} = K$.

When B goes to 0, the finite-maturity timer call price in Proposition 3 reduces to the correct final payoff $(S_0 - K)^+$, and similarly for the put in Proposition 4. The math is straightforward but very tedious, and we omit the details here.

When B goes to $+\infty$, the finite-maturity timer options become plain-vanilla European-style options. It is easy to check that C_{fin}^B and P_{fin}^B go to 0 in this case, so that Proposition 2 for timer options becomes a statement for plain-vanilla European-style options:

$$C_{\text{vanilla}} = \lim_{B \rightarrow +\infty} C_{\text{fin}} = \int_0^\infty C^{\text{BS}}(S_0, K, r, \delta, T, x) f_{\xi_T}(x) dx \quad (55)$$

$$= \left(S_0 e^{-\delta T} - K e^{-rT} \right)^+ + S_0 e^{-\delta T} \int_0^\infty (1 - F_{\xi_T}(y^2)) n(d_1(T, y^2)) dy. \quad (56)$$

In the second equation above we have performed an integration by parts. The put option price is approximated as

$$P_{\text{vanilla}} = \lim_{B \rightarrow +\infty} P_{\text{fin}} = \int_0^\infty P^{\text{BS}}(S_0, K, r, \delta, T, x) f_{\xi_T}(x) dx \quad (57)$$

$$= \left(K e^{-rT} - S_0 e^{-\delta T} \right)^+ + S_0 e^{-\delta T} \int_0^\infty (1 - F_{\xi_T}(y^2)) n(d_1(T, y^2)) dy. \quad (58)$$

The above formulas for plain-vanilla option prices are exact with no approximations made. One can also check that they satisfy the usual put-call parity for plain-vanilla options. With the approximate expression for F_{ξ_T} in equation (13), the above equations offer an alternative to Lewis's (2000) small volatility of variance expansion for the Heston model with zero correlation. In unreported numerical study, we see that the above equations with the approximated \tilde{F}_{ξ_T} are extremely accurate when η is relatively small.

Another remark is that when $\eta = 0$, our approximation reduces to the exact formula

$$C_{\text{fin}} = C^{\text{BS}}(S_0, K, r, \delta, \tau, \xi_\tau). \quad (59)$$

Similarly for the timer put. Here the exercise time $\tau \equiv T \wedge T_0$ is deterministic and is either the variance budget exceeding time $\tau^B = T_0$ or the maximum maturity T , depending on which one is smaller. The realized variance budget ξ_τ at the exercise time τ in the formula above for $\eta = 0$ is deterministic and easy to compute. To see that it reduces to this limit, it is easy to check that either C_{fin}^B or C_{fin}^T is zero, depending on whether $\tau = T$ or $\tau = T_0$. For the C_{fin}^B term, notice that when $\eta = 0$, we have $\sigma(B) = 0$, and $\mathbf{I}(a, b, s, m, 0^+, \beta) = e^{ms} N(a + bm) \cdot 1_{m < \beta}$. Therefore,

$$C_{\text{fin}}^B = C^{\text{BS}}(S_0, K, r, \delta, T_0, \xi_{T_0}) \cdot 1_{T_0 < T}. \quad (60)$$

For the C_{fin}^T term, notice that $f_{\xi_T}(x) = \delta(\xi_T - x)$ when $\eta = 0$ in equation (47) so that

$$C_{\text{fin}}^T = C^{\text{BS}}(S_0, K, r, \delta, T, \xi_T) \cdot 1_{T < T_0}. \quad (61)$$

Although we focus on finite-maturity timer options in this paper, under zero correlation, we can approximate the perpetual timer option prices by taking the limit of $T \rightarrow +\infty$ in the

above propositions for finite-maturity timer options. The limit is very easy to take and the result is given below

$$C_{\text{perp}} \approx S_0 \cdot \mathbf{J}(a_+, b, -\delta, \mu(B), \sigma(B)) - K \cdot \mathbf{J}(a_-, b, -r, \mu(B), \sigma(B)), \quad (62)$$

$$P_{\text{perp}} \approx K \cdot \mathbf{J}(-a_-, -b, -r, \mu(B), \sigma(B)) - S_0 \cdot \mathbf{J}(-a_+, -b, -\delta, \mu(B), \sigma(B)), \quad (63)$$

where a_{\pm} and b are given in Proposition 3, and the function \mathbf{J} is given by

$$\mathbf{J}(a, b, s, m, \Sigma) = e^{ms + \Sigma^2 s^2 / 2} \cdot N\left(\frac{a + b(m + s\Sigma^2)}{\sqrt{1 + b^2 \Sigma^2}}\right). \quad (64)$$

The perpetual timer option prices above are given in closed form which resembles the Black-Scholes formula, and is therefore extremely fast to compute. It is also very similar to the perpetual timer option approximation in Li and Mercurio (2013). Indeed, setting $\rho = 0$ in the formula in Li and Mercurio (2013) and performing a Taylor series expansion on η shows that the approximation above agrees with the formula in Li and Mercurio (2013) to first-order in η . It is also easy to check that the above formula has the attractive property that it reduces to the correct limits under various special parameter values. We omit the details here. For actual implementation of perpetual timer option approximation, we recommend Li and Mercurio (2013) because the formula there is applicable for any value of ρ .

One final remark is that the current method for finite-maturity options can be roughly thought of as a sophisticated moment matching method, where the first two moments of τ^B are matched very accurately. Because in the Heston model, the first two moments of ξ_T is known explicitly, an alternative method is to use these moments to approximate a normal density, and then use this normal density to approximate the timer option price. In unreported analysis, we see that this approximation is not as accurate as the current approximation, especially when T is large. The reason is that the current approximation is more than just simple moment matching. As T becomes large, the current approximation has an additional property that it satisfies the pricing PDE to first order in η . As η approaches 0, not only are the moments matched more and more accurately, but also the whole distribution as well.

2.4 Approximating Timer Option Prices for Nonzero ρ

Li and Mercurio (2013) provides an accurate and extremely fast approximation for pricing perpetual timer options under general stochastic volatility models. In particular, the method works for any correlation coefficient ρ . However, the method cannot be easily extended to finite-maturity timer options.

In developing the approximation for finite-maturity timer options in the last subsection, a crucial step is Proposition 2, where we reduce the pricing problem to two univariate expecta-

tions. This is feasible because $\rho = 0$. For nonzero ρ , by Ito's lemma, we have

$$S_t = \tilde{S}_t e^{(r-\delta)t} e^{-\frac{1}{2}(1-\rho^2) \int_0^t V_u du + \sqrt{1-\rho^2} \int_0^t \sqrt{V_u} dW_u^\perp}, \quad (65)$$

where W_u^\perp is a standard Brownian motion perpendicular to W_u^V , and

$$\tilde{S}_t \equiv S_0 e^{-\frac{1}{2}\rho^2 \int_0^t V_u du + \rho \int_0^t \sqrt{V_u} dW_u^V}. \quad (66)$$

Proposition 2 can be generalized to give the following

$$C_{\text{fin}}^B(\rho) = \mathbb{E}_0 \left[C^{\text{BS}}(\tilde{S}_{\tau^B}, K, r, \delta, \tau^B, (1-\rho^2)B) \cdot 1_{\tau^B < T} \right] \quad (67)$$

$$= \mathbb{E}_0 \left[C^{\text{BS}}(S_0 e^{-\frac{1}{2}\rho^2 B} e^{\rho \int_0^{\tau^B} \sqrt{V_u} dW_u^V}, K, r, \delta, \tau^B, (1-\rho^2)B) \cdot 1_{\tau^B < T} \right], \quad (68)$$

$$C_{\text{fin}}^T(\rho) = \mathbb{E}_0 \left[C^{\text{BS}}(\tilde{S}_T, K, r, \delta, T, (1-\rho^2)\xi_T) \cdot 1_{\xi_T < B} \right] \quad (69)$$

$$= \mathbb{E}_0 \left[C^{\text{BS}}(S_0 e^{-\frac{1}{2}\rho^2 \xi_T} e^{\rho \int_0^T \sqrt{V_u} dW_u^V}, K, r, \delta, T, (1-\rho^2)\xi_T) \cdot 1_{\xi_T < B} \right]. \quad (70)$$

To emphasize the ρ dependence, we have explicitly written these prices as functions of ρ . The proof of the above equations is very similar to that of Proposition 2. Conditioning on the filtration generated by V_u and integrating out the randomness of W_u^\perp gives the pricing formulas above.

By taking the limit $T \rightarrow +\infty$, the finite-maturity timer option becomes a perpetual timer option, and we get from Equations (67) and (69) that

$$C_{\text{perp}}^B(\rho) = \lim_{T \rightarrow +\infty} C_{\text{fin}}^B(\rho) + \lim_{T \rightarrow +\infty} C_{\text{fin}}^T(\rho) = \mathbb{E}_0 \left[C^{\text{BS}}(\tilde{S}_{\tau^B}, K, r, \delta, \tau^B, (1-\rho^2)B) \right], \quad (71)$$

where $C_{\text{perp}}^B(\rho)$ is the price of a perpetual timer option with variance budget B when the correlation is ρ . This result was used in Bernard and Cui (2011) to perform Monte Carlo simulations for perpetual timer options. Similarly, we can take the limit $B \rightarrow +\infty$ in Equations (67) and (69). The finite-maturity timer option becomes a plain-vanilla European call option in this case and we get

$$C_{\text{vanilla}}^T(\rho) = \mathbb{E}_0 \left[C^{\text{BS}}(\tilde{S}_T, K, r, \delta, T, (1-\rho^2)\xi_T) \right], \quad (72)$$

where $C_{\text{vanilla}}^T(\rho)$ denotes the plain-vanilla European-style call option price with correlation ρ and maturity T . This is the mixing formula in Hull and White (1987) when $\rho \neq 0$. The above two equations will be used later in developing an approximation for finite-maturity timer option prices.

The above representation is useful in Monte Carlo simulation as it reduces the simulation to a univariate diffusion process. Instead of separating into two components, it is easier for Monte Carlo purpose to use the following

$$C_{\text{fin}}(\rho) = \mathbb{E}_0 \left[C^{\text{BS}}(\tilde{S}_\tau, K, r, \delta, \tau, (1-\rho^2)\xi_\tau) \right]. \quad (73)$$

This equation is easily seen by summing up $C_{\text{fin}}^B(\rho)$ in Equation (68) and $C_{\text{fin}}^T(\rho)$ in Equation (70). In the numerical analysis, we use the above representation to simulate the timer option prices as benchmarks for the approximation we develop in this paper. Notice that all quantities needed above can be computed by just simulating the instantaneous variance process V_u . For each simulated sample path, we just need to record τ , ξ_τ , as well as $\int_0^\tau \sqrt{V_u} dW_u^V$.

Equations (68) and (70) can be used to simulate the two components $C_{\text{fin}}^B(\rho)$ and $C_{\text{fin}}^T(\rho)$ of the finite-maturity timer option price. This turns out to be useful when we examine the accuracy of our finite-maturity timer option price approximation. Similarly, Equations (71) and (72) can be used to simulate the perpetual timer option price $C_{\text{perp}}^B(\rho)$ and the European-style plain-vanilla call price $C_{\text{vanilla}}^T(\rho)$. For all four quantities, only the instantaneous variance process needs to be simulated since the randomness from W_u^\perp has been integrated out.

For the purpose of developing an approximation, however, the existence of the term $\int_0^\tau \sqrt{V_u} dW_u^V$ in \tilde{S}_τ in Equations (68) and (70) prevents us from reducing the two expectations to univariate ones. Therefore, we take a different approach by building upon Li and Mercurio (2013) to price finite-maturity timer options under nonzero ρ .

The approach we take is very much in the same spirit of the method of matched asymptotic expansions in solving ordinary differential equations involving two time scales. The starting point of a matched asymptotic expansions is to find an independent variable in the problem. The idea is to find valid solutions for the problem when the independent variable is very large or very small. One then somehow pastes these two solutions up in a smooth way that preserves the asymptotic behavior. For an introduction to method of matched asymptotic expansions, see Chapter 4 of Nayfeh (2000).

Such an independent variable is clearly present in the problem of pricing finite-maturity timer options, which can be taken to be either the maximum maturity T or the variance budget B . For definiteness, we take the independent variable to be T . Notice that a finite-maturity timer option behaves much like a plain-vanilla option when the target exercise time implied by the variance budget B is much larger than T . On the other hand, it behaves like a perpetual timer option when T is much larger than the target exercise time implied by B . Therefore, a simple approach is to paste these two asymptotic behavior together to approximate the ρ dependence of finite-maturity timer option prices. The separation of timer option prices into two components due to hitting B and hitting T allows us to do this pasting easily.

Notice that there exists a measure $\hat{\mathbb{Q}}^\rho$ such that we can separate the product inside the expectation in $C_{\text{fin}}^B(\rho)$ in Equation (68):

$$C_{\text{fin}}^B(\rho) = \mathbb{E}_0 \left[C^{\text{BS}}(\tilde{S}_{\tau^B}, K, r, \delta, \tau^B, (1 - \rho^2)B) \cdot 1_{\tau^B < T} \right] \quad (74)$$

$$= \mathbb{E}_0 \left[C^{\text{BS}}(\tilde{S}_{\tau^B}, K, r, \delta, \tau^B, (1 - \rho^2)B) \right] \cdot \hat{\mathbb{E}}_0^\rho [1_{\tau^B < T}], \quad (75)$$

where $\hat{\mathbb{E}}_0^\rho$ is taken under the measure $\hat{\mathbb{Q}}^\rho$, defined by the Radon-Nikodym derivative

$$\frac{d\hat{\mathbb{Q}}^\rho}{d\mathbb{Q}} = \frac{C^{\text{BS}}(\tilde{S}_{\tau^B}, K, r, \delta, \tau^B, (1 - \rho^2)B)}{\mathbb{E}_0 \left[C^{\text{BS}}(\tilde{S}_{\tau^B}, K, r, \delta, \tau^B, (1 - \rho^2)B) \right]}. \quad (76)$$

This Radon-Nikodym derivative is well-defined since it is strictly positive, and integrates to 1 under the risk-neutral measure \mathbb{Q} .⁴ For an introduction on Radon-Nikodym derivative and measure change, see Royden (1988) or Folland (1999). Therefore, by Equation (71), we have

$$C_{\text{fin}}^B(\rho) = C_{\text{perp}}^B(\rho) \cdot \hat{\mathbb{E}}_0^\rho [1_{\tau^B < T}]. \quad (77)$$

Similarly, we can separate the two terms in the expectation for C_{fin}^T by taking the expectation of $1_{\xi_T < B}$ under a different measure $\check{\mathbb{Q}}^\rho$:

$$C_{\text{fin}}^T(\rho) = \mathbb{E}_0 \left[C^{\text{BS}}(\tilde{S}_T, K, r, \delta, T, (1 - \rho^2)\xi_T) \right] \cdot \check{\mathbb{E}}_0^\rho [1_{\xi_T < B}] = C_{\text{vanilla}}^T(\rho) \cdot \check{\mathbb{E}}_0^\rho [1_{\xi_T < B}], \quad (78)$$

where the new measure $\check{\mathbb{Q}}^\rho$ is defined by the following Radon-Nikodym derivative:

$$\frac{d\check{\mathbb{Q}}^\rho}{d\mathbb{Q}} = \frac{C^{\text{BS}}(\tilde{S}_T, K, r, \delta, T, (1 - \rho^2)\xi_T)}{\mathbb{E}_0 \left[C^{\text{BS}}(\tilde{S}_T, K, r, \delta, T, (1 - \rho^2)\xi_T) \right]}. \quad (79)$$

The characterization results in Equations (77) and (78) are theoretically very pleasing because they reduce the pricing of finite-maturity timer options to the evaluation of cumulative distribution functions of τ^B under measure $\hat{\mathbb{Q}}^\rho$ and that of ξ_T under measure $\check{\mathbb{Q}}^\rho$. In practice, however, the Radon-Nikodym derivatives are complicated, and it is very difficult to characterize the two new measures to compute these cumulative distribution functions. Our approach therefore is to approximate these two probabilities. The starting point is the following trivial observation:

$$C_{\text{fin}}(\rho) = C_{\text{perp}}^B(\rho) \frac{C_{\text{fin}}^B(\rho)}{C_{\text{perp}}^B(\rho)} + C_{\text{vanilla}}^T(\rho) \frac{C_{\text{fin}}^T(\rho)}{C_{\text{vanilla}}^T(\rho)}. \quad (80)$$

From Equations (77) and (78), we see that the ratios in Equation (80) are probabilities:

$$\hat{\mathbb{E}}_0^\rho [1_{\tau^B < T}] = \frac{C_{\text{fin}}^B(\rho)}{C_{\text{perp}}^B(\rho)}, \quad \check{\mathbb{E}}_0^\rho [1_{\xi_T < B}] = \frac{C_{\text{fin}}^T(\rho)}{C_{\text{vanilla}}^T(\rho)}. \quad (81)$$

Recall that $\mu(B)$ is the deterministic time to exercise the perpetual timer option when η is 0. When $T \ll \mu(B)$, the finite-maturity timer option becomes similar to a plain-vanilla European-style option, so we have

$$\hat{\mathbb{E}}_0^\rho [1_{\tau^B < T}] \approx \hat{\mathbb{E}}_0^0 [1_{\tau^B < T}] \approx 0, \quad \check{\mathbb{E}}_0^\rho [1_{\xi_T < B}] \approx \check{\mathbb{E}}_0^0 [1_{\xi_T < B}] \approx 1. \quad (82)$$

⁴Notice that here we are dealing with two random variables $C^{\text{BS}}(\tilde{S}_{\tau^B}, K, r, \delta, \tau^B, (1 - \rho^2)B)$ and τ^B , rather than two stochastic processes. The measure \mathbb{Q} in the Radon-Nikodym derivative should be interpreted as the restriction of the risk-neutral measure to the probability space generated by these two random variables. Only real analysis is used here, and no stochastic calculus.

On the other hand, when $T \gg \mu(B)$, the finite-maturity timer option becomes similar to a perpetual timer option, so we have

$$\hat{\mathbb{E}}_0^\rho [1_{\tau^B < T}] \approx \hat{\mathbb{E}}_0^0 [1_{\tau^B < T}] \approx 1, \quad \check{\mathbb{E}}_0^\rho [1_{\xi_T < B}] \approx \check{\mathbb{E}}_0^0 [1_{\xi_T < B}] \approx 0. \quad (83)$$

This suggests that when $T \ll \mu(B)$ or $T \gg \mu(B)$, we can make the following approximations:

$$\hat{\mathbb{E}}_0^\rho [1_{\tau^B < T}] \approx \hat{\mathbb{E}}_0^0 [1_{\tau^B < T}], \quad \check{\mathbb{E}}_0^\rho [1_{\xi_T < B}] \approx \check{\mathbb{E}}_0^0 [1_{\xi_T < B}]. \quad (84)$$

These approximations become exact in the two limits $T \rightarrow 0$ and $T \rightarrow +\infty$. The interesting question is how accurate they are when T is reasonably different from $\mu(B)$, but the difference is not extremely large. To analyze this, we perform a Monte Carlo study on these ratios $C_{\text{fin}}^B(\rho)/C_{\text{perp}}^B(\rho)$ and $C_{\text{fin}}^T(\rho)/C_{\text{vanilla}}^T(\rho)$ as functions of ρ in these two limiting cases. We use the Heston model with $\kappa = 2$, $\theta = 0.09$ and $\eta = 0.250$. We use 5 values of ρ : -0.5 , -0.25 , 0 , 0.25 and 0.5 . We let $V_0 = B = 0.087$, implying a $\mu(B)$ of about 1 year. We use two values of T , 0.5 years and 2 years, corresponding to two cases: T is relatively smaller than $\mu(B)$, and T is relatively larger than $\mu(B)$. All four quantities needed in these two ratios are computed through Monte Carlo simulation. We use 4 million sample paths with a time step of two hours. The results are plotted in Figure 1. The two left subplots are the probabilities $\hat{\mathbb{E}}_0^\rho [1_{\tau^B < T}]$ for $T = 2$ years and $T = 0.5$ years, respectively. The two right subplots are the probabilities $\check{\mathbb{E}}_0^\rho [1_{\xi_T < B}]$ for $T = 2$ years and $T = 0.5$ years, respectively. The three curves in each subplot correspond to three different values of the strike price K : 90 (dotted line), 100 (solid line), and 110 (dashed line). Monte Carlo simulation is performed to compute these probabilities. We see that these probabilities have weak ρ dependence when the difference of T and $\mu(B)$ are relatively large.

The above analysis shows that when the difference of T and $\mu(B)$ is reasonably large, the following is a good approximation

$$C_{\text{fin}}(\rho) \approx C_{\text{perp}}^B(\rho) \frac{C_{\text{fin}}^B(0)}{C_{\text{perp}}^B(0)} + C_{\text{vanilla}}^T(\rho) \frac{C_{\text{fin}}^T(0)}{C_{\text{vanilla}}^T(0)}. \quad (85)$$

By design, in the limits of $T \ll \mu(B)$ or $T \gg \mu(B)$, this approximation approaches the correct corresponding limits and becomes exact.

The approximation in Equation (85) turns out to be accurate even when T and $\mu(B)$ are of comparable magnitude. This might be at first sight very surprising, because the probabilities $\hat{\mathbb{E}}_0^\rho [1_{\tau^B < T}]$ and $\check{\mathbb{E}}_0^\rho [1_{\xi_T < B}]$ have very strong ρ dependence individually in this case, as depicted by Figure 2, which plots the two probabilities $\hat{\mathbb{E}}_0^\rho [1_{\tau^B < T}]$ and $\check{\mathbb{E}}_0^\rho [1_{\xi_T < B}]$ again, but this time with $T = 1$ year. We see that $\hat{\mathbb{E}}_0^\rho [1_{\tau^B < T}]$ changes from 0.27 to 0.70 as we vary ρ from -0.5 to 0.5 . Similarly, $\check{\mathbb{E}}_0^\rho [1_{\xi_T < B}]$ changes from 0.70 to 0.27 in this range of ρ .

The reason that the approximation in Equation (85) is still very accurate when $T \approx \mu(B)$ despite the strong ρ dependence of the probabilities is the following. Notice that when $\eta = 0$,

the perpetual timer option is exercised at a fixed time $\mu(B)$, and is therefore equivalent to a plain-vanilla option. In turn, if η is small, for any ρ and any B , we have

$$C_{\text{perp}}^B(\rho) \approx C_{\text{vanilla}}^{\mu(B)}(\rho). \quad (86)$$

Comparing Equations (85) and (80), we see that in this case, in order for the approximation to work well, we do not need the individual probabilities to have weak dependence on ρ . Rather, when $C_{\text{perp}}^B(\rho) \approx C_{\text{vanilla}}^T(\rho)$ or equivalently $T \approx \mu(B)$, as long as the sum

$$\hat{\mathbb{E}}_0^\rho [1_{\tau^B < T}] + \check{\mathbb{E}}_0^\rho [1_{\xi_T < B}] \quad (87)$$

has weak dependence on ρ , the approximation in Equation (85) would be very accurate because of Equation (86). Essentially, when $T \approx \mu(B)$, the plain-vanilla option price and the perpetual timer option price are very close, and as long as the sum of the two probabilities is relatively constant as a function of ρ , it does not matter what value we assign to each of the two probabilities.

Unfortunately, it is theoretically very challenging to analyze the sum of these two probabilities under general ρ . Notice that the probabilities are taken under two different measures $\hat{\mathbb{Q}}^\rho$ and $\check{\mathbb{Q}}^\rho$. These two measures are related to the risk-neutral measure \mathbb{Q} through the two Radon-Nikodym derivatives defined earlier. We caution that even in the special case $\rho = 0$, the measures $\hat{\mathbb{Q}}^0$, $\check{\mathbb{Q}}^0$ and \mathbb{Q} are still all different. To establish the weak dependence of the sum in Equation (87) on ρ when $T \approx \mu(B)$, we resort to Monte Carlo analysis again. We still use the Heston model with the same parameters, and plot the sum in Equation (87) as a function of ρ for $T = 1$ year. The result in Figure 3 clearly shows the weak dependence of the sum on ρ when $T \approx \mu(B)$.

In view of this analysis, the approximation in Equation (85) should *not* be thought of as simply freezing the individual probabilities $\hat{\mathbb{E}}_0^\rho [1_{\tau^B < T}]$ and $\check{\mathbb{E}}_0^\rho [1_{\xi_T < B}]$ at their corresponding $\rho = 0$ values. Rather, it is crucial that we replace these two probabilities with their $\rho = 0$ counterparts *simultaneously*. When T is comparable to $\mu(B)$, the bias induced by overestimating one of the two probabilities is offset to a great extent by the underestimating of the other probability. In summary, the accuracy of the approximation is safeguarded by the fact that it is exact in the limits $T \ll \mu(B)$ and $T \gg \mu(B)$, and that sum in Equation (87) is stable when $T \approx \mu(B)$. Since both sides of Equation (85) are smooth functions of T , it is understandable that the approximation should be accurate for other values of T , by continuity.

We assume that $C_{\text{vanilla}}^T(\rho)$ can be readily computed in closed form or by methods such as analytical approximation, Fourier inversion, etc. The quantities $C_{\text{fin}}^B(0)$ and $C_{\text{fin}}^T(0)$ can be computed by Proposition 3. The perpetual prices $C_{\text{perp}}^B(\rho)$ and $C_{\text{perp}}^B(0)$ can be computed very reliably using the approximations in Li and Mercurio (2013). We remark also that the weights $C_{\text{fin}}^B(0)/C_{\text{perp}}^B(0)$ and $C_{\text{fin}}^T(0)/C_{\text{vanilla}}^T(0)$ do not in general sum up to 1.

The finite-maturity timer puts under nonzero ρ can be approximated in a similar fashion:

$$P_{\text{fin}}(\rho) \approx P_{\text{perp}}^B(\rho) \frac{P_{\text{fin}}^B(0)}{P_{\text{perp}}^B(0)} + P_{\text{vanilla}}^T(\rho) \frac{P_{\text{fin}}^T(0)}{P_{\text{vanilla}}^T(0)}. \quad (88)$$

Besides being accurate for $T \gg \mu(B)$, $T \ll \mu(B)$, and $T \approx \mu(B)$, the above approximations are exact when $\rho = 0$ by construction. They are also trivially exact when $\eta = 0$ because in this case, ρ is irrelevant for the pricing of both finite-maturity and perpetual timer options, as well as plain-vanilla options. Therefore, by continuity, the approximation should also be reasonably accurate for small η .

In actual implementation, the above matched asymptotics approximation contains two sources of errors. One error is the intrinsic error due to the approximation by matching the ρ behavior asymptotically for $T \ll \mu(B)$ and $T \gg \mu(B)$. The other error comes from approximating the quantities involved in the equations. In the Heston model, since the perpetual price approximation in Li and Mercurio (2013) is very accurate, and the vanilla prices can be computed very accurately using Fourier inversion, the main source of error is the computation of C_{fin}^B and C_{fin}^T , or P_{fin}^B and P_{fin}^T . As we have commented earlier, these quantities are relatively inaccurate compared to other quantities due to the fact that the quantile $F_{\xi_T}(B)$ is not as accurately approximated. If the weights can be computed accurately, the above approximation should be very accurate when the maximum maturity is much larger than the target exercise time implied by the variance budget, or vice versa. The matched asymptotic correlation behavior safeguards the cases with medium maximum maturity values. As we will see in the numerical analysis below that despite of the simplicity, this matched asymptotic approximation is reasonably accurate.

2.5 A Comparison of Different Analytic Methods

Now that we have derived the approximation for the finite-maturity timer option price, we make some comparison with other direct analytic methods. This will allow us to see more clearly the advantages of the proposed method over other methods, as well as its limitations.

At first glance, it seems that the finite-maturity timer option price can be obtained through the knowledge of the joint moment generating function of (V_T, ξ_T) , which is known in some models such as the Heston model. This is unfortunately not the case. While this joint distribution is useful for obtaining European-style option prices because of fixed maturity T , for perpetual timer options we need the joint distribution of τ^B and $\int_0^{\tau^B} \sqrt{V_u} dW_u^V$, where τ^B is the random time for ξ_t to hit B .⁵ For finite-maturity timer options, we need the joint distributions of τ and $\int_0^\tau \sqrt{V_u} dW_u^V$, where $\tau = \min(\tau^B, T)$. This is a completely new problem

⁵One way to see this is through Equation (71). We do not need to consider the random variable S_{τ^B} because we can perform a conditioning on the filtration generated by V_u .

from the joint distribution of (V_T, ξ_T) and is mathematically very challenging, even for the perpetual timer option case.

As far as we know, the only model in which the joint distribution of τ^B and $\int_0^{\tau^B} \sqrt{V_u} dW_u^V$ has been worked out explicitly is the Heston model (see C.X. Li (2013)). The Heston model is very special since in this model $\int_0^{\tau^B} \sqrt{V_u} dW_u^V$ can be expressed as a simple linear function of V_{τ^B} and τ^B . C.X. Li (2013) utilizes a time-change technique and relies crucially on known results for Bessel processes. The joint density of τ^B and V_{τ^B} is given by a complex Laplace inverse transformation in Proposition 3 of C.X. Li (2013), where the integrand involves modified Bessel functions with complex arguments. Special attention needs to be paid to the winding of the Bessel argument. The final pricing formula in Equation (5.3) of their paper is also fairly complicated. We emphasize again that the method only works for the Heston model and perpetual timer options. The author also only considers the zero dividend case. Nonzero dividend makes the analysis more complicated because the combination $S_{\tau^B} e^{-\delta\tau^B}$ appears after we condition on the filtration generated by V_u . This combination involves two correlated random variables, which increases complexity of the analysis. This situation differs from the European-style plain-vanilla options case, where it is easy to take into account a nonzero dividend rate because there $e^{-\delta T}$ is just a deterministic number.

To our knowledge, the only other analytic method that has been proposed for *finite-maturity* timer option prices under nonzero interest rate besides the current paper is Liang et al (2011). This paper utilizes the known results on path integrals of the Morse and Kratzer potentials in quantum field theory to study the Heston and 3/2 models. The method works for both perpetual and finite-maturity timer options. The authors only consider the zero dividend case. The math is extremely difficult. The handling of finite-maturity feature in Liang et al (2011) is highly nontrivial and requires the ability to Fourier inverse transform the marginal propagator of the stock price (Equation (35) in their paper).

In Table 1, we summarize the applicability of the three methods: C.X. Li (2013), Liang et al (2011), and the method in this paper. While C.X. Li (2013) only deals with perpetual timer options and technically is not a competitor of our approximation method, we include it because it is closely related. We notice that our method is the only method that can be applied to nonzero dividend case without any additional effort, and the only method that can be applied to models beyond the Heston model or the 3/2 model.⁶ In this sense, rather than being competitors, our method and the methods in C.X. Li (2013) and Liang et al (2011) are complementary to each other. The limitation of our method is also clear: it is an approximation derived under the assumption of a small eta in the diffusion function. It is an empirical question

⁶While it should be possible to extend C.X. Li (2013) to finite-maturity timer options, or to extend both Liang et al (2011) and C.X. Li (2013) to nonzero dividend case, the amount of math or physics knowledge required and the amount of effort required is nontrivial.

as to how small the parameter η needs to be in individual models, and numerical analysis needs to be carried out for any production level implementation.

In Table 2, we compare the implementation aspects of the three methods, as well as their performance. While it is highly nontrivial to implement both C.X. Li (2013) and Liang et al (2011), it is fairly easy to implement our approximation method. The complications in the former two methods include: high-dimensional numerical integration, integration in complex plane where one needs to pay special attention to branch cuts and poles, oscillatory integrand which can potentially be slowly decaying, and the involvement of special functions such as modified Bessel function, Airy function, gamma function, and confluent hypergeometric function. As for performance, the method in C.X. Li (2013) requires about 60 seconds to price one perpetual timer option with a percentage of about 0.05%. Due to the high dimension of integration involved, to reach higher accuracy would require considerably more time. No information on speed is provided in Liang et al (2011), but it is conceivable that it would take somewhat longer time because of the added complexity to handle the finite-maturity feature. To contrast, the approximation method in Li and Mercurio (2013) takes about 10^{-4} seconds for perpetual timer options and in this paper the method takes about 10^{-2} seconds for a finite-maturity option.⁷ In terms of accuracy, both C.X. Li (2013) and Liang et al (2011) report an average percentage error of about 0.05%. These papers only test a few parameter values and do not systematically study the error as a function of T or η . The errors of our approximation will be reported in more detail in the numerical analysis section. The average percentage error for finite-maturity option prices is about 0.5%. The error becomes considerably smaller when η is small or when T is large compared to $\mu(B)$.

One additional attractive property of our method manifests itself when we need to value a far-way-from-the-money timer option for which the strike K is much smaller or much larger than the current stock price S_0 . It is well-known that the usual Fourier inversion for European-style options becomes inaccurate numerically due to the oscillatory integrand. It is conceivable that similar difficulty will also arise in the timer option case with complex inversion methods. To contrast, when dividend rate is zero, our approximation approaches the correct no-arbitrage limit prices when $K \rightarrow 0$ or $K \rightarrow \infty$, therefore providing robustness for the method.

3 Numerical Analysis

Below we examine the accuracy of our pricing approximation numerically. We do this by first looking at the performance for the zero correlation case, and then for the nonzero correlation

⁷The approximation method has an additional attractive property in that $\mu(B)$ and $\sigma(B)$ do not depend on K , S_0 , or T . Therefore, if one needs to value many finite-maturity timer options with different S_0 , K or T (one possible situation is to compute greeks using finite difference), we just need to value $\mu(B)$ and $\sigma(B)$ once.

case. For both cases, we examine the performance for both timer calls and puts.

Table 3 reports the performance of our approximation for timer call options under the Heston model with zero correlation. Parameters used here are the same ones as in Liang, Lemmens and Tempere (2011): $V_0 = 0.087$, $\kappa = 2$, $\theta = 0.09$, $r = 0.015$, $S_0 = 100$, and $\delta = 0$. The variance budget B equals V_0 , implying a target exercise maturity of about 1 year. We vary the value of the maximum maturity T to be 0.5, 1.0, 1.5, 2.0, and 10 years, and volatility of variance coefficient η to be either 0.125, 0.250, or 0.375, ranging from small to relatively large. Monte Carlo prices are simulated based on Proposition 2 using 4,000,000 sample paths with a time step of every 2 hours and the terminal accumulated variance as the control variate. The unreported standard errors from simulation are all in the order of 10^{-4} or smaller. Each Monte Carlo simulation takes about 20 to 40 minutes for the desired accuracy. The approximated prices are computed using Proposition 3 with the quantities $\mu(B)$ and $\sigma(B)$ given in Example 1. The percentage errors are computed using the Monte Carlo prices as benchmarks.

With the given parameters, when $\eta = 0.125$, the random budget exceeding time τ^B is approximately normally distributed with $\mu \doteq 0.9969$ and $\sigma \doteq 0.1255$. When $\eta = 0.250$, we have $\mu \doteq 1.0445$ and $\sigma \doteq 0.2510$. When $\eta = 0.375$, we have $\mu \doteq 1.1239$ and $\sigma \doteq 0.3765$. The range of maximum maturities we use therefore covers situations where the cutoff is very stringent to where it is very loose. In terms of the timer call price decomposition, the relative magnitude of C_{fin}^B and C_{fin}^T therefore ranges from very small to very large. This wide range of T allows us to look at the accuracy of our approximation more comprehensively.

As we see in Table 3, our approximation is fairly accurate in general. For the relatively small value of $\eta = 0.125$, our approximation is extremely accurate, with percentage errors all in the order of 0.01%. We also see that a maximum maturity of 2 years or more puts effectively no constraint on the timer option. The approximation is still very accurate for the medium value of $\eta = 0.250$. The relatively large percentage errors occur when η is large and the maximum maturity is stringent, that is, when $\eta = 0.375$ and $T < 2.0$. When $\eta = 0.375$ but the maximum maturity is less stringent, we see that the percentage errors get much smaller. We also remark that in practice, when a timer option contract is first initiated, the maximum maturity is usually fairly loose, with typical maximum maturity 3 or 6 times the target exercise maturity. However, a timer option on an existing trade book could have a stringent maximum maturity cutoff as time progresses.

The effect of η on the timer call option prices is complex. If we focus on the timer call options with T less than 2.0 years, we see from Table 3 that the timer option prices with the same K and T are generally decreasing as η increases. This is because a large η shortens the effective exercise time, which usually implies a smaller option price. However, when T equals 2 years so that the maximum maturity cutoff becomes loose, we see the timer call option prices are increasing in η . This phenomenon is very similar to knock-out barrier call options where

a large volatility gives more optionality to exercise the option, but also increases the risk that the option becomes knocked-out.

Table 4 reports the performance of our approximation for timer put options under zero correlation. The parameters used are exactly the same ones as in Table 3. In contrast to Table 3, Table 4 shows that the timer put prices are all decreasing in η regardless of the values of T and K . The general patterns of pricing errors are the same as in Table 3 so we do not discuss the details here. The approximation is again extremely accurate when the maximum maturity cutoff is less stringent, and when η is not large.

One question is whether the accuracy of our approximation is due to the particular parameter vector we have used. In unreported analysis, we have performed extensive sensitivity analysis. The following is a partial set of what we have done.

1. We change the interest rate r to be either 0 or 3%.
2. We change the value of B to be twice of or half of the value we use right now.
3. We change the value of V_0 to be either 25% or 4 times the value we use right now.
4. We examine a much wide range of strikes, from $K = 50$ to $K = 150$.

In virtually all those cases, the accuracy of our approximation in terms of percentage errors remains the same. The only exception is when the option is deeply out-of-the-money, so the percentage error gets large, but the absolute error is still small. That the approximation works for a wide range of strikes makes it more useful in real applications.

Tables 5 and 6 report the performance of our approximation under the Heston model with nonzero correlation for finite-maturity timer call and put options, respectively. The parameters used are still the same ones as in Liang, Lemmens and Tempere (2011). Same as in Liang, Lemmens and Tempere (2011), we consider two values of nonzero ρ , namely -0.5 and 0.5 . We consider three values of T : 0.5, 1.0, and 1.5 years. The maturities T are chosen to cover the three cases that T is smaller than, roughly equal to, and larger than the target exercise time implied by B . Again, three values of η are considered, ranging from small to relatively large. The benchmark prices are obtained from Monte Carlo using Equation (73) with 4,000,000 sample paths and a time step of every 2 hours. Not surprisingly, with nonzero ρ , the accuracy of Monte Carlo decreases. However, the results are still very accurate, with standard errors all in the order of 10^{-3} or smaller.

Before we look at the accuracy of our approximation, it is interesting to point out that the finite-maturity timer option price depends on η in a more complicated way for nonzero correlation than for zero correlation. For example, when $\rho = -0.5$ and $T = 0.5$, the simulated in-the-money ($K = 90$) timer call option price decreases as we increase η , but the out-of-the-money ($K = 110$) timer call price increases. For another example, when $\rho = -0.5$ and $T = 1.5$, the out-of-the-money ($K = 110$) timer call option price first increases, and then decreases. Our

approximation in general follows the Monte Carlo prices closely and also exhibits those complex η dependencies.

In both Tables 5 and 6, we see that our simple approximation is fairly accurate, especially when η is small, or when T is large. We see that for all entries with $\eta = 0.125$ or $T = 1.5$, the percentage errors are all smaller than -0.50% for both calls and puts. For $\eta = 0.250$, we see that the errors are larger for $T = 1.0$, where the constraints due to T and B are of comparable magnitudes. The accuracy decreases significantly for $\eta = 0.375$. One reason is that C_{fin}^B and C_{fin}^T needed in the approximation become less accurate for large η . For example, in Table 5, when $\rho = 0.5$ and $\eta = 0.375$, the percentage errors for the three different strikes are -0.50% , -1.11% , and -2.41% . Had we used the values of C_{fin}^B and C_{fin}^T from Monte Carlo, the three errors would have been -0.01% , -0.06% , and -0.19% . This suggests that if accurate values of C_{fin}^B and C_{fin}^T are available through other means, then the performance of the matched asymptotic approximation can be improved, especially for relatively small or relatively large T .

4 Conclusion

In this paper, we develop an approximation for pricing finite-maturity timer options under general stochastic volatility models. The method works for both timer calls and timer puts. For zero correlation between the Brownian motions, we use the results in M. Li (2013) to approximate the probability densities we need and obtain analytic expressions for the timer option prices. For nonzero correlation, we use a linear combination approximation which matches the asymptotic correlation behavior when the maximum maturity is much larger than the target exercise time implied by the variance budget, or vice versa.

Numerical study using the Heston model shows that the approximate timer option prices under zero and nonzero correlations are fairly accurate. The approximation is more accurate when the volatility of variance is small or when the maximum expiry is not tightly binding.

There are many future directions one can go from this paper. It would be very useful to further increase the accuracy of the current approximation while keeping its analytic tractability, especially for relatively large values of volatility of variance and stringent maximum expiry. Also, our approximation for nonzero correlation leaves room for improvements. It is also useful to explore other applications of the approximation technique we have developed, such as other financial products with timer features. We leave these to future research.

References

- Ahn, D.H., and B. Gao. 1999. A parametric nonlinear model of term structure dynamics. *Review of Financial Studies* **12**, 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 R. Lee. 2010. Hedging variance option on continuous semimartingales. *Finance and Stochastic* 14, 179–207.
- Folland, G.B. 1999. *Real Analysis: Modern Techniques and Their Applications*. Wiley.
- Hawkins, D., and S. Krol. 2008. Product overview: Timer options. Lehman brothers Equity Derivatives note.
- 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.
- Lee, R. 2012. Timer options for risk-controlled variance exposure. Global Derivatives USA 2012 conference presentation.
- Lewis, A.L. 2000. *Option Valuation Under Stochastic Volatility*. Finance Press, Newport Beach, California.
- Li, C.X. 2010. Managing volatility risk. Doctoral Dissertation, Columbia University.
- Li, C.X. 2013. Bessel processes, stochastic volatility, and timer options. *Mathematical Finance*, forthcoming. DOI:10.1111/mafi.12041
- Li, M. 2013. Derivatives pricing on integrated diffusion processes: A general perturbation approach. Working paper. Bloomberg LP.
- Li, M., F. Mercurio. 2013a. Closed-form approximation of timer option prices under general stochastic volatility models. Working paper. Bloomberg LP.
- 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.
- Nayfeh, A.H. 2000. *Perturbation Methods*. Wiley-Interscience.

- Neuberger, A. 1990. Option pricing: A non-stochastic approach. Working paper. London Business School.
- Romano, M., and N. Touzi. 1997. Contingent claims and market completeness in a stochastic volatility model. *Mathematical Finance* 7(4), 399–412.
- Royden, H. 1988. Real Analysis. Pearson.
- Saunders, D. 2010. Pricing timer options under fast mean-reverting stochastic volatility. Working paper.
- 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).
- Toft, K.B. 1996. On the mean-variance tradeoff in option replication with transactions costs. *Journal of Financial and Quantitative Analysis* 31(2), 233–263.

Table 1: **Applicability of Different Analytic Methods**

	C.X. Li (2013)	Liang et al (2011)	This paper
Applicable models	Heston only	Heston and 3/2 only	Any Heston-like models in principle. Requires small η
Option types	Perpetual only	Perpetual and finite-maturity	Perpetual and finite-maturity
Dividend	Only zero dividend case treated. Extending to nonzero dividend nontrivial	Only zero dividend case treated. Extending to nonzero dividend nontrivial	Any dividend rate value allowed
Extensibility	Difficult to extend to other models since relying crucially on known math result of Bessel processes	Difficult to extend to other models since relying crucially on solvability of the Morse and Kratzer potentials	M. Li (2013) extends the method to Stein-Stein type stochastic volatility models

Table 2: **Implementation and Performance of Different Analytic Methods**

	C.X. Li (2013)	Liang et al (2011)	This paper
Complexity of final formula	High	High	Low
Ease of code implementation	Nontrivial	Nontrivial	Fairly easy
Dimension of numerical integration	3 (Method can only handle perpetual)	3 (Finite-maturity involves extra high-dimensional integration)	No integration for perpetual. 1 for finite-maturity
Property of integrand	Complex. Oscillatory.	Complex. Oscillatory.	Real. Non-oscillatory. Bounded
Special functions in integrand	Modified Bessel function. Hyperbolic functions. Airy function (analytic continuation involved)	Modified Bessel function. Hyperbolic functions. Gamma function and confluent hypergeometric function (3/2 only)	Product-log function for Heston. None for 3/2
Other considerations	Abate-Whitt algorithm with empirical damping factor and truncation. Need to properly handle winding numbers. Small B and far-away-from-the-money K pose numerical issues	No implementation details provided. Presumably the same complications as in C.X. Li (2013) also present here. Small η and far-away-from-the-money K pose numerical issues	Method approaches correct limits when option is far-away from money, where inversion methods might have difficulties
Computational time	About 60 seconds per perpetual timer	Information not provided. Presumably similar to C.X. Li (2013) for perpetual timers due to similar degree of complexity and longer for finite-maturity	About 10^{-4} seconds for perpetual (slightly slower than Black-Scholes). About 10^{-2} seconds for finite-maturity
Average percentage error	Around 0.05% (can only handle perpetual). Only tested two η values.	Around 0.05%. Only tested one η and T . Error dependence on η or T not studied	Around 0.05% for perpetual and 0.5% for finite-maturity when $\eta = 0.375$ in Heston. Error depends on η and T

Table 3: Accuracy of Timer Call Price Approximation for Zero Correlation ρ

Parameters used here are the same ones as in Liang, Lemmens and Tempere (2011): $V_0 = 0.087$, $\kappa = 2$, $\theta = 0.09$, $B = 0.087$, $\xi = 0$, $r = 0.015$, $S_0 = 100$, and $\delta = 0$. We vary the value of maximum maturity T and volatility of variance coefficient η . Monte Carlo prices are simulated based on Proposition 2 using 4,000,000 sample paths with a time step of every 2 hours and the terminal accumulated variance as control variate. The unreported standard errors are all in the order of 10^{-4} or smaller. The approximated prices are computed using Propositions 3.

T	K	$\eta = 0.125$			$\eta = 0.250$			$\eta = 0.375$		
		Approx	MC	Error	Approx	MC	Error	Approx	MC	Error
0.5	90	14.3518	14.3500	0.01%	14.3162	14.3236	-0.05%	14.1972	14.2718	-0.52%
	100	8.6915	8.6894	0.02%	8.6356	8.6436	-0.09%	8.4768	8.5579	-0.95%
	110	4.8922	4.8901	0.04%	4.8461	4.8541	-0.16%	4.7033	4.7851	-1.71%
1.0	90	17.3675	17.3649	0.01%	17.0573	17.0596	-0.01%	16.7334	16.7550	-0.13%
	100	12.1068	12.1038	0.03%	11.7531	11.7546	-0.01%	11.3766	11.3978	-0.19%
	110	8.1948	8.1919	0.04%	7.8432	7.8459	-0.03%	7.4690	7.4938	-0.33%
1.5	90	17.6277	17.6275	0.00%	17.6527	17.6299	0.13%	17.6137	17.5384	0.43%
	100	12.3947	12.3945	0.00%	12.4128	12.3882	0.20%	12.3544	12.2728	0.67%
	110	8.4786	8.4784	0.00%	8.4906	8.4667	0.28%	8.4238	8.3452	0.94%
2.0	90	17.6277	17.6277	0.00%	17.6662	17.6635	0.02%	17.7265	17.6995	0.15%
	100	12.3947	12.3947	0.00%	12.4276	12.4252	0.02%	12.4783	12.4512	0.22%
	110	8.4786	8.4786	0.00%	8.5052	8.5032	0.02%	8.5455	8.5204	0.29%
10.0	90	17.6277	17.6278	0.00%	17.6662	17.6637	0.01%	17.7305	17.7193	0.06%
	100	12.3947	12.3948	0.00%	12.4276	12.4255	0.02%	12.4827	12.4732	0.08%
	110	8.4786	8.4786	0.00%	8.5052	8.5035	0.02%	8.5498	8.5422	0.09%

Table 4: Accuracy of Timer Put Price Approximation for Zero Correlation ρ

Parameters used here are the same ones as in Liang, Lemmens and Tempere (2011): $V_0 = 0.087$, $\kappa = 2$, $\theta = 0.09$, $B = 0.087$, $\xi = 0$, $r = 0.015$, $S_0 = 100$, and $\delta = 0$. We vary the value of maximum maturity T and volatility of variance coefficient η . Monte Carlo prices are simulated based on Proposition 2 using 4,000,000 sample paths with a time step of every 2 hours and the terminal accumulated variance as control variate. The unreported standard errors are all in the order of 10^{-4} or smaller. The approximated prices are computed using Propositions 4.

T	K	$\eta = 0.125$			$\eta = 0.250$			$\eta = 0.375$		
		Approx	MC	Error	Approx	MC	Error	Approx	MC	Error
0.5	90	3.6793	3.6775	0.05%	3.6455	3.6511	-0.15%	3.5349	3.6001	-1.81%
	100	7.9443	7.9422	0.03%	7.8904	7.8964	-0.08%	7.7410	7.8115	-0.90%
	110	14.0703	14.0682	0.01%	14.0264	14.0322	-0.04%	13.8939	13.9642	-0.50%
1.0	90	6.0963	6.0943	0.03%	5.8233	5.8289	-0.10%	5.5219	5.5522	-0.55%
	100	10.6944	10.6920	0.02%	10.3819	10.3871	-0.05%	10.0305	10.0614	-0.31%
	110	16.6412	16.6389	0.01%	16.3350	16.3416	-0.04%	15.9884	16.0237	-0.22%
1.5	90	6.2921	6.2920	0.00%	6.2588	6.2461	0.20%	6.1520	6.1147	0.61%
	100	10.9107	10.9106	0.00%	10.8640	10.8505	0.13%	10.7302	10.6909	0.37%
	110	16.8461	16.8461	0.00%	16.7870	16.7751	0.07%	16.6372	16.6051	0.19%
2.0	90	6.2921	6.2921	0.00%	6.2677	6.2690	-0.02%	6.2251	6.2222	0.05%
	100	10.9107	10.9106	0.00%	10.8738	10.8757	-0.02%	10.8100	10.8097	0.00%
	110	16.8461	16.8461	0.00%	16.7960	16.7987	-0.02%	16.7104	16.7148	-0.03%
10.0	90	6.2921	6.2921	0.00%	6.2678	6.2695	-0.03%	6.2274	6.2348	-0.12%
	100	10.9107	10.9107	0.00%	10.8738	10.8764	-0.02%	10.8125	10.8236	-0.10%
	110	16.8461	16.8463	0.00%	16.7960	16.7995	-0.02%	16.7126	16.7275	-0.09%

Table 5: Accuracy of Timer Call Price Approximation for Nonzero ρ

Parameters used here are the same ones as in Liang, Lemmens and Tempere (2011): $V_0 = 0.087$, $\kappa = 2$, $\theta = 0.09$, $B = 0.087$, $\xi = 0$, $r = 0.015$, $S_0 = 100$, and $\delta = 0$. We vary the value of instantaneous correlation ρ and volatility of variance coefficient η . Three values of maximum maturity T are used. Monte Carlo prices are simulated based on equation (73) using 4,000,000 sample paths with a time step of every 2 hours and the terminal accumulated variance as control variate. The unreported standard errors are all in the order of low 10^{-3} or smaller. The approximated prices are computed using the linear combination approximation.

ρ	K	$\eta = 0.125$			$\eta = 0.250$			$\eta = 0.375$		
		Approx	MC	Error	Approx	MC	Error	Approx	MC	Error
$T = 0.5$										
-0.5	90	14.4296	14.4286	0.01%	14.4664	14.4734	-0.05%	14.4072	14.4973	-0.62%
	100	8.6748	8.6739	0.01%	8.6055	8.6101	-0.05%	8.4366	8.5172	-0.95%
	110	4.7757	4.7750	0.01%	4.6195	4.6191	0.01%	4.3889	4.4388	-1.12%
0.5	90	14.2685	14.2693	-0.01%	14.1484	14.1535	-0.04%	13.9478	14.0177	-0.50%
	100	8.7059	8.7067	-0.01%	8.6637	8.6737	-0.12%	8.5100	8.6056	-1.11%
	110	5.0035	5.0041	-0.01%	5.0603	5.0758	-0.31%	4.9898	5.1128	-2.41%
$T = 1.0$										
-0.5	90	17.4176	17.4437	-0.15%	17.1582	17.2215	-0.37%	16.8805	16.9877	-0.63%
	100	12.1074	12.1391	-0.26%	11.7432	11.8188	-0.64%	11.3494	11.4771	-1.11%
	110	8.1446	8.1683	-0.29%	7.7223	7.7779	-0.71%	7.2700	7.3645	-1.28%
0.5	90	17.3141	17.3394	-0.15%	16.9412	16.9940	-0.31%	16.5479	16.6375	-0.54%
	100	12.1047	12.1335	-0.24%	11.7554	11.8151	-0.51%	11.3837	11.4810	-0.85%
	110	8.2429	8.2668	-0.29%	7.9541	8.0100	-0.70%	7.6432	7.7422	-1.28%
$T = 1.5$										
-0.5	90	17.6572	17.6590	0.01%	17.7127	17.7036	0.05%	17.7069	17.6824	0.14%
	100	12.4284	12.4302	-0.01%	12.4748	12.4665	0.07%	12.4175	12.3985	0.15%
	110	8.5127	8.5144	-0.02%	8.5480	8.5367	0.13%	8.4527	8.4199	0.39%
0.5	90	17.5981	17.5938	0.02%	17.5914	17.5712	0.12%	17.5090	17.4679	0.24%
	100	12.3608	12.3574	0.03%	12.3499	12.3335	0.13%	12.2837	12.2467	0.30%
	110	8.4443	8.4419	0.03%	8.4323	8.4191	0.16%	8.3876	8.3551	0.39%

Table 6: Accuracy of Timer Put Price Approximation for Nonzero ρ

Parameters used here are the same ones as in Liang, Lemmens and Tempere (2011): $V_0 = 0.087$, $\kappa = 2$, $\theta = 0.09$, $B = 0.087$, $\xi = 0$, $r = 0.015$, $S_0 = 100$, and $\delta = 0$. We vary the value of instantaneous correlation ρ and volatility of variance coefficient η . Three values of maximum maturity T are used. Monte Carlo prices are simulated based on equation (73) using 4,000,000 sample paths with a time step of every 2 hours and the terminal accumulated variance as control variate. The unreported standard errors are all in the order of low 10^{-3} or smaller. The approximated prices are computed using the linear combination approximation.

ρ	K	$\eta = 0.125$			$\eta = 0.250$			$\eta = 0.375$		
		Approx	MC	Error	Approx	MC	Error	Approx	MC	Error
$T = 0.5$										
-0.5	90	3.7567	3.7565	0.01%	3.7943	3.8047	-0.27%	3.7353	3.8258	-2.37%
	100	7.9275	7.9271	0.01%	7.8604	7.8668	-0.08%	7.7014	7.7711	-0.90%
	110	13.9541	13.9535	0.00%	13.7980	13.8010	-0.02%	13.5676	13.6181	-0.37%
0.5	90	3.5956	3.5949	0.02%	3.4794	3.4817	-0.07%	3.2973	3.3441	-1.40%
	100	7.9587	7.9575	0.01%	7.9184	7.9272	-0.11%	7.7736	7.8573	-1.06%
	110	14.1819	14.1803	0.01%	14.2422	14.2546	-0.09%	14.1913	14.2899	-0.69%
$T = 1.0$										
-0.5	90	6.1447	6.1711	-0.43%	5.9174	5.9879	-1.18%	5.6547	5.7831	-2.22%
	100	10.6954	10.7253	-0.28%	10.3734	10.4484	-0.72%	10.0063	10.1389	-1.31%
	110	16.5886	16.6133	-0.15%	16.2043	16.2708	-0.41%	15.7669	15.8925	-0.79%
0.5	90	6.0449	6.0689	-0.40%	5.7153	5.7638	-0.84%	5.3561	5.4331	-1.42%
	100	10.6920	10.7219	-0.28%	10.3830	10.4482	-0.62%	10.0350	10.1430	-1.06%
	110	16.6916	16.7140	-0.13%	16.4549	16.5063	-0.31%	16.1825	16.2705	-0.54%
$T = 1.5$										
-0.5	90	6.3216	6.3196	0.03%	6.3183	6.3197	-0.02%	6.2417	6.2573	-0.25%
	100	10.9444	10.9423	0.02%	10.9263	10.9289	-0.02%	10.7954	10.8152	-0.18%
	110	16.8803	16.8782	0.01%	16.8433	16.8453	-0.01%	16.6581	16.6784	-0.12%
0.5	90	6.2625	6.2627	0.00%	6.1982	6.1964	0.03%	6.0533	6.0482	0.08%
	100	10.8768	10.8779	-0.01%	10.8008	10.8050	-0.04%	10.6577	10.6688	-0.10%
	110	16.8119	16.8140	-0.01%	16.7296	16.7369	-0.04%	16.6078	16.6189	-0.07%

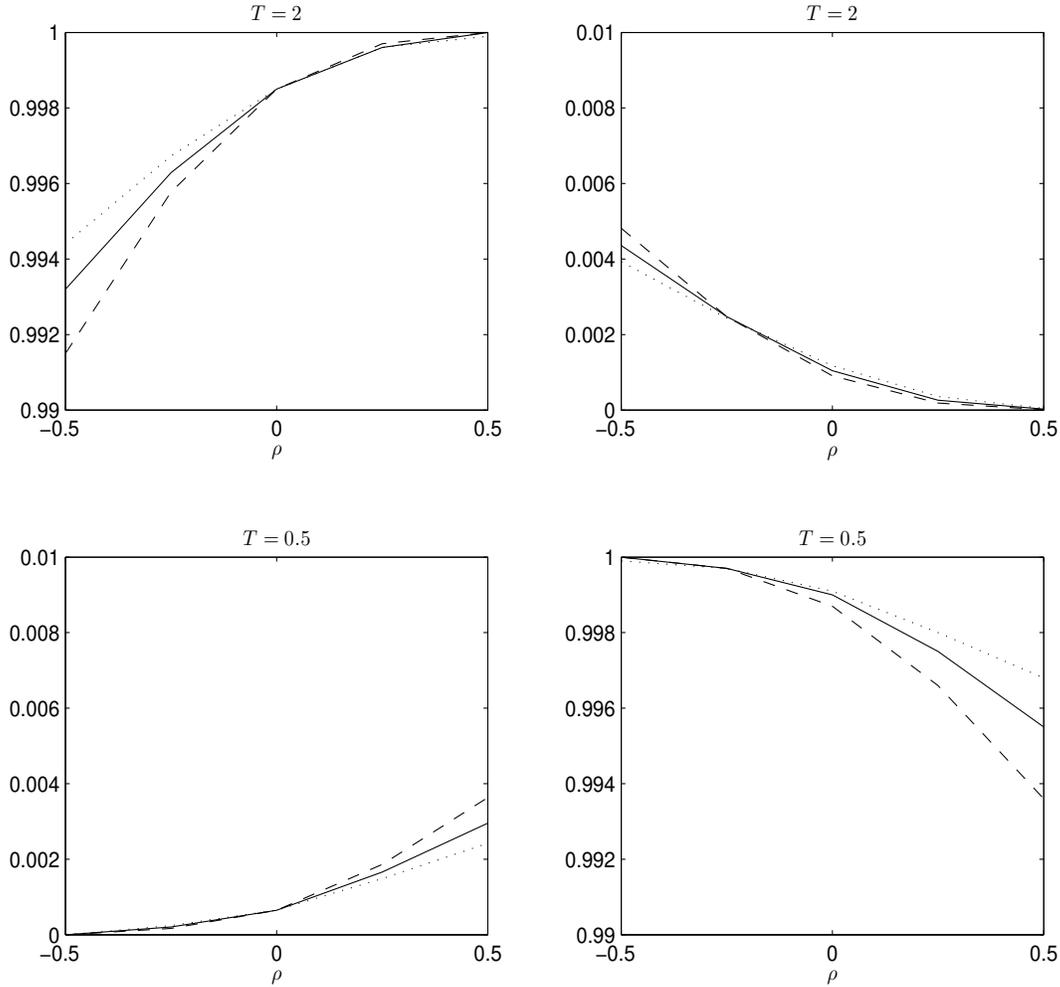


Figure 1: **Probabilities as functions of ρ when difference of T and $\mu(B)$ is relatively large.** The Heston model is used with the following parameters: $S_0 = 100$, $r = 0.015$, $\kappa = 2$, $\theta = 0.09$, $B = V_0 = 0.087$, $\eta = 0.250$, corresponding to $\mu(B)$ of about 1 year. The two left subplots are the probabilities $\hat{\mathbb{E}}_0^\rho [1_{\tau^B < T}]$ for $T = 2$ years and $T = 0.5$ years, respectively. The two right subplots are the probabilities $\mathbb{E}_0^\rho [1_{\xi_T < B}]$ for $T = 2$ years and $T = 0.5$ years, respectively. The three curves in each subplot correspond to three different values of the strike price K : 90 (dotted line), 100 (solid line), and 110 (dashed line). Monte Carlo simulation is performed to compute these probabilities. We see that these probabilities have weak ρ dependence when the difference of T and $\mu(B)$ is relatively large.

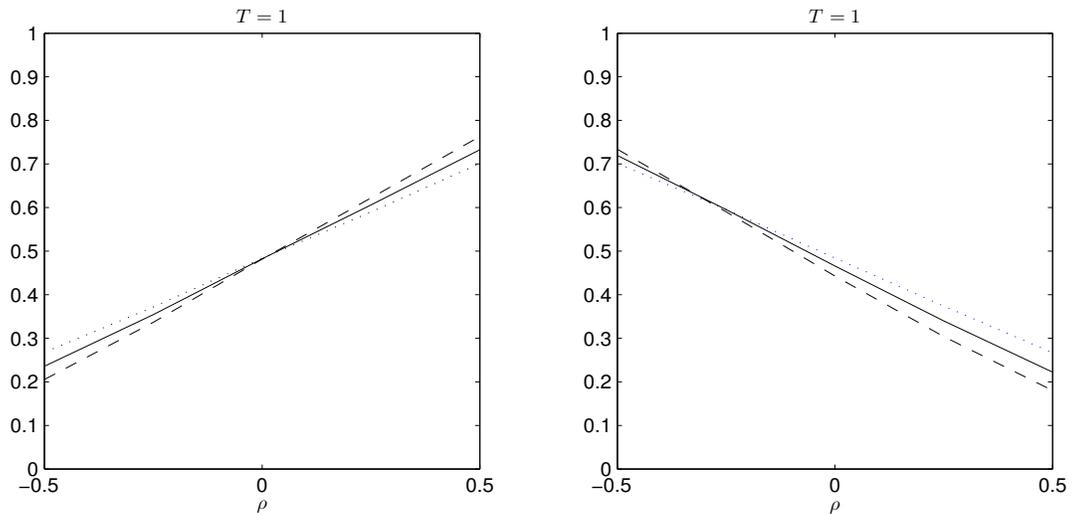


Figure 2: **Probabilities as functions of ρ when $T \approx \mu(B)$.** The Heston model is used with the following parameters: $S_0 = 100$, $r = 0.015$, $\kappa = 2$, $\theta = 0.09$, $B = V_0 = 0.087$, $\eta = 0.250$, corresponding to $\mu(B)$ of about 1 year. We use $T = 1$ year. The left subplot is the probability $\hat{\mathbb{E}}_0^\rho[1_{\tau^B < T}]$ and the right subplot is the probability $\check{\mathbb{E}}_0^\rho[1_{\xi_T < B}]$. The three curves in each subplot correspond to three different values of the strike price K : 90 (dotted line), 100 (solid line), and 110 (dashed line). Monte Carlo simulation is performed to compute these probabilities. We see that these probabilities now have relatively strong ρ dependence when $T \approx \mu(B)$.

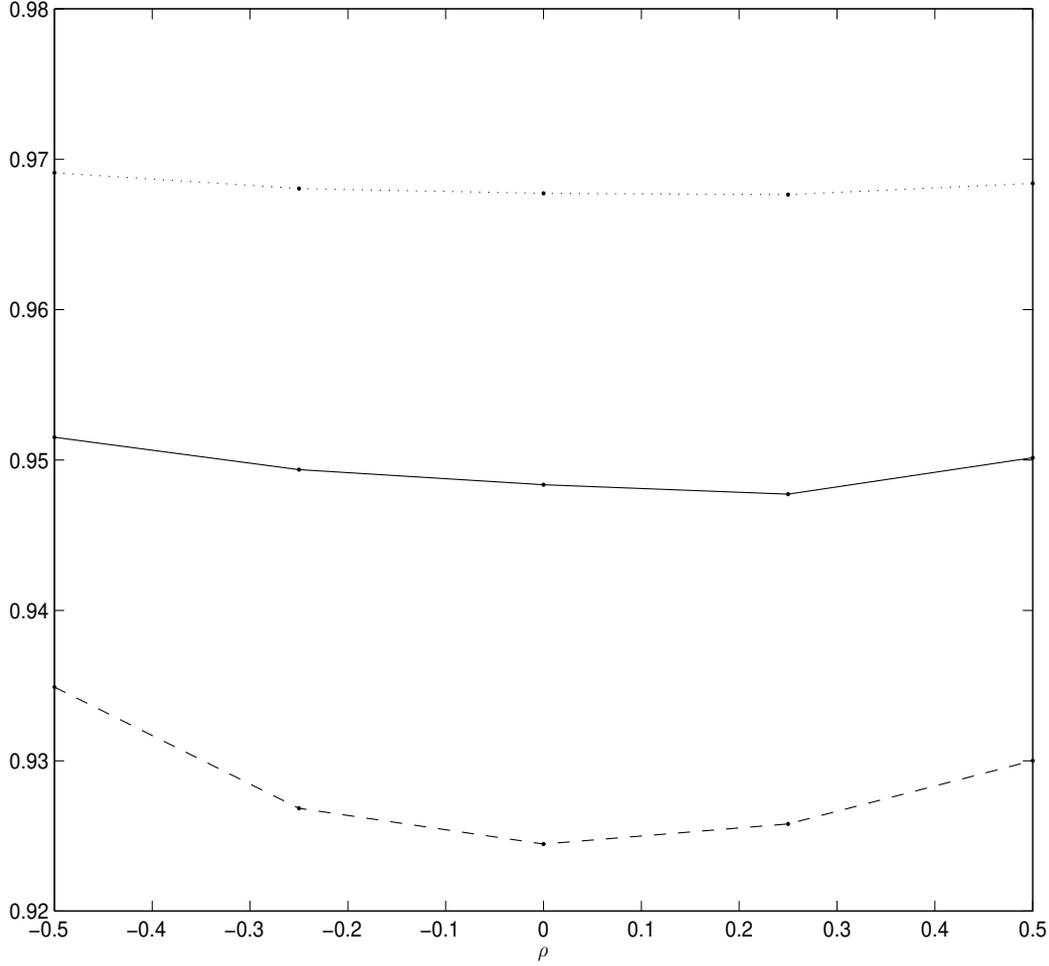


Figure 3: **Sum of probabilities $\hat{\mathbb{E}}_0^\rho [1_{\tau^B < T}]$ and $\check{\mathbb{E}}_0^\rho [1_{\xi_T < B}]$ as a function of ρ when $T \approx \mu(B)$.** The Heston model is used with the following parameters: $S_0 = 100$, $r = 0.015$, $\kappa = 2$, $\theta = 0.09$, $B = V_0 = 0.087$, $\eta = 0.250$, corresponding to $\mu(B)$ of about 1 year. The value of T is set to be 1 year. The three curves correspond to three different values of the strike price K : 90 (dotted line), 100 (solid line), and 110 (dashed line). Monte Carlo simulation is performed to compute the probabilities. We see that the sum of the two probabilities have very weak ρ dependence when $T \approx \mu(B)$, even though each probability depends strongly on ρ as seen in Figure 2.